PROCEEDINGS OF

PSAM — II

Volume 2

(Sessions 37 - 72)

An International Conference Devoted to the Advancement of System-based Methods for the Design and Operation of Technological Systems and Processes

San Diego, California, U.S.A.

March 20-25, 1994

PROCEEDINGS OF

PSAM — II

An International Conference Devoted to the Advancement of System-based Methods for the Design and Operation of Technological Systems and Processes

San Diego, California, U.S.A.

March 20-25, 1994

Editors

G.E. Apostolakis and J.S. Wu
University of California, Los Angeles

Contributing Editors

D. Bley, PLG, Inc.
S. Guarro, Aerospace Corp.
V. Ho, PLG, Inc.
R. Mulvihill, PRC, Inc.
D. Orvis, APG, Inc.
N. Sankaran, Unocal Corporation
N. Siu, Idaho National Engineering Laboratory
M. Stamatelatos, Scientech, Inc.

FOREWORD

These volumes contain papers that were presented at the PSAM II Conference. PSAM I was held in February, 1991 and its success led to the present Conference and hopefully many more.

The purpose of the PSAM Conferences is to provide a forum for the presentation of scientific papers covering both methodology and applications of system-based approaches to the design and the effective, safe operation of technological systems and processes. These include nuclear plants, chemical and petroleum facilities, defense systems, aerospace systems, and the treatment and disposal of hazardous wastes. The objective is to share experiences to the benefit of all industries.

We would like to comment on the production of these proceedings, specifically the printing and binding. Unfortunately, the publishers we had contracted to publish the proceedings did not fulfill their obligation to have them available in time for the Conference. Due to our strong belief that the proceedings should be handed out at the Conference, we have done our best to produce these volumes in the very short time available.

The Editors

SPONSORED BY THE

Society for Risk Analysis

European Safety and Reliability Association

COSPONSORED BY THE

American Nuclear Society

American Society of Mechanical Engineers

IEEE Reliability Society

SPONSORING ORGANIZATIONS

Aerospace Corporation

APG, Inc.

ASCA, Inc.

Electric Power Research Institute

Idaho National Engineering Laboratory

Kazarians & Associates

Los Alamos National Laboratory

PLG, Inc.

PRC, Inc.

SAIC

Sandia National Laboratories

Scientech, Inc.

Tenera

Unocal Corporation

University of California, Los Angeles

University of Maryland

US Department of Energy

GENERAL PROGRAM CHAIRMAN

Michael G. Stamatelatos, Scientech, Inc.

ASSOCIATE GENERAL CHAIRMAN

Doug Orvis, Accident Prevention Group

SENIOR ADVISORY BOARD

B.J. Garrick, PLG, Inc., Chairman

A. Amendola, Commission of European Communities
A.H-S. Ang, University of California, Irvine
A.T.D. Butland, SRD, United Kingdom
J.E. Fitzgerald, Jr., US Dept. of Energy
V.L. Grose, Omega Systems Group
D-Y. Hsia, Institute of Nuclear Energy Research, Taiwan
V. Joksimovich, Accident Prevention Group
S. Kondo, University of Tokyo
J. Meltzer, The Aerospace Corporation
N.C. Rasmussen, Massachusetts Institute of Technology
L. Ybarrondo, Scientech, Inc.

CONFERENCE LOCAL ORGANIZING COMMITTEE

A.A. Dykes, PLG, Inc., Chairman

D. Bley, PLG, Inc.
S.Guarro, The Aerospace Corporation
D. Henneke, TENERA
V. Ho, PLG, Inc.
M. Kazarians, Kazarians & Associates
A. Mosleh, University of Maryland
R. Mulvihill, PRC, Inc.
D. Orvis, APG
N. Sankaran, UNOCAL Corporation
N.O. Siu, INEL
M. Stamatelatos, Scientech, Inc.

ORGANIZING STAFF

D. Bell, UCLA
M. Hanna, UCLA
J. George, UCLA
C. Garrett, UCLA
T. Paulos, UCLA

TECHNICAL PROGRAM COMMITTEE

CHAIRMAN

G. Apostolakis, University of California, Los Angeles, USA

ASSOCIATE CHAIR

J.S. Wu, ASCA, Inc.

MEMBERS

D. Aldridge, USA
H.P. Alesso, USA
A. Amendola, Italy
D.R. Anderson, USA
B. Ayyub, USA
R.A. Bari, USA
J.H. Bickel, USA
V.M. Bier, USA
D. Bley, USA
E.J. Bonano, USA
S. Book, USA
B. Bream, USA
D. Brooks, USA
M.L. Brown, UK
B. Buchbinder, USA
R.J. Budnitz, USA
P.C. Cacciabue, Italy
A. Camp, USA
R.M. Cooke, Netherland
R.A. Cox, UK
D. Croucher, USA
G.E. Cummings, USA
D. Cunha, USA
K. Dahlgren, Sweden
J. Devooght, Belgium
A.A. Dykes, USA
T. Eng, USA
D. Frangopol, USA
M.V. Frank, USA
R. Friedman, USA
R.R. Fullwood, USA
W.J. Galyean, USA
D.I. Gertman, USAL.H.J. Goossens, Netherlands
S. Guarro, USA
M. Haim, Israel
G.W. Hannaman, USA
J.C. Helton, USA
D. Henneke, USA
S. Hirschberg, Switzerland
V. Ho, USA
E. Hollnagel, UK
K. Inoue, Japan
K. Jamali, USA
D.H. Johnson, USA
P. Kafka, Germany
G.D. Kaiser, USA
D.M. Karydas, USA
W.E. Kastenberg, USA
M. Kazarians, USA
D.L. Kelly, USA
K. Klein, USA
R.F. Lavelle, USA
C. Lavine, USA
L. Lee, USA
S. Lydersen, Norway
M. Maharik, Israel
D. Majumdar, USA
G. Mancin, Italy
M. Marseguerra, Italy
T. Matsuoka, Japan
M. Modarres, USA
D.A. Moore, USA
A. Mosleh, USA
R. Mulvihill, USA
K. Murphy, USA
D. Okrent, USA
L.F. Oliveira, Brazil
N. Ortiz, USA
D. Orvis, USA
I.A. Papazoglou, Greece
C.K. Park, Korea
G.W. Parry, USA
M.E. Pate-Cornell, USA
P. Prassinos, USA
C. Preyssl, Neth.
D. Rice, USA
L.K. Rudolph, USA
B. Sagar, USA
N. Sankaran, USA
J.C.H. Schuler, Netherland
N.O. Siu, USA
D. Stack, USA
M. Stamatelatos, USA
J-P. Sursock, USA
O. Svenson, Sweden
B. Thompson, U.K.
W. Tosney, USA
L.P. Vestrucci, Italy
J. Wreathall, USA
J. Ziagos, USA

CONTENTS

037 Uncertainties and Sensitivities in Physical and Chemical Phenomenology--Explosions and Combustion

Chair: S. Kaplan

Development of Fast-Running Thermal and Structural Response Models for Probabilistic Analyses of Complex Systems
A.S. Benjamin, N.N. Brown (SNL)

The Ignition Temperature of Solid Explosives Exposed to a Fire
J.R. Creighton (LLNL)

Development and Validation of a Vapour Clouds Explosion Blast Model
G. Opschoor, A.C. van den Berg, W.P.M. Mercx (TNO Prins Maurits Lab, Netherlands)

Post-Mortem Risk Modeling of the Mexico City Disaster
G.N. Pettitt, J.D. Harms, J.L. Woodward (DNV Technica)

038 Risk Assessment - Space and Missile Systems

Chair: S. Guarro, The Aerospace Corporation

An Assessment Methodology for SRB Mission Failure Probability
F. Hsu, M.A. Azarm (BNL)

The Development of Risk Profiles for Space and Missile Launch Range Safety Decisions
L.L. Philipson (ACTA)

Risk Assessment Framework for the CASSINI Planetary Mission
S. Guarro (Aerospace Corp.); B. Bream (NASA Lewis)

Risks and Issues in Fire Safety on the Space Station
R. Friedman (NASA Lewis)

039 Expert Judgment

Chair: L.H.J. Goossens, Delft U. Technol.

Method of Expert Judgment Applied to Waste Landfill Performance
Lj. Rodic, L.H.J. Goossens (Delft U. Technol.)

The Acute Toxicity of Hazardous Materials and Expert Judgment
L.H.J. Goossens, R.M. Cooke (Delft U. Technol.); F. Woudenberg, P. van der Torn (Municipal Health Serv., Rotterdam)

A Latent Variable Model in Uncertainty Analysis
A.M.H. Meeuwissen, R.M. Cooke (Delft U. Technol.)

040 Benefits of Understanding Human Errors in Complex Systems

Chair: J. Wreathall, J. Wreathall & Co.

Some Requirements for Designing and Managing Reliable Complex Systems
K.H. Roberts (UC Berkeley); M. Grabowski (Rensselaer Polytechnic U.)

Incorporation of Human Factors into Process Hazard Analysis
D.B. McCafferty, K.A. Borows (DNV Technica)

041 Methods for Improved Operations of Nuclear Units

Chair: K. Jamali, USDOE

The Savannah River Site Component Generic Data Base: Applicable to a Wide Variety of Facilities
C.H. Blanton (Westinghouse Savannah Rvr.); S. A. Eide (LANL)

Risk Based Regulation Configuration Control for the James A. Fitzpatrick Nuclear Power Plant
C.J. Everline, M.G. Stamatelatos, F.M. Quinn (SCIENTECH)

Application of Risk-Based Methods for Inspection of Nuclear Power Plant Systems
S.M. Wong, W.E. Gunther (BNL); J.W. Chung, W.D. Beckner (USNRC)

Assessment of Existing Plant Instrumentation for Severe Accident Management
A.J. Horn, D.E. True (ERIN Eng. & Res.); S. Oh (EPRI)

042 Time Dependent Effects

Chair: D.L. Kelly, INEL

Survey and Evaluation of Aging Risk Assessment Methods and Applications
D. Sanzo, P. Kvam (LANL); G. Apostolakis, J.W. Wu, T. Milici, N. Ghoniem, S. Guarro (ASCA)

Development of High Temperature Metallic Melting Processes Related to the Detritiation of Exhausted Control Rods
G. Modica, H. Dworschak, F. Mannone (CEC-JRC, Ispra); R. Nannicini (ENEA)

Use of Sensitivity Analysis in Selection of Upgrades for Low Seismic Capacity Components for the Advanced Test Reactor
S.T. Khericha, S.A. Atkinson (INEL); D.M. Henry (Tenera); M.K. Ravindra (EQE Intl.)

043 Risk Communication to Professionals

Chair: A.S. Heger, U. New Mexico

How Shall We Measure the Reliability Performance of a Safety System?
P. Hokstad (SINTEF)

The Use of Risk Contours to Clarify Risk Communication
T.E. Fenstermacher (PLG); J.T. Kopecek (Union Oil of Calif.)

Improving NRC Staff Uses of Risk Assessment
M.A. Cunningham, P. Baranowsky, W. Beckner, P. Rathbun (USNRC)

044 Uncertainties and Sensitivities in Physical and Chemical Phenomenology--Radiological Source Terms

Chair: E.J. Bonano, Beta Corp. Intl.

Uncertainties and Sensitivities in Aerosol Calculations - An Analysis of a FIPLOC-M Application to the Experiment VANAM-M2
E. Hofer, B. Krzykacz, J. Langhans, G. Weber (GRS, FRG)

Analysis of Source Term Uncertainty Issues for LWRs
K. Kuramatsu (JAERI)

Source Term Analysis for High Level Waste Storage Accidents
S. Medhekar, J. Kindinger (PLG); D. MacFarlane (LANL)

045 Interactive Fault Detection and Diagnosis--Applications

Chair: M. Modarres, U. Maryland

Applications of Pattern Recognition Techniques to Online Fault Detection
R.M. Singer, K.C. Gross, R.W. King (ANL)

Principal Factor Analysis of the Reactor Coolant Pump System
W.S. Grenzebach (Applied Mgmt. Consultants); T.J. Marx (Marx Soc. Sci. Res.)

FlowSim/FlowRisk - A Code System for Studying Risk Associated with Material Process Flows
A.M. Kaufman (LLNL)

046 Methods for Uncertainty Analysis (I)

Chair: J. Devooght, U. Libre Bruxelles

A New Approach to the Uncertainty Evaluation in Risk Assessment
P. Cecchella, M. Mazzini (U. Pisa)

Uses of Zero-One Sampling in Probabilistic Risk Assessment
A. Camp (SNL)

Coherent Sampling of Multiple Branch Event Tree Questions
A.C. Payne Jr., G.D. Wyss (SNL)

Analysis of Extreme Top Event Frequency Percentiles Based on Fast Probability Integration
B. Staple (SNL); F.E. Haskin (U. New Mexico)

047 Enhancing Safety Culture in Modern Industry

Chair: J. Philley, NUS

Safety Culture: What is it, and How can it be Developed?
D.S. Barnes, S.A. Brearley (AEA Consultancy Services, UK)

Safety Culture and its Reflection in Job and Organizational Design: Total Safety Management
G. Grote, C. Künzler (ETH-Zentrum, Switzerland)

Assessing Organizational Impact on Reactor Safety
W.G. He (PLG)

Concept Safety Review of a Sociotechnical System
G. Wells, C. Phnag (U. of Sheffield, England)

Managing Human Induced Risk Involving Technological Systems and Processes
T.G. Ryan, L.T. Ostrom, D.I. Gertman (INEL)

048 Risk Assessment of Industrial Facilities

Chair: N. Sankaran, UNOCAL

Quantitative Risk Assessment of the Los Alamos National Laboratory TA-55 Plutonium Facility
M.K. Sasser, D. Sanzo (LANL)

Quantitative Risk Analysis of a Butane Storage Facility
L.F.S. Oliveira, J.D.A. Netto, R.A. Pinto (PRINCIPIA Ltda., Brazil); J.C.A. Lima (PETROBRAS, Brazil)

049 External Events

Chair: S.T. Khericha, INEL

Site-Wide Seismic Risk Model for Savannah River Site Nuclear Facilities
S.A. Eide, R.S. Shay (LATA); W.S. Durant (Westinghouse Savannah Rvr.)

Tornado Impact Analysis for Use in Quantitative Risk Assessments
J. Shah (Four Elements Ltd.)

Aircraft Crash Analysis of the Proposed Sizewell B Containment Vessel
Y.F. Al-Obaid (PAAET, Kuwait)

050 Risk Management in DOE (III)

Chair: R.R. Fullwood, BNL

Fully Integrated Risk Management
R.J. Doyle, K. Beierschmitt, (Battelle Pantex)

051 Waste Isolation Pilot Plant--Regulations, Methodology, Modeling and Results

Chair: D.R. Anderson, SNL

Organization of Performance Assessments Conducted for the Waste Isolation Pilot Plant
J.C. Helton (Arizona St. U); D.R. Anderson, M.G. Marietta (SNL)

Using Data and Information to Form Distributions of Model Parameters in Stochastic Simulations of Performance of the Waste Isolation Pilot Plant (WIPP)
M.S. Tierney (SNL)

Modeling System Used for the Performance Assessment for the Waste Isolation Pilot Plant (WIPP)
M.G. Marietta, M.S. Tierney (SNL); J.C. Helton (Arizona St. U.)

Regulatory and Institutional Considerations in Performance Assessment for the Waste Isolation Pilot Plant (WIPP)
D.R. Anderson, M.G. Marietta (SNL)

052 Software Dependability (II)

Chair: R. Mulvihill, PRC

The Dynamic Flowgraph Methodology: A Methodology for Assessing Embedded System Software Safety
C. Garrett, M. Yau, S. Guarro, G. Apostolakis (UCLA)

Determining the Prime Implicants for Multi-State Embedded Systems
E. J. Shields, G. Apostolakis, S.B. Guarro (UCLA)

053 Expert Judgment in Assessment Studies

Chair: M.A. Cunningham, USNRC

An Application of ISPA: Sensitivity Analysis on Complex Engineering Problem
P. McClure, S. Clement, S. Ashbaugh (SAIC)

The Analysis of Competing Failure Modes
R.M. Cooke (Delft U. Technol.)

Practical Problems in Aggregating Expert Opinions
J.M. Booker, R.R. Picard, M.A. Meyer (LANL)

An Application of Expert Opinions Approach in Combining Measurement Results
M. Haim, M. Jaeger, D. Tzidony (RAFAEL, Israel)

054 HRA, the Second Generation (I)

Chair: P.C. Cacciabue, JRC

Some Thoughts on the Requirements for a Second Generation Human Reliability Assessment Process
A.J. Spurgin

The Need for, and a Proposed Structure of, a Second Generation HRA
G.W. Parry (NUS)

055 Instrumentation and Control Reliability

Chair: R. Roberts, Ultramar

Reliability-Based Design of a Safety-Critical Automation System - A Case Study
C.W. Carroll (Boeing); W. Dunn (USC); L. Doty, M. Hulet (NASA); M.V. Frank (Safety Factor Assoc.)

Safeguards Actuation Fault Tree Analysis
M. Cepin, M. Kozuh, B. Mavko (J. Stefan Inst.)

Reliability Assessment of Onsite Emergency Safety Buses at Operating Nuclear Power Plants
G. Martinez-Guridi, M.A. Azarm (BNL)

Optimal Voting Logic of Hydrocarbon Gas Detectors Based on Field Data
L. Bodsberg (SINTEF)

056 Risk Communication

Chair: R.A. Bari, BNL

Risk Assessment - Working with Your Client
K.G. Murphy, J.E. Fitzgerald, Jr. (DOE)

057 Resource Planning and Value-Impact Applications

Chair: K. Klein, USDOE

Environment, Safety & Health Risk Management: Program Planning and Budget Decision-Making
F.E. Tooper (USDOE); J.L. Von Herrmann (CYCLA)

Control Operator Overload: Identification and Prevention
D.A. Strobhar (Beville Eng.)

Probabilistic Cost-Benefit Analysis of Enhanced Safety Features for Strategic Nuclear Weapons at a Representative Location
D.R. Stephens, C.H. Hall, G.S. Holman, K.F. Graham, T.F. Harvey, F.J.D. Serduke (LLNL)

058 Examples of PRA in Design Decisions

Chair: B. Buchbinder, NASA

Examples of the Use of PSA in the Design Process and to Support Modifications at Two Research Reactors

D.H. Johnson, D.C. Bley, J.C. Lin (PLG); C.T. Ramsey, M.A. Linn (ORNL)

Risk Management and Corrective Actions

R.A. Dykes, A.A. Dykes (PLG); J.D. Blodgett (Hughes)

A Risk and Decision Analysis for Choosing Wind Tunnel Turbine Blade Strategies

M.V. Frank, S.A. Epstein (Safety Factor Assoc.)

Investigation Into the Development of A Train-Level Probabilistic Risk Assessment

C.L. Smith, R.D. Fowler, L.M Wolfram (INEL)

059 Use of PRA Models in Assessing Technologies and Device Applications

Chair: M. Zentner, Westinghouse Hanford

The Use of Robotics and Probabilistic Risk Assessment in Environmental Restoration

L.E. Cindel, G. Apostolakis (UCLA)

Risk Evaluation of Medical and Industrial Radiation Devices

E.D. Jones (LLNL), R.E. Cunningham, P.A. Rathbun (USNRC)

Integrated Systems Performance Assessment

C.N. Amos, S.D. Clement (SAIC)

060 HRA, the Second Generation (II)

Chair: E. Hollnagel, Computer Resources Intl.

Reliability of Cognition, Context and Data for a Second Generation HRA

E. Hollnagel (HRA, UK); P.C. Cacciabue (CEC/JRC, Italy)

Elements Of A Model of Operator Problem Solving and Decision Making in Abnormal Conditions

S-H. Shen, C. Smidts, A. Mosleh (U. Maryland)

Human Reliability Analysis Using Simulated Human Model

Y. Fujita, I. Yanagisawa (MAPI, Japan); H. Sakuda (NUPEC/IHF, Japan)

061 Transportation Risk (I)

Chair: T. Altenbach, LLNL

The Cost of Human Error Intervention
C.T. Bennett, W.W. Banks, E.D. Jones (LLNL)

Application of the DOE and NRC Nuclear Safety Policy to Transportation Operations
C.Y. Kimura (LLNL), G.M. Sandquist (U. Utah)

Control of Risk from Explosives Stores Within Mod (UK)
J. Connor, D. Hewkin (Ministry of Defence, UK)

062 Environmental Risk Management--Waste

Chair: C. Peabody, USDOE

The Role of Multimedia, Multi-Pathway Exposure Assessments in Risk-Based Regulations at California Hazardous Wastes Sites
T.E. McKone (LLNL), E.G. Butler, J.J. Wong (Calif. Dept. Toxic Substances Control)

Risk Analysis for New Nuclear Waste Sites - Will It Generate Public Acceptance?
H. Inhaber (Westinghouse Savannah River)

063 Methods for Uncertainty Analysis (II)

Chair: A. Camp, SNL

Development of Smart Searching Algorithms for Vulnerability and Uncertainty Analysis in Probabilistic Risk Assessments
A.S. Benjamin (SNL)

Importance and Sensitivity Analysis in Assessing the System Reliability
Y.-F. Wu (Taiwan Power Co.), J.D. Lewins (Cambridge U.)

Multiple Weight Stepwise Regression
J. Atkins (UC Berkeley); J. Campbell (SNL)

A Software System for Probabilistic Uncertainty and Sensitivity Analysis of Results from Computer Models
B. Krzykacz, E. Hofer, M. Kloos (GRS)

064 Risk Based Prioritization

Chair: G.E. Cummings, LLNL

Prioritizing Risk-Reducing Activities: Integrating Risk Management Into Business Planning
D. Brooks (Arizona St. U)

An Integrated Resource Management System and the Risk-Based Prioritization Model in Use at Martin Marietta Energy Systems
M.I. Sparks (Martin Marietta)

Risk-Based Prioritization: A Comparison of Approaches
R.S. Strait (LLNL)

065 Meaning and Interpretation of Probability

Chair: A. Mosleh, U. Maryland

The Meaning and Use of Probability in Probabilistic Safety Assessment
G.W. Parry (NUS)

Is Probability of Frequency Too Narrow?
H.F. Martz (LANL)

On Bayesian Approaches of Parametric and Modeling Uncertainty
J.S. Wu (ASCA)

On The Objectivity of Subjective Theory of Probability for PSA Applications
A. Mosleh (U. Maryland)

066 Cognitive Modeling and Simulation in Human Reliability Analysis

Chair: Y. Fujita, Mitsubishi Atomic Pwr. Ind.

A Simulation Study of Errors of Commission in Nuclear Accidents
K-S. Hsueh (NUS); L. Soth (Commonwealth Ed.); A. Mosleh (U. Maryland)

Simulating Operator Cognition for Risk Analysis: Current Models and CREWSIM
V.N. Dang (MIT); N.O. Siu (INEL)

Quantification of Multiple Error Expressions in Following Emergency Operating Procedures in Nuclear Power Plant Control Room
A.P. Macwan, P.S. Wieringa (Delft U. Technol.); A. Mosleh (U. Maryland)

067 Comparative Risk Assessment of Complex Technological Systems (I)

Chair: S. Haddad, New South Wales Govt.

Severe Accidents in Comparative Assessment of Energy Sources: Current Issues and Actual Experience Data

S. Hirschberg, C. Parlavantzas (P. Scherrer Inst., Switzerland)

An Analysis of Electricity Generation Health Risks: A United Kingdom Perspective

D.J. Ball, L.R.J. Roberts, A.C.D. Simpson (U. East Anglia, England)

External Costs as an Indicator for Comparative Risk Assessment of Different Energy Systems - A Case Study for Comparing the Nuclear and the Coal Fuel Cycle

W. Krewitt, P. Mayerhofer, R. Friedrich, A. Greβmann (U. Stuttgart, Germany); M. Dreicer (CEPN, France)

The Use of Multi-Attribute Analysis for Comparing the Health and Environmental Impacts of Different Fuel Cycles - A Feasibility Study for the Comparison Between Coal and Nuclear Fuel Cycles

M. Dreicer, P. Kuentzmann, T. Schneider (CEPN, France); W. Krewitt (U. Stuttgart, Germany)

068 Risk Assessment Failure Data Estimation

Chair: C. Fleming, PLG

Advanced Test Reactor Large LOCA Frequency Estimation

S.T. Khericha (INEL); S.A. Eide (LATA)

Component Failure Rate Comparison and Recommendations for Various Processing Fluids

L.C. Cadwallader, J.N. Wilkinson (INEL); S.A. Eide (LATA)

Application of Reliability-Centered-Maintenance to PRA Fault Tree Analysis for a BWR HPCI System

M.A. Feltus, Y.A. Choi (Penn. St. U)

069 Waste Risk Management

Chair: B. Thompson, H.M. Insp. Pollution

The HMIP Research and Development Programme for Post-Closure Risk Assessment

B. Thompson (HMIP, UK)

Risk Analysis In Regulation and Risk Communication
S.M. Stearn (HMIP, UK)

Validation: What Should a Regulator Look for in a Safety Case?
R.D. Wilmot, D.A. Galson (Galson Sciences, England)

Aspects of Review of a Proponent's Post-Closure Safety Assessment on Behalf of the Regulator
T.J. Sumerling (Safety Assessment Mgt.); D. Read (WS Atkins Sci. & Technol.)

070 Reactor PSA Studies

Chair: M. Hitchler, Westinghouse Savannah Rvr.

Probabilistic Safety Analysis of Ignalina NPP, RBMK-1500
G. Johansson (IPS); E. Shiverskly (IPE); E. Söderman (ES-Konsult)

Results and Insights from the Low Power and Shutdown Level 1 PRA of the Grand Gulf Nuclear Power Station
D.W. Whitehead, B.D. Staple, S.L.Daniel (SNL); J. Darby (Sci. & Eng. Assoc.); J. Yakle, S. Miller, J. Forester (SAIC)

Comparative Evaluation of the Performance of Advanced Semi-Passive Pressurized Water Reactors in the Loss of Offsite Power Event
M. Ouyang, M. Golay (MIT)

071 Accident Precurssors

Chair: V.M. Bier, U. Wisc.-Madison

Accident Sequence Precursor Methodology for Maritime Safety
L.H.J. Goossens (Delft U. Technol.); C.C. Glansdorp (Marine Analytics)

Developing Organizational Safety Information System for Monitoring Potential Dangers
M. Tamuz (U. Texas)

The Performance of Precursor-Based Estimator
V.M. Bier, W. Yi (U. Wisc.)

072 Aging Analysis: Methodology & Applications

Chair: D.H. Johnson, PLG

Sensitivity and Uncertainty Analyses in Aging Risk-Based Prioritizations
M. Hassan, S. Uryas'ev (BNL); W.E. Vesely (SAIC)

Bayesian Aging Analysis of Hanford Waste Storage Tank Leak
W.G. He, S.B. Rao (PLG); D.R. MacFarlane (LANL)

Lifetime Prediction and Physical Asset Management
R. Menich (Failure Analysis Assoc.)

037 Uncertainties and Sensitivities in Physical and Chemical Phenomenology--Explosions and Combustion

Chair: S. Kaplan

Development of Fast-Running Thermal and Structural Response Models for Probabilistic Analyses of Complex Systems

A.S. Benjamin, N.N. Brown (SNL)

The Ignition Temperature of Solid Explosives Exposed to a Fire

J.R. Creighton (LLNL)

Development and Validation of a Vapour Clouds Explosion Blast Model

G. Opschoor, A.C. van den Berg, W.P.M. Mercx (TNO Prins Maurits Lab, Netherlands)

Post-Mortem Risk Modeling of the Mexico City Disaster

G.N. Pettitt, J.D. Harms, J.L. Woodward (DNV Technica)

055 Instrumentation and Control Reliability

Chair: R. Roberts, Ultramar

Reliability-Based Design of a Safety-Critical Automation System - A Case Study
C.W. Carroll (Boeing); W. Dunn (USC); L. Doty, M. Hulet (NASA); M.V. Frank (Safety Factor Assoc.)

Safeguards Actuation Fault Tree Analysis
M. Cepin, M. Kozuh, B. Mavko (J. Stefan Inst.)

Reliability Assessment of Onsite Emergency Safety Buses at Operating Nuclear Power Plants
G. Martinez-Guridi, M.A. Azarm (BNL)

Optimal Voting Logic of Hydrocarbon Gas Detectors Based on Field Data
L. Bodsberg (SINTEF)

RELIABILITY-BASED DESIGN OF A SAFETY-CRITICAL AUTOMATION SYSTEM - A CASE STUDY

C. W. Carroll[1], W. Dunn[2], L. Doty[3], M. V. Frank[4], and M. Hulet[3]

[1]Boeing Aerospace Operations, NASA Ames Research Center, MS 213-26, Moffett Field, CA 94035

[2]University of Southern Colorado, NASA Ames Research Center, MS 218-6, Moffett Field, CA 94035

[3]National Aeronautics and Space Administration, NASA Ames Research Center, MS 218-7, Moffett Field, CA 94035

[4]Safety Factor Associates, 1410 Vanessa Circle, Encinitas, CA 92024

ABSTRACT

In 1986, NASA funded a project to modernize the NASA Ames Research Center Unitary Plan Wind Tunnels, including the replacement of obsolescent controls with a modern, automated distributed control system (DCS). The project effort on this system included an independent safety analysis (ISA) of the automation system. The ISA approach followed a tailoring of the risk assessment approach widely used on existing nuclear plants with one important exception - the analysis would be performed early in the design cycle so that uncovered high risk scenarios could be identified and mitigated by design change alone. The purpose of the ISA was to evaluate the completeness of the hazard analyses which had already been performed on the Modernization Project. The ISA identified several new hazards but more significantly identified many hazard causes which had been neither identified nor controlled. The ISA also noted previously identified hazard controls which were not appropriate for the hazard. The most important by-product of the ISA was the effect it had on the hazard control philosophy for the UPWT Modernization Project. Due to its top-down nature, the ISA identified simple, high level controls which would control hazards at the hazard level rather than at the cause level. The Project had been implementing controls at the cause level which resulted in the extensive use of redundant components and an extremely complex design. In addition, the use of redundancy left many single point failures with no controls. Also, the implementation of controls at the cause level made it far more important to identify all causes. By controlling hazards at the hazard level, it is not as critical to identify all causes because even if new causes are identified, the high level control will control them. The use of high level controls allowed the redundant components to be eliminated and resulted in a safer, more reliable, and less costly design. An additional by-product of the ISA was the detailed design review which it provided. The ISA provided an Independent Design Review and timely, detailed, design recommendations to the Project Engineers, which included the identification of specific design errors, omissions, and recommended corrections. The tailoring of the nuclear industry oriented risk assessment approach to the automation system and its role in reliability-based design of the automation system is the subject of this paper.

INTRODUCTION

During the 1950s, NASA Ames Research Center brought on line a very high power (216,000 hp) complex of three pressurized wind tunnels capable of testing large aircraft and spacecraft models in a speed range of 0.5 to 3.5 Mach. From startup operation to the present time, the tunnels have been a principal aerodynamic testbed for virtually every commercial, military/strategic, and military/fighter aircraft produced in the United States. Spacecraft work has covered a range of vehicles beginning with the Mercury Orbiter on through the Space Shuttle. Officially known as the Ames Unitary Plan Wind Tunnels (UPWT), the complex was declared a National Historical Landmark in 1985 in recognition of its gained and ongoing stature as an indispensable national resource in aeronautics development.

In 1986, NASA funded a project to modernize these tunnels including the replacement of obsolescent controls with a modern, automated distributed control system (DCS). The project effort on this system included an independent safety analysis (ISA) of the automation system. The objective of the ISA was to perform an independent hazard analysis of the UPWT Modernization Project Automation system. Several hazard analyses of the Automation system had already been performed by the Modernization Project including Hazard Analyses (HA) and Failure Modes and Effects Analyses (FMEA). The main purpose of this analysis was to identify any new critical or catastrophic hazards which were not identified by the HA and FMEA. To achieve this objective, a top-down analysis approach was used to contrast the bottom-up approach of the HA and FMEA. The method used for the top-down analysis was a tailored Probabilistic Risk Assessment (PRA) method as employed in nuclear power plant applications.

A key difference between a PRA performed on existing nuclear plants and the ISA was that the ISA was performed early in the design cycle. The ISA was initiated on the 30% design and continued through the 90% design. Therefore, as the ISA identified problems, the design was modified to correct those problems. In this manner, the System Safety Analysts and the Design Engineers worked concurrently to develop design solutions which would adequately control the hazards.

UPWT FACILITY AND DISTRIBUTED CONTROL SYSTEM DESCRIPTION

The Unitary Plan Wind Tunnel facility consists of three closed circuit test sections, one transonic and two supersonic, which cover a Mach number range of 0.5 through 3.5. Airflow is provided to any of the test sections by one of two compressors which are driven by 4 electric main drive motors. Air speed is controlled by compressor speed, tunnel cross section and shape, and by a moveable row of vanes on the compressor called Inlet Guide Vanes. The tunnel cross section is varied by moving flexible tunnel walls with mechanical jacks. Tunnel pressure is controlled by a network of ducts, valves, and compressors which is collectively called the Make Up Air system. These compressors, main drive motors, Make Up Air system, and all supporting subsystems are called the Auxiliaries.

The Automation portion of the UPWT Modernization Project incorporates the use of a Distributed Control System and is split into two general areas: Tunnel Automation and Auxiliaries Automation.

Tunnel Automation involves automation of the Test Conditions System (TCS) and Model Positioning System (MPS). The control of test conditions and model positioning in these tunnels currently utilizes a significant amount of open loop manual control for parameters such as Mach number and tunnel pressure.

The TCS controls the test environment and interfaces operating personnel with the tunnel processes. The TCS will receive test condition setpoints from the operators for such parameters as Mach number and tunnel pressure. The TCS will then determine the necessary setpoints for the subsystems and coordinate the control of the subsystems. For instance, in order to achieve the desired Mach Number, the TCS will vary the angle of the Inlet Guide Vanes and move the tunnel walls to the appropriate position. In order to achieve the desired tunnel pressure, the TCS will open and close the appropriate Make Up

Air system valves.

The MPS controls and monitors the Sting Model Support System (SMSS) and Turntable System (TTS) to pitch, yaw, or roll the model with respect to the windstream.

Auxiliaries Automation consists primarily of automating the Make-Up Air System, the compressor and main drive bearing lube systems, the cooling water system, and the main drive interface system.

TAILORING OF RISK ASSESSMENT APPROACH AND RESULTS

The purpose of the ISA was to evaluate the completeness of the hazard analyses which had already been performed on the Modernization Project. The PRA analysis technique was chosen because it is a top-down, comprehensive analysis technique which would provide a thorough comparison to the bottom-up approach used in the HA and FMEA. However, the PRA analysis technique was tailored in order to meet the requirements of the ISA and to reduce the budget and time required to perform the analysis.

The ISA followed the PRA technique of developing initiating events and associated fault trees but did not include formal scenario development and quantification of pivotal events. A scenario defines the possible responses (pivotal events) of a facility to an initiating event. Scenarios can be expressed in Event Sequence Diagrams and Event Trees. An Event Tree summarizes all of the Event Sequence Diagrams for a particular initiating event and is a convenient form for quantification. As explained in Dunn et al[1], event tree development and pivotal event modelling and quantification were attempted but it became obvious that there was insufficient budget and time to quantify complete scenarios. However, at the same time it became obvious that scenario quantification was not necessary because it was possible to design controls at the initiating event level and thereby ensure that the probability of the initiating event was at or less than the acceptable level of risk. By controlling the scenario at the initiating event, it is guaranteed that the risk of any resulting worst case scenario would also be at or below the acceptable level of risk.

The development of the Initiating Events Logic Diagram (IELD) was the first and most important part of the ISA. The IELD identified all initiating events which would be analyzed. Therefore the completeness of the analysis was determined by the completeness of the IELD.

The first step in development of the IELD was to determine the damage states of interest. Since the objective of the analysis was to identify hazards which could cause critical or catastrophic facility damage or severe personnel injury, these were the damage states of interest. An Initiating Events Logic Diagram (IELD) was developed for each of three damage states of interest:

1. severe personnel injury or death
2. critical or catastrophic facility damage
3. model damage

The development of the IELD consisted of identifying all initiating events which could potentially lead to the damage state. Development of the IELD's required a thorough evaluation of all Modernization Project subsystems and failure modes and an understanding of the potential consequences of a failure occurring. All of the systems and subsystems were evaluated to determine which failures had the potential to lead to any of the three damage states. These failures were then grouped into initiating events and included in the appropriate IELD.

The next step in the ISA was the preparation of comprehensive fault trees for the initiating events. Development of the fault trees required a detailed review of hundreds of drawings and procedures and interviews with tunnel operators, mechanics, and engineers. The fault trees included all causes and existing hazard controls which could mitigate the effects of the causes. The Project System Safety requirement was that no single point failure would result in a critical or catastrophic hazard. Therefore, if the fault tree for each initiating event was complete and all identified causes were controlled, the initiating event had an extremely low probability of resulting in the damage state.

The fault trees identified nine new hazards which had not been identified by the Project HA and FMEA. Three of these hazards had a catastrophic severity rating, two had a critical

severity rating and the remaining four had a marginal severity rating.

However, more significant than the number of new hazards identified, was the fact that the fault trees identified many uncontrolled causes which had not been identified by the Project HA and FMEA. The fault trees also identified controls that were on the Project hazard reports which were inappropriate and/or inadequate for controlling the hazard. An example of an inappropriate hazard control was the use of end of travel limit switches on the Inlet Guide Vanes (IGV) to prevent inadvertent Mach number increase. Mach number can exceed the test allowable well before the IGV's reach their end of travel. Therefore significant model damage can occur well before the limit switches are tripped.

Although formal scenario development and pivotal event quantification were not performed, quantification of initiating events was performed selectively to assist in design engineering tradeoffs. Specifically, quantification was used in the ISA to quantify the risk associated with using component redundancy as a hazard control. The Project hazard control philosophy was originally based on applying elaborate low level controls and using dual redundant components. The design incorporated standby redundant sensors, controllers, I/O, etc. The risks associated with using low level redundancy are numerous. The first is that there are numerous common cause failure modes. The second is that the standby spare could have undetected latent faults. In addition, it becomes imperative to identify and remove all single points of failure. The result of this low level hazard control approach was a complex system which would be difficult to design, build, test and maintain. There was concern that the use of dual redundancy was not adequate to reduce the risk to an acceptable level. To illustrate the residual risk associated with using redundancy, two initiating events were quantified:

1. inadvertent Mach Number increase
2. inadvertent model motion

The results of the quantification clearly illustrated that a design which used redundant components for hazard control would have higher probability of failure than a design which used high level controls and retained only a few simple low level controls[1]. Based on this, the Project changed its hazard control philosophy to incorporate high level controls instead of low level redundancy and high level controls were identified and implemented wherever possible.

Due to the top down nature of the ISA, it was easy to identify hazard controls which could be employed at a high level (initiating event level) rather than at the component level. The use of high level controls, with some simple low level controls, greatly simplified the design and at the same time made it safer and more reliable.

The most significant of these high level controls was the incorporation of independent safety monitors for the control of hazards associated with inadvertent Mach number/dynamic pressure (Mach/Q) increase, inadvertent model motion, and tunnel flexible wall (flexwall) overstress. The safety monitor is a computer with its own dedicated software, sensors, I/O, and hardware which is independent from the primary control system.

The Mach/Q Safety Monitor calculates Mach number and dynamic pressure (Q) independently from the primary controller. The Mach/Q calculations of the safety monitor and primary controller are then compared and if either one exceeds the Mach/Q test limit or if there is a discrepancy between the calculations, the system is shutdown. This comparison detects latent failures in both the primary controller and independent monitor. This single Mach/Q monitor replaced 12 lower level controls.

The Model Safety Monitor (MSM) calculates the model position independently from the primary controller. The model position calculations from the MSM and primary controller are then compared and if there is a discrepancy between them, the Sting Model Support System (SMSS) is shutdown. This comparison detects latent failures in both the primary controller and independent monitor.

The Flexwall Safety Monitor compares the flexwall jack positions to known safe settings and initiates a flexwall drive stop if the jacks are out of tolerance.

CONCLUSIONS AND EXTENSION TO OTHER AUTOMATION APPLICATIONS

The ISA successfully met and exceeded its objective by evaluating the adequacy of the Project HA and FMEA and in addition, providing a wealth of additional safety and detailed design information.

The most significant aspect of the ISA was that due to its top down nature, as required by the PRA analysis technique, it produced an evaluation of the functionality of the design by evaluating multiple systems and interfaces together. The top down analysis technique provided valuable insight into numerous aspects of the design and hazard control which was not provided by the bottom up approach of the Project HA and FMEA. First of all, the ISA identified the use of high level controls which resulted in the removal of unnecessary component redundancy, simplified the design, improved the safety and reliability of the design, and reduced the cost of the design. The use of high level controls also makes it less critical to identify all low level causes since the high level control should control all low level causes. It is normally not cost effective to try to identify all low level causes. Secondly, the fault trees clearly identified missing hazard controls and directed the choice of appropriate hazard controls.

This analysis demonstrated that valuable results can be gained without scenario development and pivotal event quantification. The majority of the analysis concentrated on the development of initiating events and fault trees rather than on scenario development. Quantification can be used selectively to analyze specific initiating events and aid in design engineering trade-offs.

In addition to a hazard analysis, the ISA provided an Independent Design Analysis. Due to the detailed analysis required to prepare fault trees for each initiating event, numerous design issues were identified in addition to safety issues. Due to the top down nature of the ISA, the functionality of the design was reviewed by analyzing multiple systems and their interfaces together. The ISA identified design issues such as: design mistakes, missing or incomplete interfaces between systems and subsystems, referencing mistakes between drawings, logic mistakes, and device labeling mistakes. Detailed design recommendations were made to the Project engineers as each initiating event was analyzed. Since the ISA was performed concurrently with the design, the design changes were made more easily and at less cost than if they had been uncovered during construction or testing. In short, risk assessment activities led to a reliability-based design of the safety critical Unitary systems.

Since the completion of the analysis work described in this paper, the tailored PRA approach has been applied to several other automation projects at Ames. In addition to the analytical effort, System Safety engineers have actively worked with User and Project engineers in assuring that both high and (simple) low level controls are defined, prior to design, during systems requirements definition. This effort has resulted in refined definition of how PRA can be applied to virtually any automatic control system in the requirements or early design stages. This generalized approach, which includes quantitative modelling of software failure modes, is described in these proceedings[1].

ACKNOWLEDGEMENT

This work was supported in part by NASA Cooperative Agreements NCC2-373 and NCC2-609.

REFERENCES

1. W. Dunn, M. Frank, S. Epstein, L. Doty, 1994, Risk Assessment and Management of Safety-Critical, Digital Industrial Controls - Present Practices and Future Challenges, these proceedings.

SAFEGUARDS ACTUATION FAULT TREE ANALYSIS

Marko Čepin, Mitja Kožuh, Borut Mavko,

"Jožef Stefan" Institute
Reactor Engineering Division
Jamova 39, 61000 Ljubljana, Slovenia

INTRODUCTION AND SHORT SYSTEM DESCRIPTION

The basic aim of reactor safety is to prevent the release of radionuclides. This applies to normal operation as well as to accident conditions. For accident conditions reactor is equipped with special safety systems which are designed with high requirements on availability and would in this case prevent accident or at least mitigate its consequences. In Krško Nuclear Power Plant, two loop PWR, 632MWe (net), the Engineered Safety Features is one group of such a systems. The Engineered Safety Features are used to provide protection against the release of radioactive materials in the event of a loss of a coolant accident or a steam line break accident. They function to maintain the reactor in a shutdown condition to provide sufficient core cooling to prevent or to limit the extend of fuel damage and to insure the containment integrity. They need appropriate signals for start when abnormal conditions occur. They get those signals from Engineered Safety Features Actuation System which analysis is presented here.

System Description

Engineered Safety Features Actuation System (Figure 1) has to sense abnormal situations and initiate the operation of necessary Engineered Safety Features. System uses different plant parameters which are measured and compared with the safety limits. It has proper and timely initiating signals which are supplied by the process sensors, transmitters and logic components. Redundancy is assured because it is intended that single failure will not prevent the actuation of the system. When safety limits are exceeded, the system has to send signals through analog and digital portion (Figure 2) of circuitry to start appropriate safety systems, which perform their mission.

Analog portion of circuitry [1] consists of three or four redundant channels. For the system actuating the logic 2/3 or 2/4 is required. Four channels are required where the safety and control functions of the same parameter are needed. Initiating signal is actuated on low pressurizer pressure (121.7 bar), low steam lines pressure (41.4 bar) or high containment pressure (Hi-1 0.28 bar, Hi-2 1.03 bar, Hi-3 1.59 bar).

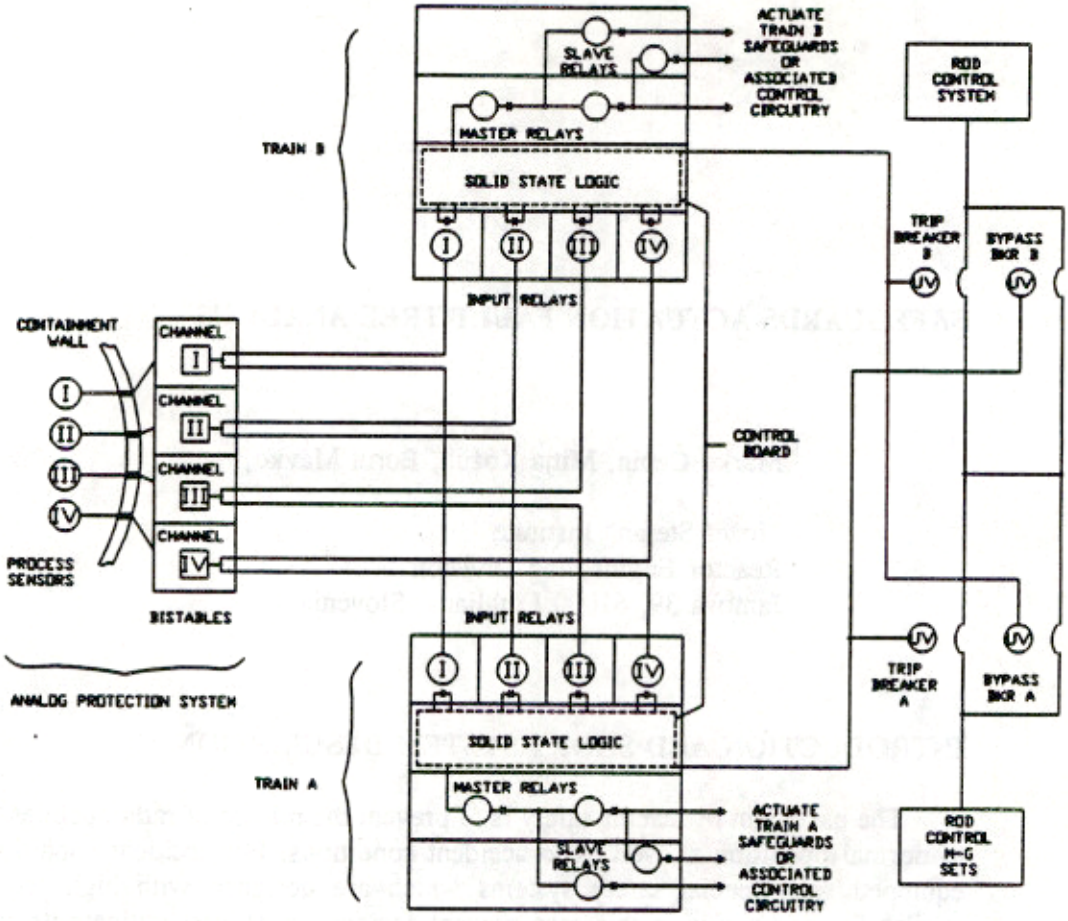


Figure 1. Engineered Safety Features Actuation System Block Diagram

Pressure transmitters are sending electrical signal which represents measured pressure to loop power supply, which has two functions. It gives information to computer and indicators which are connected through isolation amplifier and converts signal to unit signal and sends it to the lead lag card. Then the signal is led to bistable, which compares the measured with its calibrated value and if it is exceeded, changes the output to zero voltage. Outputs from bistables are led to the input relays on both trains. Namely, at bistables the transformation of analog to digital signal is achieved [2]. The redundancy in digital portion of circuitry is from here on assured by two independent trains of which each is capable to actuate required equipment.

SYSTEM ANALYSIS

System analysis was performed first to estimate the unavailability and to determine the key components of the system and second to prepare input for front line plant systems. System modelling was performed by fault tree technique [3]. Engineered Safety Features Actuation System was split to subsystems in a way that each subsystem represents signal path from relevant slave relay through master relay and all signal path to the measuring point (Table 1), pressure transmitter for example. For each subsystem one unique fault tree

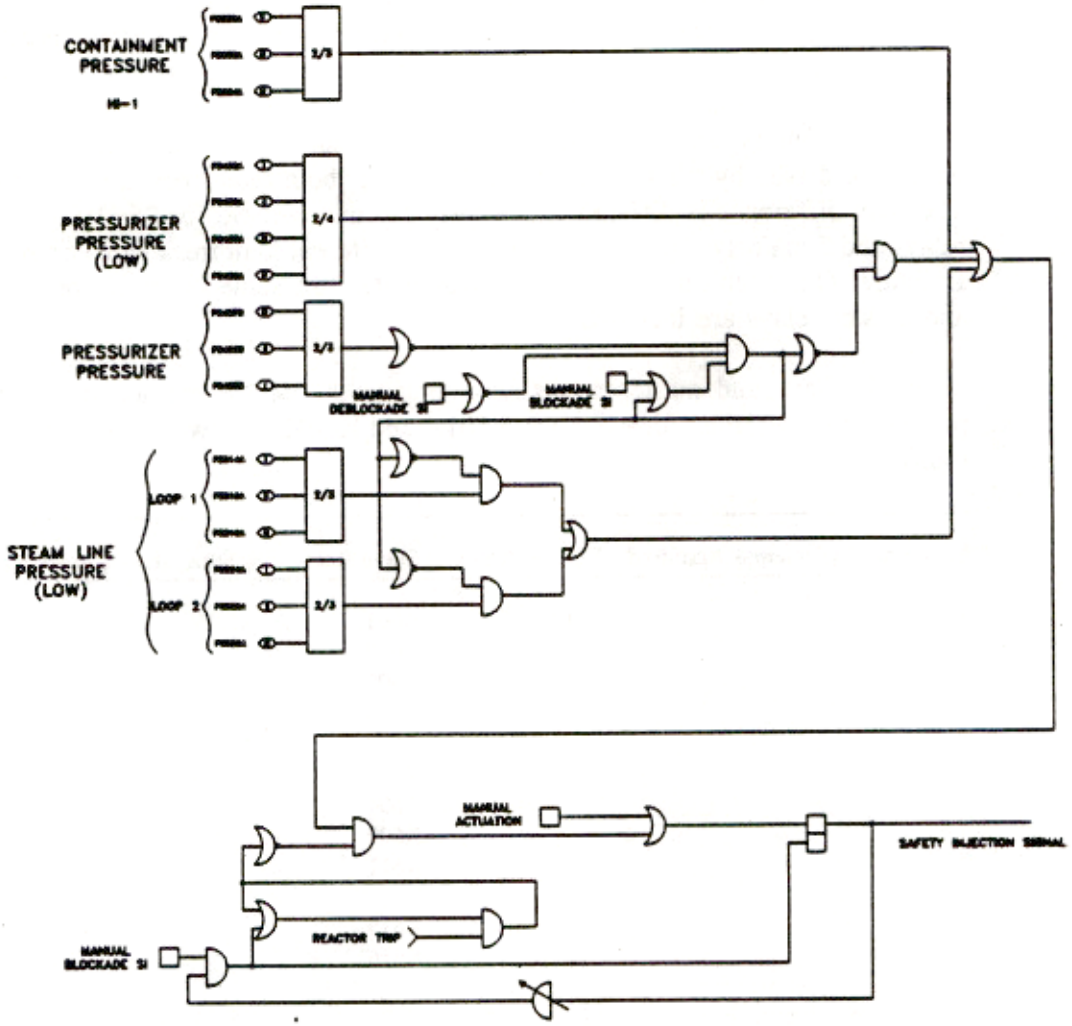


Figure 2. Safety Injection Signal Path / Digital Portion

was developed. Top event in each fault tree is defined as system incapability to actuate required Engineered Safety Features. In the fault trees are not considered events and failure modes causing reactor trip or spurious safety injection for example, which actually are two of the Engineered Safety Features functions. That is because the aim of the analysis was to estimate those different events and impacts which causes system failure in sense not to actuate required equipment. Component data are partly generic [4, 5] and partly specific for Nuclear Power Plant Krško. Advanced computer code used for evaluation was Risk Spectrum PSA [6]. System testing is included in analysis and it was foreseen that it has a relative large impact on system unavailability. Human errors considered in analysis are only those which at least potentially causes increase of system unavailability but not those causing unnecessary system actuating or reactor trip.

Analysis as Support System Analysis

Engineered Safety Features Actuation System is a support system in Nuclear Power Plant what means that is shared by more than one front line system. Its subsystem fault trees can be linked to front line system fault trees mainly to the Engineered Safety

Features Fault Trees, Safety Injection or Containment Ventilation Isolation for example. In Table 1 all subsystems with appropriate slave and master relays are presented. Master relay has function to drive slave relays which operate contacts of various pumps, motor operated valve contactors, etc. Actually for each slave relay unique fault tree was developed. But as can be seen from Table 1 for example for slave relays K601 and K603, which are driven by the same master relay K501, both fault trees are practically identical. The only difference is different slave relay tag number. So for 25 slave relays, 9 master relays and 5 main system functions we have 5 different fault trees with completely different structure. Other fault trees are with approximately the same structure, only different slave and master relays are included.

Table 1. Slave and master relays according to the equipment actuation (every relay is present on both trains, for example: K501A and K501B are two master relays one on each train)

Equipment (Systems) Actuated	Master Relay	Slave Relay	Subsystem Unavailability
Safety Injection	K501	K601 K603 K608 K610	5.19E-4
	K521	K602 K604 K609 K611	
	K525	K615 K616	
Phase A Containment Isolation	K502	K622 K624 K630	5.16E-4
	K522	K623 K629 K631	
Phase B Containment Isolation	K506	K625 K626 K632 K633	5.17E-4
Containment Ventilation Isolation	K503	K607 K614	5.16E-4
Containment Spray Actuation	K505, K519	K643 K644 K645	4.06E-4

ANALYSIS RESULTS AND DISCUSSION

The analysis of Engineered Safety Features Actuation System was performed as support system analysis because the main aim of the analysis was to prepare appropriate input for Level 1 of Probabilistic Safety Assessment for Krško Nuclear Power Plant.

Fault trees of Engineered Safety Feature Actuation System can be easily linked to Engineered Safety Features Fault Trees as support systems to perform their analysis. Nevertheless the quantitative and qualitative analysis of each Engineered Safety Features Actuation System fault tree representing each subsystem, which is path from slave relay to the measuring point, was done. Qualitative analysis which includes minimal cut sets and quantitative analysis which includes component fractional contribution, risk decrease and increase factors and subsystems unavailabilities were performed. The most important components are identified in Table 2. Matematical base and explanation of formulas for all the factors calculating are in the literature [6].

Analog channels testing does not appear at important first minimal cut sets and it does not appear in important first components ranged by fractional contribution and risk decrease and increase factors, so we conclude that its impact to system unavailability is low.

The impact of train testing where digital parts of the system are tested is greater. Basic event which represents the train testing is ranged second at minimal cut sets and it is high ranked also in table of fractional contribution, risk decrease and increase factors (Table 3).

Subsystems unavailabilities are presented in Table 1 where the list of slave and master relays and their dependencies is shown. Minimal cut sets (Table 2), fractional contribution, risk decrease and increase factors (Table 3) are presented only for the first of the fault trees containing signal path from slave relay K601A (A ... train A).

Conclusion

To get all the dependencies it is necessary to model the system as low as possible, but after the analysis is done, we can now reduce the model due to large importance of the digital part of the system. The most important components are master, slave relays and testing of the digital part of the system.

Table 2. Fault tree analysis results - minimal cut sets for subsystem containing slave relay K601A and relevant signal path

No	Unavail.	%	Minimal Cut Sets	
1.	2.18E-4	41.97	ESJC601A	Slave relay K601A failure
2.	1.34E-4	25.87	ESLT000A	Train A testing
3.	1.11E-4	21.37	ESJC501A	Master relay K501A failure
4.	4.75E-5	9.15	ESAK118A	Loss of power supply 118V AC Train A
5.	4.54E-6	.87	ESAK48VA	Loss of power supply 48 V
6.	3.70E-6	.71	ESLC518A	Card SAF OUT A518A failure
7.	7.20E-8	.01	ESZG000A	Human error no manual actuation
			ESLB516A	Card SAF OUT A516A failure
8.	7.20E-8	.01	ESZG000A	Human error no manual actuation
			ESLB517A	Card SAF OUT A517A failure
9.	7.20E-8	.01	ESZG000A	Human error no manual actuation
			ESLB206A	Card A206A failure
10.	7.20E-8	.01	ESZG000A	Human error no manual actuation
			ESLB216A	Card P/O UNIV A216A failure
11.	7.20E-8	.01	ESZG000A	Human error no manual actuation
			ESLB313A	Card P/O UNIV A313A failure
12.	4.54E-9	.00	ESZG000A	Human error no manual actuation
			ESAK15VA	Loss of power supply 15 V
13.	4.54E-9	.00	ESZG000A	Human error no manual actuation
			ESAK24VV	Loss of power supply 24 V

Table 3. Fault tree analysis results - fractional contribution, risk decrease and increase factors for subsystem containing slave relay K601A and relevant signal path

No.	Basic event	Fractional Contribution
1	ESJC601A	4.20E-01
2	ESLT000A	2.59E-01
3	ESJC501A	2.14E-01
4	ESAK118A	9.14E-02
5	ESAK48VA	8.73E-03
6	ESLC518A	7.12E-03
7	ESZG000A	7.10E-04
8	ESLB206A	1.39E-04
9	ESLB517A	1.39E-04
10	ESLB216A	1.39E-04

Top event : K601A
Description: Slave Relay K601A Failure
Mean unavailability = 5.19E-04

No.	Basic event	Risk decrease factor
1	ESJC601A	1.72E+00
2	ESLT000A	1.35E+00
3	ESJC501A	1.27E+00
4	ESAK118A	1.10E+00
5	ESLC518A	1.01E+00
6	ESAK48VA	1.01E+00
7	ESZG004A	1.00E+00
8	ESZG005A	1.00E+00
9	ESAK15VA	1.00E+00
10	ESAK24VV	1.00E+00

No.	Basic event	Risk increase factor
1	ESJC601A	1.93E+03
2	ESAK118A	1.93E+03
3	ESJC501A	1.93E+03
4	ESLC518A	1.93E+03
5	ESLT000A	1.93E+03
6	ESAK48VA	1.93E+03
7	ESLB206A	2.92E+00
8	ESLB517A	2.92E+00
9	ESAK15VA	2.92E+00
10	ESAK24VV	2.92E+00

REFERENCES

1. Nuclear Power Plant Krško, Process Instrumentation and Control, System Drawings, WEC, Pittsburgh, Pennsylvania, (1977)
2. Solid State Protection System, WEC, Baltimore, (1979)
3. N. H. Roberts, W. E. Vesely, D. F. Haasl, F. F. Goldberg, Fault Tree Handbook, NUREG-0492, US NRC, Washington, (1981)
4. IAEA-TECDOC-478, Component Reliability Data for Use in Probabilistic Safety Assessment, Vienna, Austria, (1988)
5. IEEE Guide to the Collection and Presentation of Electrical, Electronic, Sensing Component and Mechanical Equipment Reliability Data for Nuclear Power Generating Stations, The Institute of Electrical and Electronic Engineers, New York, 1983
6. Risk Spectrum PSA, Users Manual, Relcon Teknik, Sweeden, (1992)
7. A. E. Green, Safety Systems Reliability, National Centre of Systems Reliability, Warrington, United Kingdom, 1983
8. M. Kožuh, Đ. Vojnović, Common Cause in Probabilistic Safety Assessment, "Jožef Stefan" Institute, IJS-DP-5488, Ljubljana, Slovenia, (1989)
9. M. Kožuh, B. Mavko, Đ. Vojnović, Reactor Protection System Analysis, University of Ljubljana, "Jožef Stefan" Institute, Ljubljana, Slovenia, (1988)
10. M. Čepin, Engineered Safety Features Actuation System Probabilistic Safety Analysis at different Initiating Events, Thesis (in Slovene), University of Ljubljana, "Jožef Stefan" Institute, Ljubljana, Slovenia, (1992)
11. M. Čepin, M. Kožuh, Engineered Safety Feature Actuation System Unavailability Analysis, First Meeting of the Nuclear Society of Slovenia, Proceedings, Pages 237-242, Bovec, Slovenia, (1992)

RELIABILITY ASSESSMENT OF ONSITE EMERGENCY SAFETY BUSES AT OPERATING NUCLEAR POWER PLANTS*

G. Martinez-Guridi and M. A. Azarm

Brookhaven National Laboratory
Department of Advanced Technology
Upton, NY 11973

INTRODUCTION

The reliability of the supply of electricity to the electrical systems, and, in particular, to the onsite emergency safety buses, also known as the Class 1E buses, is important to the safe operation of a nuclear power plant. The power supply to the Class 1E buses comes from either or both of the following sources: 1) offsite sources (switchyards, power lines) by means of a transformer typically referred to as the Startup Transformer, and 2) the unit's main generator by means of a transformer which is usually called Auxiliary Transformer.

The main objective of this study is to evaluate the reliability of electrical circuit power paths from offsite sources to the Class 1E buses for interconnecting power circuit designs that are representative of operating nuclear power plants in the USA. Three main tasks are undertaken to conduct this study:

1. Simplified Electrical Single Line Diagrams (SESLD). A survey of operating nuclear power plants is carried out to identify designs as representative of the total population of plants. For each of the selected designs, a SESLD is developed.
2. Reliability Models. Using the SESLD, detailed fault tree models are developed for qualitative and quantitative evaluations of the reliability of each design of interconnecting power circuits.
3. Reliability Data. The relevant reliability data is surveyed to obtain the most recent and accurate values for use in the reliability models.

The first task, consisting of two subtasks: plant categorization, and the development of representative SESLDs, has been completed. The remaining two tasks are in progress. Therefore, this paper focusses principally on the work performed on Task 1.

---

* Work performed for the U.S. Nuclear Regulatory Commission

PLANT CATEGORIZATION

The objective of the first subtask is to identify a set of representative plants from the spectrum of operating nuclear power plants in the United States and to develop the associated SESLDs. To select these plants, all plants were classified according to a certain set of categories.

NUREG/CR-3992[1] proposes three main categories to classify the designs of offsite power systems and interconnecting power circuit paths. Category I deals with the independence of offsite power sources to the nuclear power plant. Category II classifies the power sources for the Class 1E buses during normal plant operation. Finally, Category III classifies the automatic manual transfer schemes for the Class 1E buses when the normal source of ac power fails and when the backup sources of offsite power fail.

General Design Criterion Number 17 (Appendix A of 10 CFR Part 50) requires, as a minimum, that there are two physically independent circuit paths from the offsite power sources to the plant's onsite power distribution system. This criterion permits a single common switchyard for both circuit paths. The categorization scheme outlined above does not address *physical* independence between power circuit paths connecting the offsite transmission network to the plant onsite power distribution system, but rather, addresses the degree of *electrical* independence provided by one switchyard, two or more switchyards and a third power circuit path.

This study classifies all the operating nuclear power sites in accordance with these categories. Reference 2 gives the resulting classification for every domestic operating nuclear power site, but is not reproduced here for lack of space. It should be noted that some of the sites were already categorized by Reference 1. For all the remaining sites, the information contained in the corresponding Updated Safety Analysis Report (USAR) was used to characterize the site.

ELECTRICAL SINGLE LINE DIAGRAMS

Simplified Electrical Single Line Diagrams (SESLD) are developed that are representative of the designs at operating nuclear power plants in the United States. To select representative designs, the categorization scheme described above is used.

First, we note that this study is mainly concerned with the availability of electrical power to the Class 1E buses. Therefore, Category I is of lesser importance to our objectives and only categories II and III are considered. Category I, "Independence of offsite power sources to the nuclear power plant", affects the availability of electrical power to the Class 1E buses, but to a lesser degree than Categories II and III.

From the results of the categorization described, we conclude that the current designs can be assigned to a certain combination of categories II and III. For example, the most common design is the following combination of categories:

> Category II, subcategory 2: During normal operation, the Class 1E buses are supplied by a preferred offsite power source .
>
> Category III, subcategory 1: If the normal ac power source fails, there are no automatic transfers and one or more manual transfers to preferred or alternate offsite power sources.

Table 1 summarizes the results for grouping all the current domestic operating electrical designs accounting for all possible combinations of categories II and III. "U" in the table means that a combination (design) is currently being used. "NA" means

Table 1. Grouping of all possible designs according to Categories II and III.

II - Power Sources for the Class 1E buses during normal operation ↓	III - Transfer Schemes when the Normal and Backup Sources of Offsite Power Fail			
	1 No automatic transfer and one or more manual transfers	2 One automatic transfer but no manual transfers	3 One automatic transfer. If this source fails, may be one or more manual transfers	4 One automatic transfer. If this source fails, there is an automatic transfer to another source
1 Main Generator	NA	U	U	U
2 Offsite Source	U	U	U	NA

that there is, at most, one plant with that configuration. As shown in Table 1, there are six representative combinations of Categories II and III (configurations) covering the majority of the current designs.

A Simplified Electrical Single Line Diagram (SESLD) is developed for each of these representative designs. Figure 1 shows the SESLD for the most common design: Category II, subcategory 2 and Category III, subcategory 1.

RELIABILITY MODELS

Task 2 consists of the development of mathematical reliability models. The reliability models employed are fault trees constructed using the Simplified Electrical Single Line Diagrams (SESLD) from Task 1. In some cases, the fault trees implement a greater level of detail than that given in the SESLDs. For example, the fault trees model the relays controlling the transfer mechanisms (opening and closure of circuit breakers). A fault tree was specifically developed for each of the six representative designs for which SESLDs were derived.

In constructing the fault trees, the components are considered to be subject to two main types of failure modes: passive and active. A passive failure mode makes the component inoperable and, therefore, interrupts the flow of power supply at that point; an example is a stuck open circuit breaker. These failures are also called non-electrical disturbances. The main contribution of an active failure mode is that it requires the actuation of other equipment to avoid the propagation of a failure. For example, a circuit breaker that fails to isolate a short circuit in a load requires another circuit breaker to open. Errors of coordination also fall in this category. The active failure modes are called electrical disturbances in the fault trees.

The level of comprehensiveness and resolution of the fault trees is enhanced by including components and failure modes not commonly considered in similar studies of Electrical Distribution Systems, for example: transfer schemes (failure), relays (common-cause failures), circuit breakers (common-cause failures and erroneous coordination), and loads (short circuit).

RELIABILITY DATA

Task 3 develops accurate values for use in the reliability models. Consistent with the level of resolution (detail) in the fault tree analysis, the reliability data analysis is structured to build upon the available generic data and extend beyond what is typically used in conventional PRAs. The starting point for constructing a generic reliability database for electrical equipment is IEEE-500[3]. The IEEE-500 data then were

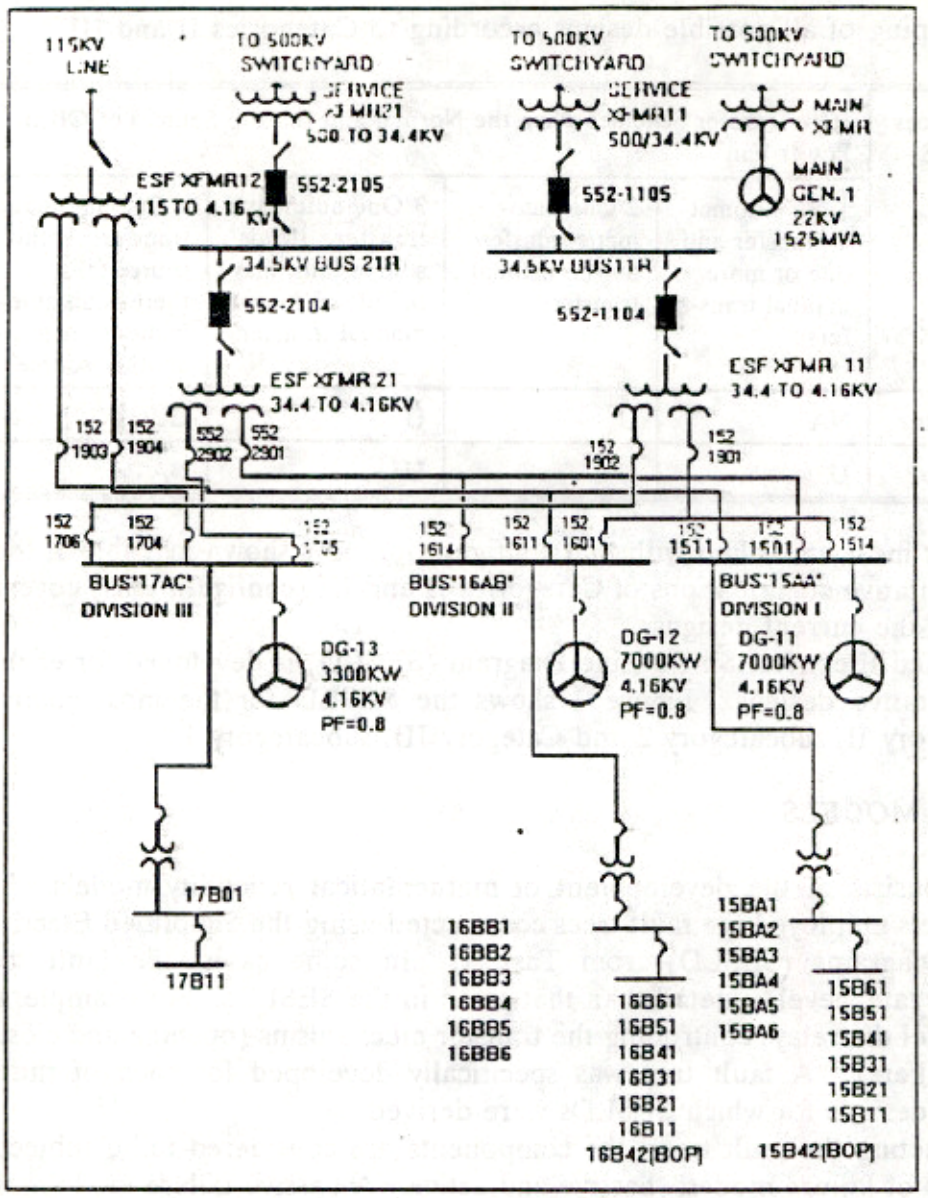


Figure 1. SESLD for Category II: 2, and Category III: 1.

compared and combined with NUCLARR[4] data which more closely represent the generic reliability data in nuclear industry. Finally, generic reliability databases, based on actual data from Nuclear Power Reliability Data System (NPRDS) for six plants were constructed and combined with the previous generic sources.

In addition, we are working on several special issues of concern to the reliability evaluation of onsite electrical distribution systems:

1. Special design reviews of various transfer mechanisms to allow proper fault-tree modelling and reliability evaluation, based on the constituent relay logic for automatic transfer and human error for manual transfer.
2. Analyses of reliability data to estimate the probability of a breaker mis-coordination and the failure in a breaker coordination scheme, based on the failure of protective relays.
3. Evaluation of reliability data to determine the frequency of shorts (single-phase or three-phase) in large loads, specifically large inductance motors.

4. Evaluation of preventive and corrective maintenance downtime for onsite electrical components, based on plant-specific data collected through two site visits.

The results of the reliability data analyses will be used to evaluate the reliability of onsite power for the six representative designs when incorporated to the fault-tree models. In addition, the generic data base developed here can be used as a basis for plant-specific reliability studies, and potentially may be incorporated to the NUCLARR system.

CONCLUSIONS

A comprehensive study to evaluate the reliability of the electrical supply to the emergency safety buses at operating nuclear power plants is being undertaken. At present, the main results obtained are those related to the first phase: Simplified Electrical Single Line Diagrams. A generic classification scheme was applied to the population of domestic operating nuclear power plants. Six representative electrical system designs were identified. A Simplified Electrical Single Line Diagram and a fault tree was developed for each design. The probabilistic data, for the components comprising the onsite electrical system, are being evaluated and updated to more accurately reflect the current reliability of these components. An evaluation of the reliability of each representative design provides a generic overview of the reliability of the onsite emergency safety buses at operating NPPs.

ACKNOWLEDGMENTS

The authors gratefully acknowledge the support provided by Frank Ashe of the U.S. Nuclear Regulatory Commission.

REFERENCES

1. Battle, R. E. , "Collection and Evaluation of Complete and Partial Losses of Offsite Power at Nuclear Power Plants", NUREG/CR--3992, Oak Ridge National Laboratory, February 1985.
2. Azarm, M.A. & Martinez-Guridi, G., "Reliability Assessment of Interconnecting Electrical Power Circuit Paths between Offsite Supplies and Onsite Power Emergency Safety Buses at Operating Nuclear Power Plants, Task 1: Simplified Electrical Single Line Diagrams", Brookhaven National Laboratory Technical Report L-2505-4-93, January 1993.
3. The Institute of Electrical and Electronics Engineers, Inc., "IEEE Guide to the Collection and Presentation of Electrical, Electronic, Sensing Component, and Mechanical Equipment Reliability Data for Nuclear-Power Generating Stations", IEEE Std. 500-1984, December 1983.
4. Gertman, D., Gilmore, W., Galyean, W., Groh, M., Gentillon, C., Gilbert, B. & Reece, W., "Nuclear Computerized Library for Assessing Reactor Reliability (NUCLARR): Summary Description", Vol. I, Revision 1, NUREG/CR-4639, EGG-2458, US Nuclear Regulatory Commission, Washington DC, May 1990.

OPTIMAL VOTING LOGIC OF HYDROCARBON GAS DETECTORS BASED ON FIELD DATA

Lars Bodsberg

SINTEF Safety and Reliability
7034 Trondheim
Norway

INTRODUCTION

Hydrocarbon gas detectors are installed to limit the consequences of a gas leakage on an offshore production installation. The detectors should provide quick and reliable detection of a hydrocarbon gas leakage and automatically initiate ignition preventing measures and production shutdown

Hydrocarbon detectors are analogue devices which initiate an alarm signal when the concentration of hydrocarbon gas exceeds a specific level of the lower explosion level (LEL). Two alarm set-points are normally used. For instance, a single gas detector may first give an alarm at 20% LEL. A second alarm at 60% LEL may automatically initiate safety actions such as a production shutdown. However, an oil company may also choose to initiate automatic production shutdown at 20% LEL. In addition to varying the set-point value for automatic safety action, the user may also employ redundant gas detectors as a means to improving the reliability of the gas detection system.

The set-point and the voting logic of redundant gas detectors is of great importance from a safety as well as production regularity point of view. There is, however, an inherent conflict between safety performance and production downtime. Field data from offshore petroleum production show that a lot of unexpected process shutdowns are due to spurious activations of gas detectors, i.e. shutdowns in which there is no abnormal gas concentration of hydrocarbon gas present in the process area. These shutdowns occur especially during maintenance activities on the installation. A set-point or voting logic which is optimal from a safety point of view (i.e. high probability of responding to a gas leak) may be unacceptable from a production regularity point of view due to high false alarm rates (e.g. the detector responds without any hydrocarbon gas in the area).

This paper presents field experience of gas detector reliability. Furthermore, guidance is provided on optimal voting logic of hydrocarbon gas detectors on an offshore production installation based on a quantitative reliability and life cycle cost analyses. Main conclusions are provided concerning the use of high set-point (60% LEL) vs low set-point (20% LEL) as criteria for automatic shutdown. For redundant detectors, also a combination of high and

low set-points are evaluated (e.g. 2oo2 - LH, two detectors generates a trip signal if at least one detector reads 60% LEL and the other detector reads at least 20% LEL).

QUANTIFICATION MODEL

Guidance on optimal voting logic and set-point are based on quantification of the following variables[1,2]:

CSU = Critical Safety Unavailability, i.e. the probability that the detector (configuration) has an unknown (dormant) failure, which prevents the detector (configuration) from initiating an automatic shutdown by the occurrence of a hydrocarbon gas leakage. The probability that the gas detector (configuration) is unavailable due to a known failure or testing, is not included in the CSU.

STR = Spurious Trip Rate, i.e. the rate of nuisance shutdowns initiated by gas detector failures. Shutdowns that are initiated deliberately, because it is known that gas detector(s) are unavailable are not included (these depends on operating strategy).

LCC = Life Cycle Costs related to gas detector. There are three main contributions to LCC:

- Primary investment
- Operation and maintenance costs
- Unavailability Costs (e.g. loss of production)

MODEL INPUT DATA

Data for failure rates of gas detectors are presented in Table 1. Failure data are given both for the detector head and the associated control card. During calculations, failure rates for these components are added to give the total value for a gas detector. Further, data are given both for low set-point and high set-point, and calculations are carried out for both these cases.

The given failure rates are split into physical and functional failures[1,3]. A functional failure occurs when the detector is unable to perform its intended function, without a physical degradation being present. Further, the rates are split according to the failure modes

- Fail-To-Operate (FTO), relevant for calculation of CSU
- Spurious Operation (SO), relevant for calculation of STR and LCC
- Non-Critical failures (NC), relevant for calculation of maintenance cost and LCC

Table 1. Failure data for hydrocarbon gas detector and detector control card (Failure rates per 10^6 hrs).

Module	Physical failure				Functional failure		
	λ_{Total}	λ_{FTO}	λ_{SO}	λ_{NC}	λ_{FTO}	λ_{SO}	TIF
Gas detector							
low set-point	11	0.5	0.5	10.0	0.5	4.5	0.002
high set-point	11	1.0	0.05	10.0	1.0	0.5	0.01
Control cards[1]							
low set-point	2	0.1	0.1	1.8	0.05	0.2	≈ 0
high set-point	2	0.2	0.01	1.8	0.1	0.02	≈ 0
Total							
low set-point	13	0.6	0.6	11.8	0.55	4.7	0.002
high set-point	13	1.2	0.06	11.8	1.1	0.52	0.01

[1]Failure rate per channel

All failure rates refer to failures that are either detectable by testing or are immediately revealed. In addition, there are some functional FTO failures, not revealed in testing (e.g failures due to lack of sensor selectivity or improper location). These are quantified by a Test-Independent Failure (TIF) probability which is defined as the probability that the detector is unable to initiate an automatic shutdown by the occurrence of an actual gas leak, even when the detector is fully tested and is apparently fault-free.

The failure date presented in Table 1 has been derived by reviewing oil company data files and records and by arranging workshop with technical experts from oil companies and suppliers of fire and gas detection systems. The estimated TIF probability is only based on expert judgement as it has been very difficult to record field data on this failure category. It should be noted that the failure rate of gas detectors may differ by a factor of 10, dependent on detector location. Physical failures of gas detectors are mainly related to poisoning of detector head, corrosion of sinter, filter blockage and electrical failures. Functional failures are related to inadequate by-pass of detectors during test and maintenance.

The modelling of dependent (common cause) failures is essential for the reliability quantification. This modelling is based on the specification of a 'multiplicity distribution', giving the fraction of failures involving exactly one detector (p_1), fraction of failures involving simultaneous failure of exactly 2 detectors due to a common cause (p_2) and fraction of failures involving failure of exactly 3 detectors due to a common cause (p_3). Separate multiplicity distributions are given for:

1. failures that are detectable in tests (or immediately revealed), being relevant for the failure rates of Table 1 (p_1= 0.95 and p_2= 0.05 in duplex system; p_1= 0.90, p_2= 0.09 and p_3= 0.01 in triplex system)

2. failure not detectable in tests (having a larger fraction of dependent failures), being relevant for the TIF probabilities of Table 1 (p_1= 0.80 and p_2= 0.20 in duplex system: p_1= 0.70, p_2= 0.22 and p_3= 0.08 in triplex system)

Input data for LCC calculations are given by Lydersen and Aarø[2].

EFFECT OF LOW VS HIGH SET-POINT

Having the option to choose between high and/or low set-point as a criteria for automatic shutdown, a lot of different configurations can be defined. The most relevant voting logics for simplex, duplex and triplex systems are presented in Figure 1. In total 12 cases are listed there, including one-out-of-three (1oo3) for high set-point, and 3oo3 for low set-point (the other combinations for 1oo3 and 3oo3 are probably less relevant). For 2oo2 and 2oo3 the possibility of trip by high set-point for at least one sensor and low set-point for the other is included as an option. Thus, for duplex systems, all five possible trip criteria are investigated.

Sensitivity analyses have demonstrated that there is a considerable uncertainty in the absolute numbers. However, Figure 1 does give a good picture of the relative performance of the different configurations. Some of the configurations seem rather irrelevant, either because they have very high CSU or because they have a very high LCC (due to a very high STR). According to Figure 1:

- The following voting logics have the highest CSU: 2oo2 - HH, 2oo2 - LH, 1oo1 - H.

  2oo2 - LH have about the same CSU as 1oo1 - H, but LCC is somewhat higher. Thus there seem to be no reason to replace 1oo1 - H by 2oo2 - LH. The voting 2oo2 - HH gives a significant improvement in safety but also have a higher LCC than 1oo1 - H.

- The following voting logics have the highest LCC: 1oo2 - L.

  In general, all triplicated solutions have a rather high LCC.

- Focusing on a high safety (at the same time avoiding configurations with the highest LCC), the following voting logics are recommended: 2oo3 - LL, 1oo3 - H.

  Observe that, from the given input data 2oo3 - LL is uniformly better than 1oo3 - H (i.e. both with respect to CSU and STR). However, input data are too uncertain to draw precise conclusions on this point.

- Focusing on a low LCC rate (at the same time having an acceptable CSU), the following voting logics are recommended: 2oo2 - LL, 3oo3 - LLL, 2oo3 - HH.

  The figures given here indicate that 2oo2 - LL is usually better than the triplicated solutions. The safety is better, and unless there is a high cost per shutdown, also LCC is lower.

- Other voting logics, giving a reasonable balance between CSU and STR are: 1oo2 - H, 2oo3 - LH.

  The figures indicate that the 1oo2 - H solution should be preferred to 2oo3 - LH, as it gives a better safety at a lower LCC.

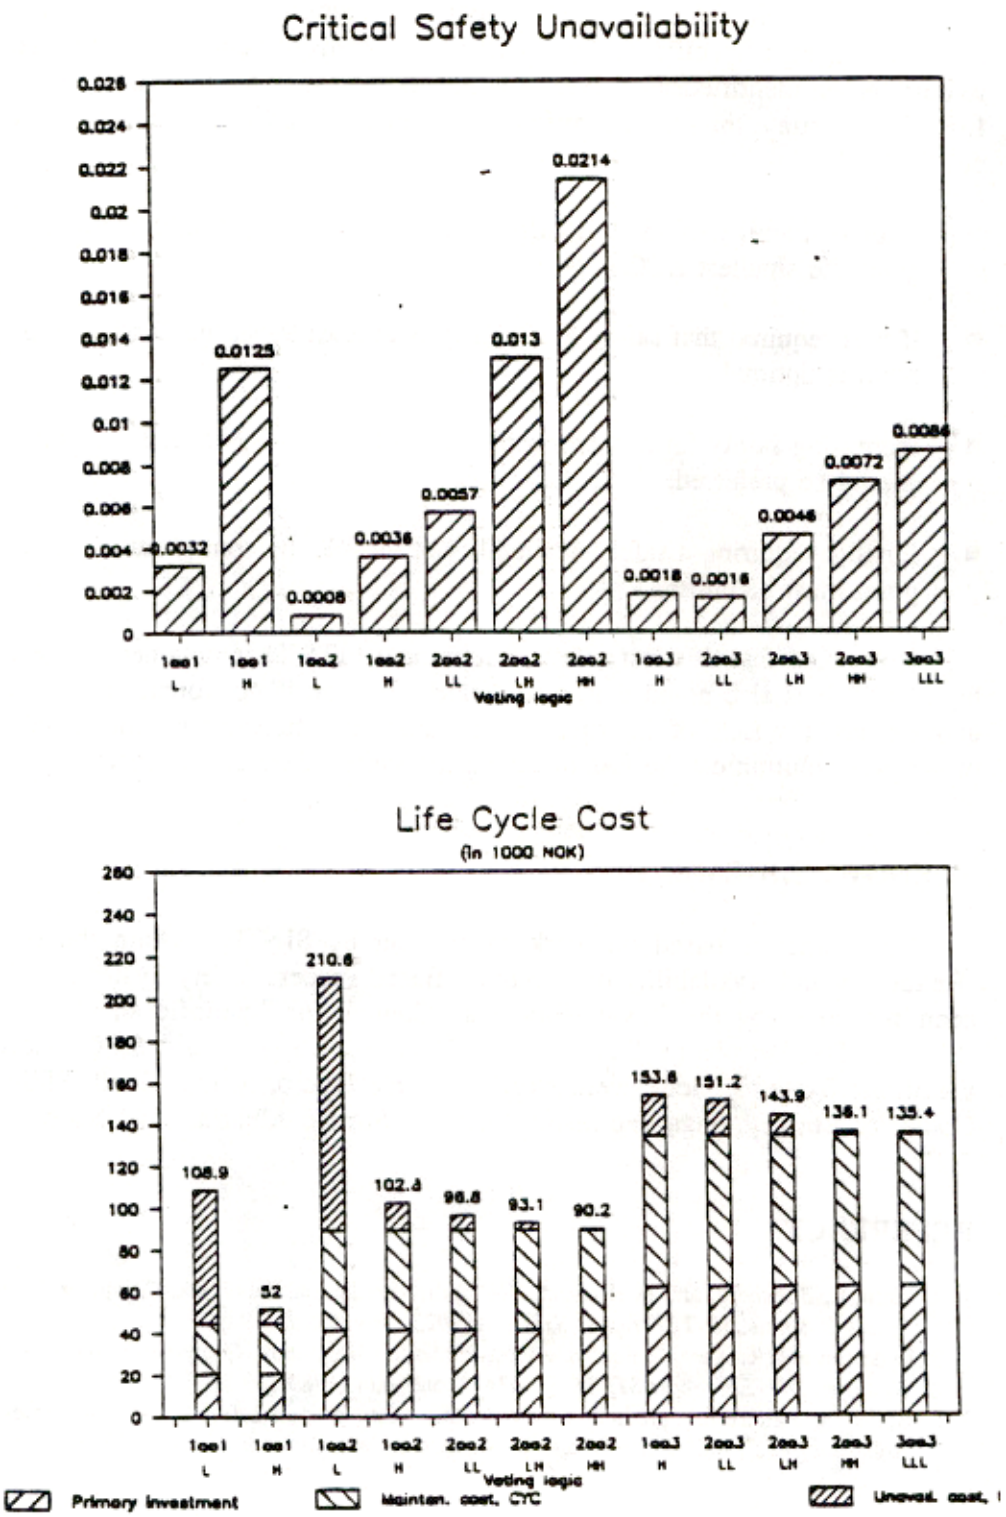


Figure 1. CSU and LCC for various voting logics, using high and/or low set-point.

CONCLUSIONS

A main conclusion is that the "optimal" configuration will depend on the chosen safety level. Requiring a specific safety availability, a configuration that is optimal with respect to cost can be identified. However, it is necessary to pay for increased safety by increased LCC. To illustrate this, the quantification gives the following results (assuming low costs per trip):

- If an availability of 98 % is sufficient, the 1oo1 voting (high set-point) is optimal, i.e. gives the smallest LCC.
- If it is required that safety availability is at least 99%, the 2oo2 voting with low set-point is optimal.
- Increasing safety further, requiring an availability of 99.5%, 1oo2 with high set-point should be preferred.
- Finally, requiring a safety availability of 99.9%, the configuration 1oo2 with low set-point must be chosen.

It is noted that this last safety requirement (99.9%) is obtained at the cost of a very high LCC. It is also noted that, a triplicated solution is not found optimal in none of the above cases. The LCC of the triplicated solutions are often so high, that usually a duplicated or simplex solution can be found, giving the same safety at less costs.

ACKNOWLEDGEMENT

This paper is based on work carried out by SINTEF within the research project "Reliability and Availability of Computer-Based Process Safety Systems". The project has been sponsored by the Royal Norwegian Council for Scientific and Industrial Research (NTNF), BP Petroleum Development Norway, EB Industry & Offshore, GP-Elliott Electronic Systems, Honeywell, Norfass, Norsk Hydro, Norske Shell, Phillips Petroleum Company Norway, Saga Petroleum, Siemens, Simrad Albatross and Statoil.

REFERENCES

1. R. Aarø, L. Bodsberg, and P. Hokstad, "Reliability Prediction Handbook, Computer-Based Process Safety Systems, *SINTEF report STF75 A89023*, Trondheim (1989).
2. S. Lydersen and R. Aarø, "Life Cycle Cost Prediction Handbook, Computer-Based Process safety Systems, *SINTEF report STF75 A89024*, Trondheim (1989).
3. L. Bodsberg, Failure rate prediction for computer-based process safety system, *in*: "Safety and Reliability '92", K.E. Petersen and B. Rasmussen, ed., Elsevier, London (1992).

056 Risk Communication

Chair: R.A. Bari, BNL

Risk Assessment - Working with Your Client
K.G. Murphy, J.E. Fitzgerald, Jr. (DOE)

RISK ASSESSMENT - WORKING WITH YOUR CLIENT

Kenneth G. Murphy[1] and Joseph E. Fitzgerald, Jr.[2]

[1]Risk Analysis Team Manager, Risk Analysis and Technology Division
[2]Deputy Assistant Secretary for Safety and Quality Assurance
Department of Energy
Washington, DC 20585

OBJECTIVE

Determining client needs is a vital first step in planning a successful risk assessment. This paper recommends a number of steps you can take to assure a thorough understanding of a client's needs and expectations and discusses ways to get the most from your risk assessment efforts. Methods for addressing the following six questions are also discussed:

1. Have you identified your client and the assessment's stakeholders?
2. Have you determined your client's needs?
3. Have you decided what to do to clarify and amplify your client's needs?
4. Have you kept in touch with your client while performing the study to assure that it is progressing in the right direction?
5. Have you assured that documentation of the results will be understood by your client and the stakeholders?
6. Have you fully explored potential uses of the analyses' intermediate products?

DEPARTMENT OF ENERGY (DOE) RISK ASSESSMENTS

The authors have nearly 40 years of combined experience as federal program managers directing technical programs, including risk assessments. Their experience as clients to risk analysts in government laboratories and private industry is the basis for the paper. The following review of risk assessment activities within DOE is included to provide additional background information regarding risk assessments.

The DOE complex includes hundreds of major facilities located at more than a dozen sites across the United States. Because facilities within the DOE complex

range from unique experimental nuclear reactors to complex chemical plants to research and development laboratories of many kinds, many different forms of risk assessment are required. Simple hazards analyses are conducted at some laboratories to identify occupational safety concerns that must be addressed to protect research personnel; at facilities that are more complex and contain substantial hazards, full-scoped probabilistic risk assessments are conducted to assure that the public is adequately protected.

An important characteristic of the DOE, as compared with other agencies, is that it owns and manages its facilities. DOE is not a regulatory agency, like the Nuclear Regulatory Commission (NRC), which has focussed regulatory responsibilities limited to nuclear safety as it concerns the public. DOE is responsible for production, for worker and public safety, and for environmental protection at its facilities. Thus, the Department's decisions must balance the needs of production with occupational safety and with public safety and environmental impact criteria.

Risk-based technologies used by DOE are derived not only from the commercial nuclear industry, but also from the chemical industry. A variety of risk management technologies are used to evaluate the many diverse attributes that must be considered to arrive at sound decisions. Because DOE can be thought of as a microcosm of the larger national industrial sector, it offers an ideal and challenging arena for the development and application of risk-based tools.

The Department is fortunate to have access to some of the most experienced risk analysts in the nation. Each of the national laboratories has strong teams of risk assessment analysts. These teams were created to support the NRC's research on the risk of commercial nuclear plants. Recently, however, the support of these teams has begun to shift to DOE, as the Department finds need for risk assessment services.

Although DOE has conducted full-scoped probabilistic risk assessments for certain of its reactors, most applications require less complex assessment approaches, such as Hazard and Operability Studies, which may use some form of semi-quantitative, risk-based assessment that roughly categorizes the frequency and consequence of postulated accident scenarios. Thus, for DOE, a risk assessment of an engineered system may mean anything from a full-scoped probabalistic risk assessment to a simple, semi-quantitative analysis.

FOCUS ON THE CLIENT

The fundamental purpose of questions related to dealing with the needs of the client is to assure that the contemplated risk-based study you are going to conduct is properly planned; that the objectives will meet the client's needs; that the information available, the risk-based tools used, and the assumptions made will satisfy the objectives; and that the final documentation of the results will satisfy the client as well as other identified users. Satisfying the client requires a concerted effort to understand, as fully as possible, the client's needs and motivations in requesting the risk-based study. Such effort is hampered by many factors, one of which involves the degree of access to the client.

Have you identified your client and the assessment's stakeholders?

There are three basic categories of clients: immediate supervisors (i.e., clients with whom you have daily or weekly contact), middle managers (i.e., clients with whom you meet once very few months), and top managers (i.e., clients with whom you have very limited contact). The strategy for understanding the client's needs is different for

each of these categories. Therefore, you must, first, specifically identify who the client is, since sometimes this is not made clear.

While asking the necessary questions to define your client, you should expand the questioning to identify all important stakeholders who have an interest in your study and attempt to understand how each stakeholder might use the results of your work. In her first press conference in May of 1993, DOE Secretary O'Leary stated that all DOE environment, safety and health documents should be prepared with an eye toward public release, indicating that she views the public as an important stakeholder when it comes to DOE environment, safety and health issues. This mandate from the Secretary clearly demonstrates the importance of identifying all the stakeholders and potential users of your study.

Have you determined your client's needs?

Even the simplest direction, such as when and where to have a meeting, is subject to misunderstanding. We have all experienced the embarrassment of walking into an empty meeting room and discovering an error in time, place, or date. When the subject is risk assessment, the potential for misunderstanding is greatly increased. Each analyst has his or her own internal definitions of what a risk assessment is, what risk means, and what the products of a risk assessment should be. The client may have concepts that are similar, but not identical, to those of the analyst. Even with daily contact and extensive discussions between analyst and client and even when both are experienced in the conduct and use of risk assessment, the chances of misunderstandings occurring are very high. Thus, clear communication and strong interaction between analyst and client are necessary. Where frequent interaction is not possible, the analyst must develop a special strategy that ensures that all meetings with the client are as effective as possible in reaching common agreement and understanding on the analyses' objectives and products.

Let's say you have met with the client for the first time and have written a list of objectives for the contemplated study. You shake the client's hand and leave. Two very useful assumptions should be made at this stage: (1) you probably misinterpreted the client's needs when you wrote down the objectives; and (2) the client probably did not explain fully or adequately his or her needs.

Have you decided what to do to clarify and amplify your client's needs?

There are many pitfalls of potential misunderstanding between analyst and client. The definition of terms always presents problems, and the definition of risk is nearly always misunderstood. The idea that risk is the product of frequency and consequence is generally agreed upon; but, beyond this, agreement ends. Risk can be averaged over a specific time period or can be defined as instantaneous. Also, risk can be measured for the most exposed individual or the average individual or can be summed over a given population. In addition, risk may be defined as basic units of dollars lost or life lost or as another measure per unit of time.

One example of potential misunderstanding is a case in which an analyst was ready to start a risk assessment of a facility that used chlorine when he happened to ask the plant manager whether he preferred the use of monetary loss or health effects as the basic impact parameter. The plant manager flinched at this and stated very emphatically that his most serious concern was the possibility that the odor of chlorine might be sensed beyond the plant boundary and that he was ready to spend a good deal of money to protect his neighbors from the slightest inconvenience. Had the analyst left the meeting without clarifying this objective, the study's results would not

have met the client's needs. But the analyst corrected his approach. The study was conducted using the odor threshold of chlorine as the impact parameter and a once-a-year frequency criterion to assure that the odor threshold would not be exceeded more than once a year. Needless to say, the resulting recommendations required substantial improvements in the mitigative systems that controlled the frequent leakage events that result when transferring chlorine from trucks to storage tanks. The newly installed mitigative systems to satisfy these stringent criteria were indeed impressive.

With such potential for misunderstanding, the analyst's very next step after the first meeting with the client should be to draft a concise report that defines the analyst's interpretation of the client's needs and discusses how the needs will be satisfied. Table 1 provides suggestions for the contents of this report, which we shall call the "Analysis Plan."

Table 1. Requirements for the "Analysis Plan" report.

Purpose

- Clearly state the purpose and summarize all major objectives.
- List the major results and products.
- Explain the benefits and limitations.
- Identify all potential users and stakeholders and discuss how they will benefit from the results and products of the study.

Scope

- Define the analytical boundaries of the process
  - the external boundary, i.e., the structures and systems to be included and excluded
  - the limit of resolution, i.e., the smallest system elements to be included and excluded.
- Define the process' operating regimes to be included and excluded.
- Define the types of accident initiators to be included and excluded.
- List other major assumptions and scoping limitations.
- Summarize the study's cost assumptions and identify key factors that might impact cost.

Methodology

- State the overall methodology, the type of study, and whether it is qualitative, semi-quantitative, or quantitative.
- Define all technical terms.
- Define the state of process information and the need to develop additional design, process behavior, and phenomenological information.
- Define the type of hazards to be analyzed and how they're to be identified and screened.
- Explain how the accident scenarios are to be developed and evaluated.
- Define the impact parameters
  - the targets, e.g., equipment, worker, public, environment
  - the impacts, e.g., dollar loss, health effects
  - the risk, e.g., define qualitative measures, semi-quantitative measures, individual or collective risk measures.
- Define how uncertainties will be treated and evaluated.
- Provide examples of how the study's results will be presented.

The most vital aspect of the Analysis Plan is a carefully crafted statement of purpose with further explanation, as required, to completely capture the analyst's interpretation of the client's needs. Complete satisfaction of all needs might not be entirely possible, given the quality of the information available and the analyses' cost limitations - the analyst must plainly discuss these limitations.

The draft Analysis Plan, and associated presentations to the client, become vehicles for further analyst-client interaction. The analyst should use the report and presentations to bring the client up to speed as to the capabilities of the selected risk assessment-risk management techniques, their benefits and limitations, and the influence uncertainties might have on the results. The analyst and client should question and discuss each element in the Analysis Plan to the point where both believe they have arrived at mutual understanding. The result should be an updated Analysis Plan that forms the basic understanding between analyst and client. The Analysis Plan is then used as the starting point of the study.

Have you kept in touch with your client while performing the study to assure that it is progressing in the right direction?

Any major risk study always results in the unexpected. Important elements may be missing from the plan. For instance you may not be able to predict how a system variable or component will behave under accident conditions; or, an unexpected system weakness that may impact the results of the study may be identified.

When you are well into the study, basic questions dealing with the adequacy of the methodologies being used or the meaningfulness of the expected results will be raised. The analyst needs to plan for periodic meetings with the client to assure that all such questions are discussed and any changes in direction approved by the client. The analyst should also expect refinements in the study's objectives as the client becomes more aware of the direction the study is taking and should assume that the client is on a learning curve and that mid-course direction changes are likely. Don't let too much time elapse between meetings with the client; otherwise, you may be wasting time going in the wrong direction.

Have you assured that documentation of the results will be understood by your client and the stakeholders?

Major risk analyses lose much of their impact when the most significant results are poorly documented. This is best illustrated by those analyses having final reports that emphasize risk numbers instead of engineering insights. Risk experts continue to say that risk numbers are the least useful results of such analyses; yet, many analyses place the numbers up front and the engineering insights at the end.

There are many examples of well-documented analyses. The analyst would serve the client well by reviewing the best-documented analyses and by understanding and using the best of the current techniques to summarize risk-based results. Also, since the field of risk communication is advancing, the analyst should become familiar with these advances in order to improve documentation and communication of the analyses' results.

As a very minimum, the purpose, scope, and methodology descriptions found in the Analysis Plan should be updated and reproduced in the final report. The executive summary should contain a concise summary of these descriptions, along with an explanation of the essence of the study in terms of important advances in knowledge and identification of the processes' strengths and weaknesses.

Have you fully explored the potential uses of the analyses' intermediate products?

It is constantly surprising to discover how useful intermediate products of a risk assessment can be and to learn that, to some, such products may be the study's most important products. The analyst should be sensitive to this possibility and should assure that potential users of such information are made aware of its existence.

A good example of a useful intermediate product is the redrawing of system piping and instrumentation drawings (P&IDs). Plant P&IDs are not adequate to effectively display critical safety functions and their interrelationships. The situation is much the same as that of a high school trigonometry class. We all remember the thrill of finding the answer to a particularly tough trigonometry problem after redrawing the problem many times, and finally discovering the one configuration that leads the eye to the solution. In the same way, we redraw P&IDs of a complex process, attempting to display the critical systems, their functions, and interrelationships effectively.

One of the most useful products of the K-Reactor risk assessment performed by Los Alamos National Laboratory was an integrated systems schematic. The schematic was an 11-inch by 17-inch multicolor drawing that showed the piping and components of a dozen primary and secondary reactor systems. Concentrating so much information on a single sheet of paper was a difficult job that required special creativity in information organization. The drawing was vital to the risk analysts in the quest to understand the workings of the reactor. It is not an exaggeration to state that this one product, by far, was the most sought after product of the risk assessment. System engineers stood in line to get copies and jealously guarded them. Why? Because the drawing displayed important information in a new and effective way.

There are many similar examples of risk assessment intermediate products that have importance outside of the study itself. Consider the case of understanding the impacts of electrical supply. Electrical engineers are in the habit of reviewing the electrical line diagrams to find the answers to what-if questions regarding how electrical power can be supplied to components under accident conditions. As with the P&IDs, this can be a time-consuming problem requiring the study of numerous drawings. There is usually no effort to shortcut this type of review; therefore, answers are not immediate and always require the review of multiple electrical line drawings.

In probabilistic risk assessments, the dependency upon electric power under numerous initial conditions must be understood. Probabilistic risk assessment analysts have developed diagrammatic approaches to summarize this information: so-called electrical dependency charts. Putting these charts in the hands of plant electrical engineers has made their work easier and allowed them to further identify subtle weaknesses in plant operation under off-normal conditions.

CONCLUSIONS

Start off on the right foot when planning a risk-based study. Understand your client's needs. Develop an Analysis Plan, discuss each element of the plan with your client. Arrive at a common understanding concerning every element. This initial time and trouble expended by you and your client will pay dividends later on when the study is in full swing. You will also be assuring the best possible chance that the client will be happy with your work and with the study's results.

057 Resource Planning and Value-Impact Applications

Chair: K. Klein, USDOE

Environment, Safety & Health Risk Management: Program Planning and Budget Decision-Making

F.E. Tooper (USDOE); J.L. Von Herrmann (CYCLA)

Control Operator Overload: Identification and Prevention

D.A. Strobhar (Beville Eng.)

Probabilistic Cost-Benefit Analysis of Enhanced Safety Features for Strategic Nuclear Weapons at a Representative Location

D.R. Stephens, C.H. Hall, G.S. Holman, K.F. Graham, T.F. Harvey, F.J.D. Serduke (LLNL)

ENVIRONMENT, SAFETY & HEALTH RISK MANAGEMENT:

PROGRAM PLANNING AND BUDGET DECISION-MAKING

Frank Tooper[1] and James L. von Herrmann[2]

[1]U.S. Department of Energy
[2]CYCLA Corporation

INTRODUCTION

The Department of Energy (DOE) is confronted with a management challenge common to many public agencies and private corporations: How to improve its environmental, safety, and health programs within a tightly constrained budget environment.

DOE owns and operates hundreds of basic and applied science laboratories and weapons material production facilities throughout the United States, including eight national laboratories. The operations at DOE labs and production facilities involve the use of many extremely hazardous radiological and toxic substances. In addition, many of the facilities in which these operations take place were built forty or fifty years ago. Internal and external assessments over the last few years have uncovered a wide variety of environmental, safety, and health deficiencies at DOE facilities. DOE has responded to the identified problems by developing and promulgating new environmental, safety and health standards and each facility has developed numerous corrective action plans to address deficiencies and comply with the new regulations.

However, at the same time that DOE is endeavoring to significantly upgrade its safety and health programs, it is being subjected, as are all government agencies, to severe budget constraints. There is simply not enough money and qualified staff to implement all of the environmental, safety and health improvements at once.

In response to this situation, the Secretary of Energy in November 1991 directed that program planning and budgeting be done:

"within an atmosphere of budgetary constraint [where] some activities will need to be eliminated, reduced in scope, or stretched out in time to accommodate new and higher priority initiatives and enhancements to ongoing programs."

In response to the Secretary's request, the Office of Environment, Safety, and Health, (EH) has coordinated the development and implementation of a risk-based DOE ES&H Management Planning Process that defines, prioritizes, and allocates budget resources to the most important ES&H activities.

Unlike many other risk-based planning processes, the ES&H Management Plan is directly linked to the Departmental budget process. Planning information generated through the process is used to support management decision-making during the Department's internal budget review process, and an ES&H "crosscut" budget derived directly from the ES&H Management Plan is submitted to the Office of Management and Budget (OMB) and to then to Congress in support of the Department's budget request.

THE DOE RISK-BASED PLANNING AND BUDGETING PROCESS

DOE's Environment, Safety and Health (ES&H) Management Plan is produced through a combination of top-down guidance and bottom-up analysis and decision-making. It is a **continuous, risk-based, resource-constrained, management process** designed to improve the way that DOE and its contractors utilize available resources to manage ES&H risks. The process is designed to help managers produce and communicate integrated information necessary for S&H strategic planning and operational management. Both the development and output of the ES&H Management Plan are directly tied to the Departmental budgeting process. The major steps of the Plan development process are:

- Strategic Planning Guidance is provided by the Secretary of Energy and key ES&H issues are identified by the Assistant Secretary for Environment, Safety, and Health.

- Headquarters managers provide direction to operating facilities, including budget targets to be used in the preparation of the field planning data.

- Field elements identify ES&H activities/programs, prepare Activity Data Sheets (ADSs) to document these activities, and enter the ADSs into the ES&H Management Planning Database System.

- ADSs are ranked using the DOE ES&H Risk-Based Priority Model (RPM). The ranking is reviewed by successive levels of management and other planning factors (such as precedence and coupling relationships between ADSs, strategic factors, etc.) are applied to adjust rankings.

- Available resources are allocated to the activities. The ranked listing of ADSs is used to develop three budget scenarios ("decrement", "target", and "program planning") for the Department's Internal Review Budget (IRB) process. Less important activities that cannot be accomplished within the prescribed budget constraints are either reduced in scope or deferred to subsequent years.

- Internal budget decisions are reflected in the ES&H Management Plans and an ES&H "crosscut budget" is prepared for OMB.

- Final OMB allowances are reflected in the ES&H Management Plans and a revised ES&H "crosscut budget" is prepared to accompany the Department's Congressional budget request.

- Budget decisions and planning implications are communicated back to the operating facilities, and the updated ES&H Management Plans provide the baseline for the next annual planning and budgeting cycle.

THE DOE RISK-BASED PRIORITY MODEL (RPM)

Environmental, Safety, and Health issues and the activities proposed to resolve these issues are evaluated and ranked using the DOE Risk-Based Priority Model (RPM). The RPM is a simple multi-attribute utility model designed to support DOE management decision-making. Issues can be ranked based on the risk associated with the current situation (i.e. the level of consequences that might occur if the identified problem is not mitigated and the likelihood of experiencing those consequences). Activities are ranked based on their risk-reduction potential (i.e. the difference in expected risk before and after implementation of the activity).

The DOE RPM has been designed specifically to be *a management tool* that supports the management decision environment addressed by the ES&H Management Plan:

- Implementation at hundreds of facilities located throughout the United States ranging from small solar energy laboratories to large nuclear weapons facilities;

- A large number of issues and activities that must be processed in a relatively short period of time and maintained current throughout the year;

- Issues and activities of many different types ranging from specific Occupational Safety improvements to broad improvements in conduct of operations;

- A wide spectrum of risk ranging from issues from high likelihood/low consequence occupational safety issues through low likelihood/high consequence nuclear safety issues;

- A relatively low level of quantitative risk assessment within the Department.

The RPM supports but does not replace the expertise of field and headquarters staff and management. The RPM is viewed as one tool within an overall management process that structures and focuses management decision-making. The RPM does not attempt to encompass all factors or provide all information required to determine the ultimate priority of an issue or activity. The ES&H planning process is designed to allow various factors important to the setting of priorities (e.g. strategic considerations concerning the expected lifetime of the facility, the level of uncertainty in the risk evaluation, or project management considerations such as precedence relationships among activities) to be incorporated into the final adjusted ranking. The RPM is not intended to provide a concise quantification of risk, but structures management experience and knowledge (and any available quantitative risk data) into a defensible and traceable relative ranking of ES&H issues and activities.

THE DOE ES&H MANAGEMENT PLANNING DATABASE SYSTEM

A powerful PC database system with a detailed User's Manual has been developed to support the production, analysis, and utilization of the ES&H Management Plans. The ES&H Database System has been designed using the FoxPro development software and is distributed to all DOE facilities as a run-time application. The user interface consists of menus, windows, dialogues, and other features that make it easy for the user to communicate with the Database System.

In addition to direct input and update of data, it provides powerful but easy-to-use filtering, indexing (sorting), and output reporting capabilities. The Database System includes an RPM "Scoring Screen" for each issue or activity record and a "Resource Screen" that captures budget costs of an Activity in a variety of resource categories. The database system supports efficient risk-management and resource allocation by presenting activities in risk-ranked order, displaying budget cut-off lines and scenarios, and allowing automatic deferral of less important activities that fall "below the line" to later years. The database system is comprised of four major modules:

1) **Assessments**, in which users can enter records and provide information about specific assessments (evaluations, etc.) that spawned ES&H issues;

2) **Issues**, in which users can enter records and provide information about specific ES&H issues identified in the assessments. Many issues can be linked to an assessment in the database.

3) **Activities**, in which users can enter records and provide information about specific ES&H activities defined to resolve the identified issues. Activities can be linked to issues within the database and through issues to assessments.

4) **Milestones**, in which users can enter records and provide information about the specific tasks or deliverables associated with an Activity and the due dates for these accomplishments. By linking the Milestones to Activities (and thereby to Issues and Assessments), this module allows the database to serve as a powerful,

integrated Commitment Management and Tracking System.

ACCOMPLISHMENTS AND FUTURE PLANS

The ES&H Management Planning Process is now in its third full year of implementation. During its first two years of implementation, the following has been accomplished:

- Operating facilities, field elements, and headquarters organizations all have a much improved basis for communicating and discussing the scope, priorities, costs, and schedules of ES&H activities.

- Headquarters senior managers and the Chief Financial Officer have greater assurance that ES&H risks are being effectively managed within the Department and that Departmental funds devoted to ES&H are being applied in a cost-effective manner to the most important activities.

- Facility managers have a much improved understanding of the scope of required ES&H activities and resources devoted to these activities.

- Facility managers have a substantially improved mechanism for identifying and allocating constrained overhead funds to the highest priority activities and communicating these planning decisions.

- Field Offices have a much better appreciation of the scope and priorities of contractor ES&H activities, allowing more effective and focussed oversight.

- Headquarters line managers have a greater ability to ensure that ES&H funds are being effectively applied and that Departmental ES&H policies and priorities are being addressed in the field.

- Line organizations and oversight groups have a structured, defensible, and traceable mechanism for discussing and negotiating compliance schedules that reflect realistic budget assumptions.

During the current planning and budgeting cycle, the ES&H Management Planning Process will continue to broaden its focus and extend its benefits within the Department by:

- Providing better integration with Departmental programs and systems that define root causes of ES&H problems and define cost-effective options for issue resolution;

- Providing better integration with Departmental programs and systems to enhance and ensure accountability for ES&H performance;

- Extend risk-based ES&H management processes to support Departmental efforts to improve overhead management and cost-control.

CONTROL OPERATOR OVERLOAD: IDENTIFICATION AND PREVENTION

David A. Strobhar

Beville Engineering, Inc.
925 Congress Park Drive
Dayton, Ohio 45459

INTRODUCTION

The control operator has become an increasingly critical element in the safety and operation of most processing facilities. The use of distributed control systems has both increased the operator's span of control and focused the control responsibilities onto him. No longer is an entire crew watching the board and ready to help out. The modern control operator is on his own to respond to plant disturbances.

Assurance that the control operator will not become overloaded during an upset condition and will be able to perform his duties is a critical issue. The distributed control system allows an almost limitless amount of information to be presented to the operator, information that must be processed or overlooked. The distributed control system also allows the operator's span of control to be increased, with proximity of the unit or size of the control boards no longer a factor in how much a control operator can handle.

In order to assure that the operator will not be overloaded, it is necessary that operator loading be measured and that the factors contributing to overload be identified and eliminated from the system. This paper summarizes data on control operator performance from field studies at more than thirty processing plants and references it to theoretical principles in human factors engineering.

NATURE OF HUMAN PERFORMANCE

Human performance has very few absolutes. A given level of performance in a man-machine system can be obtained through various combinations of the following:

. Selection (Who is chosen to do the tasks?)
. Training (What are the skill and knowledge levels of the individual?)
. Motivation (How well do they want to do, are their circadium rhythms disrupted?)
. Supervision (Have the goals been well established, individual empowered?)

- Task allocation (Are the tasks appropriate for human control?)
- The interface between the man and the machine/process (workstation, alarms, displays)

Similarly, improvement in performance can be obtained by improving any one variable. Or, poor characteristics of one variable can be compensated by superior characteristics of another.

There are, however, absolutes in human characteristics, particularly in the human information processing system. Since the role of a board operator is primarily that of information processing, these characteristics become key factors in the cause and prevention of operator overload. Two of the key characteristics are human short term memory limitations and mental workload demands. Information processing tasks will place demands on both short term memory and mental workload. This is colloquially described as workload. When the demands become excessive, the probability of operator error increases dramatically, and the human sub-system, the operator, becomes overloaded.

Human short term memory limitations make the human information processing system capacity limited. Short term memory is limited to about seven chunks of information.[1] A person cannot retain in his short term memory, or consciousness, much more than seven distinct chunks of information. While this is an often overstated principle, it has profound implications in board operator performance. Since distributed control can present almost limitless amounts of information, the operator's short term memory can quickly become overloaded. This phenomenon is particularly evident in the processing of alarm information.

A fluid catalytic cracking unit at a small refinery experienced a failure of one of the slide valves separating the reactor from the regenerator. The resulting incident at one point resulted in 300 alarms actuating in a ten minute period. The information on the alarms was shown on a CRT screen sequentially by time of actuation. Needless to say, the operator could not process any more alarm information at this point and the alarm system had simply become a nuisance. Fortunately, the operator knew the problem. However, it would be unlikely that the operator could have detected a second independent problem using the alarm system. This type of example is typical of DCS operators, who routinely say that the alarm system is useless during an upset.

Even if the capacity limitation of short term memory is not overloaded, there must be sufficient mental resources to act upon the information. This is the second key consideration of human information processing. Demands on mental resource requirements are theorized to be task dependent and vary with the individual. The task characteristics are the complexity of the task, the time constraints required to perform the task, and the stress of the task (or consequences of error).[2]

A prime attribute of human involvement in systems is that people rarely shut down in an overloaded situation. We continue to perform, but our error rate increases dramatically. High mental resource demands reflect this phenomenon. Tasks requiring over 40% mental workload can be handled by a person. However, the longer that level of mental resource demand is maintained, the greater the probability of operator error.[3]

The configuration of the distributed control system has a major interaction with short term memory and mental workload. If information is spread out over multiple display pages, the operator may have to hold the information in short term memory while accessing another display, increasing the potential for a capacity bottleneck. Display paging uses mental resource, reducing the availability for other applications and increasing the potential for operator error. As is evident, calculation of human error rates independent of the design of the distributed control system will be fraught with errors.

OPERATOR WORKLOAD ANALYSIS

Determination of operator workload must be done for both steady state and off-normal operation. The demands of each are different, but both are important. The importance of evaluating off-normal situations (i.e., upsets) is obvious. Steady state evaluation is important in that it is the precursor for the upset situation, it provides insight into upset characteristics that are not easy to obtain in an upset, and it reflects cumulative demands on the board operator.

Steady State Analysis

Determination of board operator loading for steady state analysis requires determination of time spent on job related tasks and the number of tasks performed. The board job in most processing plants is highly reactive, creating a system-paced job. Simple examination of time on job related activities fails to capture the nature of the job. Likewise, demands other than those imposed by the process (maintenance requirements, outside operator interactions, administrative, SPC) must be accounted for in the overall operator workload.

Beville Engineering determines the steady state operator loading through calculation of a busyness factor, reflecting both task duration and frequency. Figure 1 shows a plot of board operators' steady state demands, reflecting the task duration and task frequency.

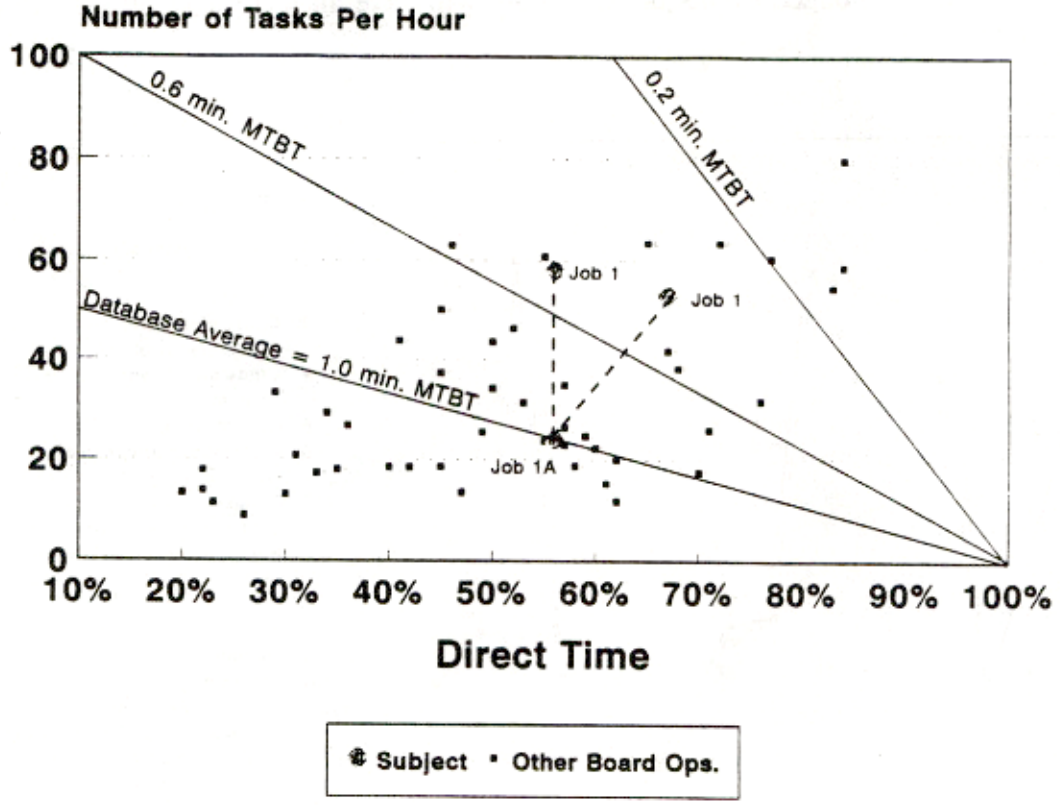


Figure 1. Steady State Operator Workload

The plot shows the results of almost fifty data samples of refinery board operators. The vertical plot is the number of tasks per hour the operators performed. The horizontal plot is the percentage of the sample time (four hours) spent directly on job related tasks. The diagonal lines represent equivalency curves reflecting mean time between tasks (MTBT) for varying numbers of tasks and levels of direct time. The lower the MTBT, the busier the operator. Increasing workload is reflected in increasing direct time, increasing numbers of tasks, and in the resulting decrease in MTBT.

The point labeled Job #1 in the figure was analyzed because management felt the operator to be overloaded. As the figure clearly shows, management's perception was valid. However, a subsequent analysis of the data indicates that much of the loading resulted from an alarm actuation rate four times the industry average and a communication rate with outside operators also significantly above the industry average. Altering the alarm actuation rate and communication rate would change the loading to that identified as Job #1a.

Off-Normal Analysis

The use of task duration and frequency is insufficient for off-normal analysis, both because the duration has a ceiling effect (there are no breaks in an upset) and the data are obtained through observations which aren't amenable to the unplanned nature of upsets. Therefore, a different approach has to be utilized.

Beville uses the Subjective Workload Assessment Technique (SWAT), developed by the U.S. Air Force, to measure off-normal workload. SWAT requires the rating of off-normal events by operators which are then mapped to a scale of mental workload demands. Figure 2 shows the result of such an analysis for board operators at a major refinery. The figure shows the mental workload demands for 16 events, with the operators rating each event for working with an experienced crew and an inexperienced crew.

As can be seen from the figure, Tasks 4, 5, 6, 7 and 14 utilize more than the 40% mental workload. Prolonged events of this nature will have an increasing probability of operator error.

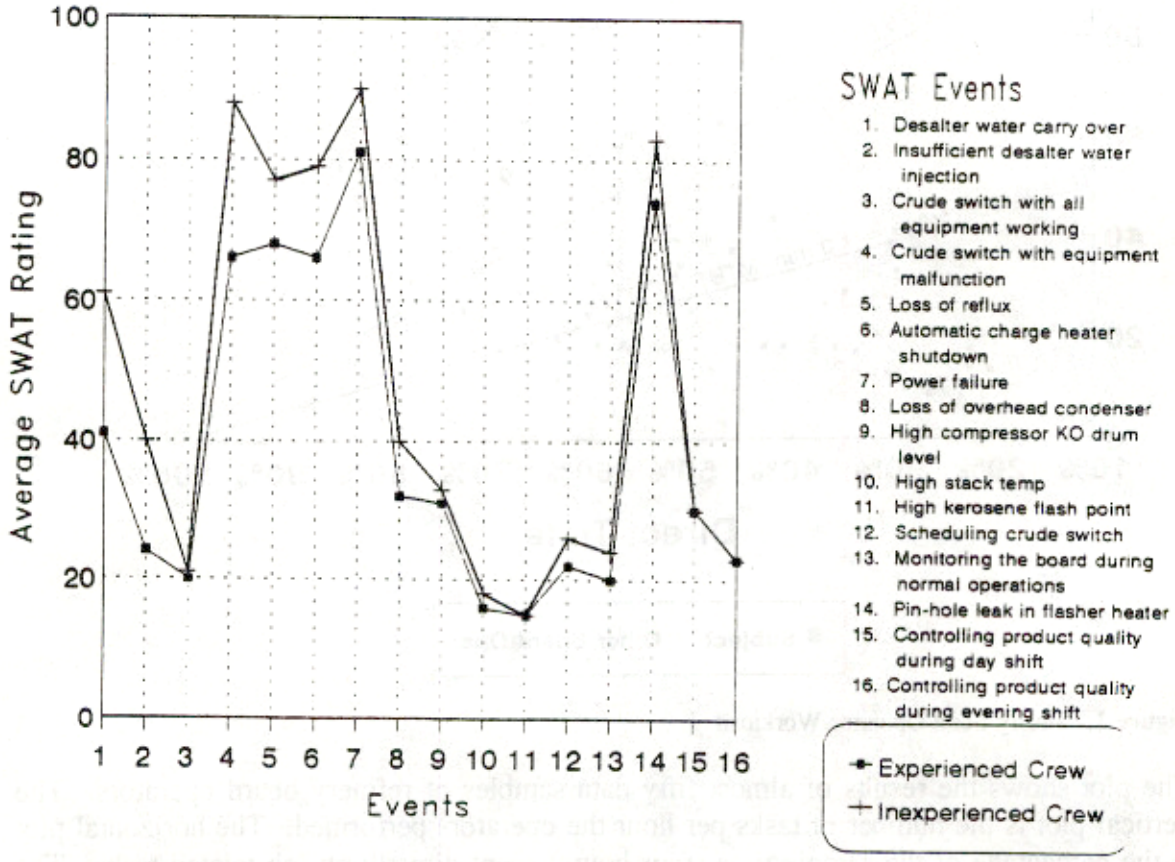


Figure 2. Example of Upset Workload Analysis

Inspection of the figure also shows the impact of experience or training on workload demands. Tasks 1, 2, and 4 all imposed significantly less workload on the board operator

when an experienced or well trained crew was working the unit.

DESIGNS TO REDUCE WORKLOAD

If an operator overload situation is possible, what can be done to minimize the probability or lower the workload? The impact of enhanced training was quite clear from the SWAT profile. However, configuration of the operator-process interface offers great potential to reduce workload demands. In particular, properly structuring the alarm and display system is a prime opportunity to reduce the potential for the operator to become overloaded.

Enhancing alarm presentation is simple in concept. First, ensure that each alarm prompts a unique operator action. This means that there should be no alarms for which the operator takes no action and that there are not multiple alarms that all prompt the same behavior. Simple application of this principle to a variety of processing plants has reduced the number of points that can alarm by 70%. The result is a ratio of alarms to controllers of less than 3:1. Second, provide qualitative information on alarm importance by system. Alarms need to be prioritized and the operator provided indication on the status of all alarms on various systems in the plant. This usually requires creation of common trouble alarms for all the alarms in each system, indicating the priority of the alarms by color and position. This type of qualitative presentation allows the operator to process large numbers of alarms with little demand on the information processing system.

Improvement of the display system can be done through ensuring that the display system is compact. A compact display system will reduce the amount of paging and load on short term memory demands. Creating a compact display system requires that the display be built for a qualified operator. A qualified operator knows the hardware he is controlling; therefore, extraneous information on what is in a pipe, or the number of trays in a tower, or the number of exchangers in a exchanger bank need not be shown on a graphic display. The resulting savings in space can be used for real information, with again a potential reduction in the number of displays by 70%.

Enhancing the alarm and display system reduces workload for both normal and off-normal operation. Normal operation is improved due to a reduction in the number of tasks the operator must perform, i.e., fewer alarms to respond to and less pages that must be accessed. Upset response is improved because with less pages more time is available to correct the problem, less mental resources are utilized, and the complexity of the problem is reduced due to proper grouping of the information. The alarm status indicators allow chunking of the alarm information so that the alarm system is actually useful to the operator during an upset.

SUMMARY

Overloading the control operator is a justified concern in the processing industry. Inherent limitations in the human information processing system create the potential for operator overload and error. Examination of operator task frequencies and durations can indicate when steady state loading is becoming excessive. Excessive mental workload for off-normal situations can be ascertained through mental workload assessment techniques such as SWAT. Redesign of the operator-process interface, the alarms and the displays, has significant potential to reduce the probability of an operator becoming overloaded.

NOTES

1. G.A. Miller, The magical number seven, plus or minus two: some limits on our capacity for processing information, *Psychological Review*, 63 (1956)
2. G.B. Reid, C.A. Shingledecker, and F.T. Eggemeier, Application of conjoint measurement to workload scale development, in *Proceedings of the Human Factors Society-25th Annual Meeting* (1981)
3. G.B. Reid, and H.A. Coole, Critical SWAT values for predicting operator overload, in *Proceedings of the Human Factors Society-32nd Annual Meeting* (1988)

PROBABILISTIC COST-BENEFIT ANALYSIS OF ENHANCED SAFETY FEATURES FOR STRATEGIC NUCLEAR WEAPONS AT A REPRESENTATIVE LOCATION*

Douglas R. Stephens, Charles H. Hall, Garry S. Holman,
Keith F. Graham, Ted F. Harvey and Franklin J. D. Serduke

Lawrence Livermore National Laboratory
P.O. Box 808, L-85
Livermore, CA 94551

INTRODUCTION

We carried out a demonstration analysis of the value of developing and implementing enhanced safety features for nuclear weapons in the U.S. stockpile. We modified an approach that the Nuclear Regulatory Commission (NRC) developed in response to a Congressional directive that NRC assess the "value-impact" of regulatory actions for commercial nuclear power plants. Because improving weapon safety shares some basic objectives with NRC regulations, i.e., protecting public health and safety from the effects of accidents involving radioactive materials, we believe the NRC approach to be appropriate for evaluating weapons-safety cost-benefit issues.

Impact analysis includes not only direct *costs* associated with retrofitting the weapon system, but also the expected costs (or economic risks) that are avoided by the action, i.e., the *benefits*. The categories of impacts that we addressed are listed in Table 1.

Table 1. Impacts addressed by this study.

Retrofit costs	Avoided costs	
• Research and development (R&D)	• Physical and radiation damage	• Environment decontamination and cleanup
• Production	• Property loss	• Natural resource damage
• Implementation	• Health costs	• Emergency response
• Other life-cycle costs	• Evacuation and relocation	• Litigation and claims
	• Debris removal and disposal	• Property condemnation
	• Losses of business and wages	• Medical and environment monitoring
	• Hearings, public information	• R&D

* Work performed by Lawrence Livermore National Laboratory under the auspices of the U.S. Department of Energy under contract number W-7405-ENG-48.

RETROFIT COSTS

Retrofit costs are costs of implementing enhanced safety features in strategic nuclear weapons. These include the required research and development, testing and engineering (RDT&E) of the retrofitted weapon by DOE, and any RDT&E for modifications of the delivery system by the DOD. Most production costs for a completed weapon system would be incurred by DOE. The DOD may also incur some production expenditures for software or hardware changes or system testing. Other life-cycle costs, such as transportation, would be incurred by both agencies.

Credible retrofit costs were obtained from recent government studies. These are summarized in Table 2 as expected values with upper and lower bounds. We assumed a Drell recommended retrofit[1*] for 1000 chemical high explosive (CHE) weapons.

Table 2. Estimated retrofit costs ($M) of safety features. The lower and upper bounds delimit about a 70% confidence range, i.e., one standard deviation.

Operation	Lower bound	Expected value	Upper bound
DOE R&D	400	500	600
DOE production	300	700	1000
DOD implementation	10	100	1000
Rounded Totals	700	1300	2600

AVOIDED COSTS

Enhanced safety features would decrease the likelihood both of a nuclear weapon accident and incurring the high costs of its aftermath. Among the benefits of enhanced safety are the avoided-accident costs resulting from reduction of the risk of such accidents. We considered three accident scenarios: (1) Pu release by a fuel fire, (2) Pu release by nonnuclear high-explosive (HE) detonation, and (3) unintended nuclear detonation. For a given accident, the "actuarial costs" (expected costs) are a product of three terms: **(consequence) × (mitigation cost) × (accident probability)**. In our analysis, these costs and consequences have been divided into three components: onsite economic risk (e.g., risk to a military base where the accident was assumed to occur), offsite economic risk (e.g., risk to the general public), and government operations (e.g., accident response, investigations, hearings, public information, and R&D).

Accident Consequence

The accident site was assumed to be a military base in a mixed urban-rural area. We assumed that the nuclear weapons accident occurred at a single, specific location on the base. State-of-the-art atmospheric transport codes were employed for consequence estimation. The Lawrence Livermore National Laboratory (LLNL) KDFOC3 code2 was used to estimate fallout from the nuclear detonation. Plutonium dispersal by detonation and by fire were assessed by the LLNL ADPIC model.3 The cloudrise model for the HE detonation was a version of the ERAD model4 developed at Sandia National Laboratories (SNL). For fuel-fire buoyancy, we used a version of the Briggs model.5

The accidental nuclear detonation was assumed to have a yield of 1 kt. For both the HE-detonation and fuel-fire dispersal scenarios we assumed that 10 kg of Pu were involved. In the nuclear detonation case, we assumed 100% of the Pu was lofted in aerosolized form and that 20% of the Pu was in the respirable size range. Uncertainties in these quantities are believed to be factors of two to three. Uncertainties in the source terms for fuel fire acci-

---

* The term Drell-1 is used for current state of the art safety improvements including ENDS (enhanced nuclear detonation safety), fire-resistant pit (FRP), and insensitive high explosive (IHE). The term Drell-2 refers to safety enhancements beyond the current state of the art , e.g., systems which separate the HE from the fissile material.

dents are larger. The lofted fraction ranged from 0.5% to 20% with a median value of 2%. Corresponding respirable fractions ranged from 0.0025% to 1% with a median of 0.05%.

We used source-term variations with the probabilistic distributions of wind direction to produce the radioactive-material dispersal consequences and their associated uncertainty ranges. As an example, Fig. 1 shows the probabilistic population exposure for a fuel-fire dispersal accident.

Consequences of an accident involving nuclear weapons can vary from moderate, in the case of a fuel fire accident, to extreme for an accidental nuclear detonation. For example, onsite and offsite cleanup areas based upon a draft EPA screening level[6] of 0.2 $\mu Ci/m^2$ can range up to a few tens of square kilometers for fuel fire accident, a few hundred to over a thousand square kilometers for HE detonation, and one to several thousands of square kilometers for a nuclear detonation. Likewise, the estimated number of latent cancer fatalities varies from less than one for a fuel fire, as much as several tens for HE detonation, and up to several thousand for a nuclear detonation. Early deaths from a 1 kt nuclear detonation are in the range of one-to-several thousand. These numbers are highly scenario dependent; there are less serious and more serious scenarios.

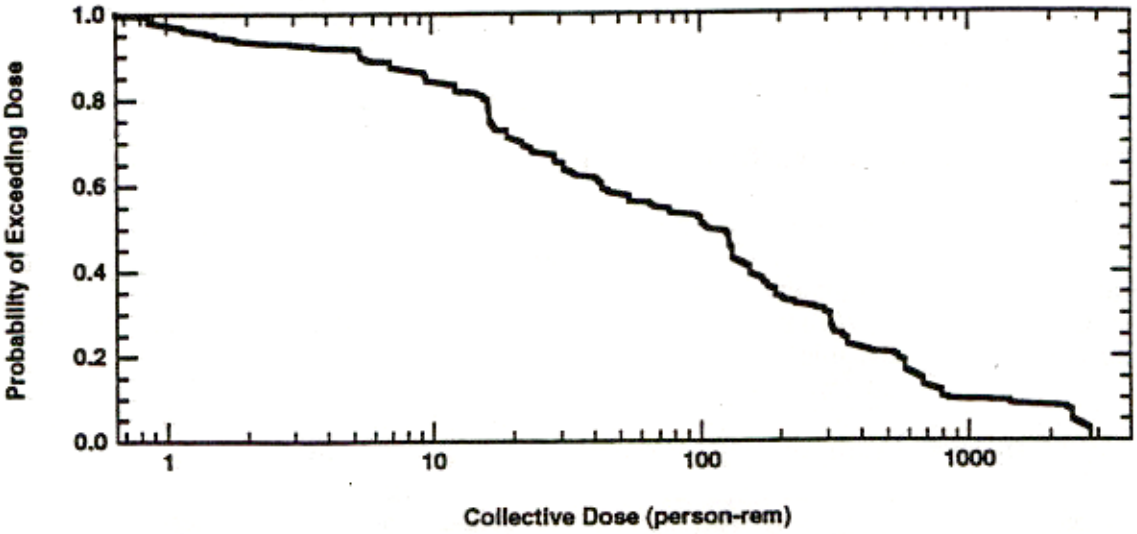


Figure 1. Probabilistic population exposure in person-rem for a fuel-fire accident. Explicit uncertainties included in the analysis were the wind directions and the respirable fraction of the Pu aerosol.

Mitigation cost

Our methodology for estimating avoided-accident costs was to establish a number of credible accident scenarios and determine their consequences in terms of health and affected areas. We used a population data base to estimate the probable number of injured individuals. We identified a large number of damage categories in which payments are being made for expenses, claims, and settlements. We used costs associated with actual, recent, natural and man-made disasters as a guide to estimating the relevant cost ranges of an accident. Some of the costs considered include major passenger-aircraft disasters, major hurricane and earthquake costs, Three Mile Island litigation, asbestos and Agent Orange litigation, Love Canal and Times Beach settlements, Fernald, Alyeska, the Exxon Valdez, and Rocky Mountain Arsenal costs, the World Trade Center costs, and the Clark AFB replacement estimates. Accident costs are based on the predicted consequences, current environmental costs, and litigation costs; offsite economic risks are the principal contributors.

The total impact of a single HE-detonation Pu dispersal incident is estimated to be in the neighborhood of $80 billion. Similar estimates for a fuel-fire Pu-dispersal incident indicate an impact of the order of $3 billion. The impact of an accidental nuclear detonation is approximately $200 billion. These values are specific to the assumed accident site, the amount of plutonium involved, details of the accident, the weather, and whether or not the affected population could be warned or sheltered. Different conditions would result in different impacts.

Accident Probability

The lifetime probability of an accident is the annual frequency of an accident multiplied by the expected weapon-system lifetime. We employed results from 12 nuclear weapons risk assessments to derive a range of Pu release probabilities. We based our estimates on a system of 1000 weapons with a life cycle of 30 years. Representative accident frequencies are shown in Table 3. Pu release probabilities for air transport, ground transport, storage, and deployment are bounded by the ranges and fall within approximately 70% confidence ranges of accident frequencies per weapon. Nuclear detonation estimates were obtained from expert opinion.

Table 3. Estimated annual probability of an accident per weapon used in this study.

Accident	Frequency (yr^{-1})
Fuel Fire	3×10^{-6} to 3×10^{-8}
CHE Detonation	10^{-6} to 10^{-8}
Nuclear Detonation	10^{-9} to 10^{-13}

Actuarial Costs Avoided (Benefits)

Both the implementation costs and avoided-accident costs are large for a possible nuclear weapons retrofit. We have also demonstrated that accidental HE detonation (for a nuclear weapon containing *conventional* HE and not the safer, insensitive HE) is the dominant contributor to actuarial risk when compared to accidents resulting from fire or nuclear detonation (See Table 4).

Table 4. Best estimates for frequency-weighted avoided costs ($M).

Avoided Costs	Fuel fire	CHE detonation	Nuclear detonation
On-site	3	5	0.0003
Off-site	30	700	0.004
Government operations	1	1	0
Rounded Totals	30	700	0.004

Lognormal distributions were assigned to the retrofit costs shown in Fig. 2. The much smaller uncertainties in the retrofit costs are evident. A cumulative distribution of net savings was determined as a difference of the actuarial avoided-cost and retrofit-cost distributions. It is shown in Fig. 3. The *mean* value of the savings is a few billion dollars and depends on the upper limit imposed on the costs of a dispersal incident. Interestingly, the *median* value of the savings is a net loss of $0.4B.

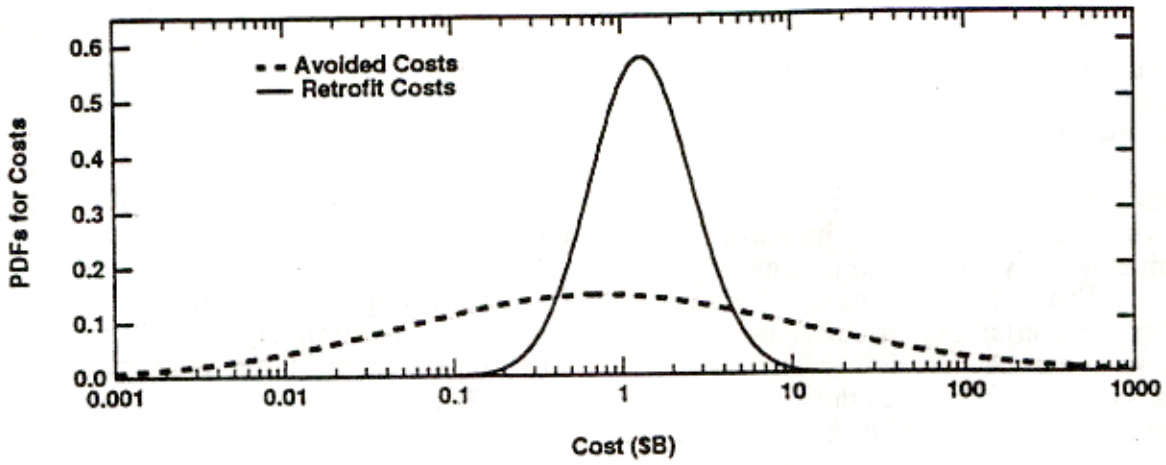


Figure 2. Distributions of potential retrofit costs and actuarial-avoided costs for replacing CHE with insensitive high explosive (IHE) in a 1000 weapon system.

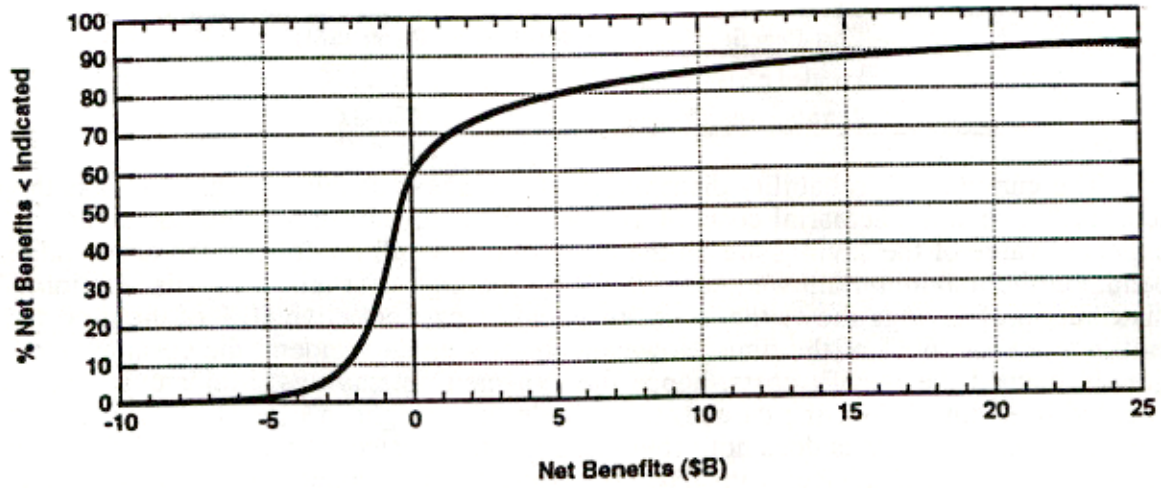


Figure 3. Distribution of future net savings of retrofitting 1000 nuclear weapons with IHE.

DIFFICULT-TO-QUANTIFY FACTORS

Considerations that are not easily quantifiable are significant and possibly dominant in decisions such as the one analyzed here. Political issues will probably play a major role. Nuclear weapon safety improvements require nuclear testing, which some view as an impediment to curtailment of proliferation. An important issue is the probability of technical success. Even after the weapon has been proven, uncertainties in the configuration of the DOE Complex-21 weapons productions facility lead to uncertainties in the ability to implement the retrofit. The possibility of reduced military effectiveness for a new, safer, weapon system and new security issues should also be considered.

Public reaction to an accident involving a nuclear weapon could lead to adverse affects on national defense, *e.g.*, a loss of deterrence or a temporary stand down of the weapon system while it is investigated and perhaps modified to minimize future accidents. It is hard to overstate the consequences that a serious accident could have for national security.

These unquantified values may overshadow the quantified impacts. Thus, our preliminary results suggest that considerations other than a simple cost-benefit tradeoff will influence and possibly control safety retrofit decisions.

CONCLUSIONS

We have estimated the benefits of a retrofit of CHE nuclear weapons with safer IHE. Retrofit costs and avoided-accident costs were estimated by extrapolating from historical data. These were combined with estimates of accident frequencies and with accident consequences that were estimated using state-of-the-art atmospheric dispersal codes. An HE detonation accident results in an impact on the order of $80 billion. When the expected cost is computed as an actuarial cost, the expected cost is reduced to less than $1 billion. This is within a factor of two of its retrofit costs. Other assumptions, scenarios, and locations could dramatically alter these findings.

The overall cost-benefit value is the difference in large and uncertain numbers. The large uncertainties in our analysis include a factor of approximately 10 in many accident consequences and another factor of at least 10 in accident probabilities. We estimate at least a factor of 15 when all the uncertainties in the actuarial costs avoided are considered. As an example Table 5 shows the 70% confidence ranges of costs and benefits.

Table 5. Results of the cost benefit analysis on whether to replace the CHE with IHE in a nuclear weapon system.

Cost-Benefit	Range ($M)
Avoided costs	50–10,000
Retrofit costs	700–3,000

The cumulative probability distribution of savings is shown in Fig. 3. This is the difference between the actuarial costs of avoided accidents and retrofit implementation. The expected value of the savings are on the order of a few billion dollars; the exact value depends on the maximum allowed cost (cutoff on the high-cost tail of the distribution). The distribution of savings shows that retrofits would be cost effective 40% of the time and result in a net loss 60% of the time. However, for any given accident, the costs avoided can greatly exceed the retrofit costs, hence the expected net savings. Major costs would be avoided in about 10% of the cases (off the scale to the right in Fig. 3). We suggest that a pure cost-benefit analysis does not provide a definitive answer to whether or not we should retrofit CHE nuclear devices with IHE. The decision must depend on other factors. These difficult-to-quantify factors that involve political considerations, the risk of loss of deterrence, and others need more thorough scrutiny.

REFERENCES

1. Appendix F-39, *in:* "Report of the Panel on Nuclear Weapon Safety of the House Armed Services Committee," S.D. Drell, chairman, Government Printing Office, Washington (1990).
2. T. Harvey and F.J. D. Serduke, "Fallout Model for System Studies," Lawrence Livermore National Laboratory, Livermore, UCRL-52858 (1979).
3. R. Lange, ADPIC—A Three-dimensional particle-in-cell model for the dispersal of atmospheric pollutants and its comparison to regional tracer studies, *J. Appl. Meteor.* (320) (1978).
4. B. Boughton and J. DeLaurentis, "Description and Validation of ERAD: An Atmospheric Dispersion Model for High Explosive Detonations," Sandia National Laboratory, Albuquerque, SAND92-2069 (1992).
5. G. Briggs, Plume rise and buoyancy effects *in:* "Atmospheric Science and Power Production," USDOE, DE84-005177, Washington, D.C., (327) (1984).
6. Persons exposed to transuranium elements in the environment, Federal Register **42**, (230) (60956) (1977); also, "Interim Recommendations on Doses to Persons Exposed to Transuranium Elements in the General Environment," U.S. EPA, draft, 1987.

058 Examples of PRA in Design Decisions

Chair: B. Buchbinder, NASA

Examples of the Use of PSA in the Design Process and to Support Modifications at Two Research Reactors

D.H. Johnson, D.C. Bley, J.C. Lin (PLG); C.T. Ramsey, M.A. Linn (ORNL)

Risk Management and Corrective Actions

R.A. Dykes, A.A. Dykes (PLG); J.D. Blodgett (Hughes)

A Risk and Decision Analysis for Choosing Wind Tunnel Turbine Blade Strategies

M.V. Frank, S.A. Epstein (Safety Factor Assoc.)

Investigation Into the Development of A Train-Level Probabilistic Risk Assessment

C.L. Smith, R.D. Fowler, L.M Wolfram (INEL)

EXAMPLES OF THE USE OF PSA IN THE DESIGN PROCESS AND TO SUPPORT MODIFICATIONS AT TWO RESEARCH REACTORS

David H. Johnson,[1] Dennis C. Bley[1], James C. Lin,[1]
Charles T. Ramsey,[2] and Mark A. Linn[2]

[1]PLG, Inc.
4590 MacArthur Boulevard, Suite 400
Newport Beach, CA 92660-2027

[2]Oak Ridge National Laboratory*
FEDC Building
P.O. Box 2009
Oak Ridge, TN 37831-8218

ABSTRACT

Many, if not most, of the world's commercial nuclear power plants have been the subject of plant-specific probabilistic safety assessments (PSA). A growing number of other nuclear facilities as well as other types of industrial installations have been the focus of plant-specific PSAs. Such studies have provided valuable information concerning the nature of the risk of the individual facility and have been used to identify opportunities to manage that risk. This paper explores the risk management activities associated with two research reactors in the United States as a demonstration of the versatility of the use of PSA to support risk-related decision making.

INTRODUCTION

The use of probabilistic safety assessment techniques to support design decisions at two special-purpose nuclear reactors is explored. The two research reactors considered are the High Flux Isotope Reactor (HFIR) and the Advanced Neutron Source (ANS). The HFIR is operating at the Oak Ridge National Laboratory (ORNL), and the ANS is planned to be constructed at Oak Ridge. The versatility of the use of PSA to support decision making is

---

*Portions of this manuscript were prepared by the Oak Ridge National Laboratory, Oak Ridge, Tennessee, 37831-6285, managed by Martin Marietta Energy Systems, Inc., for the U.S. Department of Energy under contract DE-AC05-84OR21400.

demonstrated by its application to these two facilities, which are in two very different stages of life.

HFIR: AN EXISTING RESEARCH REACTOR

The first example involves the use of PSA in finalizing the design of modifications that were proposed to address a specific class of postulated scenarios at the HFIR, an existing research reactor. The class of scenarios of concern involved overpressurization or overcooling of a radiation-embrittled vessel.

The HFIR is an 85-MW flux trap reactor that has operated for over 20 years. The peak thermal flux in the flux trap is 5×10^{15} n/cm^2-sec. The core (17.5 inches in diameter and 24 inches in height) is contained in an 8-foot-diameter pressure vessel that is about 19 feet in height. In 1986, the HFIR was shut down due to vessel embrittlement concerns. Uncertainty existed about whether embrittlement of the pressure vessel had progressed too far to support continued safe operation. These concerns were addressed through a probabilistic fracture mechanics analysis.[1] Prior to restarting the reactor, virtually every safety aspect of the operation of HFIR was reviewed. The HFIR Probabilistic Risk Assessment[2] was one of the safety analyses conducted prior to restart. The HFIR PRA was used to refine proposed design changes as well as to identify additional opportunities for safety enhancement.

To lessen the likelihood of overpressurization and overcooling of the pressure vessel, design changes were proposed to trip the main secondary pumps in response to a reactor scram and to install emergency depressurization capability. Tripping the secondary pumps would prevent overcooling, while the emergency depressurization valves would provide an additional line of defense by depressurizing the vessel on sensed low inlet temperature.

The original concept for tripping the secondary pumps included the automatic starting of a single, small standby secondary pump. This small secondary pump would provide a measure of cooling to the heat exchangers and, in addition, continue to provide cooling to the primary pump upper motor bearings (one of the functions of the secondary cooling system). If secondary cooling flow is lost, a single annunciator alarms in the control room to alert the operators. At the time, the existing procedures told the operators to identify the one pump experiencing high upper bearing temperature and to shut down that one pump. The procedures, written prior to the proposed design modification, did not explicitly consider the possibility that a single event such as the failure of the standby secondary pump to start could cause the upper bearings for all three running primary pumps to overheat. Information from the pump manufacturer indicated that bearing damage would occur if the pumps were not tripped within several minutes.

Early results of the PSA indicated that failure of the standby secondary pump followed by failure of the operators to trip the primary pumps, even with enhanced training and procedures, was a significant contributor to core damage frequency. In fact, such scenarios were about one-quarter of the total core damage frequency from internal events. Plant engineers and operators revisited the proposed design change following the interim PSA results and identified a simple change that virtually eliminated this family of scenarios. A direct trip of the primary pump motors on high upper motor bearing temperature was installed.

The proposed air-operated emergency depressurization valves were to "fail open" on loss of instrument air. This is certainly the "safe" mode if only vessel embrittlement is to be considered. The HFIR is, however, rather sensitive to pressure changes due, in part, to the fact that its primary system is "water solid." Rapid pressure changes could impact the fuel/coolant heat transfer process as the critical heat flux decreases with lower pressures.

The original design of the emergency depressurization system would have resulted in a loss of instrument air event to initiate effectively a small LOCA.

The PSA was used to consider the two alternatives (fail open or fail closed) for the final design of the emergency depressurization valves. The results clearly indicated that the "fail closed" alternative was the preferred choice.

The HFIR PRA also identified other opportunities for safety enhancement. The addition, for example, of a check valve in a pool cleanup line was identified as an effective way to prevent a loss of the ultimate heat sink during moderate-to-strong earthquakes. The very closely spaced nature of the fuel plates makes flow blockage events a concern. The PSA identified the potential blockage scenarios, which, in turn, were used to make physical and procedural changes in the plant to lower the likelihood of blockage.

ANS: A RESEARCH REACTOR IN THE ADVANCED CONCEPTUAL DESIGN PHASE

The Advanced Neutron Source (ANS) reactor, proposed for construction at ORNL, is being designed to be the most intense source of continuous thermal and sub-thermal neutrons currently in existence. The ANS, which may replace HFIR, will produce a peak thermal flux greater than 8×10^{15} n/cm^2-sec and operate at 330 MW.

Although the ANS resembles the HFIR in that it uses closely spaced, plate-type, highly enriched fuel elements, the ANS contains many design features that the HFIR does not. The ANS will contain two liquid deuterium cold sources and a hot source to provide long wavelength or short wavelength neutrons, respectively. The ANS uses heavy water for primary cooling and moderation and light water for secondary and component cooling. The ANS has a full containment, whereas the HFIR uses a confinement strategy. Most interesting, the ANS has a number of passive safety design features. These include natural circulation cooling capabilities, passive accumulators on the hot legs of each primary loop to lessen the impact of a loss of coolant or loss of pressure control, water-filled cells around much of the primary and reflector cooling piping to lessen the impact of breaks within these cells, and limited volume air cells around piping to limit the amount of water lost for specific breaks. Active safety enhancements include redundant power sources with fast transfer capability and redundant and diverse scram systems.

The results of a PSA were recently published[3] based on the conceptual design of the ANS. This PSA was restricted in scope and is considered a "limited-scope" or "Phase 1" PSA. As such, it takes a broad look at the scenarios that comprise the risk of fuel damage at ANS. A limited-scope PSA typically makes simplifying conservative assumptions and considers only a limited amount of "recovery" actions to refine the results. A limited-scope PSA can provide significant insights with a fraction of the effort of a full-scope PSA and, at the same time, can provide a focused basis for a future extension to a full-scope PSA.

The Phase I Level I PSA of the ANS conceptual design examines the contribution to plant risk and unavailability from internal event initiators. The primary goal of the PSA was to determine approximation of the core damage frequency and to show if the proposed core melt risk limitation conceptual design goal of 1×10^{-5} events resulting in core damage per year could be met with the existing conceptual plant configuration. A secondary goal stems from the ability to accomplish the first goal. This secondary goal was to identify areas in the design that contribute significantly to plant risk and/or unavailability.

The top five contributors to fuel damage in percent of the total estimated mean core damage frequency are as follows:

Flow Blockage	83%
Undetected Defective Fuel	26%
Large LOCA Near Core	0.5%

ATWS	$<< 0.5\%$
Station Blackout	$<< 0.5\%$

The only significant contributors to the plant risk are flow blockage and a large loss of coolant accident very near the core. The lack of impact from other potential contributors is attributable to the many redundancies and safety features in the plant design. For most transients, the only requirements for successful mitigation are (1) a reactor scram by one of the two redundant and diverse scram systems, and (2) at least one primary flow path available for natural circulation cooling. The results of the sequence quantification indicate that the core damage frequency goal of less than 1×10^{-5} per year may not be met with the conceptual design, due to reasons not directly related to the systems design. Therefore, efforts are focusing on ways to decrease the estimated likelihood of flow blockage scenarios and evaluating the cost-risk benefit of some of the redundant systems. The current flow blockage and undetected defective fuel models are based largely on the HFIR PSA. A plant-specific review of the ANS flow blockage frequency is now underway. Various availability studies, some of which are discussed below, are also being performed to address RAM issues and to optimize design and cost.

The purpose of an ANS Three versus Four Loop Study was to evaluate the unavailability of power operation due to two specific alternative plant designs: (1) the four 33.3% capacity loop conceptual design where one loop is maintained in standby during plant operation, and (2) a four 25% capacity loop design where all loops are operating at the same time. Fault tree methodology similar to that used in the ANS PSA was used to model each configuration. Time-dependent availabilities were calculated using BRAVO 2.0[4] from JBF Associates, Inc., and failure data from the ANS PRA. Repair times were taken from Reference 5 but were adjusted upward to allow for xenon poisoning as specified below.

In this ORNL study, fault tree models for each design option questioned the response to loss of a single loop and considered the likelihood of multiple-loop failures. Due to modeling complexities, the four 33% loop fault tree modeled the availability of only 100% power operation, whereas the four 25% loop tree modeled availability of power operation at either 100% or at reduced power (approximately 80%). Because of the short time before xenon poisoning prevents startup and the length of time required to repair a scram-producing failure, the repair time, unavailability determination was postulated to have a minimum value of 3 days to allow for xenon decay prior to restart. The results of the model quantification indicated that, from the standpoint of unavailability (0.015 versus 0.025), unreliability (0.78 versus 0.92), and expected number of failures (1.5 versus 2.5), the four 33.3% loop design is preferable to the four 25% loop design.

A second availability study[6] was performed by Reliability and Performance Associates (RAPA) based on the analysis framework provided by the ANS PSA. This study examined the possible plant operating scenarios for each cooling configuration based on loop availability using a success state matrix. The RAPA analysts examined HFIR operating data for the main cooling pumps from May 1990 until June 1993 to obtain loop availabilities for the HFIR primary system. Using this data in combination with simple reliability equations for the loop configuration, the outcome was that, as loop availability decreases, the time spent at power levels below full power increased greatly for the four 25% loop design over the four 33% loop design, primarily due to the opportunity to maintain full power using the spare loop with the later design.

Both studies indicate that high availability can be obtained with either system. However, as unavailability increases, the time spent at less than full power or shutdown greatly increases with the four 25% loop design, while there is only a modest increase with the four 33% loop design. Given that the four 33% loop results reflect 100% power operation, the basis exists for maintaining this design.

An additional study is being conducted by ORNL and RAPA to investigate the contribution of the cold neutron source to plant unavailability. Preliminary results suggest that, with the current design, unanticipated scrams based on monitored cold source parameters may preclude the plant from reaching its conceptual design availability goal of 80%.

The ANS PSA and its ancillary unavailability analyses show great potential to identify sources of risk and unavailability, as well as to identify and quantify the effects of candidate plant design changes.

OBSERVATIONS

The results of the application of PSA techniques to two different facilities have been summarized. The first facility was a nuclear reactor with an extensive operating history that was undergoing extensive safety enhancements and review. The second facility was also a nuclear reactor that incorporates several new concepts and is in the advanced conceptual design stage. In both cases, PSA provided a meaningful framework to identify and measure the safety impact of potential changes.

REFERENCES

1. R.D. Cheverton, J.G. Merkle, and R.K. Nanstad, "Evaluation of HFIR Pressure-Vessel Integrity Considering Radiation Embrittlement," Oak Ridge National Laboratory, ORNL/TM-10444 (1987).
2. D.H. Johnson, D.R. Buttemer, S. Kaplan, et al., "The High Flux Isotope Reactor Probabilistic Risk Assessment - Analysis of the Risk From Internal and External Events," prepared for Martin Marietta Energy Systems, Inc., Oak Ridge National Laboratory, PLG-0829 (1991).
3. Phase I Level I Probabilistic Risk Analysis of the Advanced Neutron Source Conceptual Design, Draft (1993).
4. V.H. Guthrie, "BRAVO 2.0 — A Windows 3.1 Program for Quantitative Risk Assessment," JBF Associates, Inc. (1992).
5. V.M. Bier, "Database for Probabilistic Risk Assessment of Light Water Nuclear Power Plants — Maintenance Data," PLG-0500, Vol. 3, Rev. 1 (1989).
6. Reliability and Performance Associates (RAPA), "Primary Coolant Loop Availability Analysis of the High Flux Isotope Reactor in Support of Advanced Neutron Source Reactor," prepared for Martin Marietta Energy Systems, Inc., Oak Ridge National Laboratory (1993).

RISK MANAGEMENT AND CORRECTIVE ACTIONS

Robert A. Dykes,[1] Andrew A. Dykes,[1] and Jimmie D. Blodgett[2]

[1]PLG, Inc.
4590 MacArthur Boulevard, Suite 400
Newport Beach, CA 92660-2027

[2]Hughes Missile Systems Company
P.O. Box 85715, MZ 84-6130
San Diego, CA 92186

INTRODUCTION

During a typical full-scale development phase of a missile system, the approach to corrective actions is to usually tackle the most obvious problem and resolve it quickly to keep the development program on track. Engineers are generally fully confident that this corrective action will solve the problem and that there will be no additional problems. However, this is rarely the case. In addition, there are usually several possible corrective actions that can be implemented across major subsystems. The problem then is threefold:

1. How can corrective actions be assessed regarding their impact on end-product reliability?
2. Which corrective actions can provide the greatest improvement in system reliability at the lowest cost?
3. How can we evaluate the corrective action before actually implementing it?

When a missile has significant initial reliability problems, the full-scale development and flight tests experience a continuous stream of failures, resulting in a number of possible corrective actions. Attempts to assess the impact of these corrective actions on the performance of the missile based on ground test results rather than on actual performance have failed in the past. This is due to the limited ability to fly the missile (if it is an expensive one-shot device) and the inability to integrate diverse ground test data into a meaningful quantitative perspective. To resolve this problem, a methodology using Bayes' theorem can be used to incorporate the ground testing program into the missile reliability assessment. The methodology is based on a reliability database that is used to collect and assess the impact of all of the available success and failure test data on the potential for mission success. The observed failures are then accounted for in light of the effectiveness of each corrective action for preventing their reoccurrence.

The use of Bayes' theorem allows comparisons of the existing missile configuration to proposed configurations resulting from corrective actions using all of the evidence available. This comparison can be accomplished at both the subsystem and system levels. When compared against the cost of the corrective action, it provides management with a valuable tool in selecting cost-effective corrective actions that can have a great positive impact on system reliability. This paper describes the methodology of how this can be accomplished.

ASSESSING FAILURES

Generally, only those failures that occur during testing in which the equipment is expected to function properly are counted for mission reliability assessments. (Failures during initial manufacturing checks may be addressed as part of the corrective action process, and may be indicative of overall process control but would not typically be addressed directly.). These failures were then categorized by impact on mission.

1. Catastrophic — Guaranteed loss of mission.
2. Critical — Potential loss of mission.
3. Launch Abort — Occurs only in captive carry.
4. Minor — Affects only logistical reliability.
5. Other — Not counted for reliability, but helps to identify corrective actions.

All Criticality 1 and 2 failures are subjected to a numerical assessment of the conditional probability for loss of mission (LOM) as a result of the observed effects of the failure. This LOM likelihood is then used to convert the severity of the observed failure effects into an equivalent LOM estimate. When all of the observed catastrophic and critical failures are summed, one obtains the observed equivalent mission losses for those failures. When considered with the total number of tests in which the failures accumulate and the configuration of the units undergoing tests, this number can then be used in the assessment of the missile mission reliability. For example, a catastrophic failure would be assigned a conditional LOM of 1.0. Conversely, a slightly out-of-tolerance reading on a bias of the gyroscopic might be assigned an LOM of 0.1 since the terrain matching capability of the guidance system would likely compensate for the slight drift. The advantage of this approach is that it is no longer necessary to take an "all or nothing" approach when including failures into the reliability assessments. Of course, the rationale for each criticality assignment and LOM assessment must be documented.

CORRECTIVE ACTION EFFECTIVENESS

This evaluation is the expression of the confidence of the corrective action team that the observed failure mode will not occur again. This confidence is based on evidence regarding the

- Root cause of the problem.
- Measures (design changes or process changes) taken to prevent its reoccurrence.
- Experience verifying that those measures were indeed successful.

An evaluation of 1.0 means that it is absolutely positive that there is no possibility of the observed failure of ever occurring again, and that the fix has not introduced any failure modes that did not exist prior to the fix. This belief must be based on a thorough understanding of the root cause, a direct correlation of the corrective action to the elimination of that cause, and a comprehensive verification that the root cause has been eliminated with no additions of other failure modes. Evaluations between 0 and 1 reflect a lesser degree of confidence based on the product of the assessments of the three factors

involved. These assessments should recognize that measures that require continued human vigilance, such as improved training and procedures, should produce a lower degree of confidence than design changes that physically remove the root cause of the failure.

DATABASE STRUCTURE

To evaluate the potential reliability growth that can be achieved by a corrective action, one must be able to relate configurational data with test data, failure information, and the resultant corrective actions. Figure 1 illustrates how typical relational links among the various databases are indicated by the arrows. The objective is to produce a closed loop system, so that the test program results can be correlated to the evolution of the design that results from the corrective actions.

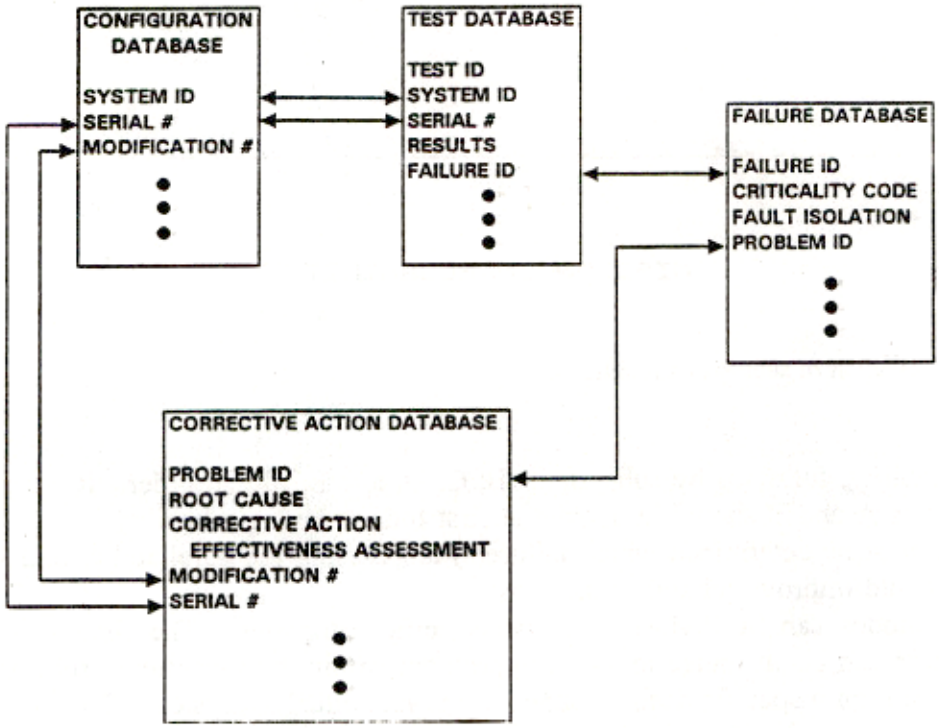


Figure 1. Relational database for managing reliability improvement.

ASSESSING POTENTIAL FOR RELIABILITY GROWTH

Once the database is developed, the probability that the failure will impact the mission is estimated.[1,2] The expected corrective action effectiveness is also estimated, and comparisons can be made between multiple failure modes. The question is, How does one select the subsystem that is having the greatest current impact on reliability?

The missile is expected to achieve a desired flight reliability specification. This is divided into reliability allocations for its various subsystems. Using the information in the database described above, and comparing the actual reliability performance of the subsystems and their contribution to the overall reliability, the subsystems that have the greatest impact are easily identified.

Figure 2 shows this relationship. The system curve expresses our confidence regarding missile flight reliability based on current evidence. Curves A to E are the associated reliability curves for the various major subsystems. Subsystem A contributes the greatest

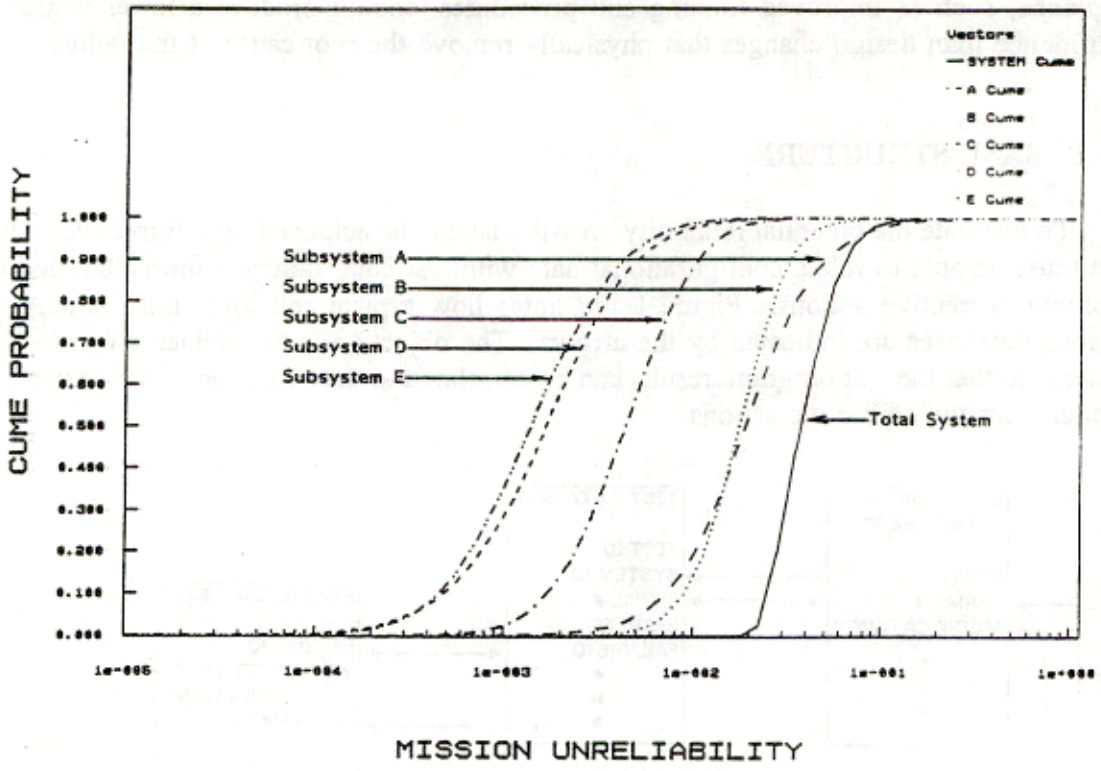


Figure 2. Contribution of subsystems to reliability.

to the unreliability, followed by subsystems B, C, D, and E, in that order. It is apparent that subsystems A and B should be looked at first for improving reliability.

Now it must be determined which failure(s) are causing the unreliability and which, if resolved, could improve reliability the most.

Failure modes can be divided into two specific categories. The first is inherent failures. These can be attributed to a design problem within the subsystem itself, and can be resolved through a specific design change to the individual subsystem. The other type of failures is associated with quality/workmanship and process control.

Figure 3 shows the relationship of the two types of failures. The difference between the two curves might be considered the "cost of quality." When it comes to improving the reliability, quality/workmanship failures may be the easiest and least costly to fix. Many times, the subsystem can be significantly improved enough to cause a major increase in the reliability of the missile.

If all of the quality/workmanship problems are resolved and the reliability of the subsystem is still causing a negative impact on the overall missile flight reliability, the inherent problems must be fixed. Redesign requires the engineers to completely analyze the subsystem to determine why it is not functioning as expected. There may be several possible alternative solutions suggested that require further investigation. Once implemented, a requalification test program is required, which includes design margin requirements, single point failure analysis, and stress testing. This process can be extremely costly.

Once the high priority failure modes have been identified, corrective actions can be analyzed and individually evaluated using the techniques discussed above. Several corrective actions can be rapidly evaluated using this methodology and the most effective ones retained for additional cost-benefit analysis. Although this analysis is normally conducted for a single subsystem at a time, in reality, corrective actions for several subsystems can be examined at the same time. When working with limited dollars or late

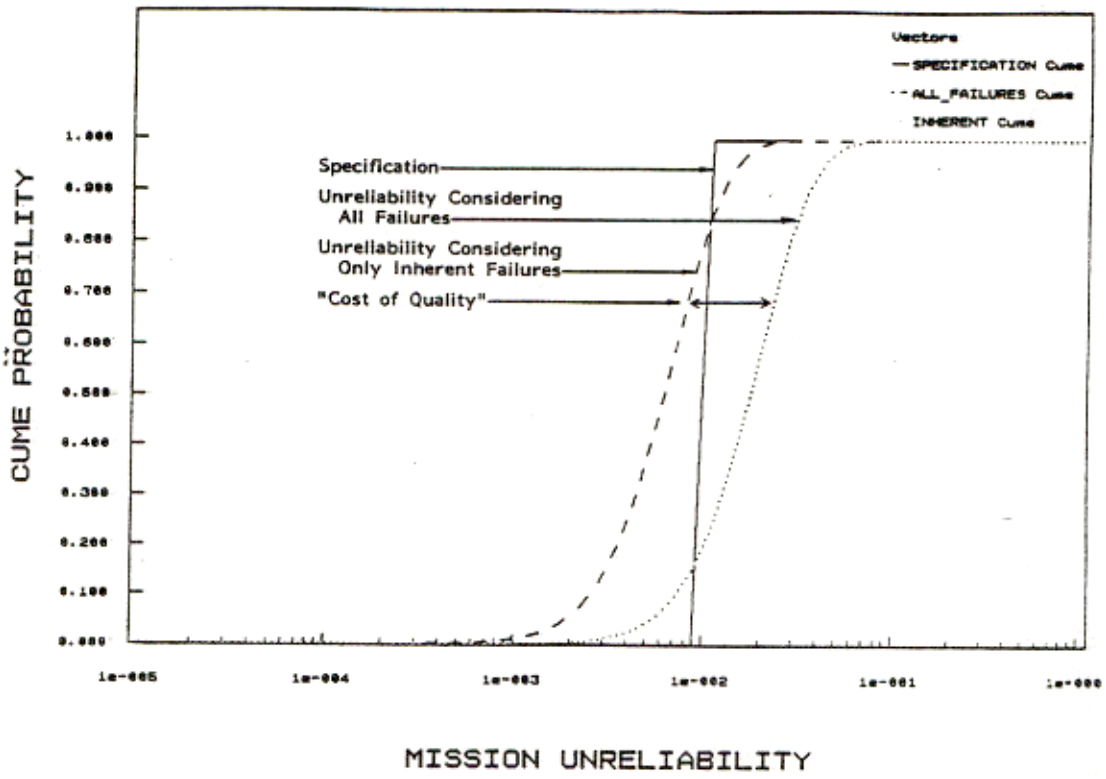


Figure 3. Relationship of failures.

in the development program, this is a useful tool to maximize the reliability gain for the dollar invested.

When all of the evaluations have been completed, a rank order of the effectiveness of the corrective actions can be listed. By building a table with the corrective action and its associated costs, management can rapidly see how much reliability can be gained for the dollar spent. An example of this can be seen in Table 1.

Table 1. Cost comparison of possible corrective actions.

Corrective Action	Reliability Improvement (percent)	Cost (dollars)	Marginal Cost per Percentage (dollars)
A	5	200,000	40,000
B	15	750,000	50,000
C	12	360,000	30,000
D	2	100,000	50,000

Since this was a Bayesian analysis, new test results can be used to verify that the subsystem is achieving the estimated reliability. Also, test results of new subsystems as they complete initial testing can be compared to what was expected to verify that the expected gains are obtained before full design implementation. This is especially valuable since some corrective actions can result in a negative reliability growth due to unforeseen interactions between other subsystems.

Figure 4 shows the effect of a major subsystem design change. Distribution 1 is the experienced failure rate prior to installing it in the missile. This was unacceptable to management, so a major design change was directed. The expected improvement is reflected in distribution 2. Notice that the distribution is to the left of the current reliability, showing that an improvement is expected; however, it is broader, showing that there is increased uncertainty that the desired result will be achieved. Distribution 3 is distribution 2 updated with additional test data. In this case, the subsystem has shown a definite improvement, bearing out the corrective action decisions by management.

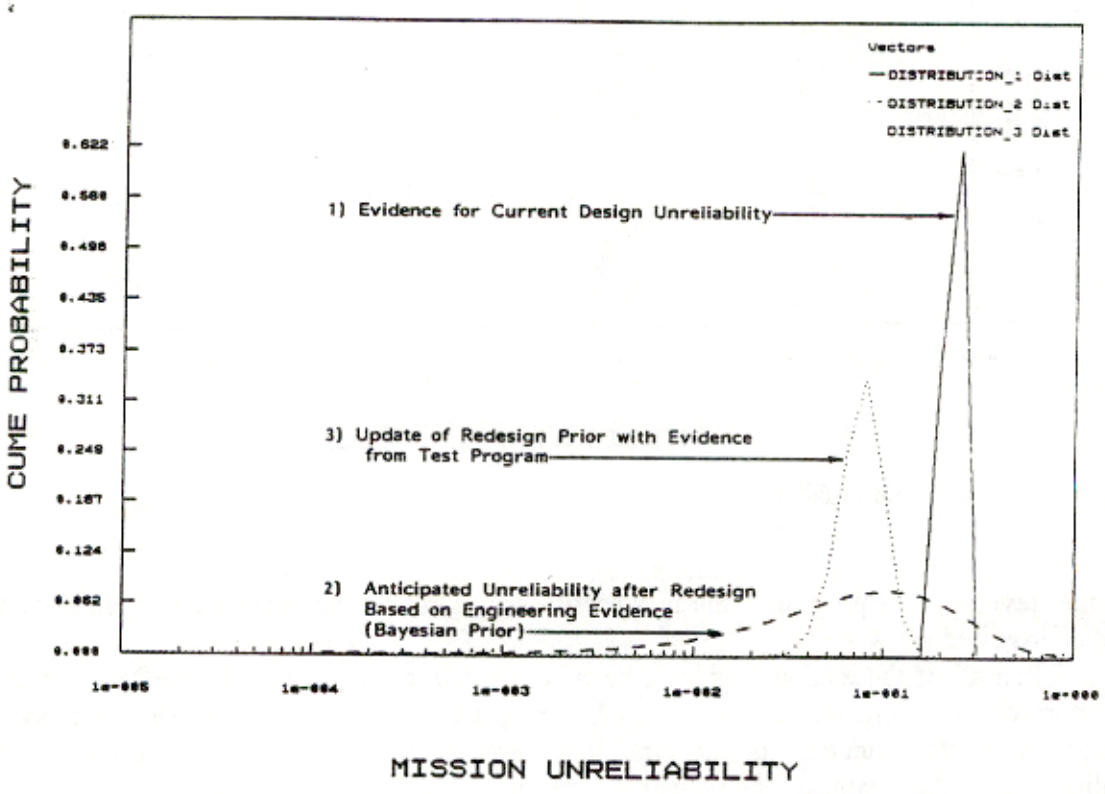


Figure 4. Effect of a major design change.

SUMMARY

The use of Bayes' theorem in missile reliability analysis can be very beneficial. The dynamics of the Bayesian analysis with its confidence curves provides engineers and management with a quantitative value to (1) determine the impact of quality/workmanship failures against design or inherent failures, (2) allow the analyst to isolate the impact of individual or multiple failures, and (3) provide a tool to do cost-benefit analysis for multiple corrective actions to determine where the greatest reliability could be gained for the dollar invested.

REFERENCES

1. S. Kaplan, et al., "Bayesian Methodology for Assessing Reliability during Product Development," *Annual Reliability and Maintainability Symposium*, Los Angeles, California (1990).
2. R.A. Dykes, et al., "An Application of Bayes' Theorem to Management Decision Making," to be presented at *Annual Reliability and Maintainability Symposium*, Anaheim, California (1994).

A RISK AND DECISION ANALYSIS FOR CHOOSING WIND TUNNEL TURBINE BLADE STRATEGIES

Michael V. Frank, P.E., Ph.D. and Steven A. Epstein

Safety Factor Associates Inc.
1410 Vanessa Circle
Encinitas, CA 92024

INTRODUCTION

The purpose of the turbine blade trade study was to aid NASA Ames decision makers in deciding among many alternative compressor and blade combinations for the 3-stage compressor of the Unitary Plan Wind Tunnel (UPWT) at Ames Research Center. This paper can only summarize what was a complex multi-criteria risk and decision analysis.[1]

The compressor was designed and built in the 1950's with aluminum blades. The compressor and blade design was such that the ramp up to operational speed was through several severe resonance frequencies. Consideration of composite blades along with modifications to the compressor arose because of their resistance to damage (e.g. pitting), freedom from the annoying resonances, and less potential for damage to other equipment. These potential advantages, however, were weighed against the large redesign and capital costs of the modification. The essential question in this decision was, will the advantages outweigh the cost? We were asked to help the UPWT staff at Ames develop an analysis as input for the decision because of our experience in analyzing risks in the face of uncertainties.

The objectives in structuring the analysis were 1) to discover and quantify the significant differences among alternatives and 2) rank the alternatives. We did this by carrying out a three step evaluation:

1. Development of a decision hierarchy that identified the criteria against which the options were evaluated and identified the alternative compressor and blade options. Projected life cycle and safety were the most relevant criteria for the decision-maker at Ames.

2. Performance of a risk analysis of each option with respect to each criterion. In this context risk is defined as the probability of achieving specified cost and safety where both are uncertain quantities. Safety was defined as the frequency with uncertainties of catastrophic damage of the wind tunnel. An analysis was performed for each option.

---

[1] M.V. Frank and S. A. Epstein, "Turbine Blade Trade Study: A Risk and Decision Analysis", Safety Factor Associates, Inc. Report for Ames Research Center, February, 1993

3. Using the results of the risk analyses, rank the options with respect to each criterion (i.e. rank with respect to cost and separately rank with respect to safety) and combine criteria such that an overall ranking that includes all criteria is obtained (i.e. rank with respect to cost and safety together).

In carrying out this evaluation the fundamental viewpoint was probabilistic. For example, if we could actually run each blade and compressor option for a period of many years, we could know fairly precisely what each option would cost. We do not have this luxury nor are we prescient beings. We, therefore, relied on "gedanken" experiments to predict costs. The results of these experiments, however, were inherently uncertain because of less than perfect understanding of costs, the models we used to predict costs, and the input data for these models about future costs.

OPTIONS, CRITERIA, AND A DECISION HIERARCHY

The goal was to choose an option that produces the safest and most cost effective compressor operation. The options were defined by the UPWT staff and we assisted in grouping and characterizing them as follows:

- Option 1 represented continued use of the current aluminum blades, three stage compressor, testing practices, maintenance practices, and refurbishment practices. Operational startup and shutdown resonance stress reduction practices had been developed over the years since 1958 when the compressor first became operational. In addition, the rate at which the hub turned (i.e. rpm) was varied during steady state run to avoid resonance regions. Aluminum blades become pitted because of hits from high velocity fragments and debris in the tunnel. A substantial amount of operating costs are incurred by changing out pitted blades and grinding them to remove the pits.

- Option 2 included replacing current blades with new high aspect ratio aluminum blades when the current blades wear out. It also included addition of motorized stators and revised internal guide vanes (IGVs). It was assumed that startup and shutdown resonance stress reduction as above would be practiced but the motorized stators and IGVs would minimize the need for rpm changes during operation.

- Option 3 included replacing current blades with new high aspect ratio composite blades. This option also included motorized stators and revised IGVs. It was assumed that no resonance avoidance activities would be required with this option and that blade grinding would not be performed.

- Option 4 included replacing current blades with new high aspect ratio composite blades and also purchase of motorized stators, revised IGVs, new hubs, and other modifications to enhance ease of operation.

Ames management wanted to address these options with respect to various factors. The factors included blade life, blade/hardware replacement cost, life cycle cost, lost productivity time, energy usage, maximum Mach number capability, management attention, catastrophic failure probability, catastrophic failure damage, and recovery period after

catastrophic failure. We proposed two decision criteria that were independent, were easy to communicate, and encompassed the concerns of Ames managers.

- Life Cycle Cost (excluding time lost from catastrophic failures)
- Catastrophic Blade Damage Frequency

Separation of catastrophic blade damage from life cycle cost recognized that risk averse decision-makers often consider this attribute more important than the expected value of lost revenue from catastrophic failures would indicate. Therefore, life cycle cost with uncertainties and catastrophic blade damage frequency with uncertainties were calculated for each option

Life cycle cost was defined as the levelized dollars per year of recurring and nonrecurring costs over the lifetime of the facility. Because life cycle cost cannot be precisely known in advance for postulated options, they were characterized using probability distributions expressing uncertainty in the levelized dollars per year. The total calculated life cycle cost included cost of energy usage, maintenance costs, inspection and blade change out, operations costs, lost productivity owing to resonance avoidance, design-development-testing-evaluation costs, material costs, and cost of money.

Catastrophic blade damage frequency was defined as the number of catastrophic failures per year. General causes of catastrophic failure evaluated in this study included human error, spontaneous equipment (hardware and software) failure, and design errors. This analysis used the past catastrophic failure history of other tunnels combined with that of the Unitary by applying Bayes' Theorem for forecasting the probable range of future catastrophic failure frequencies caused by the compressor and blades. We have, admittedly, incomplete historical incidents of wind tunnels and we did not attempt to quantify the error owing to the incomplete history.

The decision problem may be illustrated by a hierarchy[2] in Exhibit 1. At the top level of the diagram we depict the goal of the decision analysis: safe and cost effective operation of the 3-stage compressor. The bottom level of the diagram depicts the options (sometimes called strategies) for attaining that goal; and the middle level depicts the criteria whose distribution of values were calculated for each option.

RISK ANALYSIS FOR LIFE CYCLE COST

A risk analysis (i.e. an analysis that accounted for uncertainties and variabilities) of the life cycle cost of each option was performed as follows.

We developed a cost model consisting of simple equations that included all cost factors mentioned above. Uncertainty and variability distributions were developed for each of the parameters and variables in the equations. In general the uncertainty in material costs were state-of-knowledge uncertainties derived from experience with previous purchases at Ames. Personnel costs were from the rate schedule of support contractors at Ames with uncertainty associated with the number of hours applied to each task. Maintenance cost uncertainties for Options 3 and 4 were taken from the variability of such costs in other facilities. Energy cost uncertainties were based on a statistical analysis of 20 months of energy usage and megawatt hour prices for the Unitary wind tunnel.

We calculated the life cycle costs by solving the equations with input probability density functions using Monte Carlo simulation. The equations were programmed in EXCEL 4.0 running under Windows 3.1 and the Monte Carlo calculations were performed using CRYSTAL BALL, an add-in to EXCEL.

---

[2] T. L. Saaty, The Analytic Hierarchy Process, RWS Publications, Pittsburgh, 1990

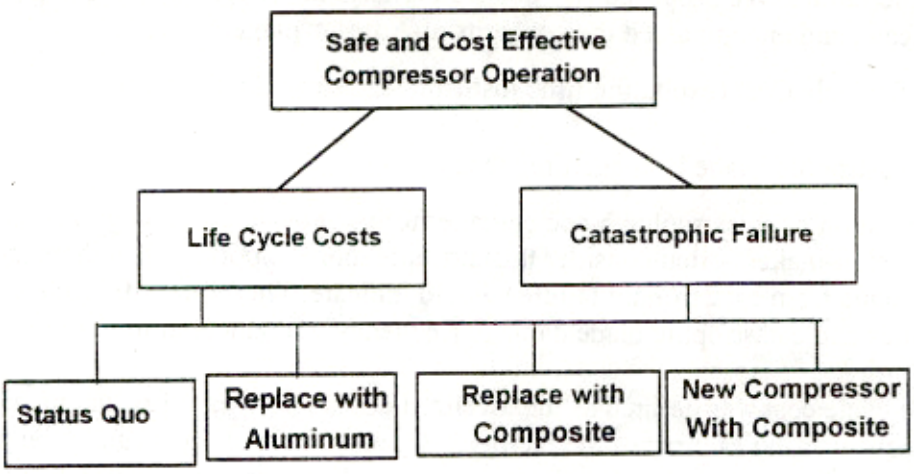


Exhibit 1. Decision Hierarchy

Sensitivity calculations were also performed to account for uncertainties in the model itself. Perhaps the largest modeling uncertainty was the year that the current blades are to be replaced. Blades are replaced when they wear out and time to wear out is a sensitive function of local blade stress levels which are not well known. Therefore, calculations were performed assuming replacement of all current blades in 1995 and replacement of all current blades in 2005.

Resultant probability density functions over life cycle cost are summarized by their mean values as follows (in millions of dollars per year): Option 1 - 4.1; Option 2 - 4.0; Option 3 - 3.4; Option 4 - 4.0. Not only was Option 3 the least expensive, but the spread of its cost uncertainty was also the least. The spread of the resultant probability density functions were dominated by the variability in energy consumption and cost per megawatt hour, which is a function of energy consumption.

Energy costs are the most significant variable. Options 3 and 4, which save the most energy, have the lowest life cycle costs even considering uncertainties. At 95% confidence, energy cost savings are large enough to compensate for (i.e.. pay back) lump sum capital outlays of Option 3 (over Option 1) in less than 10 years and large enough to compensate for lump sum capital outlays of Option 4 (over Option 1) in less than 15 years. The ranking of the options is insensitive to large variations in purchase costs, labor hours, and labor rates because of the dominance of energy costs.

RISK ANALYSIS FOR CATASTROPHIC FAILURE FREQUENCY

We also developed estimates of the compressor blade induced catastrophic failure frequencies via a Bayesian approach using compilations of catastrophic failure incidents from wind tunnels around the world. These were compiled by personnel at Ames Research Center for two general causal categories. The first is catastrophic blade failure owing to damage caused by equipment failure or design flaws. These manifest themselves as seemingly "spontaneous" failures. The second is catastrophic blade failure owing to damage caused by foreign objects left in the tunnel by human error.

Design flaws and equipment failures appear to occur either early in operational life or after many years of operation. Early incidents tend to be caused by (a) operational characteristics that were not known, (b) design mistakes, or (c) inadequate quality control/configuration management. Incidents that occur later in life tend to be caused by components that have worn out. Human errors do not exhibit a time bias.

The following table summarizes the data that resulted from this compilation.

Number of Wind Tunnels Surveyed	50
Number of Design Flaw/Equipment Failures	17
Number of Foreign Object Caused Incidents	10
Number of runs per year per tunnel	25 to 100
Operating years per tunnel	17 to 38

Unfortunately the compiled records were not explicit enough to suggest differences between metal and composite blades, although both of these tended to have more frequent severe accidents than wood blade tunnels. This information coupled with generic human error rates formed the information basis for developing two prior distributions. Unitary specific data for updating the prior distributions is: 1 design flaw failure in 36 years and 0 foreign object failures in 36 years.

The resultant updated distributions were summed (again using Monte Carlo simulation) to produce a total probability density function for catastrophic failure that is characterized by a distribution that is approximately lognormal in shape with a mean value of 1 failure in 36 years and with a 90% uncertainty interval of about one order of magnitude. There was not sufficient discrimination in the data to be able to distinguish among the options.

RANKING THE OPTIONS USING BOTH CRITERIA

The essence of the Analytic Hierarchy Process is to (1) determine the importance of each option (at the third level) with respect to each criterion (at the second level); (2) determine the importance of each criterion (second level) with respect to the goal (first level); and (3) combine above 2 steps to obtain the rank of each option (third level) with respect to the goal (first level).

The first two steps are accomplished by constructing matrices whose elements are pairwise comparisons of each option and each criterion, respectively. The third step is accomplished by obtaining the eigenvectors of the matrices and performing a dot product of new matrices constructed by these eigenvectors.

Constructing the matrix that assessed the importance of each option with respect to life cycle cost was relatively straightforward because we had the results of the risk analysis. However, the ranking of each option with respect to the criterion of catastrophic blade damage was less straightforward because the risk analysis did not distinguish among the options. We relied, therefore, on intuition within the framework of the Analytic Hierarchy Process to construct the matrix associated with the importance of each option with respect to catastrophic blade damage frequency.

We first constructed a four by four matrix whose elements are a pairwise weighting of the importance of each option with respect to life cycle cost. We then constructed a four by four matrix whose elements are a pairwise weighting of the importance of each option with respect to catastrophic failure frequency. The ranking of each option with respect to each criterion is recovered by obtaining the eigenvector or the maximum eigenvalue of each matrix. The eigenvectors are then normalized and concatenated into a two by four matrix.

It is now left to define the matrix that ranks the importance of life cycle cost (LCC) against catastrophic blade damage frequency (CDF) with respect to the goal of safe and cost effective blades. Rather than presume to understand the thoughts of the decision maker in this matter or limit his perspective to a single relative importance, the author's believe

that a more illuminating procedure is to calculate what is coined, herein, as a *Decision Trajectory*. We developed a decision trajectory by varying the quantity **i** in the following matrix over the entire range of possible rankings (9 to 1/9) in the analytic hierarchy process.

Goal	**LCC**	**CDF**
LCC	1	**i**
CDF	**1/i**	1

The decision trajectory is locus of points obtained by varying the eigenvector of the dot product of this matrix with the above two by four matrix over all values of **i**. As illustrated in Exhibit 2, this provides a graphical depiction of the ranking of each strategy with respect to the overall goal across the spectrum of relative importances of cost and safety.

Option 3 is always superior to the others. As safety becomes more and more important relative to cost, the rankings of Option 1 converges with Option 2 because they are weighted identically with respect to safety. A similar phenomenon occurs with Options 3 and 4. When cost is the most important factor, the ranking of options coincides with the ranking with respect to life cycle cost only. As safety approaches being of equal importance with cost, the Option 4 becomes preferable to Option 2. This is because consideration of safety (reduced catastrophic blade damage frequency) is influential enough to overcome the slight cost disadvantage Option 4 has with respect to Option 2 .

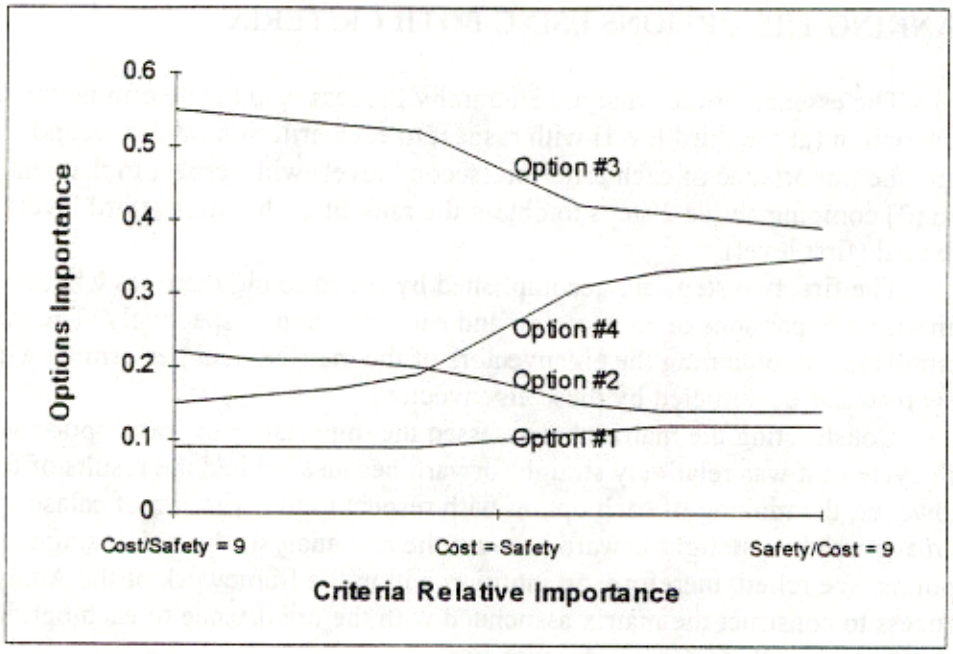


Exhibit 2. Decision Trajectories

Our recommendation to Ames management was to pursue Option 3. This is the course of action they took.

ACKNOWLEDGMENTS

The authors wish to thank Laura Doty and Ken Mort of Ames Research Center for their moral support during this analysis and Ken Mort for the material help in development of the uncertainty distributions. Thanks also goes to Robin Orans of Ames for her considerable efforts at obtaining the needed data for this study.

INVESTIGATION INTO THE DEVELOPMENT OF A TRAIN-LEVEL PROBABILISTIC RISK ASSESSMENT

Curtis L. Smith, Richard D. Fowler, Leon M. Wolfram

Idaho National Engineering Laboratory
EG&G Idaho, Inc.
Idaho Falls, Idaho 83415

INTRODUCTION

EG&G Idaho, Inc., located at the Idaho National Engineering Laboratory (INEL), investigated the development, use, and limitations of train-level probabilistic risk assessment (PRA) databases. These train-level databases are intended to allow the Nuclear Regulatory Commission to investigate the impact on plant core damage frequency (CDF) when a train or important train-related component is taken out of service or is failed. These train-level databases were loaded into the Safety Analysis Programs for Hands-on Integrated Reliability Evaluations (SAPHIRE) version 4.0 package,[1] a suite of PRA analysis tools that is also being developed at the INEL. A detailed discussion of the results for the train-level PRA investigation can be found in the report *Development and Use of a Train-Level Probabilistic Risk Assessment*.[2]

The intent of the train-level PRA development program was to construct a PRA database that required minimal personal involvement to utilize and would accurately measure the effect on plant CDF if a train is taken out of service or is failed. The basic premise was that this development could be achieved by truncating system fault tree models at the train level. That is, the models would not be expanded to include basic events below the top gate of a train, with the possible exception of including a few of the important train-related components (e.g., pumps and motor-operated valves).

In general, a PRA is used to define and evaluate the types of accident sequences that can occur, their frequency of occurrence, and their consequence. Normally, system level failures are modeled through fault trees using a top-down approach (i.e., system function fails is the top event of the fault tree). Unfortunately, not all trains are modeled in the PRA so they can be failed by failing one important component. Some trains may require that multiple components be failed before emulating failure of the associated train. Also, it is necessary to consider that a component (which would cause train failure) could be used for more than one train, system, or operating mode. Consequently, if the component is

not actually failed, it should not be used to fail its associated train. These problems result in the need for a PRA database specifically developed for modeling train-level failures along with component-level failures.

The term "train-level PRA" is used to denote a PRA that can be used to evaluate not only component failures but failures of trains. For a PRA to be considered a train-level PRA it must: (a) contain models that allow for train level failures (the modeling can not be stopped at the system level), (b) be constructed such that a determination can be made as to where a train begins and ends, and (c) be structured such that an analyst can fail trains, not just components in trains. While most nuclear power plant PRAs satisfy the first two conditions, the third condition may not be satisfied by many PRAs and their corresponding analysis tools. However, the SAPHIRE 4.0 software tools do have this capability.

To direct the construction of the train-level PRA models, several programmatic requirements were developed to describe the desired outcome of the train-level PRA research. These requirements include: (a) the databases be user friendly and require minimal personal involvement to utilize, (b) the databases would require a minimal amount of analysis time, including the time needed to add cut set factors such as recovery factors, (c) the databases allow investigations of both important component as well as train failures, (d) the train-level models provide an accurate representation of the plants and produce the associated CDF, (e) the methodology will be applicable to plant PRA or IPE models, and (f) the databases be developed using the SAPHIRE analysis programs.

TRAIN-LEVEL PRA DEVELOPMENT METHODS

Five train-level development methods were hypothesized. These methods are listed from the simplest (Method 1) to the most complex (Method 5).

- **Method 1.** Revise the PRA fault trees so simplified train-level models replace the detailed component-level fault trees. This method was not judged to be very useful for event analysis since it lacks component modeling.

- **Method 2.** Truncate the logic modeling of the original PRA fault trees at the train-level so the trains are treated as developed events[a]. The train-level developed events would then replace any component cut sets that may have appeared in the original cut set list.

- **Method 3.** Truncate the logic modeling of the original PRA fault trees at the train level so that trains are treated as developed events. Then, via the failure propagation feature in SAPHIRE, tie important components to the train so component level failures can result in train level failures. The train-level developed events would still replace any component cut sets, but component failures could cause its associated train to fail

---

[a] The SAPHIRE tools treat top event gates and transfer gates in fault trees as developed events. The term "developed event" should not be confused with the familiar PRA term "undeveloped event." A developed event is specifically meant to define the top event of an event tree or the top gate of a fault tree (or sub-tree) in the SAPHIRE package. But, in SAPHIRE, a developed event does show up in the list of basic events, which signifies that it is something more than just a fault tree gate.

via the failure propagation feature. The failure propagation method is specific to SAPHIRE and may not be available in other PRA software.

- **Method 4.** Truncate the logic modeling of the original PRA fault trees at the train level so the trains are treated as developed events. Also, put failures of major components at the same level (in the logic model) as the train failures.

- **Method 5.** Keep the logic modeling of the original PRA fault trees, but modify the train-level logic gates so the trains *may* be treated as developed events. Using this method allows a train to be failed, thereby removing its component cut sets, but still allows the rest of the PRA cut sets to be intact.

In order to compare the differences between the potential train-level methods, two of the methods (Methods 2 and 5) were used to develop separate databases using the Peach Bottom Unit 2 NUREG-1150 PRA.[3] The Method 2 database is referred to as PBTRAIN while the database for Method 5 is referred to as PBTLPRA.

COMPARISON OF TWO TRAIN-LEVEL DATABASES

Both the PBTRAIN and PBTLPRA databases are based on the Peach Bottom Unit 2 PRA. But, the two databases handle the concept of a train-level PRA differently. The PBTRAIN database truncates (by using the SAPHIRE process flags on the developed events) the fault tree logic models at the top gate of each train, starting from the bottom of a fault tree. The PBTLPRA database uses the same bottom-up approach to clarify where in the logic model the train top gates appear, but the database does not truncate the logic at that point. Instead, the PBTLPRA database allows failing a particular train by truncating the logic at a point in the fault trees that represents the failed train. A comparison of the two train-level databases, PBTRAIN and PBTLPRA, is presented in Table 1.

To fail a train using a developed event, the developed event that represents the train's top gate in the fault tree would be set to a TRUE house event (a TRUE house event in the SAPHIRE analysis tools is a logical true expression, i.e., a guaranteed failure) with the process flag set to an "X." Within the SAPHIRE analysis tools, setting a developed event's process flag to an "X" signifies that the logic *below* the developed event should be ignored and that the parameters (e.g., probability of failure) specified for the developed event should be used in place of the logic.

One of the most important advantages of the PBTLPRA database is that it provides results comparable to the original PRA for the calculation of the overall plant CDF. To illustrate this difference between using the PBTRAIN database and the PBTLPRA database, the 18 dominant sequences for each database were compared to the dominant sequences for the original Peach Bottom PRA. The summation of the dominant sequences for the PRA yielded a yearly CDF of 1.9×10^{-4}. The summation of the dominant sequences for PBTLPRA yielded a yearly CDF of 2.0×10^{-4}, while for PBTRAIN the yearly CDF was 2.6×10^{-6}. The total CDF for the PBTRAIN database is approximately two orders of magnitude lower than the original PRA results while the PBTLPRA database reasonably agrees with the CDF reported in the PRA. Several of the dominant sequences for the PBTRAIN database are too low by over three orders of magnitude. These differences between the original PRA and the PBTRAIN database can be attributed primarily to oversimplification of the PRA models.

CONCLUSIONS

While the train-level databases have specific advantages and disadvantages, they each have particular problems or limitations which became apparent during the development and use of the databases. Some of the problems or lessons learned are listed below.

- Dependencies, both for components and support systems, were not handled very well when simplifying the PRA models at the train level.
- Using the PBTRAIN database for event analysis would be limited to only evaluating a few front-line systems which are not buried under other train-level developed events.
- Even though the total number of accident sequence cut sets in PBTRAIN is over 10 times less than the total number of accident sequence cut sets in PBTLPRA, a large number of accident sequence cut sets will still appear during the event analysis.
- It may be necessary to generate recovery actions for use with a train-level event analysis or to duplicate the original PRA results. Unless applicable recovery factors are added to the cut sets, the analysis results could be erroneous and misleading.

Given the lessons learned, it is not recommended that the PBTRAIN database be used for any real-world risk analysis. Conversely, the PBTLPRA database provides appropriate results for an event analysis and could be used for a variety of risk analysis activities. The authors believe that for many PRAs or IPEs, simplifying or using generic logic models could result in erroneous and highly misleading results similar to those seen in the PBTRAIN simplified database. Further work is necessary to fully quantify the potential differences between full PRA models and simplified models.

Limitations to the recommended method and use of the PBTLPRA database are the length of time needed for analysis and the lack of preservation for accident sequence recovery factors. Version 5.0 of SAPHIRE has a module that allows for the automatic preservation of recovery factors. Additionally, SAPHIRE 5.0 is able to handle larger event trees and provides sequence cut set generation 2-5 times faster than version 4.0. It is expected that these improvements will make the PBTLPRA-type database and methodology suitable for train-level event analysis.

This work was supported by the U.S. Nuclear Regulatory Commission, Office of Nuclear Regulatory Research, Under DOE Idaho Field Office Contract DE-AC07-76ID01570.

REFERENCES

1. K. D. Russell et al., *Integrated Reliability and Risk Analysis System (IRRAS), Version 4.0*, NUREG/CR-5813, January 1992.

2. C. L. Smith, R. D. Fowler, and L. M. Wolfram, *Development and Use of a Train-Level Probabilistic Risk Assessment*, NUREG/CR-5976, April 1993.

3. A. M. Kolaczkowski et al, *Analysis of Core Damage Frequency: Peach Bottom, Unit 2 Internal Events*, NUREG/CR-4550, Volume 4, Revision 1, August 1989.

Table 1. A comparison of the two Peach Bottom train-level databases.

Database feature	PBTRAIN	PBTLPRA
Provides results comparable to the original PRA when used for event analysis	No (results may be off by orders of magnitude)	Yes
Breaks the system fault trees into train-level developed events	Yes	Yes
Significantly reduces the overall number of sequence cut sets (compared to the original PRA)	Yes (by over a factor of 10)	No
Time required to analyze accident sequences for event analysis	Short (sequences are requantified)	Long (sequences are regenerated and requantified)
Analyst can effectively fail components along with trains in the model	No	Yes
Crossties and interdependencies are handled correctly in the logic models	No	Yes
An uncertainty analysis can be easily performed for the event analysis	No	Yes (either Monte Carlo or Latin Hypercube are available)
Recovery factors are automatically included in the sequence generation	No (automatic recovery factors not available in SAPHIRE 4.0, but are included in SAPHIRE 5.0)	No (automatic recovery factors not available in SAPHIRE 4.0, but are included in SAPHIRE 5.0)
Nested trains and support systems failures are modeled adequately for event analysis	No	Yes

059 Use of PRA Models in Assessing Technologies and Device Applications

Chair: M. Zentner, Westinghouse Hanford

The Use of Robotics and Probabilistic Risk Assessment in Environmental Restoration
L.E. Cindel, G. Apostolakis (UCLA)

Risk Evaluation of Medical and Industrial Radiation Devices
E.D. Jones (LLNL), R.E. Cunningham, P.A. Rathbun (USNRC)

Integrated Systems Performance Assessment
C.N. Amos, S.D. Clement (SAIC)

THE USE OF ROBOTICS AND PROBABILISTIC RISK ASSESSMENT IN ENVIRONMENTAL RESTORATION

Laura E. Cindel and George Apostolakis

University of California
Mechanical, Aerospace, and Nuclear Engineering Department
38-137 Engineering IV
Los Angeles, CA 90025-1597

ABSTRACT

There is not much historical information when combining the technology of Robotics, with the methodology of Probabilistic Risk Assessment, for the application of Environmental Restoration and Waste Management. This purpose of this paper is to be an introduction to the topics above and will include: some of the key issues in each, what work has been performed in these areas, and a simple example of the combination.

INTRODUCTION

The Department of Energy (DOE), Environmental Restoration and Waste Management (EM) [1], has the responsibility of locating, identifying, and cleaning up the hazardous waste created by defense production endeavors of the past. As the magnitude and scope of the cleanup task are revealed, the implementation of advanced technologies, such as Robotics, seem a logical and beneficial solution in the reduction of further risk to the environment and the public.

Across the United States there is an estimated 120 sites which in some way are contaminated with nuclear and/or hazardous waste. Each of these sites need to be mapped out and the waste has to be sampled, characterized, treated, repackaged and disposed of in a safe and rapid manner [1]. The waste will be placed in newer storage facilities within the next twenty-six years and the environment will be restored to a safe state.

Robots and/or robotics systems are used in more than eighteen countries for countless applications within a very diverse group of industries. Robotics research has increased in recent years due, in part, to the cost of these systems becoming more competitive. In addition, robots have the added advantages of high precision, reliability,

and better quality. Some of the benefits of installing robotics systems are that they reduce occupational exposure risks associated with performing tasks in hazardous environments, fatigue associated with physically demanding tasks is reduced, and robots minimize human error associated with repetition [2].

The use of Probabilistic Risk Assessment (PRA) in industry has become more common in recent time. PRA involves utilizing the tools of statistics to analyze a system, component of a system, environment, etc. and then evaluate the risk or safety associated with using it. The methods and principles of Risk Management are tools which portray the problem/situation in a different perspective to help clarify which approach or direction to take.

The DOE, EM, Office of Technology Development Robotics Program is in the process of developing the robotic technology for environmental applications. The areas in which robotic assistance would be useful are: buried waste, decontaminating and decommissioning of nuclear facilities, facility management, sample analysis, and the minimization of waste production [3].

The combination of PRA and Robotics in the Environmental Field is a relatively new concept. Therefore, identifying areas where this combination could be beneficial is an important task. When PRA is used to analyze these systems, a deductive approach must be implemented since there is not much historical data available.

First some of the major issues involved when dealing with the above areas will be explored. Next there will be a brief summary of some of the work that has been, or is being, done. Then a simple example will demonstrate how these areas can be combined.

MAJOR ISSUES

Environmental Restoration and Waste Management (EM) is an extremely large and complex task. The most important issue that must be dealt with is that the materials and wastes that need to be cleaned up are still relatively unknown. These materials were disposed of many years ago and the constituent could have, and probably has, changed form. It is suspected that the waste includes but is not limited to: explosives, asbestos, inorganic chemicals and nuclear materials [4]. The integrity of the canister, drum, or tank is also unidentified [5], therefore lifting and transporting these receptacles may not be feasible, and the amount of ground water and soil contaminated due to receptacle leakage has not yet been calculated. It is difficult to plan far ahead, or to even devise a common plan for cleanup, since the environments are unstructured and each site is different.

Robotics is a blend of sophisticated portions of many engineering disciplines. Robots depend on controls, sensors, vision, statics, dynamics, computers, etc. [6] for existence. This makes the issues of robotics very diverse since the important problems from each of these areas must be taken into account. Autonomous robots are computer program and sensor controlled, and teleoperated robots are manually operated by a human. Computer software problems are a significant portion of the important issues of Robotics. This is due to the fact that software runs, or controls, equipment that is quite capable of harming/killing humans or damaging the environment, and the possibility of these accidents is not acceptable [6]. Humans write, and execute the software, therefore, software failure, for the most part, results from human error. Consequently reliable, and user-friendly, software is essential. Another issue is that robots have limited on board power. Also it is necessary to have a communication link. Therefore many robots use tethers. This is satisfactory for a relatively small traveling area, but for the EM application large areas of land must be traveled.

Probabilistic Risk Assessment (PRA) uses the principles of probability to analyze a system, component of a system, environment, etc., and to evaluate the risk or safety associated with it. Since human error is a large factor in robot accidents [7], PRA and the Markov method are used in the reliability analysis of a robot system with human error. There is not much historical data for EM applications. In these cases a deductive approach must be taken. PRA is useful, input, when data are sparse. It is most beneficial to implement the PRA process during the developmental stage, as time and cost can be reduced. It can also be particularly helpful in improving the designer's understanding of the importance of safety and its significance/impact on the overall system design. Since most of the innovative applications for environmental restoration and waste management are still in the developmental stage it is possible to use the PRA technique.

LITERATURE REVIEW

As stated above, there has not been much work done on the implementation of PRA to analyze the use of Robotics in the Environmental Field therefore, identifying areas where this analysis could be advantageous, or where it has been used, should be done first. The following are applications where the three have been employed together.

The robot reliability issue is an important one when placing the robot in a situation where retrieval is difficult. Therefore, the design of a fault-tolerant system is an important development. The fault-tolerant systems developed [8, 9,10,11], use rule-based and model-based reasoning to determine, and then recover from, events that were not originally planned for. Fault trees are used extensively and are the key component to this fault tolerant concept.

Goossens' [12] work demonstrates how the PRA process is useful to determine, and avoid, possible human/robot accidents. Cost-benefit analysis of robotics is also addressed, this is to display how, where, and to what degree, safety precautions can be implemented to reduce accidents and the cost associated with them.

As stated previously there is considerable uncertainty when dealing with the environment. Due to this, there is difficulty in planning for, or determining whether, the robot is in contact with another object. Spreng [13] employs hypotheses' analysis on the uncertainty model of the PRA methodology. Tests will be performed by the robot to determine if it is actually touching (not just near) another object, and what position the object is in.

Extended Petri nets [14] are another tool to aid in the detection of failure. These failures can occur in any part of the system, not necessarily just the robot. This is beneficial since the total system must be operative for a successful mission. Neural Maps [15] can detect and respond to changes in the system, even suggesting alternatives, so that the task does not have to be terminated.

ROBOTIC, PROBABILISTIC RISK ASSESSMENT, AND ENVIRONMENTAL RESTORATION AND WASTE MANAGEMENT EXAMPLE

The following is a simplified example of how to implement some of the PRA techniques to assist in the analysis of the use of a robot for a specific application of EM. The specific situation studied was the use of an autonomous robot to uncover, inspect, open, repackage, and load hazardous waste for its transportation from its previous disposal

site to another disposal location.

What is of concern is whether the robot will fail, and if it does fail, what the consequences will be. A fault tree, figure 1, was created to display many common areas which could cause the malfunction or failure of a robot. The main branches of the fault tree are: human interference, internal hardware malfunctions, external hardware malfunction, and software malfunction. Once there is a failure in a main branch and the robot fails, the consequences and outcomes must be investigated.

An event tree, figure 2, was created to look at the various outcome paths to determine if there was human harm. The initiating event of the event tree is a robot failure. There are various questions that must be asked, and these answers form the branches of the event tree. The outcome possibilities of these branches are one of the following, no human harm, latent death, or instantaneous death.

The results of the thirty-two outcomes are that thirteen of these branches resulted in instantaneous deaths, two outcomes lead to latent deaths, and seventeen outcomes resulted in a non-hazardous situation. The outcomes that resulted in latent or instantaneous deaths must then be examined. This is to see what safety features could be added to the system to minimize the probability of the occurrence of these outcomes. Although fifteen of the thirty-two outcomes resulted in death, the probability of their occurrence is small. Also, by using robots, the human is removed from the contaminated area which in turn reduced human exposure level.

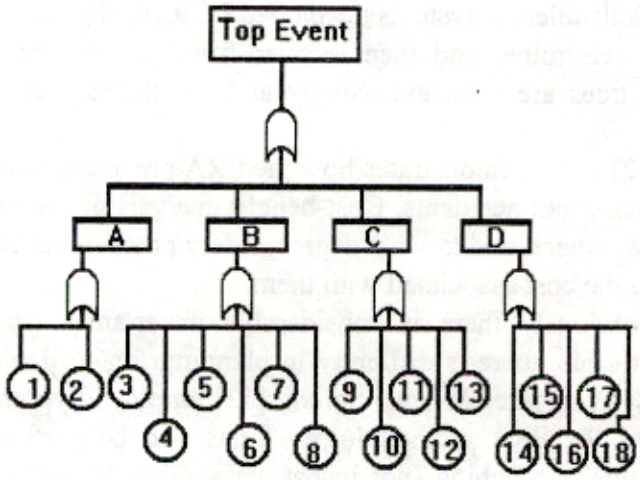


Figure 1: Fault Tree

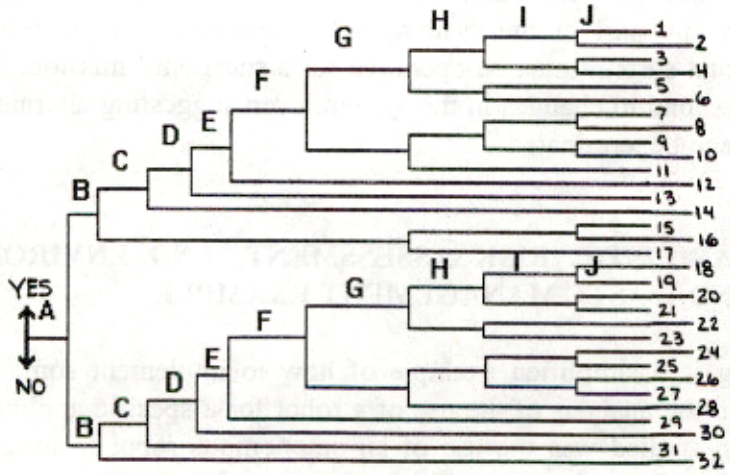


Figure 2: Event Tree

Table 1: Key for Figure 1

Symbol	Description
Top Event	Robot Malfunction or Failure
A	Human Interference
B	Internal Hardware Failure
C	External Hardware Failure
D	Software Malfunction
1	Intentional Entry
2	Unintentional Entry
3	Radiant Heating
4	Fire
5	Explosion
6	Shock & Vibration
7	Noise Interference
8	Hardware/ Software Interfacing Problem
9	Radiant Heating
10	Liquid Spray
11	Fire
12	Fumes & Vapor
13	Explosion
14	Radiant Heating
15	ROM Problem
16	RAM Problem
17	Software/ Hardware Interfacing Problem
18	Error with the Programming

Table 2: Key for Figure 2

Symbol	Description
A	Initiating Event: Robot Failure (from top of Figure 1)
B	Is there Involuntary Movement of Robot?
C	Is the Robot handling Hazardous Material?
D	Was there damage to the canister?
E	If so, was there a release of Haz. Mat. to the Environment?
F	Was the Haz. Mat. in an uncontrolled environment (site, field, not lab)
G	Did the Haz. Mat. get into the ground water or air?
H	Were there Humans in the area at time of accident?
I	Were the Humans exposed to radiation?
J	Was the Human struck by the robot?

Table 3: Outcome

1	Instant Death	12	Non- Hazardous	23	Non-Hazardous
2	Instant Death	13	Non-Hazardous	24	Instant Death
3	Instant Death	14	Non-Hazardous	25	Instant Death
4	Non-Hazardous	15	Instant Death	26	Instant Death
5	Latent Death	16	Non-Hazardous	27	Non-Hazardous
6	Non-Hazardous	17	Non-Hazardous	28	Non-Hazardous
7	Instant Death	18	Instant Death	29	Non-Hazardous
8	Instant Death	19	Instant Death	30	Non-Hazardous
9	Instant Death	20	Instant Death	31	Non-Hazardous
10	Instant Death	21	Non-Hazardous	32	Non-Hazardous
11	Non-Hazardous	22	Latent Death		

Note: Death is the worst case scenario, sickness or injury could result instead but it is still a hazardous situation.

CONCLUSION

Combining Probabilistic Risk Assessment, Robotics, and Environmental Restoration and Waste Management seems to be a very practical and advantageous idea. Although it is still a relatively new concept, there is plenty of information of each of these areas individually, and the literature search shows that there are many people working on uniting them.

REFERENCES

[1] U.S. Department of Energy, Environmental Restoration and Waste Management Five-Year Plan, Fiscal Years 1993-1997; Washington, DC., August, 1991

[2] Robotics Program Development: Applicable Lessons Learned from TMI-2, Final Report, EPRI NP-6521, September 1989

[3] H.B. Meleran, "Robots in the USA Nuclear Industry: An Overview to the End of the 20th Century", Proceedings of the 2nd ASME/JSME Joint Conference on Nuclear Engineering, March 1993, Vol. 2, pp 255

[4] B.L. Burks, B.S. Richardson, G.A. Armstrong, W.R. Hamel, D.H. Thompson, "A demonstration of Remote Survey and Characterization of a Buried Waste Site Using the SRIP Testbed", Proceedings of the 4th ANS Topical Meeting on Robotics and Remote Systems, January 1991, pp 179

[5] F. Hechendorn, C. Ward, and D. Wagner, "Remote Radioactive Waste Drum Inspection with an Autonomous Mobile Robot", Proc. of the ANS 5th Topical Meeting on Robotics and Remote Systems, April 1993

[6] J.A. McDermid, "Issues in Developing Software for Safety Critical Systems", Reliability Engineering and System Safety, 1991, Vol. 32, pp 1

[7] B.S. Dhillon, Robot Reliability and Safety, Springer-Verlag, New York, Inc., 1991

[8] M.L. Visinsky, I.D. Walker, J.R. Cavalaro, "Robotic Fault Tolerance: Algorithms and Architectures", Robotics and Remote Systems in Hazardous Environments, Prentice Hall, Inc., 1993, Chapt. 3, pp 53-73

[9] M.L. Visinsky, I.D. Walker, J.R. Cavalaro, "Layered Dynamic Fault Detection and Tolerance for Robots", IEEE International Conference on Robotics and Automation, Atlanta, GA, May 1993, Vol. 2, pp 180-187

[10] M.L. Visinsky, I.D. Walker, J.R. Cavalaro, "Expert System Framework for Fault Detection and Fault Tolerance in Robotics", Proceedings of the International Symposium on Robotics and Manufacturing, Santa Fe, NM, Nov. 1992, pp 793-799

[11] K.S. Tso, M. Hecht, N.I. Marzwell, "Fault-Tolerant Robotic System for Critical Applications", IEEE International Conference on Robotics and Automation, May 1993, vol. 3, pp 691-696

[12] L.H.J. Goossens, "Risk Prevention and Policy-Making in Automatic Systems", Risk Analysis, Vol. 11, No. 2, 1991

[13] M. Spreng, "A Probabilistic Method to Analyze Ambiguous Contact Situations", IEEE International Conference on Robotics and Automation, May 1993, vol. 3, pp 543-54

[14] S. Ramaswamy, K.P. Valavanis, "Modeling, Analysis and Simulation of Failures in a Materials Handling System with Extended Petri Nets", IEEE International Conference on Robotics and Automation, May 1993, vol. 1, pp 323-330

[15] A. Syed, H.A. Elmaraghy, N. Chagneux, "Real-Time Monitoring and Diagnosing of Robotic Assembly with Self-Organizing Neural Maps", IEEE International Conference on Robotics and Automation, May 1993, Vol. 2, pp 188-195

RISK EVALUATION OF MEDICAL AND INDUSTRIAL RADIATION DEVICES*

Edwin D. Jones

Fission Energy and Systems Safety Program
Lawrence Livermore National Laboratory
P.O. Box 808, L-632
Livermore, CA 94551-9900

Richard E. Cunningham and Patricia A. Rathbun

Office of Nuclear Material Safety and Safeguards
U.S. Nuclear Regulatory Commission
Mail Stop 6H3, OWFN
Washington, D.C. 20555-0001

INTRODUCTION

In 1991, the NRC, Division of Industrial and Medical Nuclear Safety, began a program to evaluate the use of probabilistic risk assessment (PRA) in regulating medical devices. This program represents an initial step in an overall plan to evaluate the use of PRA in regulating the use of nuclear by-product materials. The NRC envisioned that the use of risk analysis techniques could assist staff in ensuring that the regulatory approach was standardized, understandable, and effective.

Traditional methods of assessing risk in nuclear power plants may be inappropriate to use in assessing the use of by-product devices. The approaches used in assessing nuclear reactor risks are equipment-oriented. Secondary attention is paid to the human component, for the most part *after* critical system failure events have been identified.

This paper describes the risk methodology developed by Lawrence Livermore National Laboratory (LLNL), initially intended to assess risks associated with the use of the Gamma Knife,+ a gamma stereotactic radiosurgical device. For relatively new medical devices such as the Gamma Knife, the challenge is to perform a risk analysis with very little quantitative data but with an important human factor component. The method

* This work was supported by the U.S. Nuclear Regulatory Commission, Office of Nuclear Material Safety and Safeguards, under the Interagency Agreement DOE 40-550-75 with the U.S. Department of Energy.
+ The Gamma Knife is a registered trademark of Elekta Instruments, Inc.

described below provides a basic approach for identifying the most likely risk contributors and evaluating their relative importance.

The risk analysis approach developed for the Gamma Knife and described in this paper should be applicable to a broader class of devices in which the human interaction with the device is a prominent factor. In this sense, the method could be a prototypical model of nuclear medical or industrial device risk analysis.

RISK ASSESSMENT METHODOLOGY

The risk approach used is motivated by a need to have a flexible analysis framework that can incorporate both qualitative and quantitative data about human and equipment factors, and that is consistent with the ALARA principle (reducing risk to a level "as low as reasonably achievable"). The methodology should not be a rule-based methodology, but should be a systematic approach to uncovering risk at various levels of resolution for a range of levels of effort. The methodology must also be able to accommodate a variety of medical practices and devices. It thus must be empirically based and must not rely on preconceived notions of system processes. For relatively new devices, most of the operating experience data will be qualitative, i.e., anecdotal, rather than quantitative. Thus, the risk analysis must not rely on only quantitative data to be useful. It should be able to compare a range of data types and data quality. In the methodology, there must be equanimity between human and equipment factors: the methodology cannot simply be machine- or human-centered in its orientation. The notion of ALARA provides a principle of relativity for the risk analysis: to lower risks to a level as low as reasonably achievable. This can be accomplished without absolute measures of risk (which, after all, are very difficult and costly to ascertain): only relative improvements in apparent risks are needed.

After a few iterations of considering potential risk analysis methodologies, it was decided that the above criteria could best be met by the approach of developing relative rankings of risk or risk profiles. Profile analysis is a general analytic tool which has been employed since the late 1940s. In the last decade, profile analytic techniques have been applied to the evaluation of both machine failures and human errors in nuclear facilities (Seaver and Stillwell, 1983; Banks and Paramore, 1983).

The risk assessment methodology is an empirically based, systematic approach to uncovering potential risks. It consists of a hierarchy of iterated stages: (1) process and sequence identification; (2) hazards evaluation; (3) event/task analysis and data collection; and (4) relative rankings and risk profiles. The first three stages serve to systematically identify elements most likely to contribute to risk. The last stage provides a measure of the relative risk importance of each of the identified risk contributors. Even though each stage requires its own analysis techniques and tools, the stages are not independent, and each requires iterated inputs from previous and subsequent stages. In this way, the total analysis is thorough, balanced, and self-consistent. Documented and anecdotal information as well as subject matter expertise is used, as much as possible, at every stage of the methodology.

Identifying Risk Contributors

The first stage is to define the problem to be analyzed, which involves selecting the series of events or process to be evaluated. The selection requires familiarity with the system of interest, its operational requirements and functions, and the role of the human operator. A process should be selected that can be separated, by a reasonable person, into a sequence of steps. The criteria for the delineation of a sequence step is that the step must be completed successfully before the next step can be initiated.

Identifying and judging hazardous elements or conditions associated with the use of the device is a very important stage. It provides the boundary conditions or circumscription for subsequent analyses. Insufficient resources are available to analyze everything, ad infinitum, so a clear understanding of hazardous situations is necessary to direct, focus, and constrain the analysis efforts. The hazards evaluation helps identify those sequence elements that deserve further evaluation.

As part of the risk identification effort for the Gamma Knife, it was important to search for extraneous radiation fields that may affect people during system failures or abnormal operating modes (Smith et al, 1993). First, the radiation levels were checked at the intended treatment target as a function of patient positioning during a normal treatment cycle. Then fields were checked during a failure mode of the hydraulic system while a phantom was within the radiation unit. (The primary concern was extraneous radiation to the patient close to the primary sources.)

Two kinds of radiation hot spots were discovered to which a patient would be subject while being transported between the shielding door and the treatment position, but not while in the treatment position. One hot spot (approximately 8–10% of maximum dose rate) is due to transmission of the primary beams through the stainless steel of the collimating helmet. Other hot spots (approximately 1–2% of maximum) are due to inadvertent, non-attenuated transmission through misaligned collimators. These effects disappear at the treatment position, because the tungsten collimators are aligned and they prevent transmission of the primary beams.

After these determinations, Rhode Island Hospital carefully checked for radiation hot spots (with the shielding door open) on the walls of their Gamma Knife suite. They found a collimated beam coming out of each side of the open shielding door due to a systemic design flaw. The radiation outside of the shielding door is thus not purely scattered radiation. (This problem has now been successfully corrected at all U.S. Gamma Knife facilities.)

Once a process sequence is developed and the hazards are well defined and understood, a human-oriented task sequence list needs to be developed for each step of the process of interest. The first task in each list should be the initiating task for the process step, and the last task or subtask in each list needs to be completed successfully before the next step of the process can occur. Such task lists were developed for each of the Gamma Knife treatment path processes of imaging and localization, treatment planning, and patient positioning and treatment. All tasks have the characteristics of a purpose or goal, an input or stimulus, a decision or response by the operator, and a system or process change which can be fed back to the operator.

For the purpose of a risk analysis, tasks should be selected that are judged to be the most pertinent activities affecting risks associated with the device. The analyst must ascertain where errors most relevant to risk can or do occur. Then each event or task sequence must be clearly delineated. Each task should have a well-defined error and a measurable consequence associated with that error. The selection of tasks should be based on expert actual experience with the use of the device.

In the case of the Gamma Knife, data were collected from medical associations, standard-setting organizations, the manufacturer, users, and experts. A multidisciplinary team of physicians and medical physicists with expertise in teletherapy, risk assessment experts, and scientists and engineers with extensive knowledge of task and safety analyses inspected Gamma Knife units, attended acceptance tests, interviewed users, and observed patient treatments. The data were verified for accuracy, completeness, and self-consistency by the use of subject matter experts, simulations, facility walk-throughs, and observing actual practices. Data analyzed by the project team were subsequently reviewed, critiqued, and validated by medical community expert review teams.

Evaluating Risk Contributors

Once the tasks have been analyzed and selected for errors pertinent to risk, it is possible to identify those tasks associated with the highest risks. In a first-order risk analysis, likely error rates and consequences for each task are treated as independent from other tasks, and estimated as if they are associated with isolated events. However, many errors or consequences are mitigated by verification or checking procedures. Such procedures must be adequately reflected in the task list, so that final ranking schemes can logically combine probability and magnitude of errors for a concatenation of tasks, in order to accurately represent potential error propagations. Scenarios involving concatenations of tasks are examined to validate or adjust the rankings for each task, in order to ensure appropriate relative rankings,

The relative likelihood of error and the degree of consequence can be estimated by subject matter experts. Evaluations have provided encouraging support for the use of expert judgment (Comer et al, 1983). Experts are good at making relative estimates on limited scales, and relative rankings are reproducible. One may not conclude, however, that the expert judgments have predictive validity if no statistics are available for comparison or calibration. An advantage of direct numerical estimation is that it can be used to obtain estimates of uncertainty bounds.

Ranking data were collected for each task by asking relevant experts to provide their estimation of error frequencies or likelihoods and error magnitudes (dose deviations) associated with those errors. Experts were asked to make estimates based on their personal knowledge or experience. At this level of analysis, the issue was not how or why errors occurred, but how often errors have in fact occurred. Relative ratings or discrete distributions were used — absolute values and continuous distributions are desirable, but not necessary.

Comparing rankings among tasks is done by means of a "risk profile." There are various ways to compare tasks. One instructive representation is to graph the tasks by their probabilities versus their consequences, by using the mean values from the discrete data distributions. Another first-order risk profile results from multiplying the probability of error by the magnitude of consequence for each task. Such profiles clearly identified the high risk tasks: those with relatively high probability and high consequence.

The relative rankings and error distributions also can be used in computerized Monte Carlo simulations of scenarios that provide a higher level of risk analysis, because concatenations or interactions among diverse tasks can be simulated and evaluated. Relative measures are sufficient for the Monte Carlo technique, since only weighted stochastic choices are used.

In the Gamma Knife project, a Monte Carlo computer code was developed and used to simulate and evaluate the relative risks of possible error scenarios. The code was named the Monte Carlo Risk Scenario Code or MCRSC ("McRisk"). It makes full use of the developed error rate and error magnitude distributions and can model the interactions among any number of tasks, logically convolving distributions. It was used to aggregate subtasks and their error distributions, determine best- and worse-case extremes, and perform the analyses outlined below. It is important to have a tool like MCRSC for handling the uncertainties associated with human errors in the use of medical devices. The large values of uncertainty indicate that the first order risk analysis, where only the mean values were used, may not be adequate to represent interacting errors among tasks.

The Monte Carlo simulations generated a multitude of error scenarios and their associated risks. Such simulations exposed synergies or scenarios that otherwise may have gone unnoticed. MCRSC produced distributions that enabled the identification of high-risk scenarios and tasks most likely associated with the highest risk scenarios.

Uncertainty, sensitivity, and mitigation studies were performed by analyzing the changes in risk profiles that were due to variations in task or error data. Thus, evaluations were made on the risk impacts of system design changes, site and device variabilities, potential risk mitigation and management measures, and regulatory requirements.

If costs are associated with proposed changes in tasks or error rates and magnitudes, then risk cost-benefit studies are also possible, e.g., risk vs. cost utility curves can be generated. The Monte Carlo simulations can generate risk outcomes for a variety of risk mitigation strategies in such a way that minimax criteria (and hence linear programming techniques) can be applied to determine the optimal mix of strategies. The mixed strategies may involve, for instance, a combination of changes in design, procedures, and training.

SUMMARY AND CONCLUSIONS

The efforts in applying probabilistic risk techniques and methods to the use of by-product devices show promise for developing indices of risk importance and effective risk management practices. The methods provide a flexible, basic approach for identifying most likely risk contributors and the relative importance of each contributor. The risk analysis tools also provide a platform for evaluating the effectiveness of regulatory practices.

The results obtained have their limitations, however, because the work is new and innovative. The results so far are device-specific, and it remains to test the validity, reliability, and applicability of the methods to other devices. Deeper analyses may reveal shortcomings of the methodologies or lead to the development of improved analysis techniques. Finally, because of the qualitative nature of the data available for the devices studied, uncertainty bounds are not well understood.

REFERENCES

Banks, W.W., and Paramore, B., October 1983, Systems integration: a pilot task analysis of the DOE size reduction facility; EGG-REP-6440, Idaho National Engineering Laboratory.

Comer, M.K., Kozinsky, E. J., Eckel, J.S., and Miller, D.P., February 1983, A data bank conception and system description, *in:* "Human Reliability Data Bank for Nuclear Power Plant Operations," Vol. 2; NUREG/CR-2744, General Physics Corporation and Sandia National Laboratories.

Seaver, D.A., and Stillwell, W.G., March 1983, Procedures for using expert judgment to estimate human error probabilities in nuclear power plant operations; NUREG/CR-2743, Decision Science Consortium and Sandia National Laboratories.

Smith V., Verbey, L., Jones, E., and Lyman L., 1993, Consequences to the patient in the event of hydraulic unit failure; *to appear in:* "Stereotactic and Functional Neurosurgery," Vol. 61, Supplement 1.

INTEGRATED SYSTEMS PERFORMANCE ASSESSMENT

Christopher N. Amos and Steven D. Clement

Integrated Technology Applications Group
Science Applications International Corporation
Albuquerque, NM 87106

INTRODUCTION

Integrated Systems Performance Assessment (ISPA) refers to the application of knowledge-based decision analysis to the probabilistic analysis of engineered systems. ISPA is an extension of the event progression analysis methodology that was developed by Sandia National Laboratories[1] as part of the nuclear power plant risk assessment program undertaken by the U.S. Nuclear Regulatory Commission[2]. To assess the containment response of these plants to postulated severe core damage accidents, large complex event trees were developed. The EVNTRE computer code[1] was written to evaluate these trees. ISPA involves the application of EVNTRE to the assessment of engineered systems. This assessment involves the development of event tree or decision tree models for system performance. Software tools that facilitate these applications have been developed by the authors and their co-workers.

EVNTRE requires that sequences in an event tree be described through a programming language developed specifically for that purpose. For each top event, rules are developed that define the probabilities of the outcomes (i.e., branches) for that event based on the outcome of previous events in the tree (i.e., the sequence up to this point). Thus, the set of rules that specify any given tree are the framework for knowledge-based decision analysis. Consistent with the graphical representation of the model as an event tree, the rules can reference only prior events in the tree (i.e., ISPA allows only backward references). In addition to specifying the branch split fractions for any top event, the rule-based language allows the user to specify values for an arbitrary number of parameters that can be associated with each sequence. Parameters can be used to carry an estimate of cost, or other consequence, associated with events in the sequence. Rules can reference the value of the parameters at a given point (i.e., event) in the sequence to determine the event outcome. For example, the probability of a branch representing successful funding of a program can be determined based on the value of a parameter representing the program cost that has been calculated based on the previous event outcomes in the sequence. Figure 1 illustrates these capabilities schematically.

I	Event 1	Event 2	Event 3	Event 4	What is the Program Cost? COST	Is the Program Cost $2 or Less? FUNDING	SEQ. PROB.
1.00E+00	A	A	A	C	$2	Yes	5.00E-02
			9.00E-01	A 3.00E-01	$1	Yes	2.25E-01
			B	B 2.00E-01	$2	Yes	1.35E-01
				C	$3	No	9.00E-02
			A	B 5.00 E-01	$2	Yes	2.23E-01
	5.00E-01	A	1.00E-02	C	$3	No	2.23E-01
	B	1.00E-01	C	B	$4	No	4.50E-03
		B	A	A	$2	Yes	5.00E-02

ISPA RULES

Event 1 (E1)
A: 0.5
B: 0.5

Event 2 (E2)
If E1=A then
A: 1.0
B: 0.0

Otherwise
A: 0.9
B: 0.1

Event 3 (E3)
If E1=A then
A: 0.1
B: 0.9
C: 0.0

If E1=B and E2=B then
A: 1.0
B: 0.0
C: 0.0

Otherwise
A: 0.99
B: 0.0
C: 0.01

Event 4 (E4)
If E1=A and E3=A then
A: 0.0
B: 0.0
C: 1.0

If E1=A and E3=B then
A: 0.5
B: 0.3
C: 0.2

If E1=B and E2=A and E3=A then
A: 0.0
B: 0.5
C: 0.5

If E3=C then
A: 0.0
B: 1.0
C: 0.0

Otherwise
A: 1.0
B: 0.0
C: 0.0

ISPA PARAMETERS

PARAMETER	COST
A	0
B	1
C	2

Figure 1. Schematic Representation of ISPA Model

ISPA APPLICATIONS

This paper describes the general capabilities of ISPA through the presentation of three examples. These presentations are tailored to emphasize the role of ISPA in the assessment. It is not the intent of this discussion to describe the specific needs that were met through the application of ISPA or detail the results that were obtained. The three application examples are (1) comparison of two conceptual designs for a space nuclear

power source, (2) combining the analyses comprising the safety assessment report (SAR) for a nuclear process facility to provide a more integrated assessment of risk, and (3) a system for the safety assessment of aircraft operations. An additional example of applying ISPA to perform engineering sensitivity calculations is presented by McClure, et al.[3] in another paper at this conference.

In each of these cases, an ISPA model was developed that predicts the performance of the system (or systems) under consideration in terms of a number of pre-established criteria. These models could be represented graphically as an event tree. Each sequence represents a different combination of sub-system performance characteristics. Performance measures of the overall system are developed by combining the sequence-specific performance characteristics for the sub-systems and/or components. The predicted probability for each sequence indicates the relative likelihood of achieving the corresponding level of overall system performance. Performance risk (i.e., the sum of the probability-weighted performance indicators for each sequence) provides a useful assessment measure for the system.

Conceptual Designs for Space Nuclear Power Systems (SNPSs)

ISPA was used to assess two conceptual designs for a SNPS. Performance criteria were specified in advance of the ISPA model development. These criteria were developed in a hierarchy, consistent with accepted practice in decision analysis[4]. The ISPA model was developed to reflect the relationships between design features of the proposed alternative designs and their performance with respect to these criteria. For instance, the extent of development effort required to engineer the system components so that they would achieve the level of performance indicated by the designer was assumed to be directly related to technology risk (i.e., the probability that a working prototype with these specifications could be developed). In this assessment, two types of models were used to express these relationships. For some cases, such as weight of the spacecraft, mathematical expressions (e.g., simple addition of component weights) could be used to compare the characteristics of the proposed design to the performance criteria. Other relationships were more subjective. For instance, public acceptance could be evaluated based on a number of subjective assessments related to reliability, maturity of the technology employed, and reliability of the required launch vehicle given the assessed weight.

The motivation for putting these relationships among design characteristics and the assessment criteria into an ISPA model is twofold. First, there are uncertainties associated with the assessment of performance. These uncertainties result from both lack of data and the subjective aspects of the assessment. Sensitivity studies can be conducted to determine the effect of these uncertainties on the overall assessment of performance. Second, it is desirable to develop an overall "score" for each design relative to the criteria. Each criterion was assigned a weight (i.e., relative importance). The performance of each design relative to the criteria was then assessed on a numerical scale. An overall score was then calculated for each design by summing the individual "scores" for each criterion. The uncertainty in the score is also readily assessed in this approach. This assessment indicates whether it is reasonable to discriminate between the two designs based on the difference between their scores.

This same assessment model could also be applied for design optimization once a conceptual design is selected. Assessing the effect of proposed design changes on the design "score" indicates the level of design optimization achieved. The simultaneous assessment of uncertainty is useful in determining whether the design changes have a significant effect on the assessed performance. The use of knowledge-based decision analysis in this process ensures that relationships between design characteristics and

performance indicators are measured in a consistent fashion. Thus, the potential negative side effects of a design alteration intended to resolve a particular issue are readily identified.

Quantitative Assessment of Facility Safety

ISPA was applied to predict the risk of facility operations as part of the revision of a facility SAR. For many facilities, the accident (or hazard) analysis portion of the document consists of a number of separate analyses of postulated upset events (i.e., accidents). For U.S. Department of Energy (DOE) facilities, these analyses are typically probabilistic. The results obtained for each accident analyzed is the estimated frequency of occurrence and the predicted consequence. While the form of these results is consistent with those of a risk assessment, they cannot be used directly to estimate the risk of facility operation. Because the analyses were developed independently, their results are only additive if the events analyzed are completely independent. This is rarely true within a facility. There are typically common support systems such as electric power, compressed air, and ventilation. Since failures in these common systems can affect more than one of the postulated accidents that have been analyzed, the actual risk of facility operation can be significantly underestimated if the accidents are treated as being independent.

ISPA was used to integrate the independent assessment of accidents for a nuclear process facility and obtain a realistic assessment of operating risk. In addition to providing the framework for this integration, the ISPA model allowed the risk from seismic and fire initiators to be analyzed. This could not be accomplished using the individual accident models, since the common initiator would not be treated consistently and the individual accident frequency predictions were not additive.

The model was constructed by designating each postulated accident as a top event in the facility "event tree"*. Events that represent the failure of common systems were added to the front of the tree. Sequences through the set of support system events represented the possible support system states. Accident frequencies for these support states were then calculated using the cut sets that were developed from the accident fault trees. Flags in these cut sets were set by the ISPA model according to the support state for the sequence being evaluated. The values of these flags (i.e., combinations of "true" and "false") indicated the particular support state for which the cut sets were to be evaluated. Accidents that could induce other accidents were treated in the same manner (i.e., their prior occurrence in the sequence was indicated by a flag).

This relatively complex model was developed within a few weeks and used much of the existing models without substantial modification. The ISPA framework allowed the complex event tree to be developed in a straightforward fashion. Since the small event approximation was frequently not valid for situations in which cut set evaluation was necessary, a method for evaluating the cut sets more precisely was incorporated into the EVNTRE code. A more extensive discussion of this application can be found in Sharp and Brandyberry[5] and Brandyberry and Amos[6] in other papers at this conference.

Safety Assessment for Air Operations

The safety of aircraft operations can be strongly dependent on mission planning. It is frequently true that the individual responsible for planning may not be completely

---

* The model is not actually an event tree since the sequences do not necessarily indicate the progression of an accident. Some of the sequences are single accidents (i.e., one event "occurs" and it is the accident of interest). Other sequences represent multiple accidents that are either induced by a common initiator (e.g., electric power failure triggers several upsets) or as the result of accident progression (i.e., the upset induced by the initiator causes subsequent events to be initiated).

informed with respect to all factors that determine mission success. Rather than require extensive participation from a number of experts in the planning process, it was desirable to incorporate their expertise into a planning tool. An ISPA model was developed to demonstrate the potential to serve as an aid in the planning process. The demonstration model allows the mission planner to enter information about the mission plan into a personal computer. A data entry form, formatted as a questionnaire, is provided. Based on the responses to the questions, the ISPA model predicts the probability of mission success. Predictions of the potential consequences of failure, in terms of damage to the aircraft or other property, injury, or loss of life, are also developed. A completed ISPA model for this application would identify those aspects of the mission profile, as indicated by the questionnaire, that should be adjusted to reduce the risks associated with the operation. It was intended that the completed model could be used as a "sounding board" for the development of mission plans that appropriately consider the potential risks.

To develop the model, the planning process was subdivided into an number of areas: (1) the mission profile (e.g., number and type of aircraft, payloads, terrain, flight level, expected speed); (2) the amount of experience and specific training of the personnel involved; (3) the procurement process used to acquire the aircraft, pilots, and ground crews; (4) the extent of flight plan development; and (5) the expected weather conditions. Experts in the field of aviation safety were closely involved in the development of the ISPA model. In each of the areas listed above, considerations that could impact the probability of mission success were identified. For instance, if formation flying was required, the probability of a collision between aircraft will be strongly influenced by the formation flying experience of the pilots involved. If rapid refueling of the aircraft was required, the experience of the ground crews with such operations influences the probability of a refueling accident. The requirements of the mission and the type of aircraft used (e.g., fixed-wing versus helicopter) affect the potential for an accident. These relationships were incorporated into the ISPA model as Boolean expressions that determined the probability of a negative outcome given relevant combinations of the factors in the planning process that could affect their probability of occurrence. The preliminary model relied solely on expert opinion. The proposed complete model would incorporate relevant portions of the extensive air safety data base so as to provide a more reliable estimate of risk.

This example is relevant to the use of ISPA to manage risks. Identification of the risk factors in a facility, process, or undertaking is part of the risk assessment or hazard analysis process and are incorporated into the model. Aspects of the operation, or decisions, that could influence these factors are identified and their effects are modeled within the ISPA framework. This example uses a questionnaire to drive the assessment of the consequences of operational decisions. Other formats are also possible.

SUMMARY

ISPA provides a powerful tool to support risk management and decision analysis. The examples described above illustrate a number of potential applications. These applications span the range of risk management requirements for many contemporary engineering undertakings. Selection of a technology, or system design, considering both public and programmatic risks, is a requirement in many arenas. As illustrated in the SNPS example, ISPA can provide a platform for comparing candidate system designs against a fixed set of criteria. Risks (e.g., public, technology, or programmatic) associated with a particular selection are readily predicted. Once the technology, or system design has been selected, the ISPA model used in the selection process can be refined to support design optimization. As demonstrated in the facility modeling application, ISPA can incorporate existing models for failure rates and accident consequences with little modification. This

makes it a cost-effective tool in many applications. Refinement to reflect operational considerations will allow the ISPA model to be used for risk management once the process or facility has been fielded. Again, little additional development resources are required.

This cradle-to-grave perspective on ISPA applications is entirely consistent with regulatory requirements. The National Environmental Policy Act (NEPA) requires an assessment of the environmental consequences of any major federal action prior to its implementation. The comparison of alternatives that would be included in an environmental impact statement can be appropriately supported with an ISPA model. ISPA provides a tool that can be used to support the evaluation of alternatives, comparing the associated risks in a consistent and systematic fashion. For actions that require permitting under the Resource Conservation and Recovery Act or the National Pollution Discharge Elimination System, an ISPA model can provide a systematic means of meeting the risk assessment requirements for certain types of facilities. If an initial risk assessment was required under NEPA, refinement of the ISPA model for this purpose provides a cost-effective approach to meeting that requirement. This same model can be used as an aid to managing risks from the implementation of the process through to final disposition. ISPA can play a similar role in actions where the Comprehensive Environmental Response, Compensation, and Liability Act is invoked. An ISPA model can be used to predict the risk of alternative remedial actions. Once the approach to remediation has been selected based on risk minimization, the same ISPA model can be refined as a risk management tool for the implementation phase. In addition to the use of ISPA to support risk management, regulatory requirements can be incorporated into the model to assist in the identification of decisions made during the implementation of an action that have the potential to impact regulatory compliance. This capability is incorporated through the inclusion of these requirements into the ISPA framework.

In summary, ISPA has broad applications. It offers the potential for making risk assessment, and risk management, available on a routine basis for daily activities. This is accomplished by incorporating the knowledge of the risk analyst and the compliance specialist into an ISPA model that is used to assess day-to-day project activities.

REFERENCES

1. J.M. Griesmeyer and L.N. Smith. "The Event Progression Analysis Code (EVNTRE) Users' Guide," United States Nuclear Regulatory Commission, Sandia National Laboratories, NUREG/CR-5174 (1989).

2. United States Nuclear Regulatory Commission. "An Assessment of Five Nuclear Power Plants." NUREG-1150 (1990).

3. P.R. McClure, S.D. Clement, and S.G. Ashbaugh, "An Application of ISPA: Sensitivity Analysis of a Complex Engineering Problem," Proceedings of PSAM II, March 1994.

4. T.L. Saaty, "The Analytic Hierarchy Process," RWS Publications, Pittsburgh, 1990.

5. D.A. Sharp and M.D. Brandyberry,"Risk Assessment Methodology for the DWPF at the Savanna River Site", Proceedings of PSAM II, 1994.

6. M.D. Brandyberry and C.N. Amos, "Application of Accident Progression Event Tree Technology to the Savanna River Site Defense Waste Processing Facility SAR Analysis", Proceedings of PSAM II, 1994.

060 HRA, the Second Generation (II)

Chair: E. Hollnagel, Computer Resources Intl.

Reliability of Cognition, Context and Data for a Second Generation HRA
E. Hollnagel (HRA, UK); P.C. Cacciabue (CEC/JRC, Italy)

Elements Of A Model of Operator Problem Solving and Decision Making in Abnormal Conditions
S-H. Shen, C. Smidts, A. Mosleh (U. Maryland)

Human Reliability Analysis Using Simulated Human Model
Y. Fujita, I. Yanagisawa (MAPI, Japan); H. Sakuda (NUPEC/IHF, Japan)

RELIABILITY OF COGNITION, CONTEXT AND DATA FOR A SECOND GENERATION HRA

Erik Hollnagel[1], P.Carlo Cacciabue[2]

[1]Human Reliability Associates, Higher Lane,
Dalton, Parabold, Lancashire, WN8 7RP, UK.

[2]Commission of the European Communities, Joint Research Centre,
Institute for Systems Engineering and Informatics, 21020 Ispra (Va), Italy

THE IMPORTANCE OF THE CONTEXT

It has become accepted to think of Human Reliability Assessment (HRA) as taking place in a context, meaning that the assessment of human reliability cannot be done without recognising that human activity always is part of a situation thus assessing the reliability of isolated human actions, such as observing a signal or turning a switch, is of limited value for practical applications (Hollnagel, 1993, Singh, Parry, Beare, 1993; Dougherty, 1993). But the context of HRA can be understood in two different ways: either the HRA as part of a context, or the context as part of the HRA.

Considering HRA as part of a context means that the HRA takes into account its role within the wider scope of a PSA/PRA. This is in particular important for a **Cognitive Reliability Analysis** (CRA), i.e., an HRA that looks specifically at the role and reliability of human cognition. The role of the CRA is not only to provide the basis for quantitative estimates of the probability of certain events happening or of certain cognitive functions failing, but also and perhaps more importantly, to provide the basis for an assessment of the likely functioning of the joint man-machine system with regard to overall reliability and safely. This means that the HRA/CRA must produce results that are meaningful within the larger context of the PSA/PRA.

RETROSPECTIVE AND PROSPECTIVE ANALYSES

Considering context as part of a HRA/CRA means that the HRA/CRA explicitly describes the context of human cognition. The main difference between **retrospective** event analyses and **prospective** (predictive) analyses is in the existence of a context (figure 1) (Hollnagel, 1993). A retrospective analysis by definition takes place in a clear context, namely the event itself. The analysis typically moves backwards from the event to the underlying causes. During this process knowledge of the context can be used to eliminate

candidate explanations and to identify the most likely causal connections. In contrast to that, a prospective analysis must explicitly describe or develop the context.

A prospective analysis typically moves forwards from an action or an event and tries to find the possible consequences. If this propagation of effects is done without constraints, the number of possible end events very becomes unmanageable; even the simplest taxonomy will produce a combinatorial explosion.

The context is therefore needed both to limit the number of possible consequences and to give meaning to the results of the analysis. Without the context the analysis becomes a purely syntactic exercise. Specifically, the reliability of cognition cannot be described with reference only to internal cognitive functions but must acknowledge the conditions under which the person has to work - such as the interface, the procedures, the team, the working environment and organisation, etc. In particular, this kind of context description may naturally extends the CRA to include also considerations relative to group collective and co-operative decision making as well as management roles and functions.

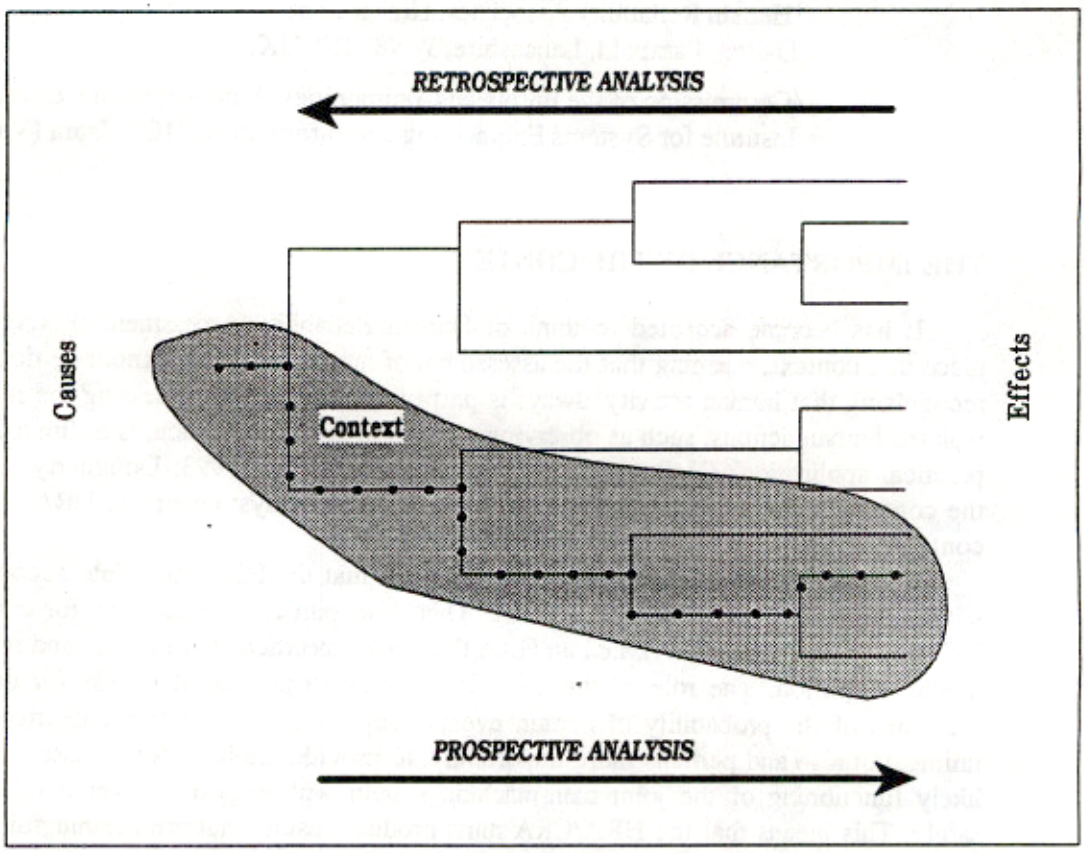


Figure 1. Retrospective and prospective analyses.

A CRA must therefore begin by establishing the context. This can be done by means of a Goals-Means Task Analysis (GMTA) supplemented by specification of the common performance modes through e.g. the Dependent Differentiation Method (Hollnagel, 1993). That provides the appropriate basis for making quantitative estimates of the probabilities of certain specific actions. On the whole, however, the aim must be to ensure that the goal of the actions can be achieved. Individual steps or actions must accordingly be seen in relation to that goal - which is another way of saying that they must be seen in a context.

THE CONTEXT OF COGNITION

This is, however, not the only perspective from which the context should be considered. Indeed, the dynamic evolution of the system, or plant, under control plays an essential role in how operators choose to respond to an incident or an unexpected event. Consequently, providing an account of this context - possibly in the form of a simulation - with its balance equations, physical variables and constraints, must be explicitly addressed as an element of a second generation approach to CRA/HRA.

As mentioned above, a GMTA is an efficient and convenient way of producing a description of the context. By the GMTA the analyst defines **which** are the important quantities affecting operator behaviour. The simulation of the system under control has then to account for these quantities either by the classical conservation equations or by parametrical or logical models, when simulation tools are not available or simpler descriptions are acceptable. Moreover, for each quantity, one has to identify **how** its evolution affects the behaviour, i.e. which values identify an abnormal situation vs. an operational transient or a simple deviation from nominal behaviour, which values characterise a specific accident or system configuration etc. These questions come prior and are of support to the evaluation of probabilities necessary for the PSA.

Even if the example of the physical variables is very representative for contextual elements affecting behaviour, there are other relevant components of the working environment that can affect the behaviour and generate erroneous performances. In general, five categories of factors can be identified, namely: Interference, Communication, Equipment/Interfaces, Workplace and Procedures (Table 1).

Table 1. External factors affecting human behaviour.

External factors	Example
Interference	Conflicting activities
Communication	Colleague, noise
Equipment/Interfaces (transient)	Spurious signal, trend, critical values
Workplace (permanent)	Ambiguous labelling
Procedures	Ambiguous text

In order to characterise these factors, the analyst has to acquire a domain knowledge by a number of activities, such as:

- a plant visit;
- a review of the information from the designer and/or the safety analyst;
- the task analysis.

By this approach, which is common to many classical human factors methods, such as THERP (Swain and Guttman, 1980), the relations existing between the actual working environment and the operators can be described in practice either through analytical forms or by logical expressions. These analytical and logical forms can be developed using techniques like the fuzzy set theory, the logic of predicates, the rule based systems, as well as simpler mathematical formulations, like polynomials or differential equations.

The contextual factors affecting human behaviour are similar to the "Performance Shaping Factors" (PSF), which are well known and widely adopted by the more classical human reliability methods. However, there is a crucial difference between the PSFs and the proposed factors: while the PSF are the result of a process of decomposition of the whole human activity into the individual actions and the associated environmental conditions affecting it, the proposed set of factors are consistent with notion of Common Performance

Modes (CPM), intended as general determinants of performance (Hollnagel, 1991), which affect the performance as a whole rather than generating a specific effect on a single action. The actual forms that the inappropriate behaviour may take are to be evaluated independently of the CPMs (see below), even if the origin of these inappropriate behaviours are to be traced back to the CPMs.

When the CPMs are identified and correlated to the context by some formulations, it is possible to associate to each of the identified CPMs a probability value. In this way, the boundary and initial conditions of the context affecting human behaviour can be appropriately defined and can then be coupled to a dynamic analysis of the human machine interaction (Cacciabue and Vivalda, 1991), able to generate at the same time:

- the environmental conditions leading to the inappropriate behaviour and
- the probability of failure to successfully perform the expected actions.

THE CONTEXT AND ERROR MODES

The manifestation of erroneous actions can be described in terms of error modes and/or phenotypes. The error modes refer to the various dimensions by which things can be wrongly done, typically using a set of categories as shown below (Table 2).

Table 2. Basic error modes.

Basic error modes	Example
Timing	Too early, too late, not at all (omission)
Duration	Too short, too long
Force	Too little, too much
Direction	Wrong direction
Sequence	Reversal, repetition, omission, intrusion
Magnitude	Too short , too far
Object	Wrong action, wrong object

The phenotypes are a more abstract way. of describing how incorrect actions can appear, using a psychological approach (Hollnagel, 1993). The simple phenotypes are: repetition, reversal, omission, delay, premature action, replacement, insertion, and intrusion. In making a prospective analysis it is important to know precisely how the erroneous actions are going to manifest themselves. This means both **which** error modes (or phenotypes) are likely and **what** the magnitude or limits are.

To tackle the last aspect first, it is obviously essential to know what the normal range of variation in executing an action is Consider, for instance, incorrect timing - either too early or too late. It depends on the process, and therefore on the context, what constitutes too early or too late. In some cases it may be a matter of seconds, in others a matter of minutes or even hours The determination of what the normal variations are, hence of what the limits of acceptability are, is a clear example of context dependence. The HRA cannot determine if an action is carried out too early, unless it is precisely known what the acceptable limits are. These limits can be derived either from a description of the process, from experience, or from a systematic GMTA combined with a process simulation. The same considerations are, of course, applicable to the other error modes In retrospective

analyses the limits are obviously known, because the consequences are known. In prospective analyses the limits must be determined as part of specifying the context.

The list of error modes (or phenotypes) is made without taking the concrete context into account. The lists simply describe the logically possible ways in which an incorrect action can manifest itself. In order to apply these categories as part of an HRA it is therefore necessary to determine which categories, i.e., which error modes, are actually possible. This can only be done if the context is known. Detailed knowledge of the nature of the process, the interface, the barriers and interlock mechanisms, etc. can be used to specify which particular subset of error modes are actually possible. If, for instance, input is provided via a touch panel display, error modes referring to magnitude, duration, and direction are probably not applicable (as the computer can be used to remove possible variation in these dimensions). If the set of possible error modes can be reduced - and even detailed further by including limit values - then the analysis can be simplified in the sense that fewer possible consequences have to be considered. In this way the context improves the quality and precision of the analysis, at the same time as the analyst's work is made easier. Instead of making "blind" combinations of semantically meaningless categories, he can now concentrate on identifying the causal connections that are meaningful vis-à-vis of the context (Cacciabue, Pedrali, Hollnagel , 1993).

These two aspects of an approach to Human Reliability Analysis could very well be fitted into the existing methodologies (as SHARP or THERP) or they could even replace some of the fundamental steps of specific methods, while retaining all the other methodological features, specially those related to the analysis of operator behaviour and the familiarisation with the actual work setting. They would strengthen the outcome of the qualitative survey of the working environment and of the analysis of the procedures for accident management. They would provide input data for evaluating the consequences of erroneous actions on the plant behaviour. They would lead to anticipation or predictive results concerning possible outcome of human machine interaction.

CONCLUSIONS

The present generation of Human Reliability Assessment methods shows some very valuable features as well as some shortcomings, mainly with reference to the methods for analysing and predicting the interaction of the operators in their socio-technical working environment. There are a number of possible ways for improvements, aiming at a more complete account of the actual human machine interaction process and at a better understanding of the causes of human erroneous behaviour.

In this work, we have focused on the **context**, and on the two ways in which the context interacts with the human, i.e., in terms of "the environment affecting the human" and of "the human responses to control the environment". The inclusion of the effects of the context in existing methods of HRA is possible and it would already grant a considerable progress towards the formulation of more complex and complete techniques. For example, the retrospective studies of real accidents, for the identification of the root cause(s) of erroneous behaviour can be carried out accounting for the contextual factors, as they have been proposed here, in combination with tables (Cojazzi, Pinola, 1994) or diagrams describing the sequence of the events (Pedrali, 1993). This analysis provides:

- the environmental ("external") and the cognitive ("internal") factors affecting the human behaviour;
- their likely combinations; and, possibly,
- the different timings at which they modify the interaction processes, resulting in the erroneous manifestation, i.e., the inappropriate action.

The consideration for the context is the first crucial step towards the formulation of really innovative methods. However, the second fundamental step forward which needs to be done consists in the coupling of the contextual effects with the actual dynamics of the interaction between the system and the human in control. In this way the overall simulation scenario is completed, as all the actors of the interactive loop, namely the humans and the machines, are simulated in parallel and in consideration of their mutual effects for the complete study of the safety of a system. The accomplishment of this second development is also being studied and will be the object of the future research.

Finally, it can be argued that the improvement of the "classical" HRA techniques can be performed in a stepwise process aiming at new approaches, which are more complete, both for the consideration for the actual context and for the inclusion of the dynamic characteristics. The completeness of the approaches that could be obtained in this way, and their integration into the PSA methodological framework would allow the full exploitation of the PSA for more realistic and accurate studies of systems with strong dependence on the human machine interaction and on the dynamic evolution of sequences.

REFERENCES

Cacciabue P.C., Vivalda C. 1991, A Dynamic Methodology for Evaluating Human Error Probabilities. In Apostolakis, G. (Ed.), Proc. of Int. Conf. on *Probabilistic Safety Assessment and Management (PSAM)*. Beverly Hills, California, 4-7 February 1991. Elsevier, New York, NY, 507-512.

Cacciabue P.C., and Hollnagel E., 1993, Human Models in Reliability and Safety Analysis of Interactive Systems. Proc. of Int. ANS/ENS Topical Meeting on Probabilistic Safety Assessment, PSA 93, Clearwater Beach, FL, Jan. 26-29, 1993. American Nuclear Society, La Grange Park, Ill., 25-31.

Cacciabue P.C., Pedrali M., Hollnagel E., 1993. Taxonomy and Models for Human Factors Analysis of Interactive Systems: An Application to Flight Safety. 2nd ICAO Flight Safety and Human Factors Symposium, Washington, DC, 12-15 April, 1993.

Cojazzi G., Pinola L., 1994, Root Cause Analysis Methodologies: Trends and Needs. *Probabilistic Safety Assessment and Management II (PSAM-II)*. San Diego, California, 24-28 March, 1994.

Dougherty E.M., 1991 Issues of Human Reliability in Risk Analysis. In Apostolakis, G. (Ed.), Proc. of the Int. Conf. on *Probabilistic Safety Assessment and Management (PSAM)*. Beverly Hills, California, 4-7, February 1991, Elsevier, New York, NY, 699-704.

Dougherty E, 1993, Context and Human Reliability Analysis. *Reliability Engineering and System Safety, RE&SS*, **41**, 25-48.

Hollnagel E, 1993, *Reliability of Cognition: Foundations of Human Reliability Analysis.* Academic Press, London.

Hollnagel E., 1991, The Phenotype of Erroneous Actions: Implications for HCI Design. In G.R.S. Weir & J.L. Alty (Eds.), *Human Computer Interaction and the Complex Systems*, pp. 73-121. Academic Press, London.

Pedrali M. 1993, Modelli e tassonomie di errori umani per l'analisi di sicurezza nell'aviazione civile. Il caso studio delll'incidente di Zurigo: ricerca delle cause e indagine prospettica. Graduation Thesis in Aeronautical Engineering, Politecnico di Milano, April 1993.

Singh A., Parry G., and Beare A., 1993, An Approach to the Analysis of Operating Crew Responses for use in PSAs. Proc. of Int. ANS/ENS Topical Meeting on Probabilistic Safety Assessment, PSA 93, Clearwater Beach, FL, Jan. 26-29, 1993. American Nuclear Society, La Grange Park, Ill., 294-300.

Swain A.D., Guttman H.E., 1980, Handbook on Human Reliability Analysis with Emphasis on Nuclear Power Plant Application. NUREG/CR-1278. SAND 80-0200 RX, AN. Sandia National Laboratories, Albuquerque, New Mexico.

ELEMENTS OF A MODEL FOR OPERATOR PROBLEM SOLVING AND DECISION MAKING IN ABNORMAL CONDITIONS

S.-H. Shen, C. Smidts, A. Mosleh

Department of Materials and Nuclear Engineering
University of Maryland
College Park, MD 20740

INTRODUCTION

The need for a new approach to modeling human interactions with systems has been recognized by researchers and practitioners particularly in the context of PRAs, where probabilistic assessment of operator performance in response to emergencies is of interest. Modeling behavior through understanding the cognitive elements of human response has gained popularity among researchers in recent years. The efforts have led to development of frameworks[1,2], taxonomies[3], evolutionary models[4], and "engineering approaches"[3].

By "engineering approach", we mean that available theoretical and empirical results together with approximations and assumptions are synthesized into a working model without requiring a basic scientific knowledge of all the elements involved. While the success of this approach is evident from the history of technology and applied sciences, its relevance and applicability to modeling human performance is debatable. Nevertheless, the temptation to try it despite the skepticism of many is strong and motivated to some extent by the desire to help in decisions that cannot wait until a comprehensive science of human behavior has been developed.

The model of this paper is best characterized as an engineering model. The basic approach is to use the available knowledge about human cognition, techniques of artificial intelligence, experimental results, actual experience, intuition and "common sense" to develop a predictive probabilistic model of operator/system interactions. The domain of applicability of the model is the relatively constrained environment where behavior is significantly influenced by high levels of training and written procedures. These constraints in fact simplify the modeling by removing "degrees of freedom" from the broader human response spectrum.

The model presented here consists of a crew task assignments and communication model, and of a model for individual operator behavior. The major components of the individual operator model are the Information Module (dealing with observation and perception of information), the Action Module (dealing with transforming intention to act into action), and the Problem Solving & Decision Making Module (which formulates and intention to act given the information forwarded by the Information Module).

The present work primarily focuses on the development of the problem solving and

decision making aspects of the model of an individual operator. In the following we will first describe the fundamental elements of the various components of the model along with a partial description of the way in which these elements interact[5].

THE OPERATING CREW MODEL

Our model of the crew behavior makes the explicit assumption that the crew is a team of individuals working on different assigned tasks and communicating with each other. Our model of the crew is rather crude because it is not the main focus of this work. However, it was recognized that approximating the behavior of a team with a model of an individual operator is not only unrealistic but also introduces more complexity in developing a generic model of cognitive processes. The present model is a partial extension of Huang et al.[6]. Despite the large diversity of crew structures observed in nuclear power plants, we simplify the problem by assuming that the basic elements of a crew are (see Fig-1): one (or several) Reactor Operators (RO) who perform the actions, a supervisor (SRO) who makes decisions, and a consultant (STA) whose main function is to advise the supervisor. Moreover, we assume that RO's do not communicate. In Fig-1, solid lines represent the "normal" channels of communication whereas dotted lines represent "passive" reception of information. Errors in communication occur when the content of the message is lost in, or distorted by, the communication channel. Communication by-passes may occur as a function of the receiver characteristics. The actual or perceived nature of the sender might affect the credibility of a message content and its weight in future decisions. Examples of a more extensive list of communication error modes are provided in Ref.-5.

MAIN COMPONENTS OF THE INDIVIDUAL OPERATOR MODEL

Fig.2 shows the main components of the individual operator model. The Knowledge Base (KB) contains information relevant to the process memorized during training or previous experience. The Information Module (IM) includes information which is provided by the plant, the procedure and/or other members of the crew and perceived by the operator. The Action Module (AM) transforms the intention into action. Intentions are generated by the Problem Solving/Decision Making (PS/DM) Module. The Working Memory (WM) Module is a short term memory unit which contains information that has been activated through the PS/DM process and gathered from the I module and/or from the KB. This representation and break-down of human cognitive activities is quite traditional and is found in most cognitive models[]. However, refinements can be brought such as allowing an additional layer of memorization, the Intermediate Memory (IM), to distinguish clearly between the current information and past information which is still retrievable because it is not yet forgotten or stored permanently (learned) in the Knowledge Base. This work does not address the way in which information is observed and perceived or the way in which an intention is transformed into an action. We refer the reader to Allport[8], Swain[9] for work related to perception and action.

The Working Memory Unit

The Working Memory module is a short-term memory unit that includes pieces of information extracted and activated by the PS/DM process from the Information Module and the Knowledge Base. The volume of information kept in WM is controlled by the limited storage capacity[9] as determined by cognitive and psychological experiments. Some pieces of information are of particular interest if we want to predict future operator behavior. These specific elements are worth separating from the mass and assembling in a new unit which we will call the Mental State (MS) of the operator. The MS itself can

be subdivided into information-related and emotion-related states (such as stress). A partial description of the MS for nuclear power plant operators is provided in Table -1.

Table 1. The Operator's Mental State.

Elements of Mental State Related to:	Information Part of Mental State (I)	Emotional Part of Mental State (Ψ)
Situation Assessment	▪Global Diagnosis ▪Specific Diagnosis ▪Local Diagnosis ▪Global Expected Plant Behavior ▪Specific Expected Plant Behavior ▪Local Expected Plant Behavior ▪Previous Specific Diagnosis	▪ Confidence in Global Diagnosis ▪ Confidence in Specific Diagnosis ▪ Confidence in Local Diagnosis ▪ Time Pressure ▪ Level of Fatigue ▪ Perception of Danger
Problem Solving	▪Current Goal or Problem ▪Previous Goal or Problem ▪PS/DM Strategy ▪Previous Goal ▪Previous PS or DM Strategy	▪ Confidence in PS/DM Strategy ▪ Confidence in Implementation of Specific Steps of the PS/DM Strategy
Crew	▪ Knowledge of respective responsibilities	▪Assessment of Other's and Own Capacity to Perform given Role

The Intermediate Memory Unit

The Intermediate Memory Unit is the other short term memory unit added for modeling convenience. It contains the operator's most recent experience, and, elements of the mental state such as the *previous diagnosis* or the *PS/DM strategy*, are directly related to its content . Information stored in this memory will be either removed after some time and hence fully disappear or be stored in the Knowledge Base as part of a learning process.

The Problem Solving/Decision Making Module

The main elements of the PS/DM module are respectively the problem which is currently being solved and the strategy which is used to solve it.

Goals-Problems. Each goal/problem uses a given PS/DM strategy. With time different goals are generated and resolved, strategies are selected and executed or eventually abandoned if unsuccessful, steps in strategies are performed, questioned and reexecuted. There exists a vertical hierarchy of goals and problems: sub-problems are generated to solve specific parts of a higher level problem. An example is the response of an operator to an initiating event that leads to a reactor trip. After the reactor trip, the operator may decide to follow the Emergency Operating Procedures (EOP). If the EOP leaves room for interpretation (underspecified diagnosis rules), then the choice between two (or more) initiators becomes by itself a sub-goal, requiring a new problem solving strategy. There exists also an horizontal hierarchy of goals related to their respective importance or priority.

In the context of accidents and abnormal situations we have identified two high level goals. The first, called *maintaining safety margin*, refers to the desire to keep the integrity of the plant by initiating safety systems when required. The second, called *search for accident root cause* speaks by itself. There could be multiple goals that the crew will want to address at the same time. These multiple goals can be multiple goals of a given individual and/or multiple goals for a crew that translate into unique goals for each member of the crew. Rational behavior implies that goals will be resolved according to their levels of priority. Transitions between goals can occur before the current goal is entirely fulfilled.

Problem Solving and Decision Making Strategies.

Strategies. A limited set of problem solving and decision making strategies seem to suffice to describe a wide spectrum of possible operator behaviors. Seven problem solving strategies have been identified and ranked in Table-2 from the less to the more cognitively demanding. We have identified one decision making strategy also showed in Table 2.

Table 2. Problem Solving and Decision making Strategies.

Reflex Response is an immediate reaction to a major threat.
Direct Matching consists in the matching of the signature of a situation (for instance the pattern of alarms on a control room board) to an immediate solution stored either in the WM or in the Knowledge Base.
Follow a Procedure (written or oral)
Logic Expansion is a cause analysis process which lists the potential causes of a symptom and then goes over this list to check for contradictory or confirming evidence, eliminate progressively non confirmed causes and converge to one (if possible) cause for the observed accident or abnormal condition.
Trial & Error which identifies possible causes for given accident symptoms (or a restricted and simplified set of accident symptoms) and tests them by implementing related actions and observing their outcome ,i.e., the plant response.
Wait & Monitor consists in delaying decision and action in the expectation of further information and feedback from the system.
Optimization between choices which is based on perceived costs and benefits (where costs and benefits can be in terms of efforts and/or money).

Transition between Strategies. Once a problem has been defined, the operator has to select a strategy to solve it. We think that this initial selection is governed by a combination of different influencing factors:

- Degree of cognitive complexity (i.e. the operator will privilege the problem solving strategies that imply less mental activity. He will prefer skill[10]-based, over rule-based behavior, and rule-based over knowledge-based behavior).
- Efficiency of the strategy given the time left to implement it (balance between accuracy and time).
- Problem solver's level of knowledge. A novice is not likely to exhibit direct matching because he has no ready made solutions for the current problem.
- Expected consequences of a given strategy (risk).
- Personal biases (training, experience). For example, past deceiving experience might discourage operators from using procedures.
- Organizational biases (rules and standards). For instance, the requirement for the use of procedures after reactor trip or safety injection.

Transitions also occur between strategies if for instance a strategy has proven to be unsuccessful in coping with the problem. This reflects on the *confidence in the current problem solving strategy* index of the mental state. Factors that influence these later transitions and hence progressively degrade the confidence level are: (1) mismatch between reality and expected system behavior (inadequacy), (2) low level of improvement exhibited by the situation (fatigue), (3) time pressure (time left for improvement).

Once the transition out of the strategy has occurred, a new strategy is selected. This selection follows the same governing principles as the selection of the initial strategy, however is excluded from the set of potential choices. This new strategy is entered with a state of mind that reflects the previous efforts and the knowledge the operator has gathered through the partial execution of the unsuccessful strategy.

For instance, an operator has started to approach a problem with a "Logical Expansion" strategy. Using this strategy, he obtains confirming evidence that the accident can not be a Loss of Coolant Accident (LOCA). Along the process, he also gathers

information about principal plant parameters. However, he fails to identify the root cause of the accident and decides to switch to the EOP. Hence, he will initiate procedure E-0. But given his initial state of knowledge (his state of mind) he is likely to skip the different verification steps which constitute the first part of E-0 and to avoid to enter in procedure E-1 (LOCA). In conclusion, the point at which the new selected strategy is entered and the way it is executed (for instance, the existence of short cuts) depends mainly on the current operator's MS.

Transitions may also occur within strategies. They then take different forms such as *back-up*[11] or *barge on* (i.e. relaxation of steps) and are initiated by a mismatch between expectancy (or in the mental state, *the expected behavior* of the plant) and real situation. It seems reasonable to assess that the operator will only resume to switch to another strategy when all possibilities within the current strategy have been exhausted.

The Knowledge Base

The knowledge base is a long term memory unit containing:

- Operator's basic physical knowledge (e.g., a break in a closed loop leads to its depressurization).
- Operator's specific mental representation of the plant such as: the functional representation, the layout of the components, a qualitative model of the plant response under given accident conditions.
- Operating biases such as the fact that instrumentation indications could be questionable.
- Memorized steps of procedures (scripts).

Representations may take several forms: neural networks (for the description of the physical plant behavior), rules: A and/or B→C, visual representations of accident patterns, scripts (procedures). The elements of the representation are accessed and eventually combined through the reasoning process. Future work needs to determine if there exists a correct representation, the way in which elements of a representation or the representations themselves are combined and the spectrum of possible operators (analogous to mathematical ones) used in these combinations. An initial list of these operators includes: operators that transform generic knowledge into specific knowledge (i.e. context dependent), specific knowledge into generic knowledge (abstraction), operators that interpolate, extrapolate, approximate, operators that allow analogies between different fields of knowledge.

Modeling Dynamic Aspects

Information becomes available in the WM through the I module. This information may be of two different kinds. It is either the answer to a question the operator elaborated during his reasoning process or it has nothing to do with the ongoing reasoning process and was passed to the WM because of its level of criticality. If this information has a high level of criticality it is likely to interrupt the ongoing process and modify the MS, the content of the IM, the goal or the relative priority of goals, the PS/DM making strategy. This interruption of the normal cycle of operations related to the dynamics of the systems we study (be it the plant or the other members of the crew), has not yet been fully studied and will be part of future developments.

CONCLUDING REMARKS

This paper has presented some of the fundamental elements of a model which when fully developed is expected to provide a predictive tool for nuclear power plant operator response in emergencies. The core of the ideas presented here relate to an important (and perhaps the most challenging) aspect of human cognitive process, namely the problem solving and decision making. The elements introduced here, and more fully explained in

Ref.-5, have been utilized as the basis for techniques useful in a number of applications. For instance, we developed a data classification scheme for systematic and exhaustive gathering of data from actual experience and applied it to several AEOD events. This enabled us to identify gaps in the interviews of operators, and also it provided qualitative input to our model enhancement. A taxonomy of errors has been developed by assuming a sequential action formation process which goes from Information gathering to Problem Solving then Decision Making and finally Action, and, by assuming that errors are likely to happen at any stage of this chain. Potential applications of the concepts and of the model include the improvement of operating conditions, for instance, by placing the predictive model in a simulation environment under different external conditions. This will help improving training, procedures and man-machine interfaces.

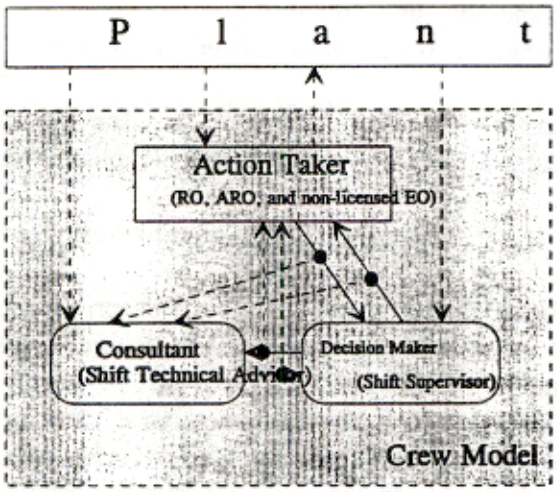


Figure 1. The Operating Crew Model

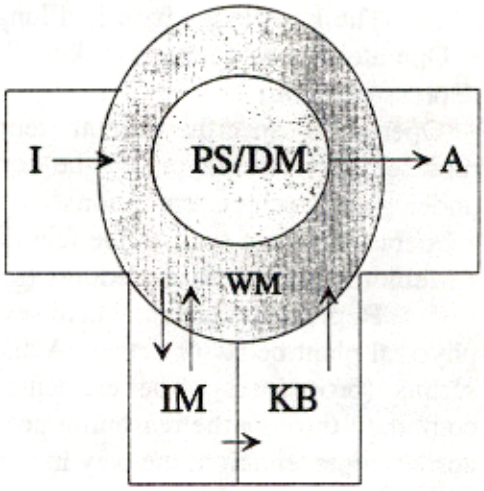


Figure 2. Main Elements of the Individual Operator Model

REFERENCE

1. D.D. Woods, E.M. Roth, H.Jr. Pople, "Modeling intention formation for human reliability assessment", Rel. Eng. and System Safety 22: 169 - 200, (1988).
2. P.C. Cacciabue et al, "A cognitive model in a blackboard architecture: synergism of AI and psychology", Reliab. Eng. and System Safety, 36: 187-197, (1992).
3. E. Hollnagel, P.C. Cacciabue, "Modeling cognition and erroneous actions in system simulation context", submitted for publication to Int. Journal of Man-Machine Studies.
4. A.P. Macwan, A. Mosleh, "A simulation-based approach to modeling errors of commission during nuclear power plant accidents: Application to PRA", PSA'93, (1993).
5. S.-H. Shen, C. Smidts, A. Mosleh, "A model for operator problem solving and decision making in abnormal conditions and its applications", MDNE-93-002 Report, (1993).
6. Y.Huang, N. Siu, D. Lanning, J. Carroll and V. Dang, "Modeling control room crew for accident sequence analysis", Grant Final Report, MIT (1991).
7. A. Newell, P.S. Rosenbloom, J.E. Laird, "Symbolic architectures for cognition" in "Foundations of Cognitive Science", Ed. Michael I. Posner, The MIT Press, Cambridge, (1990).
8. A. Allport, "Visual Attention", in "Foundations of Cognitive Science", Ed. Michael I. Posner, The MIT Press, Cambridge, (1990).
9. A.D. Swain, H.E. Guttmann, "Handbook of Human Reliability Analysis with Emphasis on Nuclear Power Plant Applications", NUREG/CR-1278, NRC, Washington, D.C., (1983).
10. J. Rasmussen, "Information Processing and Human-Machine Interaction: An Approach to Cognitive Engineering", North-Holland, New York, (1986).
11. K. VanLehn, "Mind Bugs: The Origins of Procedural Misconceptions", The MIT Press Cambridge (1990).

HUMAN RELIABILITY ANALYSIS USING SIMULATED HUMAN MODEL

Yushi Fujita,[1] Hiroshi Sakuda,[2] and Ichiro Yanagisawa[1]

[1] Mitsubishi Atomic Power Industries, Inc. (MAPI)
4-1, 2-Chome, Shibakouen, Minato-ku, Tokyo 105, Japan

[2] Institute of Human Factors, Nuclear Power Engineering Corporation (NUPEC/IHF)
17-1, 3-Chome, Toranomon, Minato-ku, Tokyo 105, Japan

INTRODUCTION

Simulated human models have increasingly been recognized as a useful vehicle in cognitive systems engineering (Roth and Mumaw, 1993). They allow us to look into context-driven interactions between human and situations which are believed to be a major determinant of human performance (Fujita, 1992; Hollnagel, 1993). They also provide us with practical merits including the repeatability and the ability to allow parametric studies which are not necessarily easy to obtain in empirical approaches.

A task analysis tool called CAMEO Task Analysis Tool (CAMEO/TAT) incorporates such a simulated human model called CAMEO (i.e. Cognitive and Action Modeling of Erring Operator) which simulates the behavior of a supervisory operator working in process control environments (Fujita et al., 1993a). CAMEO/TAT is designed to be used for evaluating task design (e.g. design of operating procedures or other work rules) in its early stages. It allows task designers to ensure that operators can carry out tasks as they expected. This evaluation is important because there could be factors which designers are not always able to fully take into account when they design tasks. Such factors which might impair the basis of task design, thereby potentially giving negative influences on operator performance, include task load, task complexity, time pressure, opportunistic change of task order, multiple task environments, negative feedback from previously made decisions or actions, operator's policies and traits, etc.

Task designers input task design and relevant plant data (e.g. data on plant configuration, specifications of individual components) in the form of operational knowledge or rules which they assume that operators will possess, and then CAMEO simulates operator behavior by dynamically interacting with a given situation (e.g. accident) and provides information on "erroneous tendency." Here, erroneous tendency refers to an internal state in which "resources" (Norman and Bobrow, 1975; Gopher and Donchin, 1986) allocated to a certain task are less than necessary. When an erroneous tendency occurs, it manifests itself as erroneous behavior or an error depending on external performance criteria. Designers can iteratively modify and evaluate task design until no erroneous tendencies are found with CAMEO/TAT.

The purpose of this paper is to detail CAMEO and also to claim that simulated human models such as CAMEO can improve human reliability analysis (HRA).

COGNITIVE AND ACTION MODELING OF ERRING OPERATOR

Framework

CAMEO takes a view "... it may often be useful to describe cognition as if it is information processing, there is no need to assume that it really is that." (Hollnagel, 1993, p.148). Based on a widely accepted model of human information processing (Wickens, 1984), CAMEO is composed of six "functional modules" (Figure-1); perception & recognition (P&R), decision-making (DM), action (ACT), working memory (WM), long-term memory (LTM), and attention resource controller (ARC). P&R has two "task modules;" auditory (P&R-A) and visual (P&R-V). P&R-A and P&R-V receives intruding auditory signals (e.g. alerting sounds) and visual signals (e.g. flushing alarm lights), respectively. In addition, PR&-V collects visual signals necessary for reasoning by DM. DM has four task modules; vigilance (DM-V), diagnosis (DM-D), response selection (DM-RS), and confirmatory evaluation (DM-CE). DM-V serves three functions; the acknowledgment of cued information (e.g. flushing alarm lights), routine vigilance, and the evaluation of feedback of control actions. DM-D identifies a failure mode. DM-RS selects responses to be implemented by ACT. DM-CE evaluates the conclusion of DM-D or DM-RS. ACT selects a set of actions corresponding to a response selected by DM-RS and implements the actions as specified. WM is a kind of pad shared by other functional modules (except LTM) for recording or communicating information. LTM stores sets of knowledge used by DM and ACT. It is this knowledge which task designers evaluate with CAMEO/TAT. ARC supplies resources to P&R, DM, and ACT.

Reasoning Processes

In every "cognitive unit cycle," ARC supplies resources to P&R, DM, and ACT (Figure-2). The amount of resources supplied to each module is determined on demand. All the modules may receive sufficient resources when the total amount of resources is larger than the sum of resources demanded. When the total amount of resources is not enough, there will be a module or modules which cannot receive sufficient resources. Resources are supplied in the order specified by "resource allocation priorities rules." Hence, less prioritized modules will not receive sufficient resources. The total amount of resources may be influenced by stress. It is reduced when no meaningful conclusion is obtained by DM for a considerable period of time. In the following several paragraphs, how each module behaves is explained.

P&R-A or P&R-V receives and stores cue signals in a buffer. Then it sends a cue to DM which carries information on a relevant task and subjects together with their scores. Here, scores refer to the importance of the relevant task and subjects. Since more than one cue signal may be received by P&R-A and P&R-V, DM may receive many cues in a cognitive unit cycle.

DM uses two task switching mechanisms to select the next task to be attended; opportunistic task switching mechanism and temporary task switching mechanism. An assumption that underlies these task switching mechanisms is that actual task sequences are dynamically determined through interactions between

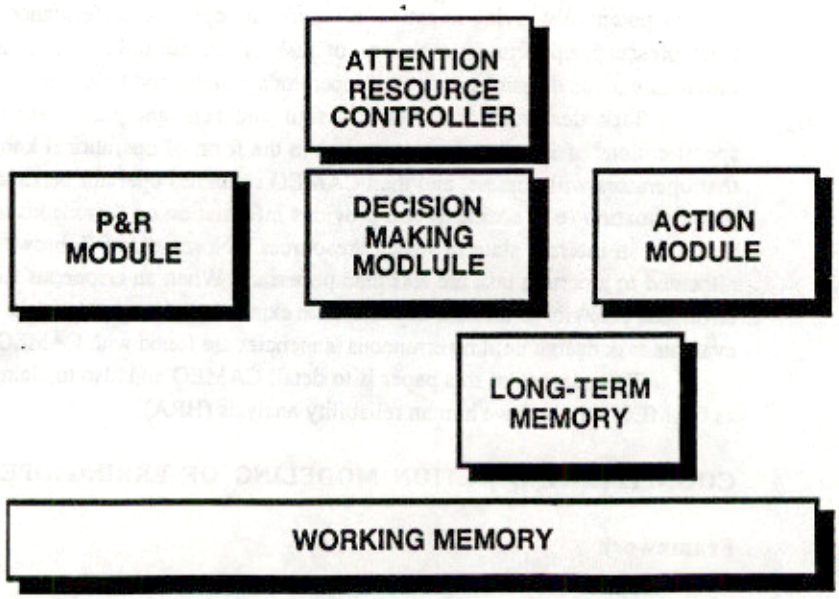


Figure-1. The framework of CAMEO is based on human information processing model.

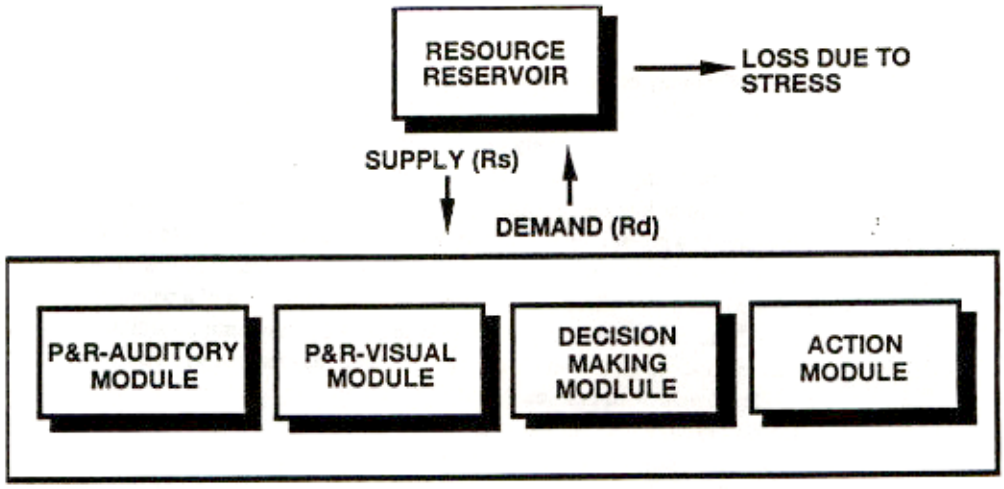


Figure-2. ARC supplies resources to P&R, DM, and ACT in each cognitive unit cycle. The amount of resources supplied to each module is determined on demand. Resources are supplied in the order specified by "resource allocation priorities rules." The total amount of resources is affected by stress.

functional requirements/constraints and operator's characteristics (e.g. job knowledge, policies, traits), rather than by formally described task structure alone. The process is assumed to be inherently interactive and opportunistic. The initial selection is made by the opportunistic task switching mechanism. It calculates a "transition priority score" for each task and selects one with the highest score as "primary task." Transition priority scores are calculated in terms of (1) scores from P&R-A or P&R-V, (2) scores from a decision made in DM itself in a previous cognitive unit cycle, and (3) scores reflecting operator's policies. Depending on situations, the first two scores may change dynamically. On the other hand the last one is assumed to be a fixed value. When a primary task is selected, the temporary task switching mechanism may be activated randomly. When it is activated, the primary task is replaced with one having the second highest score (i.e. "temporary task"). How frequent this activation may happen depends on the assumption about operator's traits. The next step is the selection of a "subject" which refers to a specific concept or task to be evaluated (e.g. LOCA, Confirm Start-up of Pump X). A subject having the highest score in the primary or temporary task is selected as one which is of current interest. The final step is the evaluation of subject which starts with the selection of a knowledge-strategy file specifying a knowledge set and a strategy. Two types of strategies are available; pattern matching and good-bad mapping (Rasmussen, 1986). It is assumed that the latter requires more resources. The initial selection is done by default. If no conclusion was obtained in previous cognitive unit cycles with the default selection, a different file may be selected. Usually, one associated with a more elaborate knowledge set or strategy is selected. A selected knowledge set is evaluated under the control of the associated strategy. Information necessary for the evaluation is collected through P&R-V. When a conclusion is obtained, a cue is generated which will be evaluated in a subsequent cognitive unit cycle in the calculation of the transition priority score. Note, if no conclusion is obtained by the temporary task, the primary task will automatically resume as a primary task in the next cognitive unit cycle.

When ACT receives a cue from DM-RS, it begins to search for the relevant action file in which a series of specific action steps to be implemented are specified. When the relevant action file is found and action steps implemented, ACT sends a cue to DM so that the feedback of control actions will be confirmed.

Error-inducing Mechanisms

When resources supplied by ARC are less than necessary, erroneous tendencies begin to appear. Each module shows specific forms of erroneous tendencies. (Note that errors may of course occur if operational knowledge or rules are incorrect even when sufficient resources are supplied. Also note that the task switching mechanisms may give considerable influences on the behavior of CAMEO. Depending on situations, however, they may work favorably or unfavorably.)

When resources are in short, less salient auditory or visual signals begin to receive less attention from P&R-A or P&R-V. This results in the decrease in the number of signals received by P&R-A or P&R-V

(Figure-3). This may cause important cue signals to be overlooked. It also happens to P&R-V that it cannot collect visual signals requested by DM. This results in the failure of updating corresponding information stored in WM, leading DM to rely on old pieces of information which could be incorrect.When a primary or temporary task being attended in DM cannot receives sufficient resources, it begins to search for a functionally similar knowledge-strategy file which it can afford with available resources. This may results in the utilization of a potentially inappropriate knowledge set, leading DM to make an incorrect judgment.

When ACT is not supplied with sufficient resources, it begins to search for a functionally similar action file which it can afford. This may results in the utilization of a potentially inappropriate knowledge set, leading ACT to implement incorrect action steps. ACT may omit conditional action steps which need to be implemented prior to goal-relevant action steps. This may happen when a similar action file cannot be found.

The mechanism that leads DM and ACT to search for an "affordable" solution overlaps with ideas that Reason (1990) calls "similarity matching" and "frequency gambling." It also accounts for a commonly observable tendency to adopt incomplete but functionally sufficient solutions as a "quick cure" when overloaded. The omission or violation of conditional steps is also a tendency that can be seen everywhere. It should be noted that operators are not free from these ubiquitous tendencies (Fujita et al., 1993b).

Simulations

Using a simple hypothetical water supply plant (Figure-4), how CAMEO interacts with simulated accidents were studied. It was confirmed that a variety of erroneous tendencies may appear as designed when resources supplied are less than demanded. The following is an example:

> "Because of leakage taking place at its outlet, Pump #1 must be isolated. Procedures specify that Tank #1 should continue to be used. This means to start Pump #2 after opening Valves #4 and #5. Though DM-RS successfully selected the correct response, ACT selected a wrong action file which lead ACT to open Valves #3 and #4. These actions are used in 'Pump Rotation' operation which is periodically done during normal operations."

This example demonstrates a possibility that rarely used unfamiliar actions may be replaced with frequently used familiar actions when resources are not enough. Depending on situations, however, erroneous tendencies may or may not manifest themselves as errors. For instance, the actions taken in this example, though deviated from procedures, does not lead the plant to a critical state, because the primary system continues to be cooled by water from Tank #2. Therefore, they are erroneous behavior, but not errors. It was also interesting to find that policies (i.e. resource allocation priorities rules, task selection policy) make considerable differences in the performance even when the amount of resources supplied are the same.

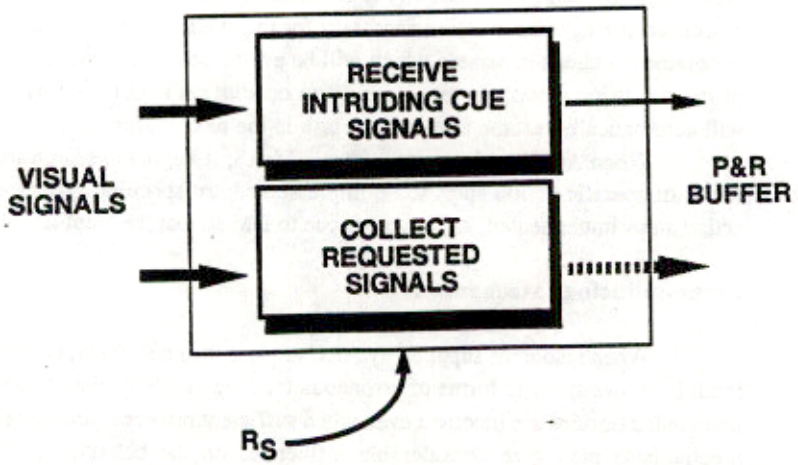


Figure-3. Error-inducing mechanism of P&R-V task module. When resources are not enough, less salient visual signals will not intrude. It also happens that P&R-V cannot collect signals which DM demanded to update the corresponding information.

HRA USING SIMULATED HUMAN MODEL

CAMEO takes a position in which it claims that the essence of human performance lies in the interactions between human and environments or situations in which he or she is embedded. It argues that human and situations cannot be treated in isolation. Persons may behave quite differently in the same situation even when they are assigned to the same tasks. The idea that humans are a mere component which executes assigned tasks only in a way specified must be abandoned: Humans are driven by concept or knowledge which is believed to have personalized structures or characteristics: Humans have their own strategies which are normally tuned to experiences or in accordance with some criteria which often involve their policies: Humans have their own traits.

The interactions between human and situations are believed to be inherently contextual, especially in environments such as process control environments. Therefore, it is necessary to look into "context" so that erroneous behavior can be understood. Human models such as CAMEO feature the ability to simulate these dynamic context-driven interactions. Human behavior are determined in an opportunistic manner through interactions among situations, knowledge, strategies, policies, and traits under the constraints of resource capacity. This approach allows us to explain why operators behaved that way, or predict dominant forms of erroneous tendencies which are likely to occur in given situations. No doubt, this is an invaluable source of information with which meaningful countermeasures can be developed.

Given that the ultimate purpose of HRA in general is to identify meaningful countermeasures against human factors deficiencies, in-depth understanding of erroneous behavior is a prerequisite. It is, however, often criticized that the present HRA fails to meet this crucial requirement. Hollnagel (1993) argues that the present HRA is not satisfactory because they are based on an "atomistic approach," and claims the need for the "whole approach." A conclusion that can naturally be drawn from this line of arguments and discussions made above is that simulated human models such as CAMEO can improve HRA. However, there is a fundamental question. As Hollnagel (1991) argues that human may not be quantifiable. It is technically possible to extend CAMEO so that it can calculate frequencies with which particular forms if erroneous tendencies may occur. This extended CAMEO may possibly be useful for HRA as a part of probabilistic safety assessment (PSA) in a sense that it provides "numbers" that can be mathematically handled.

CONCLUSIONS

CAMEO/TAT is a tool designed to be used by task designers for qualitatively evaluating cognitive difficulties associated with tasks designed by them. It incorporates a simulated human model called CAMEO.

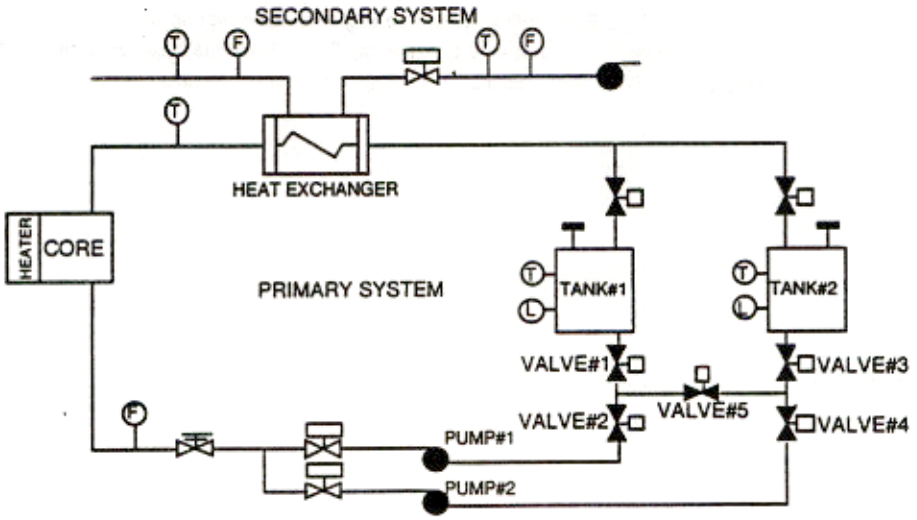


Figure-4. Schematic of a water supply plant used in simulation studies. The water in the primary side of the plant is circulated by one of two pumps which are occasionally changed for maintenance and management purposes (i.e. Pump Rotation). The heat generated in a core is transferred to the water in the secondary side of the plant through a heat exchanger.

Using a resource metaphor, CAMEO can identify erroneous tendencies by dynamically simulating context-driven interactions between human and situations. It is technically possible to extend CAMEO so that it can calculate frequencies with which particular forms of erroneous tendencies will occur. This extended CAMEO may serve as a better tool for HRA. However, the true merit of using a simulated human model such as CAMEO is that, regardless of whether it can provide quantitative estimates of human reliability or not, it allows us to look into "context." This is an indispensable step towards better understanding of cognitive difficulties that operators are likely to experience. If this step receives proper attention, HRA will contribute to find meaningful countermeasures, either as a part of PSA or as an independent evaluation.

ACKNOWLEDGMENT

The authors would like to thank the Ministry of International Trade and Industry (MITI) for sponsoring a series of projects in which CAMEO and CAMEO/TAT were developed (1990-1992). Thanks are also given to the members of the NUPEC/IHF Steering Committee for their technical supervision.

REFERENCES

Fujita, Y, 1992, Human reliability analysis: a human point of view, Reliability Engineering and System Safety, 38:71.

Fujita, Y., Yanagisawa, I., Nakata, k., Itoh, J., Yamane, N., Kubota, R., and Tani, M., 1993a, Modeling operator with task analysis in mind, in: "Proceedings of Topical Meeting on Nuclear Plant Instrumentation, Control and Man-Machine Interface Technologies," American Nuclear Society, Illinois.

Fujita, Y., Ishigaki, N., and Takaura, K., 1993b, Data, keyholes for the hidden world of operator characteristics, in preparation.

Gopher, D., and Donchin, E., 1986, Workload - an examination of the concept, in: "Handbook of Perception and Human Performance," K. Boff, L. Kaufman and J. Thomas, eds., Wiley, New York.

Hollnagel, E., 1991, What is man that he can be expressed by numbers? in: "Proceedings of the International Conference on Probabilistic Safety Assessment and Management (PSAM), " Elsevier Science Publishing, New York.

Hollnagel, E., 1993, "Human Reliability Analysis: Context and Control," Academic Press, London.Norman, D., and Bobrow, D., 1975, On data-limited and resource-limited processing, Journal of Cognitive Psychology, 7:44.

Rasmussen, J., 1986, "Information Processing and Human Machine Interaction," North-Holland, New York.

Reason, J., 1990, "Human Error," Cambridge University Press, New York.

Roth, E., and Mumaw, R. J., 1993, Cognitive engineering: issues in user-centered system design, in: "Encyclopedia of Software Engineering," J. J. Marciniak, ed., John Wiley & Sons, in press.

Wickens, C. D., 1984, "Engineering Psychology and Human Performance," Charles E. Merill Publishing.

061 Transportation Risk (I)

Chair: T. Altenbach, LLNL

The Cost of Human Error Intervention
C.T. Bennett, W.W. Banks, E.D. Jones (LLNL)

Application of the DOE and NRC Nuclear Safety Policy to Transportation Operations
C.Y. Kimura (LLNL), G.M. Sandquist (U. Utah)

Control of Risk from Explosives Stores Within Mod (UK)
J. Connor, D. Hewkin (Ministry of Defence, UK)

THE COST OF HUMAN ERROR INTERVENTION

C. Thomas Bennett
William W. Banks
Edwin D. Jones

Human Factors Engineering & Risk Assessment
Fission Energy & Systems Safety Program
Lawrence Livermore National Laboratory
University of California & Department of Energy
P. O. Box 808
Livermore, CA 94551

BACKGROUND

DOE has directed that cost-benefit analyses be conducted as part of the review process for all new DOE orders. This new policy will have the affect of ensuring that DOE analysts can justify the implementation costs of the orders that they develop.

We would like to argue that a cost-benefit analysis is merely one phase of a complete risk management program—one that would more than likely start with a probabilistic risk assessment. The safety community defines risk as the probability of failure times the severity of consequence. An engineering definition of failure can be considered in terms of physical performance, as in mean-time-between-failure; or, it can be thought of in terms of human performance, as in probability of human error.

The severity of consequence of a failure can be measured along any one of a number of dimensions—economic, political, or social. Clearly, an analysis along one dimension cannot be directly compared to another but, a set of cost-benefit analyses, based on a series of cost-dimensions, can be extremely useful to managers who must prioritize their resources.

Over the last two years, DOE has been developing a series of human factors orders, directed at lowering the probability of human error—or at least changing the distribution of those errors. The following discussion presents a series of cost-benefit analyses using historical events in the nuclear industry. However, we would first like to discuss some of the analytic cautions that must be considered when we deal with human error.

RADIOBIOLOGICAL COST-BENEFIT OPTIMIZATION

Approaches to Optimization

There have been sophisticated approaches to radiobiological, cost-benefit, analytic techniques developed for management guidance (Dunster, 1989; Gonzalez, 1983). One of the simplest forms of these techniques is shown in the following equation:

$$B=V-(P+X+Y) \quad (1)$$

where B is the net benefit of the introduction of a practice;
V is the gross benefit of the introduction of such practice;
P is the basic cost, excluding the radiation protection;
X is the cost of achieving a given level of radiation protection;
Y is the cost of the detriment of achieving the given level of protection.

There are several other more sophisticated variants of the equation presented above. However, the goal of all of these is to present information such as that shown in Figure 1.1.

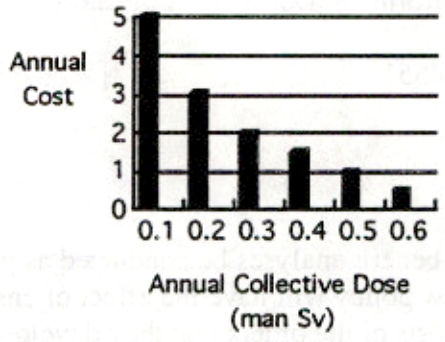


Figure 1.1. A representation of a cost-benefit profile of different design alternatives (indicated by each bar) providing different levels of radiation protection. Managers can choose cutoffs based on cost or annual collective dose. (Man Sv = radiation dose in sieverts.)

In Figure 1.1, data is provided concerning the cost of radiological protection of several hypothetical design alternatives. A management team could then make a decision based on either the cost of the different alternatives, or on the level of radiological protection that they believe should be achieved.

Level of Precision in Optimization

All the techniques discussed and alluded to in the previous section will provide exceedingly fine levels of analytic precision. There is, however, a more simplified procedure that we would like to present.

Using an extremely conservative technique, we will compare only the cost of a human error intervention program to that of constructing a facility and the revenues lost if that facility were to be destroyed by human error. In the examples that we will present, we will not factor into our equations the cost of human lives lost. We believe that because cost-

estimates of human lives are arbitrary, such estimates would act merely as mathematical constants and would not provide additional information for high-level management decisions.

In the following section, we will present a brief background on how human error should be treated statistically and brief examples of cost-benefit analyses of human error intervention programs.

THE COSTS AND BENEFITS OF HUMAN ERROR INTERVENTION

Human Error as a Random Variable

Classical human risk analysis usually considers human error as a random variable. There are, at present, too many unknown variables influencing human behavior to determine accurately what a person might do in a given setting. This uncertainty forces us to consider human error in probabilistic terms.

The definition of a random event is one that can have, *a priori*, multiple outcomes. If a process can have only one outcome, it is thought of as deterministic. Taking this argument one step further because of the uncertainty associated with human error, analysts must regard it as a random event.

Once consensus is reached that human error is a random event, or more properly, a series of random events then the nature of this stochastic process must be described. That is, the probability density functions must be specified.

We are stating, in essence, that we do not know the cause of human error. In one sense, the assertion that human error is random will lower the "comfort level" of design teams and their program managers.

On the other hand, treating human error in a probabilistic sense allows us to bring to the problem powerful statistical tools if the probability density functions are known. These statistical tools are the same mathematical techniques used in standard engineering design to reduce the uncertainty associated with mechanical failure.

As scientists, we are not quite prepared to give up the notion that human behavior is deterministic. However, a probabilistic view of human behavior provides us with the most sound analytic basis for providing management guidance.

Human Error Intervention

There is ample scientific evidence showing that much can be done to decrease the probability of human error. A known set of influencing variables affect the probability of human error. Importantly, these variables can be brought under the control of design teams and program managers.

Figure 1.2 highlights the interactive affects of operator, equipment, training, procedural, management, and environmental variables. A slight shift in the value of any of these variables can produce dramatic behavioral changes that are followed by increased human error rates.

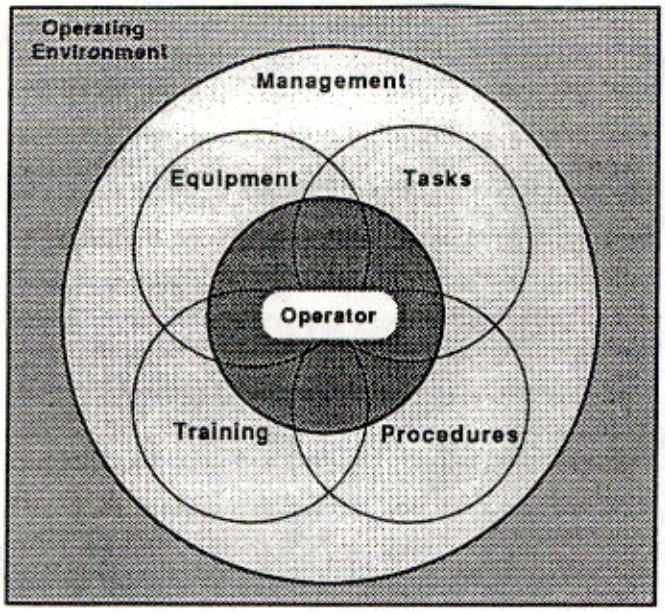


Figure 1.2. This Venn diagram shows the relationships among various factors that influence operator performance. Each must be considered during the design process.

From a probabilistic view, we can make reliable, predictive estimates of human error if the multivariate set of performance influences are well understood. This notion is supported by the following quotation:

> Some previous analyses of the Three Mile Island accident have been read to attribute it to "operator error." We reject this conclusion as being incomplete.
>
> While there is no question that operators erred... There were a number of important factors not within the operators' control... These include inadequate training, poor procedures, lack of management, misleading instrumentation, plant deficiencies, and poor control room design.
>
> (Rogovin and Frampton, 1980)

The High Cost of Human Error

The costs associated with finding and correcting human errors may be one of the largest hidden costs in the DOE budget. This has not been extensively studied by DOE yet, but an example from another industry demonstrates this point.

Prior to the forced corporate breakup, AT&T found that it was spending $100 million dollars a year locating and correcting human errors (Bailey, 1983). These costs were considered to be less than the costs that would have been incurred if these errors were allowed to be shipped to the public. Other industries have documented similar costs resulting from human error.

In the nuclear industry, we now know that the Chernobyl nuclear accident as well as the Three Mile Island incident were caused by interactive faults in equipment, procedures, training, and management. This interaction is what led directly to the series of human errors that in turn, propagated into events of severe economic, political, environmental, and social consequences.

THE COST OF FAILURE TO PROMOTE HUMAN ERROR INTERVENTION

The Three Mile Island Incident

It has been reported that Metropolitan Edison had rushed the Three Mile Island nuclear reactor into service by the end of 1978 in order to save $40 million dollars in taxes (Stephans, 1980; Medvedev, 1991). Three months later, a design-driven human error resulted in the worst commercial nuclear disaster in U.S. history.

In 1987, *US News & World Report* estimated that the direct costs of the Three Mile Island accident would exceed $3 billion. This estimate did not even consider the lost revenue which is approximately $365 million per year for every year TMI-2 does not produce power. The total loss, $3 billion plus $365 million over 40 years, will total approximately $15 billion.

DOE analysts have computed that the normal cost of a comprehensive human factors engineering design program would be no more than 1% of the total cost of a facility. If the Three Mile Island nuclear facility had cost $200,000,000 to build, the human error intervention program would have cost only $2 million.

The simple monetary cost benefit of such a human factors engineering program at Three Mile Island would have been 7,500:1.

Brookhaven National Laboratory

At BNL, a water purification system was operated by an individual who was misled by a control panel layout. In fact, the panel design violated several existing HFE standards.

In this event that occurred just a few years ago, an indicator light that correctly signaled the position of a valve control was uncorrelated with its function. During one shift, an operator observed what he thought was an overfill indication. In fact, the indication had nothing to do with an overflow condition. In an attempt to correct what he believed to be the problem, he allowed the tank to overflow. Due to its location on the roof, the overfilled tank caused the building to collapse partially. The repair costs totaled over $2 million.

It was estimated that a ten dollar label would have significantly reduced the probability of this two million dollar accident. By one simplistic analysis, the cost-benefit ratio of the human error intervention act to correct the problem would be 200,000:1.

The Chernobyl Accident

On April 26, 1989, after what turned out to be a series of miscalculations, an experiment on the graphite-moderated reactor went suddenly out of control. The result was the world's worst nuclear reactor disaster costing perhaps trillions of dollars around the world, and resulting in the evacuation of millions of people. The legacy of that experiment has been the documented increase in radiation-related illnesses.

There is little doubt that a chain of human errors—faulty procedures, followed by a series of misjudgments—was the cause of the accident (Medvedev, 1991). We will never really know what the costs of these errors will be but, if we take an extremely conservative estimate of $200 million for the cost of the facility, and $1 trillion for the aftermath, the cost-benefit ratio of a $2 million dollar human intervention program would be 50,000:1.

When the costs of such devastation are so great, it is clear that sophisticated, precise optimization algorithms are not needed to argue for human error intervention programs. We do not, however, want to leave the impression that we do not believe in the value of precision in cost-benefit optimization programs.

Such analytic tools are useful in ranking different radiological protection procedures once the decision to institute a program has been made. However, precision is not necessary in making the high-level, go-no-go decisions.

SUMMARY OF HUMAN ERROR INTERVENTION COST-BENEFIT ANALYSES

The two most costly nuclear reactor disasters were the result of human error. Almost seventy percent of the events documented in the Unusual Occurrence Report database have been traced to human error (Waters, 1993).

The analysis of human errors can be conducted using probabilistic-based methods. The error predictions that these analytic techniques produce are based on verified, empirically oriented, human factors engineering methods.

DOE analyses have shown that standard human error intervention programs could average 1% of the cost of the facility. Based on these estimates, even the most sophisticated of human error intervention programs would provide significant cost-reduction.

REFERENCES

Dunster, H. J., 1989, "Optimization and decision-making in radiological protection. A report of Committee 4 of the International Commission on Radiological Protection," *Annals of the ICRP*, Vol. 20, No. 1.

Gonzalez, A. J., Ilari, O., and Snihs, J. O., 1983, "Cost-benefit analysis in optimization of radiation protection. A report of Committee 4 of the International Commission on Radiological Protection," *Annals of the ICRP*, Vol. 10, Nos. 2/3.

Medvedev, G., 1991, *The Truth About Chernobyl*, Basic Books: New York.

Rogovin, M., and Frampton, G. R., 1980, "Three Mile Island. A report to the Commissioners Nuclear Regulatory Commission Special Inquiry Group."

Stephans, M., 1980, *Three Mile Island*, Random House: New York.

Waters, R. N., 1993, Personal communication to C. T. Bennett.

Webb, G.A.M., 1989, "Optimization and decision-making in radiological protection. A report of Committee 4 of the International Commission on Radiological Protection, *Annals of the ICRP*, Vol. 20, No. 1.

APPLICATION OF THE DOE AND NRC NUCLEAR SAFETY POLICY TO TRANSPORTATION OPERATIONS*

Chris Y. Kimura[1] and Gary M. Sandquist[2]

[1]Fission Energy and Systems Safety Program
Lawrence Livermore National Laboratory
Livermore, CA 94551

[2]Nuclear Engineering
University of Utah
Salt Lake City, UT 84112-1107

INTRODUCTION

The transportation of Department of Energy (DOE) materials and components throughout the United States and the world is a major, ubiquitous activity essential to the mission of DOE. Historically, DOE transportation activities have been performed without major impact upon the public; and in comparison with the safety record of other potentially hazardous materials, the DOE safety record has been excellent. Indeed, the risk to public health and safety from packaging, shipping and receiving DOE transported materials has been remarkably minimal.

However, the rapid de-escalation of the Cold War and the resulting international commitments to reduction in nuclear defense have resulted in a significant increase of international and domestic shipments of DOE nuclear defense materials and components to meet these commitments for arms reduction and elimination and modification of the US nuclear weapons arsenal. Current DOE policy is that transportation activities will continue to be conducted and regulated so as to ensure public health and safety.

The radioactive and hazardous materials typically associated with the transportation of nuclear defense systems and their components include tritium, uranium isotopes, plutonium isotopes, beryllium, and other selected materials[1]. These materials serve as the potential radiological and hazardous material source term in any transportation accident in which a radiological and hazardous material release occurs. The specific physical and chemical characteristics exhibited by these radioactive and hazardous materials depend upon the particular component application and shipment scenario and can vary in mass from micrograms to kilograms.

The radionuclides (e.g., tritium and the isotopes of uranium and plutonium) and hazardous materials (e.g., beryllium) released from a transportation accident in the form of particulates or gases, may establish an airborne plume that disperses downwind and can affect the public in the path of the plume from:

- an external dose caused by ionizing radiation emitted from the radioactive materials in the airborne plume or on the ground or water surface

* This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract no. W-7405-Eng-48.

• internal dose from inhalation, ingestion, or dermal absorption of the radioactive or hazardous materials deposited upon the body surface.

TRANSPORTATION ACCIDENTS

Given a transportation accident with release of material from the shipping package, there is a variety of pathways by which these radioactive and hazardous materials can be transported through the environment to expose the public. The most important pathways are atmospheric and surface water movement.

Each of the transport mechanisms has its own unique physical process for material movement and different pathways by which humans are exposed. The atmospheric pathway is of particular concern because of the speed by which the radioactive and hazardous material can move through the environment and expose the public. Within minutes these materials can be transported from the damaged containment package to the atmosphere and then expose persons within the local vicinity both internally from inhalation and externally from the airborne plume and material deposited upon the ground[1].

To provide adequate security and safety for DOE defense related materials, surface transportation within the U.S. is conducted along pre-planned and approved routes in SSTs (Safe, Secure Trailers). Typically, highway movement of these materials is made along US Interstate highways and other primary roadways. Stopping and maintenance locations are identified along these routes to provide "safe havens" for the SSTs.

The principal factors which affect all highway accident rates and also possible SST transportation accident rates include:

- driver experience and performance
- traffic and roadway conditions
- packaging and transportation equipment performance.

Driver competence and equipment performance are essentially deterministic factors which are readily quantified and assessed by actual records, past experience, and appropriate testing. However, the modeling and projection of traffic and roadway conditions, resulting accidents and types, release mechanisms and actual transport pathways constitute stochastic process and are generally best treated by probabilistic risk assessments (PRA) methods[1]. The most significant aspect associated with a given highway transportation accident involving an SST transporting nuclear defense materials is the magnitude and distribution potential release of these radioactive or hazardous materials and the consequences associated with subsequent exposures to the public.

DOE AND NRC NUCLEAR SAFETY POLICY

On September 9, 1991, the U.S. Secretary of Energy, James Watkins, issued Notice SEN-35-91 Nuclear Safety Policy[2]. This Notice declared the U.S. Department of Energy's (DOE's) nuclear safety policy and established two quantitative safety goals to limit the health and safety risks to the public associated with the DOE's operations. This notice declared that:

"The risk to an average individual in the vicinity of a DOE nuclear facility for prompt fatalities that might result from accidents should not exceed one-tenth of one percent (0.1%) of the sum of prompt fatalities resulting from other accidents to which member of the population are generally exposed. For evaluation purposes, individuals are assumed to be located within one mile of the site boundary."

"The risk to the population in the area of a DOE nuclear Facility for cancer fatalities that might result from operations should not exceed one-tenth of one percent (0.1%) of the sum of all cancer fatality risks resulting from all other causes. For evaluation purposes, individuals are assumed to be located within 10 miles of the site boundary."

The safety goals specified in SEN-35-91 are difficult to directly apply to the transportation of nuclear materials, special assemblies, and components associated with the DOE Defense Programs because these safety goals were developed principally for fixed site nuclear facilities, not mobile transportation operations. Accidents involving transportation operations, unlike operations at fixed site facilities, cannot be characterized as occurring at a particular location. These accidents may occur at any point along the transportation route.

Since DOE SEN-35-91 safety goals are difficult to apply directly to DOE transportation operations, it is appropriate to review the safety goals developed by the U.S. Nuclear Regulatory Commission (NRC) for comparison with DOE policies. In 1983, the NRC issued their Policy Statement[3] on Safety Goals for the Operation of Nuclear Power Plants. Two quantitative safety objectives were proposed as follows:

"The risk to an average individual in the vicinity of a nuclear power plant of prompt fatalities that might result from reactor accidents should not exceed 0.1% of the sum of prompt fatality risks result from other accidents to which members of the U.S. population are generally exposed."

"The risk to the population in the area near a nuclear power plant of cancer fatalities that might result from nuclear power plant operations should not exceed 0.1% of the sum of cancer fatality risks resulting from all other causes."

Both the DOE and NRC relate their safety goals to a fixed percentage (0.1%) of the sum of the prompt fatality risk resulting from accidents and the sum of the cancer fatality risk resulting from all causes. For transportation operations, since the accident location cannot be fixed, the DOE and NRC safety goals for all nuclear transportation operations could be conservatively interpreted to be 0.1% of the sum of the cancer fatality risk resulting from all causes for the entire U.S population. Thus, it is assumed here that DOE policy will ensure that the risk to the public from prompt fatalities and latent cancer deaths from DOE transportation activities is at least 1000 times less that risk to the public from other typical industrial and societal activities.

APPLICATION OF NUCLEAR SAFETY POLICY TO TRANSPORTATION OPERATIONS

Radiation doses and hazardous material exposures to the public may result in specific health effects. These effects are usually classified as chronic effects and include carcinogenic, mutagenic and tetratogenic health effects (HEs) and arise from inhalation, ingestion, dermal contact and direct exposure to radioactive and hazardous materials. For convenience, both acute and chronic fatalities are referred to as HEs.

Radiation doses and hazardous material exposures may be converted to health effects using appropriate dose conversion factors. The data base generally employed for assessing health effects (including carcinogenic and non carcinogenic effects) from hazardous materials is IRIS[4] (Integrated Risk Information System) and HEAST[5] (Health Effects Assessment Summary Tables) for radionuclides. Both data bases are prepared and maintained by the U.S. EPA. Data used here which is extracted from HEAST for inhalation exposures are 1.6 HE/MBq for Pu-239 (typical of radioactive materials in DOE shipments) and 2.4 HE/(mg/cubic meter) for beryllium.

Table 1, using official data from Vital Statistics of the U.S. and Census data[6,7], provides the U.S. total resident population as of 1 July of the year shown, the total resident deaths from all causes, total accident deaths, total cancer deaths, and cardiovascular deaths from 1960 to 1989 (the latest year from which official data is published for the death categories shown in the table). Death data for the years 1990 and 1991 are extrapolated from the previous 5 years. From these data various evaluations and conclusions can be drawn regarding the consequences of DOE policy for DOE transportation of its radioactive and hazardous materials.

It is apparent from Table 1 that DOE policy regarding prompt fatalities from all DOE transportation would limit risk potential to a total of 97 accidental fatalities per year based upon 1991 data. This risk potential is 103 if accidental fatalities are based upon the mean value of accidental deaths occurring between 1960 and 1991. The DOE transportation accidental fatality rate would be about 0.5 deaths per million population.

The DOE policy regarding latent cancer deaths from all DOE transportation would result in a total annual cancer risk potential of about 512 fatalities based upon 1991 data. This risk potential is 381 if cancer fatalities are based upon the mean value of cancer deaths between 1960 and 1991. The DOE transportation cancer fatality rate would be about 2 deaths per million population.

DOE fatalities from cancer imply a U.S. population exposure limit of (based upon 0.02 HE per person-Sv) 100 microSv above background (i.e., 3.6 mSV) if the entire population were exposed to activities associated with DOE transportation.

The acute radiation exposure level for humans without medical attention is about 4.5 Sv received within a 24 hour period. Assuming such an acute exposure scenario for prompt fatalities arising from a transportation accident, then an exposure of about 440 Sv is required to the critical population group to exhibit 97 prompt fatalities.

CONCLUSIONS

For the assumption that Secretary of Energy's Nuclear Safety Policy[2] is applicable to the transportation of DOE materials, then various consequences of this policy can be observed from the data shown in the Table 2.

DOE transportation accidents which annually resulted in 97 prompt fatalities traceable to DOE transportation activities would not be acceptable to an informed public. Certainly acute releases of 440 Sv would be readily detected and widely reported and harshly condemned.

Latent cancer fatalities of 512 persons annually might be lost in natural background radiation variations if the releases occurred in a chronic, disperse manner since 100 microSv represents only a 3% increase in average US background. However, the disclosure of any accidental release and supporting radiation monitoring activities to evaluate the consequences would undoubtedly reveal such a background increase.

From the discussion above, we have concluded that the DOE and the NRC Nuclear Safety Policy cannot be applied to the area of transportation of radioactive materials. The risk limits given by the DOE and the NRC Nuclear Safety Policy while suitable for fixed sites, are far too liberal when applied to an operation where the potential accident site may be anywhere along a transportation route.

REFERENCES

1. G.M. Sandquist et al, Risk Assessment of Radioactive and Hazardous Materials in DOE Defense Package Transport Accidents, *Radioactive Materials Transportation*, Vol. 3, No 2/3 P 121-128 (1992)

2. Secretary of Energy Notice SEN-35-91, Nuclear Safety Policy, Office of the Secretary, U.S. Department of Energy, Washington, DC, September 9, 1991

3. U. S. Nuclear Regulatory Commission, "Safety Goal Development Program", Federal Register, Volume 48, pages 10772-10781, March 14, 1983.

4 Integrated Risk Information System(IRIS), PB91-591330, U.S. Environmental Protection Agency, 1992

5 Health Effects Assessment Summary Tables (HEAST), Annual FY-1992, PB92-921100, U.S. Environmental Protection Agency, 1992

6. Vital Statistics of the United States, U.S. National Center for Health Statistics, annual.

7. Current Population Reports, Series P-25, Numbers 311, 1045, 1083, U.S. Bureau of the Census.

TABLE 1. U.S. RESIDENT DEATH STATISTICS FOR RISK ASSESSMENT

Year	US population	% chg	US deaths	Death rate	Acci-dent	Accident rate	Cancer	Cancer rate	Cardio-vascular	Cardio rate
1960	179,979,000	1.67%	1,711,982	.00950	93,806	.000520	267,627	.00149	935,668	.00520
1961	182,992,000	1.67%	1,701,522	.00930	92,249	.000500	273,502	.00149	936,218	.00512
1962	185,771,000	1.52%	1,756,720	.00950	97,139	.000520	278,562	.00150	968,809	.00522
1963	188,483,000	1.46%	1,813,549	.00960	100,669	.000530	285,362	.00151	994,747	.00528
1964	191,141,000	1.41%	1,798,051	.00940	103,843	.000540	289,577	.00151	984,261	.00515
1965	193,526,000	1.25%	1,828,136	.00950	108,004	.000560	297,588	.00154	1,000,787	.00517
1966	195,576,000	1.06%	1,863,149	.00950	113,563	.000580	303,736	.00155	1,021,188	.00522
1967	197,457,000	0.96%	1,851,323	.00940	113,169	.000570	310,983	.00157	1,012,047	.00513
1968	199,399,000	0.98%	1,930,082	.00970	114,864	.000580	318,547	.00160	1,023,399	.00513
1969	201,385,000	1.00%	1,921,990	.00950	116,385	.000580	323,092	.00160	1,013,015	.00503
1970	203,984,000	1.29%	1,921,031	.00940	114,638	.000560	330,730	.00162	1,007,984	.00494
1971	206,827,000	1.39%	1,927,542	.00930	113,439	.000550	337,398	.00163	1,017,145	.00492
1972	209,284,000	1.19%	1,963,944	.00940	115,448	.000550	345,618	.00165	1,035,146	.00495
1973	211,357,000	0.99%	1,973,003	.00930	115,821	.000550	351,055	.00166	1,037,492	.00491
1974	213,342,000	0.94%	1,934,388	.00910	104,622	.000490	360,472	.00169	1,010,926	.00474
1975	215,465,000	1.00%	1,892,879	.00880	103,030	.000480	365,693	.00170	971,047	.00451
1976	217,563,000	0.97%	1,909,440	.00880	100,761	.000460	377,312	.00173	974,429	.00448
1977	219,760,000	1.01%	1,899,597	.00860	103,202	.000470	386,686	.00176	961,539	.00438
1978	222,095,000	1.06%	1,927,788	.00870	105,561	.000480	396,992	.00179	966,120	.00435
1979	224,567,000	1.11%	1,913,841	.00850	105,312	.000470	403,395	.00180	958,282	.00427
1980	227,220,000	1.18%	1,989,841	.00880	105,718	.000470	416,509	.00183	988,545	.00435
1981	229,457,000	0.98%	1,977,981	.00860	100,704	.000440	422,094	.00184	972,992	.00424
1982	231,669,000	0.96%	1,974,797	.00850	94,082	.000410	433,795	.00187	967,868	.00418
1983	233,806,000	0.92%	2,019,201	.00860	92,488	.000400	442,986	.00189	981,098	.00420
1984	235,847,000	0.87%	2,039,369	.00870	92,911	.000390	453,492	.00192	973,418	.00413
1985	237,950,000	0.89%	2,086,440	.00880	93,457	.000390	461,563	.00194	977,879	.00411
1986	240,162,000	0.93%	2,105,361	.00880	95,277	.000400	469,376	.00195	968,240	.00403
1987	242,321,000	0.90%	2,123,323	.00880	95,020	.000390	476,927	.00197	963,611	.00398
1988	244,534,000	0.91%	2,167,999	.00890	97,100	.000400	485,048	.00198	969,400	.00396
1989	246,820,000	0.93%	2,150,466	.00870	95,028	.000390	496,152	.00201	931,838	.00378
1990	249,415,000	1.05%	2,183,925	.00876	96,666	.000388	503,268	.00202	934,917	.00375
1991	252,177,000	1.11%	2,202,994	.00874	97,162	.000385	511,753	.00203	925,825	.00367
Mean Value	(1960-91)	0.011	1,951,927	.00905	102,848	.000481	380,528	.00174	980,809	.00458

Death data from "Current Population Reports", Series P-25, Nos 311,1045, 1083, US Bureau of Census
Population data from "Vital Statistics of the United States", US National Center for Health Statistics.

TABLE 2. DOE SAFETY POLICY APPLIED TO TRANSPORTATION ACCIDENTS

Category	Fatalities		Rate (HE/million persons)	Acute release	Chronic release
Prompt fatalities	97		0.5	440 Sv	- -
Latent cancer fatalities	512	2	510 MBq (Pu-239) 210 mg/m^3(Be)		100 microSv

Projected fatalities are based upon 1991 statistic data given in Table 1

CONTROL OF RISK FROM EXPLOSIVES STORES WITHIN MOD(UK)

J Connor[1] and D Hewkin[2]

[1]Chief Safety Officer
[2]Leader ESTC Risk Assessment Study
Ministry of Defence
London SW6 1TR

INTRODUCTION

MOD(UK) has a good safety record for the storage and transport of explosives stores, supported by Quantity-Distance regulations. However, it is recognised that the level of protection afforded by these regulations differs under various circumstances, and it is frequently not easy to quantify the level of protection achieved in terms that can be readily explained to the public or compared with standards elsewhere.

Increasingly, the Health and Safety Executive, as the body charged with regulating Health and Safety in the UK, is demanding evidence that risks are managed to protect employees and the general public. MOD(UK) is not excepted from these requirements and therefore needs, in visible and defensible terms, to:-

a. Demonstrate arrangements for risk control

b. Identify sources of risk and justify their tolerability.

c. Manage a policy of continuing reduction of risk in a structured manner.

---

[1]Dr J Connor, Chairman ESTC, MOD Aquila, Golf Road, Bromley, Kent. BR1 2JB (Telephone UK (0)81-467-2600 ext 6600 - fax UK (0)81-295-0366)

[2]Dr D Hewkin TL-RAST(ESTC) Room 2006 Empress State Building, Lillie Road LONDON SW6 1TR. (Telephone UK (0)71 824-2210 fax UK(0)71-824-3592).

d. Explain and justify its safety decisions in terms which permit comparison and defence in public when necessary.

Risk Assessment has been identified as a valuable management tool which assists MOD(UK) to achieve these objectives.

The purpose of this paper is to review the extent to which the use of Risk Assessment techniques is helping MOD(UK) to carry out its necessary explosives operations including storage and transport with the minimum of risk.

MANAGEMENT AND CONTROL OF RISK

The arrangements in place within MOD(UK) to control the transport and storage of explosives stem from the responsibilities placed on the Secretary of State by the 1875 Explosives Act, and from his Health and Safety Policy Statement which places responsibility for safety on line management within each of the active Services and MOD(UK)'s Procurement Executive (PE).

Each active Service and PE has appointed its own Chief Inspector of Explosives, who issues operating licences for situations which meet standards based on recommendations of the triservice Explosives Storage and Transport Committee (ESTC). The CIE may seek advice from ESTC and others before licensing less orthodox work. Operational line management is required to operate within the licensed limits, and regular single service safety inspections backed up by triservice audits of safety arrangements are carried out to monitor safety standards. The Chief Safety Officer for MOD is required to report annually on safety performance targets and effectiveness to the Health and Safety Steering Committee.

CIEs may find specific cases where the regulations are unduly restrictive for a particular site and set of operating conditions. In such cases CIEs may consider all the relevant information including the results of a risk assessment, and, if the circumstances justify it, decide to issue a revised licence which, in their view, generates tolerable risks but offers wider scope. This procedure is subject to management reporting and safety audit controls, but permits efficient and effective achievement of Service defence objectives.

Although a CIE can be confident that the risks generated at explosives sites he licenses and controls are tolerable, site management is better placed to ensure that risks are kept as low as reasonably practicable (ALARP). A prime target for site managers seeking to minimise risks will be to reduce the quantities of explosives handled, but the development of alternative and safer procedures is also encouraged. Safety training and monitoring of safety performance also makes a valuable contribution to the stimulation of staff interest and demonstrates a commitment to safety throughout the organisation.

With the assistance of a contractor, ESTC has developed a procedure which can be used to estimate maximum risks of fatality generated by MOD(UK) explosives storage activities. Preliminary assessments at several explosives storage sites are now complete and maximum individual risks at relevant locations (exposed sites) within and adjacent to a facility can be identified.

RISK IDENTIFICATION

In order to compare estimated risks generated by a number of MOD explosives sites which conform to existing MOD Q-D regulations we have estimated the risk to a continuously exposed person in the open at the minimum permitted distance for an inhabited building. The estimates we produce are maximum values based upon worst credible case assumptions, and vary with location both within site and between site. Four estimations have been completed for each of four sites and Tables 1 - 3 show the structure of the model and the range of risk of fatality for each site.

Table 1 The components of Frequency and Consequence which contribute to Individual Risk (IR) of fatal injury in the ESTC Risk Assessment Procedure. The example shows the contributions from one Potential Explosion Site (PES) to the IR inside a house at a single Exposed Site (ES).

Frequency (yr^{-1})		Consequence (Prob Fatality)	
Store	4.7e-5	Blast	1.9e-5
Processing	(Nil)	Building Debris	8.3e-5
External Hazards	6.6e-8	Weapon Fragments	3.0e-5
Mal.Damage + Transp't	6.5e-7	Thermal Radiation	(Nil)
Communication	4.8e-5		
Total : 9.6e-5 (yr^{-1})		Aggregate : 1.3e-4	
INDIVIDUAL RISK OF FATALITY		1.3e-8 per year	

The comparatively wide range of values of risk reported in Table 3 is not too surprising when it is recalled that the Q-D rules take no account of the predicted frequency of an initiation, which for our standard situations we estimate to range from 1 in 1,000 per storehouse per year, to one in 1,000,000. There is also a difference in fatality risk contribution when the cause of fatality changes from blast to weapon fragment, building debris or thermal radiation.

TOLERABILITY

MOD(UK), through ESTC, takes a considerable interest in the continuing debate on what level of risk might be considered tolerable under particular conditions. We support the view that rigid universal risk tolerability criteria are ill advised. Those

Table 2 Total Individual Risk at an ES from all Contributing PES's.

PES	IR Contribution	PES	IR Contribution
1	3.5e-7	6	5.7e-7
2	3.6e-7	7	7.5e-7
3	4.1e-7	8	1.1e-6
4	4.3e-7	9	1.2e-5
5	4.6e-7	10	2.3e-4

(all other contributions were below 3.5e-7)

INDIVIDUAL RISK OF FATALITY FOR ES IS 2.5e-4 PER YEAR

Table 3 Comparison of Individual Risk at minimum permitted Inhabited building Distance from Four Sites which meet Quantity Distance Regulations.

Site Location	Individual Risk per year (Continuous Exposure Inside Domestic Building)	
	Highest	Lowest
A	5.0e-6	4.5e-7
B	9.3e-6	3.2e-6
C	7.5e-5	2.1e-6
D	2.3e-5	2.1e-7

responsible for making judgements on the tolerability of a given situation need to take into account all relevant factors and conditions including the need for the work to be carried out; the perception of the risks by various sectors of the public; local factors, and the influence on risk of decisions to move part or all the task elsewhere.

It is our view that for MOD's explosives risks, the CIE responsible for licensing the site is best placed to exercise such judgements. However, in reaching his decision

the CIE needs guidance on, for example, standards which have been accepted for comparable activities elsewhere, and the level of protection afforded by some of MOD(UK)'s existing regulations. Advice from ESTC is intended to reflect currently held views on risk tolerability and risk perception beyond the explosives field[1]. Our conclusion, based upon 14 site assessments so far, is that MOD explosives storage sites which meet the existing Q-D Regulations are most unlikely to generate a risk which would be considered intolerable.

Although our ideas are not fully developed, we are also trying to quantify societal risks which are generated by MOD's explosives activity. We are interested in the extent to which MOD(UK)'s existing regulations protect people from transient risks such those which are experienced by persons travelling past a hazardous site. Figure 1 outlines a credible situation.

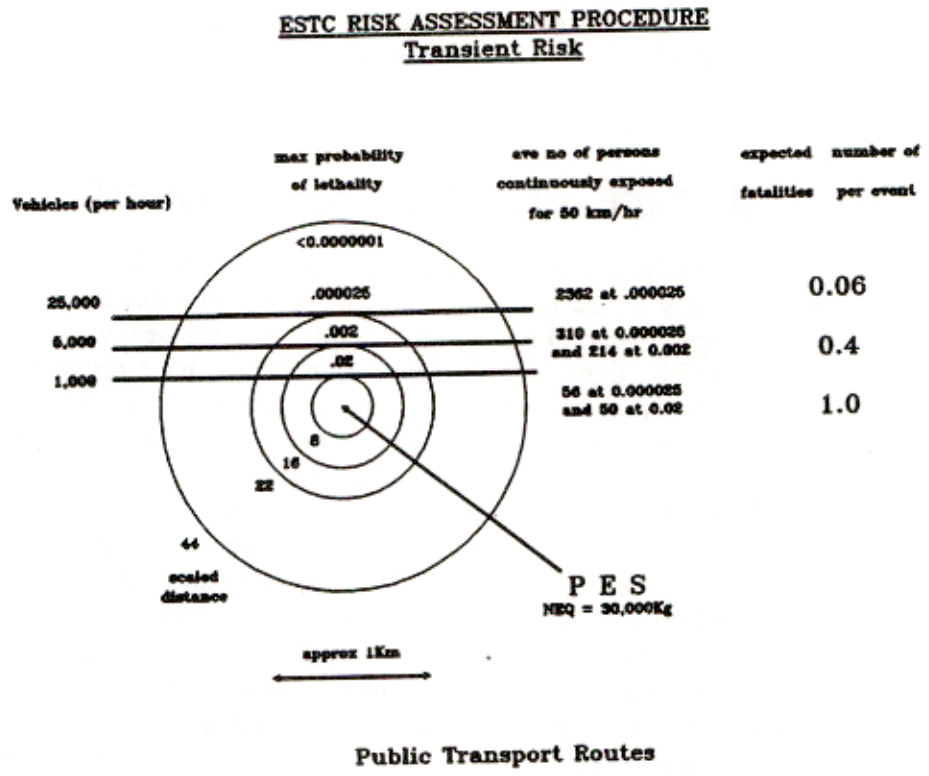


Figure 1 Given the indicated probability of fatality at each scaled distance and the vehicular flows shown, then the expected number of fatalities from an event will be << 10.

Unless structural collapse or heavy public transport is involved, the number of fatalities expected from a severe event is not normally sufficient to warrant comparison with societal risk situations. Our impression is[1] that intolerable situations do not arise until situations which involve more than 100 fatalities are predicted to occur at frequencies approaching 1 in 1,000 per year. We have not identified any situation where risks as high as this could arise from explosives storage.

RISK REDUCTION

The data output from these assessments (listed in Tables 1 and 2) can be best assimilated in the form of a spread sheet for each Exposed Site with a line for each Potential Explosion Site. The relative contributions to event frequency from components such as external or adjacent hazards, and to fatality probability from blast, building debris, weapon fragments or thermal radiation, make it comparatively straightforward to identify the PES's which contribute most to the risk and the key components within it. The method offers management guidance on which buildings and activities are contributing most to the risk, and what risk reductions can be obtained by altering the work plan. It is also possible to compare the risks generated at different sites, and to consider where most effort on improvement should lie. The procedure, although simplistic, necessitates considerable background knowledge and thought. It is not acceptable merely to allocate numbers to programme variables from a remote desk and then to rank projects entirely in terms of cost per predicted life saved per year. However, when it is completed using technically justifiable assumptions, a risk assessment provides a sound basis on which to formulate management decisions for risk reduction and to clarify the issues affecting a specific operation.

RISK COMPARISON AND THE PUBLIC PERCEPTION OF RISK

An essential feature of any method of risk control is the need to communicate findings in a way which is understood and which commands trust. This conference provides an opportunity for us to review how that might best be done. Risk analysis provides a useful structural framework on which to base discussions and to identify any areas of disagreement with a view to resolution. A report for the Royal Society[2] has identified the need to accept the differing levels of perceived risk considered tolerable for the same hazard by different groups and to reach an equitable balance. We see the solution to this conundrum as essential to the wider acceptance of risk assessment as a decision making tool. We believe our paper demonstrates that MOD(UK) has developed a capability which will enable the Department to meet its safety obligations into the next century. Risk Assessment will have an important role to play, and we have developed a relatively crude but effective and appropriate method for the analysis of the risks from explosives storage and transport. Now we must put theory into practice, and ensure that our motives can be clearly understood by those who are affected by our activities. Risk Assessment has become a valuable aid in the development of a trustworthy, transparent and defensible basis for achieving balanced decisions.

© British Crown copyright 1993/MOD : This report may be freely reproduced in whole, or in part, provided that the source is acknowledged

REFERENCES

1. Health and Safety Executive, "The Tolerability of Risk from Nuclear Power Stations" HMSO London. (Revised 1992) ISBN 0 11 886368 1

2. A Royal Society Study Group, "Risk: Analysis Perception and Management", The Royal Society, London. (1992) ISBN 0 85403 467 6

062 Environmental Risk Management--Waste

Chair: C. Peabody, USDOE

The Role of Multimedia, Multi-Pathway Exposure Assessments in Risk-Based Regulations at California Hazardous Wastes Sites

T.E. McKone (LLNL), E.G. Butler, J.J. Wong (Calif. Dept. Toxic Substances Control)

Risk Analysis for New Nuclear Waste Sites - Will It Generate Public Acceptance?

H. Inhaber (Westinghouse Savannah River)

THE ROLE OF MULTIMEDIA, MULTI-PATHWAY EXPOSURE ASSESSMENTS IN RISK-BASED REGULATIONS AT CALIFORNIA HAZARDOUS WASTES SITES

Thomas E. McKone[1], Edward G. Butler[2], and Jeffrey J. Wong[2]

[1]University of California
Lawrence Livermore National Laboratory
Livermore, CA 94551
[2]Office of Scientific Affairs
California Department of Toxic Substances Control
Sacramento CA 95812–0806

INTRODUCTION

The Department of Toxic Substances Control (DTSC), within the California Environmental Protection Agency, has the responsibility for managing the State's hazardous-waste program to protect public health and the environment. The Office of Scientific Affairs (OSA) within the DTSC provides scientific assistance in the areas of toxicology, risk and environmental assessment, training, and guidance to the regional offices within DTSC. Part of this assistance and guidance is the preparation of regulations, scientific standards, guidance documents, and recommended procedures for use by regional staff, local governmental agencies, or responsible parties and their contractors in the characterization and mitigation of hazardous-waste-substances-release sites.

CalTOX has been developed as a spreadsheet model to assist in health-risk assessments that address contaminated soils and the contamination of adjacent air, surface water, sediments, and ground water (McKone, 1993a, 1993b, 1993c). The model system includes a multimedia transport and transformation model, exposure scenario models, and efforts to quantify and reduce uncertainty in multimedia, multiple-pathway exposure models. This paper lists the objectives of a multimedia, total exposure model; provides an overview of the CalTOX model components; and describes the philosophy under which the model was developed.

OBJECTIVES OF THE CalTOX MULTIMEDIA, TOTAL EXPOSURE MODEL

In order to meet the risk management needs of the California Environmental Protection Agency, we believe it is not necessary to develop ever larger and more complex deterministic models. What is needed instead are scientifically defensive exposure models that incorporate uncertainty and population-distribution data. From a scientific perspective, the precision of environmental models is most often limited by the precision of the inputs. Under these conditions increasing the complexity of models does not necessarily increase the precision of model predictions and might often decrease them. The most simple view of the environment that could be posed in an environmental model is that of a world composed of a number of physical phases, each homogenous, well-mixed, and static. Our goal in developing the CalTOX model was to start with this view and then reluctantly, but as necessary, add compartments, heterogeneity, exposure pathways, and time dependence.

The objectives of CalTOX are to (1) provide realistic, although uncertain, ranges of potential dose estimates, (2) link exposure with multiple environmental media at hazardous waste sites and other contaminated sites, (3) provide a tool for risk-based priority setting on the basis of chemical properties and landscape properties, (4) provide a comprehensive approach for screening-level health-risk assessments, (5) provide a relatively simple and transparent model, (6) be compatible with uncertainty and sensitivity analyses, (7) provide a process for default risk assessments at small waste sites, (8) provide time-varying chemical inventories in soil layers, (9) be capable of addressing both continuous inputs and initial concentrations, and, importantly, (10) conserve contaminant mass.

Confronting Uncertainties

Regulatory toxicology and risk assessment often operate under the premise that, with sufficient funding, science and technology will provide an obvious and cost-effective solution to the problems of cleaning up hazardous waste sites. However, in reality there are many sources of uncertainty and variability in the process of human health-risk assessment. Many of these uncertainties and variabilities are not reducible. Effective hazardous waste clean-up policies are possible under conditions of uncertainty, but such policies must take the uncertainty into account. One often used method for addressing uncertainty in risk assessments is the compounding upper bound estimates in order to make decisions based on a highly conservative estimate of exposure and risk. Such an approach is contrary to accepted principles of decision making under uncertainty. The use of compounding upper bounds leaves the decision maker with no flexibility to address margins of error; to consider reducible versus irreducible uncertainty; to separate individual variability from true scientific uncertainty; or to consider benefits, costs, and comparable risks in the decision making process.

In any issue involving uncertainty, it is important to consider a variety of plausible hypotheses about the world; consider a variety of possible strategies for meeting our goals; favor actions that are robust to uncertainties; hedge; favor actions that are informative; probe and experiment; monitor results; update assessments and modify policy accordingly and favor actions that are reversible. In order to make CalTOX consistent with such an approach, it was

designed to have both sensitivity and uncertainty analyses incorporated directly into the model operation. Parameter values suggested for use in CalTOX are described in terms of mean values and a coefficient of variation in place of plausible upper values. Models are described in terms of the confidence intervals associated with model predictions. This is done to allow the users to produce more than a single number for an outcome such as a soil clean-up goal.

Chemical Classes

The are many classes of chemicals that must be addressed in environmental transport/transformation models, including organic chemicals, metals, inorganic chemicals, and radionuclides. These chemicals species can also be categorized according to the physical state in which they are introduced to the environment (gas, liquid, or solid), according to whether they dissociate in solution (ionic or nonionic) and according to the charge distribution on the molecule (polar or nonpolar). Mass potential (or fugacity) approach used in CalTOX and described below is most appropriate for nonionic, organic chemicals in a liquid or gaseous state. However, with modifications for condensation of solids on air particles, this approach can be made appropriate for solid-phase organic chemicals. Additional adjustments make possible the treatment of inorganic species, metals, and fully ionized organic species. Metals (such as mercury) and inorganic chemicals with a relatively large vapor pressure pose special problems, which are not yet addressed in the model. Special modeling problems also occur with mixed polarity, dissociating organic species, such as surfactants. The CalTOX model, in descending order of reliability, is capable of handling nonionic organic chemicals, radionuclides, fully dissociating organic and inorganic chemicals, and solid-phase metal species. With careful attention to inputs, the model can be used for partially dissociated organic and inorganic species. The model has not been designed to work with surfactants, inorganic chemicals species with high vapor-pressure-to-solubility ratios, and volatile metals such as mercury.

OVERVIEW OF THE COMPONENT MODELS OF CalTOX

There are two major model components within CalTOX and each can operate independently of the other. These two model components are the multimedia transport and transformation model, which is used to determine the dispersion of soil contaminants among soil, water, and air media and the human exposure model, which translates environmental media concentrations into estimates of human contact and potential dose. A detailed description of the multimedia transport and transformation model is provided by McKone (1993b) and a detailed description of the human-exposure model is by McKone (1993c). An identification of inputs to these models including ranges needed to represent uncertainty and variability are provided in a supplemental report describing the values and ranges of parameters to be used in the CalTOX model and issued by the California (DTSC).

The Multimedia Transport and Transformation Model

The transport and transformation component of CalTOX is a seven-compartment regional multimedia fugacity model that was developed for assessing the spread of contamination among the environmental media at a landscape unit that represents a toxic-substances-release site. It is designed to model a landscape on the order of 1 to 100 km^2 associated with a toxic-substances-release site. The CalTOX model relies on fugacity as a means of describing the mass potential of organic chemicals in two adjacent compartments. Fugacity models for organic chemicals were proposed and initially developed by Mackay (1979, 1991) and Mackay and Paterson (1981, 1982) and are now used widely in Canada, Europe, Japan, and the U.S. The fugacity approach is best suited to organic chemicals for which partitioning is related to chemical properties, such as vapor pressure, solubility, and the octanol-water partition coefficient. CalTOX has been designed to also handle inorganic contaminants such as radionuclides and metals.

The multimedia transport and transformation model in CalTOX is a dynamic model that can be used to assess time-varying concentrations of contaminants introduced initially to soil layers or for contaminants released continuously to air or water. This model assists the user in examining how chemical and landscape properties impact both the ultimate route and quantity of human contact. Using this model, we view the environment as a series of interacting compartments. The model allows the user to determine whether a substance will (a) remain or accumulate within the compartment of its origin, (b) be physically, chemically, or biologically transformed within the compartment of its origin (i.e., by hydrolysis, oxidation, etc.), or (c) be transported to another compartment by cross-media transfer that involves dispersion or advection (i.e., volatilization, precipitation, etc.).

The seven-compartment structure used in CalTOX is illustrated in Figure 1. The seven CalTOX compartments are (1) air, (2) ground-surface soil, (3) plants, (4) root-zone soil, (5) the vadose-zone soil below the root zone, (6) surface water, and (7) sediments. The air, surface water, ground-surface-soil, plants, and sediment compartments are assumed to be in quasi-steady state with the root-zone-soil, and vadose-zone-soil compartments. Contaminant inventories in the root-zone soil and vadose-soil zone are treated as time-varying state variables. Contaminant concentrations in ground water are based on the leachate from the vadose-zone soil.

It is important to note that, in contrast to many models used for assessing environmental fate, CalTOX imposes conservation of mass on the contaminated landscape unit. In addition, the model accounts systematically for gains and losses in each compartment and for the whole system in concert.

The Human Exposure Model

Human exposures to chemicals can result from contact with contaminated soils, water, air, and food as well as with drugs and consumer products. Exposures may be dominated by contacts with a single medium or may reflect concurrent contacts with multiple media. Assessment of human exposure to environmental contaminants requires translating environmental concentrations into quantitative estimates of the amount of chemical that contacts individuals within an exposed population. Potential dose, expressed as average daily dose, is the amount of material per unit of body weight per day

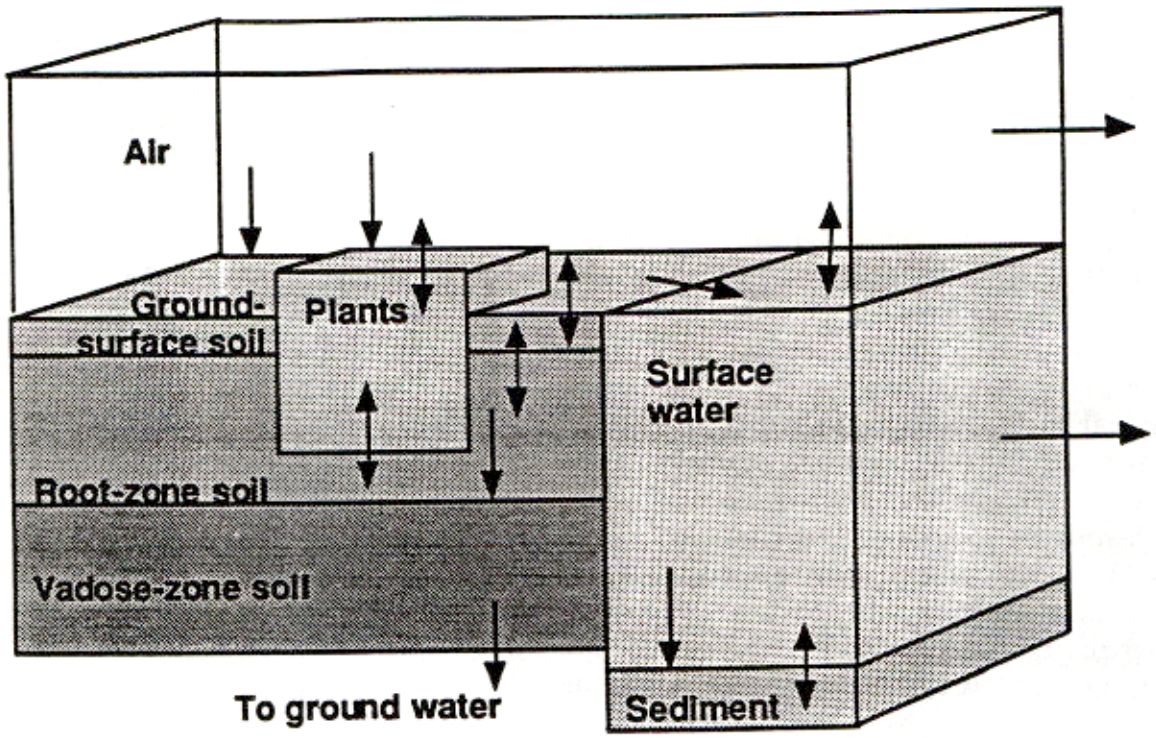


Figure 1. An illustration of mass-exchange processes modeled in the seven-compartment environmental transport and transformation model. (Ground water is not explicitly modeled in the system of equations but is used in the exposure calculations.)

(mg/kg-d) that crosses the mouth of an exposed individual by inhalation or ingestion or enters the outer layer of the skin through dermal contact. Dose models used in CalTOX are based on those described by the U.S. Environmental Protection Agency (EPA, 1989; 1992).

The nature and extent of multimedia exposures depends largely on two things—(1) human factors and (2) the concentrations of a chemical substance in contact media. Human factors include all behavioral, sociological, and physiological characteristics of an individual that directly or indirectly affect his or her contact with the substances of concern. Important factors in this regard are contact rates with food, air, water, soils, drugs, etc. Activity patterns, which are defined by an individual's allocation of time spent at different activities and locations, are also significant because they directly affect the magnitude of inhalation and dermal exposures to substances present in different indoor and outdoor environments.

The exposure analysis component of CalTOX begins with concentrations in ambient air (gas and particle phases), surface water, ground water, surface soil, and root-zone soil at a hazardous-waste-substances-release site. The exposure-assessment process consists of relating contaminant concentrations in these environmental media to contaminant concentrations in the media with which a human population has contact (personal air, tap water, foods, household dusts, soils, etc.). The average daily dose is the product of the exposure concentrations in these contact media and an intake or uptake factor that relates the concentrations to the distributions of potential dose within the population.

SUMMARY AND DISCUSSION

Ultimately, environmental samples are a set of static pictures used to characterize a dynamic world. Unless these "pictures" can be guided by an appropriate theoretical framework, they are of little value in decision making unless we have a very large set of "pictures". The goal of the CalTOX model is to maximize the amount of information obtained from each "picture". In such a system, the CalTOX model must serve as a repository for much of the current knowledge of environmental pollution and exposure processes. We envision that CalTOX can provide a "portfolio analysis" approach to human exposure at hazardous waste sites. That is, the exposure assessor should first develop a preliminary exposure matrix based on a portfolio of scenarios to match potential dose by route to contaminant-specific, multimedia dispersion in the environment. The result is a multimedia, multiple-pathway, and multiple-route assessment and gives indications of where it is most valuable to focus our resources to more fully characterize distributions of population exposure. This "value of information" approach is particularly important when there is a need to minimize the cost of providing exposure information without jeopardizing the precision required of this information to give the regulator confidence about the protection of public health.

ACKNOWLEDGMENTS

This work was performed under the auspices of the U.S. Department of Energy (DOE) through Lawrence Livermore National Laboratory (LLNL) under Contract W-7405-Eng-48. Funding was provided by the State of California Department of Toxic Substances Control (DTSC) through Contract Agreement 91-T0038.

REFERENCES

Mackay, D., 1979, Finding fugacity feasible *Environ. Sci. Technol.* 13, 1218–1223.

Mackay, D., 1991, "Multimedia Environmental Models, The Fugacity Approach," Lewis Publishers, Chelsea, MI.

Mackay, D., and Paterson, S., 1981 Calculating fugacity, *Environ. Sci. Technol.* 15:1006.

Mackay, D., and Paterson, S., 1982, Fugacity revisited, *Environ. Sci. Technol.* 16:654.

McKone, T.E., 1993a, "CalTOX, A Multimedia Total-Exposure Model for Hazardous-Wastes Sites Part I: Executive Summary," prepared for the State of California, Department Toxic Substances Control, Lawrence Livermore National Laboratory, Livermore, CA, UCRL-CR-111456PtI.

McKone, T.E., 1993b, "CalTOX, A Multimedia Total-Exposure Model for Hazardous-Wastes Sites Part II: The Dynamic Multimedia Transport and Transformation Model," prepared for the State of California, Department Toxic Substances Control, Lawrence Livermore National Laboratory, Livermore, CA, UCRL-CR-111456PtII.

McKone, T.E., 1993c, "CalTOX, A Multimedia Total-Exposure Model for Hazardous-Wastes Sites Part III: The Multiple-Pathway Exposure Model," prepared for the California Department Toxic Substances Control, Lawrence Livermore National Laboratory, Livermore, CA, UCRL-CR-111456PtIII.

U.S. Environmental Protection Agency (EPA), 1989, "Risk Assessment Guidance for Superfund, Volume I: Human Health Evaluation Manual (Part A)," Office of Emergency and Remedial Response, EPA/540/1-89/002.

U.S. Environmental Protection Agency (EPA), 1992, Guidelines for exposure assessment: Notice, *Federal Register* 57(104):22888.

RISK ANALYSIS FOR NEW NUCLEAR WASTE SITES - WILL IT GENERATE PUBLIC ACCEPTANCE?

Herbert Inhaber

Westinghouse Savannah River Co.
1991 S. Centennial Ave.
Aiken, South Carolina 29803

INTRODUCTION

Over the past decade, one of the major uses of risk analysis has been to estimate the hazards of existing or potential nuclear waste sites. Its ultimate purpose, aside from meeting regulatory requirements, is to assure local residents that the risks are small, if indeed they exist at all. A wide variety of risk communications techniques have been suggested by social scientists and others to explain the results to non-scientists. However, regardless of the specific techniques employed, the overall process often fails. Potential neighbors often become even more hostile than they were prior to the risk analysis process. As has often been pointed out by sociologists, the risk analysis itself is "captured" by facility opponents. They point out that uncertainties exist in the computations and that zero risk cannot be guaranteed for any of the project phases.

Does this mean that risk analysis should not be applied in finding new waste sites or in evaluating old ones? Few would claim that risk analysis is not necessary, but it clearly is not adequate for producing social acceptance. In the context of waste siting, risk analysis has to carry too heavy a burden. While few risk analysts would claim explicitly that the field is sufficient to solve siting problems, the claims are implicit in the process. The Yucca Mountain controversy over high-level wastes is a case in point. While millions have been spent on various risk analyses - taken in the broad sense - over the years, there is little evidence that many Nevadan minds have been swayed. Indeed, the subject is little mentioned in public debates, except by opponents who note that the risk analyses always find a non-zero risk.

It is clear that other tools, not replacing but in addition to risk analysis, must be used to reach a successful conclusion to the decades-long process. Only in this way can the nuclear industry be freed from the albatross of nuclear wastes, and go on to more productive activities.

A BRIEF HISTORY

Since the beginning of the nuclear era, it has been recognized that the wastes generated would pose a problem to public acceptance. It is sometimes believed that nuclear pioneers disregarded wastes, and that they assumed a future generation could deal with it. The record suggests otherwise. For example, David Lilienthal, the first chair of the former Atomic Energy Commission (AEC), said in a book published four decades ago that unless the problem was solved, the industry would gradually diminish in public support. In the

mid-50s, a sanitary engineer (Lieberman, 1955) from the AEC Reactor Development Division wrote,

> "One has only to consult the popular press to become acutely aware of the militant interest of the public in matters directly concerned with waste disposal and environmental sanitation..."

It is not clear when quantitative risk analysis for nuclear waste sites came into being. From the beginning, it was stressed by regulatory agencies that public health and safety would be protected in the siting and disposition process. By the mid and late 1970s, calculations were being performed, and some had a probabilistic aspect. Risk computations were performed around that time in the U.S. by the Environmental Protection Agency (Energy Daily, 1979); the AEC (Atomic Energy Commission, 1974), Science Applications Inc. (Erdmann et al., 1979), DOE (Department of Energy, 1978), and the American Physical Society (Hebel et al., 1978). Calculations were also performed in Sweden (Karn-Bransle-Sakerhet, 1978).

Inhaber (Inhaber, 1982) and others, in reviewing these studies, found that the waste risks, while non-zero, were a smaller component of total nuclear fuel cycle risk than any other aspect, such as mining or reactor operation. Since that time, many other reports have been issued, some in conjunction with the proposed characterization of a high-level waste site at Yucca Mountain in Nevada.

Waste risk studies have now been produced in many countries and with varying assumptions. None of them predict that the risk of nuclear wastes, when properly handled and disposed of either above-ground or in geological formations, would pose other than a minuscule threat to human health and safety.

DO RISK STUDIES HAVE ANY EFFECT ON THE PUBLIC?

Risk studies seem to have little effect on public attitudes, which is presumably the intended object of the studies. Surveys of the public in Nevada and elsewhere have indicated little knowledge of or even interest in these calculations.

Perhaps the most graphic example of this was seen during the Chicago convention of the American Nuclear Society in November 1992. A high official of the Department of Energy Office of Civilian Radioactive Waste Management (OCRWM) in Nevada was reviewing progress to date on characterization. He noted that all of the risk studies performed on behalf of his agency were available to the citizens of Nevada. Since OCRWM has made strenuous efforts in terms of public outreach, all of the reports noted in this paper and others issued since the 1970s would be supplied to anyone interested. Presumably at least some of the studies have been read by a number of Nevadans.

An individual in the audience then posed a question. Was the official aware of any Nevadan who had read one or more of these studies and then publicly stated that his or her opinion had been changed from anti-repository to pro-repository? The official was unable to supply the name of any such person.

This exchange suggests that while risk studies have been useful in meeting regulatory requirements, they have not, by themselves, generated much change in public opinion on these sites. The situation described above could be multiplied many times in experiences around the world.

THE PILING-ON OF RISK STUDIES

The risk studies situation is illustrated in Figure 1. In the first part, the dark sheets indicate the studies performed. Public opinion is generally unfavorable in spite of this work. Siting agencies then conclude that the solution is more studies, not a new approach. These studies are performed, and public opinion may shift slightly in favor of the proposed waste site. However, it is still a long way from overall approval.

Undaunted, the siting agency redoubles its efforts, as shown in the third part of Figure 1. More studies are produced, and public opinion is changed slightly. But the goal of general approval is still not reached.

Finally, the process collapses of its own weight, as shown in the fourth and last part of Figure 1. After years, or in some cases, decades, of delay, the process and the accumulated studies are abandoned. The process begins again with new risk studies.

NEW TOOLS TO SOLVE THE PROBLEM

Some decision-makers are aware of the seemingly endless cycle of studies, more studies, abandonment and restart. They have searched for another tool to solve the problem. Some, such as the Decision Research group in Oregon (Kunreuther et al., 1988), have proposed greater "trust" between risk analysts and the public. However, many of the proposed techniques to generate trust have been employed in the past, with few beneficial effects from the viewpoint of project sponsors.

For example, social scientists suggest that trust can be produced by having officials of a project readily accessible to any member of the public. This credo has been followed in the Yucca Mountain situation. The home phone number of the chief Department of Energy representative is listed in the phone book, available to any critic. Yet the level of opposition there remains strong, with all officials who have been elected state-wide for years being opposed to the Yucca Mountain project.

THE NEGOTIATOR APPROACH

Another approach, some distance from risk analysis itself, is that of a "negotiator". Community groups have often pointed out that the plethora of agencies concerned with siting have made it difficult for citizens to evaluate the facts. Congress, in the 1987 amendments to the Nuclear Waste Policy Act, provided for a federal negotiator for high-level wastes. David Leroy, former lieutenant governor of Idaho, was appointed to this position in 1990 (Nuclear News, 1990). His tenure lasted until June 1993.

In principle, a negotiator could reduce the conflicting signals often produced by the various bureaucracies of DOE, EPA, the Department of the Interior, and others. However, the authority given the negotiator by law is unclear; his relationship to the Secretary of Energy and the head of the civilian high-level nuclear waste office remains cloudy (Inhaber, 1991).

From the viewpoint of risk analysis, it is instructive that the negotiator made little if any use of the extensive risk studies sponsored by DOE, EPA and others. He may have brought them to the attention of the counties and Indian tribes to whom he wrote, but they apparently played a negligible part in the discussions he held.

COMPENSATION - THE ULTIMATE TOOL?

To date, the amount of nuclear industry concern has been much greater for risk analysis than any risk compensation to a potential community or state. The results suggest that the industry and the public might be better off if the proportion of concerns had been reversed.

Compensation has been a tool in siting, both in conjunction with and separate from risk analysis. Observers inside and outside the nuclear industry have said that communities should be fully compensated for the detriment caused by a waste site (Ackerman, 1977; Bacow and Milkey, 1982), but until recently no methodology for setting the appropriate amount has been devised. In consequence, compensation amounts have been selected arbitrarily, leading to rejection by affected communities.

For example, in the case of Yucca Mountain, the original amount of compensation proposed in a Congressional committee in 1987 was $100 million annually. By the time the bill emerged from Congress, the amount was whittled down to $10 or $20 million annually, depending on the stage of construction. Whatever the level of funding, it was chosen without any apparent input by Nevada or Nevadans.

In a more recent case in France (Nuclear News, 1991), a fee of about $10 million annually was to be paid to affected communities. However, as long as the fee amounts are

set by outsiders and not the affected community, there will be a tendency to reject the amounts proffered. This took place in Nevada, and is on the way to happening in France.

It is evident that compensation systems have been treated as if they are all equal. However, there are as many types of compensation systems as there are varieties of risk analysis. Just as risk analysis must be carefully constructed to yield valid results, so a compensation scheme must be devised to achieve the goals of fairness and equity. In many instances, these twin goals can be reached only with adequate community control over the entire process. The present siting system, with or without compensation, offers little or no control to communities accepting wastes. The implicit assumption is that allowing this control would result in the facility being shut down or not constructed at all.

A number of solutions to the above-noted problems have been suggested. Kunreuther and others at the University of Pennsylvania (Kunreuther, 1987) have proposed compensation-lottery systems, with the goals of ensuring equity and an open process. Others (Inhaber, 1991) have propounded auction-style systems, in which communities would, in effect, set the level of compensation they think adequate.

The reverse Dutch auction system of Inhaber is shown in the next two figures. Figure 2 shows the steps until the first (and probably only) bid is made; the latter traces subsequent steps. It is assumed that a state wishes to site a waste facility somewhere in its boundaries. Counties are relevant political jurisdictions. The principles can be applied to siting high-level waste, where the states would be the relevant political entities.

After announcement of risk and environmental ground rules, the siting authority gradually raises the bonus level. The graph shows increments of $10 million, but other amounts can be used. Since communities often do not have the technical expertise to evaluate the complex criteria, they will be allowed to hire consultants at no cost to them to perform this task. Those who do are under no obligation to bid.

The bonus level rises until a community decides to make the first bid. At this point, the process transfers to Figure 3. The proposed site is considered by regulatory agencies - EPA, NRC or state bodies. At the same time, most of the bonus (perhaps 2/3) is paid into a trust fund. Many citizens will believe that, regardless of the safeguards, their community will not receive the bonus, but the trust fund will reassure them. If the site exceeds the criteria (even by a small amount) the process is completed and the total bonus is paid. If it does not, the auction continues, starting at the previous level. Eventually, a second community will come forward. The process is completed with the minimum number of complaints of unfairness and favoritism.

The reverse Dutch auction is apparently the only system that allows for volunteerism and payment of the correct social cost to the affected community.

In the phrase used in geometrical proofs, risk analysis is regarded as necessary for siting, but not sufficient. In these new systems for siting, risk analysis is not abandoned but is no longer the sole driving force.

CONCLUSIONS

Ultimately, the approach of suggesting to citizens that they study and understand risk analyses may be fundamentally at odds with the American experience. As a professor of history (Diggins, 1987) has noted, the U.S. Constitution was based in part on the philosophy of David Hume, a Scottish philosopher. Hume argued that the aim of government should not be to exhort citizens to attain grace or virtue. In the case of risk analysis, this state would be achieved by understanding its complex mathematics and elaborate assumptions. Rather, Hume's philosophy proceeded under the assumption that man is a pleasure-seeking creature, in pursuit of material gain. The reverse Dutch auction system outlined here would be instantly recognizable to him.

REFERENCES

Bruce Ackerman, 1977, The jurisprudence of just compensation, *Env. Law* 7:510.

Atomic Energy Commission, 1974, "High-Level Radioactive Waste Management Alternatives", AEC, Washington.

Lawrence Bacow and Milkey, J., 1982, Overcoming local opposition to hazardous waste facilities: the Massachusetts Approach, *Harvard Envtl. Law Rev.* 6:265.
John Patrick Diggins, 1987, Science and the American experiment, *The Sciences* 27:28 (Nov./Dec.)
Department of Energy, 1978, "Environmental Aspects of Commercial Radioactive Waste Management", DOE, Washington.
Energy Daily, 1979, EPA study pinpoints the real villain in nuclear waste: technetium-99, 7:4 (Jan. 29).
R. C. Erdmann et al., 1979, "Status report on the EPRI Fuel Cycle Accident Risk Assessment", Science Applications Inc., Palo Alto., California.
L. C. Hebel et al., 1978, Report to the American Physical Society by the study group on nuclear fuel cycles and waste management, *Rev. Mod. Phys.* 50:S86.
H. Inhaber, 1982, "Energy Risk Assessment", Gordon and Breach, New York.
H. Inhaber, 1991, Can an economic approach solve the high-level nuclear waste problem?, *Risk: Issues Health Safety* 2:341.
Karn-Bransle-Sakerhet, 1978, "Handling of Spent Nuclear Fuel and Final Storage of Vitrified High Level Reprocessing Waste", AB Teleplan, Solna, Sweden.
Howard Kunreuther et al., 1987, A compensation mechanism for siting noxious facilities: theory and experimental design, *J. Envtl. Econ. and Mgmt.* 14:371.
Howard Kunreuther, William H. Desvousges and Slovic, Paul, 1988, Nevada's predicament, *Environ.* 30:17.
Lieberman, J.A., 1955, Disposal of radioactive wastes, *Civ. Eng.* 25:44.
Nuclear News, 1990, Negotiator nominee Leroy confirmed by Senate, 97 (Sept.).
Nuclear News, 1991, France considers bill to resume exploration, 34:101 (July).

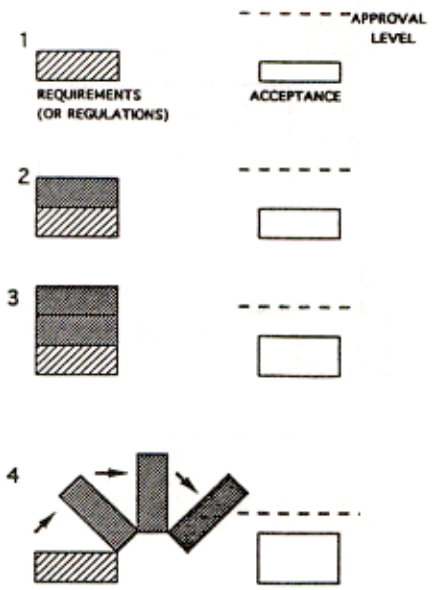


Figure 1. An analogy to the "overkill" of regulations designed to generate waste acceptance. The pile of requirements continually mounts. However, the amount of public acceptance never gets to the dotted approval level. Finally the requirements collapse of their own weight, and the process begins anew.

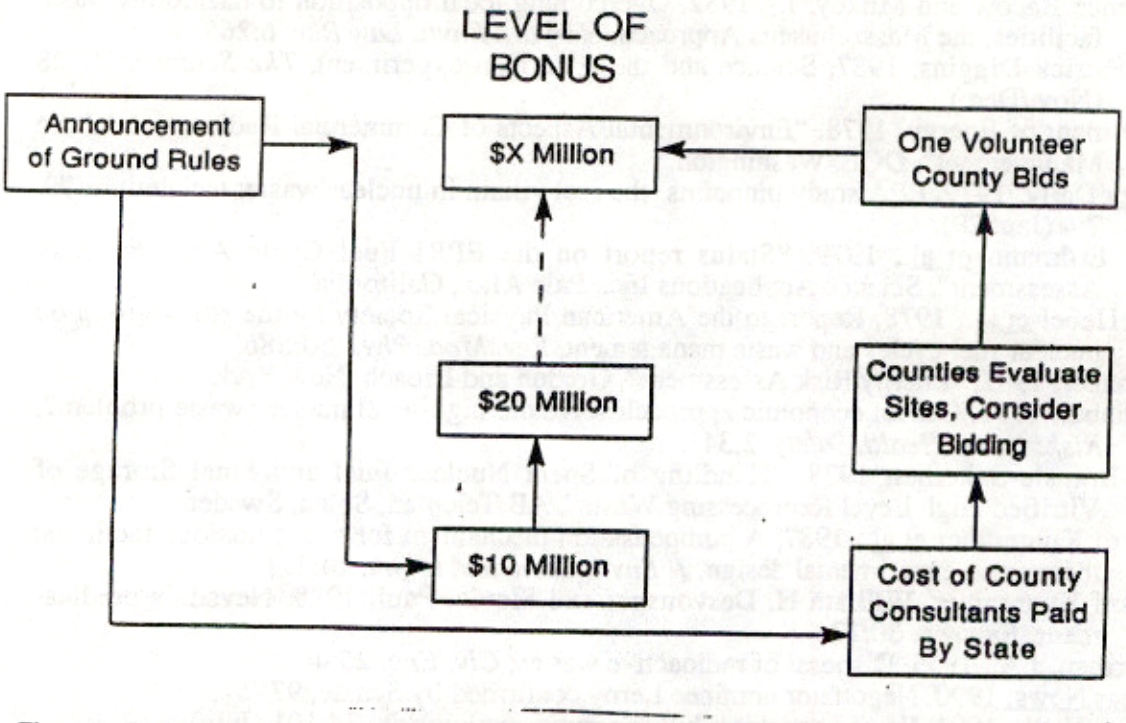


Figure 2. Pre-bid steps in the reverse Dutch auction for waste facilities. They begin with the announcement of the environmental criteria or ground rules. After each county (or other political jurisdiction) evaluates the regulations, the auction begins. The velocity of the bonus rise is set by the siting authority. As the price increases, a volunteer community will eventually bid.

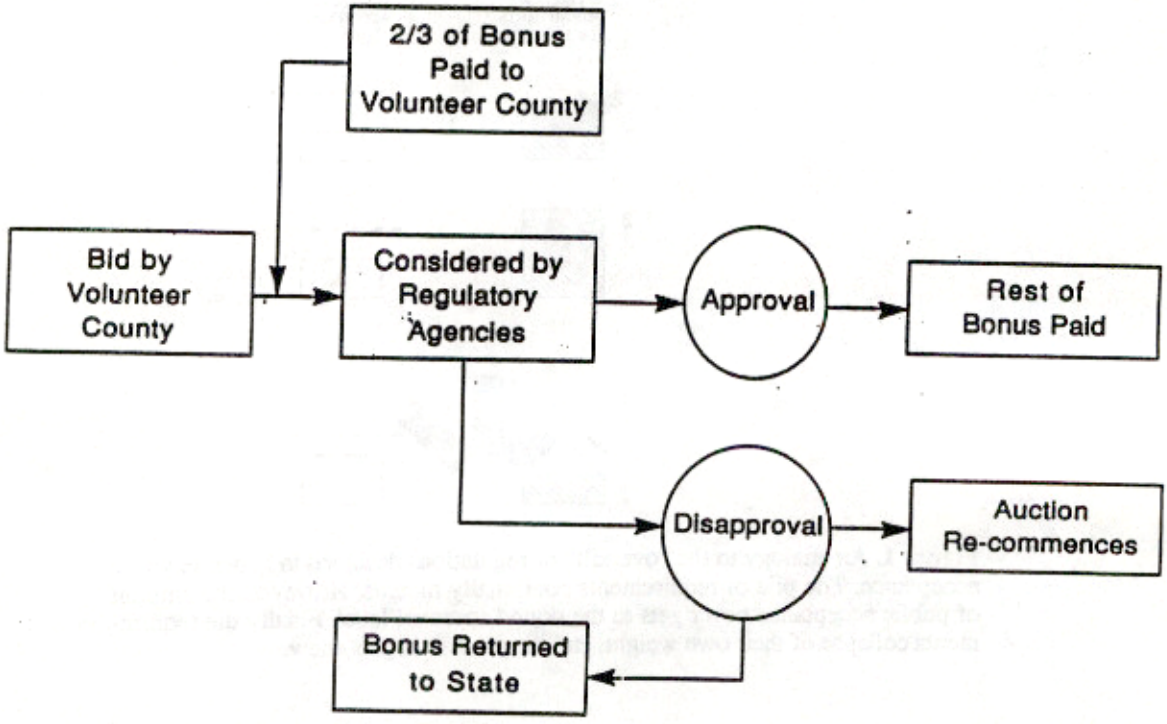


Figure 3. Post-bid steps in the reverse Dutch auction. After a bid is made, it must still pass environmental muster. The arrow to the right of the bid box ensures this. To produce confidence on the part of the bidding community that the bonus will indeed be paid in a timely manner, most of the funds is transferred to a temporary trust fund after the bid is made. If the proposed site meets the pre-announced environmental criteria, the total bonus is paid. If it does not, the auction begins again, starting at the bonus level previously achieved.

063 Methods for Uncertainty Analysis (II)

Chair: A. Camp, SNL

Development of Smart Searching Algorithms for Vulnerability and Uncertainty Analysis in Probabilistic Risk Assessments

A.S. Benjamin (SNL)

Importance and Sensitivity Analysis in Assessing the System Reliability

Y.-F. Wu (Taiwan Power Co.), J.D. Lewins (Cambridge U.)

Multiple Weight Stepwise Regression

J. Atkins (UC Berkeley); J. Campbell (SNL)

A Software System for Probabilistic Uncertainty and Sensitivity Analysis of Results from Computer Models

B. Krzykacz, E. Hofer, M. Kloos (GRS)

SAND93-1236

DEVELOPMENT OF SMART SEARCHING ALGORITHMS FOR VULNERABILITY AND UNCERTAINTY ANALYSES IN PROBABILISTIC RISK ASSESSMENTS*

Allan S. Benjamin
Special Projects Department 6411
Sandia National Laboratories
Albuquerque, NM 87150

In order to evaluate the risk inherent in a complex system, including a reasonable accounting of the many uncertainties that characterize both the environments to which the system is exposed and the responses of the system to these environments, risk analysts are forced to use statistical approaches. Some have chosen to assume that the variables have specialized distributions that lend themselves to rapid convolution. Others have resorted to Monte Carlo techniques or to similar sampling approaches. Such techniques have been used to determine the resultant distributions of various risk indices, such as the frequency of core melting in a nuclear power plant,[1] the probability of accidental nuclear detonation in a nuclear weapon system,[2] or the expectation of health effects from radioactive dispersal,[3] and to estimate the medians, means, and confidence intervals of these resultant distributions.

If the system is very well designed for safety, then simple convolution or sampling approaches may have to be augmented by intelligent searching schemes. An excellent example is the prediction of the probability of an accidental nuclear detonation for a nuclear weapon system. Under practically all situations, accidental nuclear detonation is virtually impossible because modern nuclear weapon systems are designed to preclude it. They have a series of strong links and weak links that are guaranteed to fail in a prescribed order when exposed to virtually all credible abnormal environments, such as fires, impacts, punctures, crushes, external pressures, lightning, or chemical attack. However, no engineered system is perfect, and under certain peculiar conditions involving combined environments that are spatially directed in the worst possible way, the strong links may fail before the weak links and an accidental nuclear detonation may be plausible. The challenge is to identify those conditions through an intelligent searching process and to determine whether their probability of occurrence is high enough to be of concern. (The current requirements for nuclear weapon systems stipulate that the probability of nuclear detonation must be less than one chance in a million for any credible abnormal environment or combination of abnormal environments.)

In the weapon system application, LHS generates millions of sample members, each consisting of a different set of parameter values, and we must employ intelligent searching techniques to reduce this set to a more workable number. We do this through a discriminator subprocess, which we will describe momentarily. The discriminator selects a fraction of the LHS sample members (i.e., several thousand) which are judged to possess the highest potential for producing the undesired outcome (i.e., nuclear detonation).

*This work was supported by the United States Department of Energy under Contract DE-AC04-76DP00789.

To evaluate the response of the system for the remaining sample members, we exercise a set of risk-compatible physical response codes such as TEMPRA-3D and STRESS-3D. These are fast-running thermal and structural analysis codes which are benchmarked to results obtained from more detailed finite-element codes. We describe the concepts employed in these codes in a companion paper.[6]

We combine results from TEMPRA-3D and STRESS-3D with a fault tree model which is produced by the codes SEATREE[7] and SETS/SABLE.[8] The fault tree for the weapon system application describes the pathways by which electrical energy may propagate from a source to the detonators in the nuclear explosives package. We use a cut set uncertainty analysis code, TEMAC,[9] to evaluate whether the undesired event occurs for each sample member and to analyze the principal contributing factors. For the sample members that produce the undesired outcome or come close to producing it, we perform further analysis using the more detailed finite-element physical response models together with the previously generated fault tree models. Thus we obtain a detailed analysis for those combinations of parameters that provide the highest potential for achieving the undesired outcome. The whole process, or a part of it, may be iterated as many times as necessary to assure that the accounting of potential pathways is reasonably complete.

The discriminator subprocess provides the basis for intelligent searching. Two examples of discriminator subprocesses are illustrated in Figure 2. In the first example, we develop a much simpler representation of the system by employing only a few representative cut sets to describe the pathways to the undesired outcome and a few algebraic equations to describe the physical responses of the components. We select the representative cut sets and formulate the algebraic equations by analyzing results obtained with the more detailed fault tree and risk-compatible physical response models for a select number of cases. We then apply the simpler representation to determine, for each of the original millions of sample members, a figure of merit loosely interpreted as a "closeness to the undesired outcome". More specifically, the figure of merit records the margin by which the most critical race in each cut set is won or lost. Sample members having the highest figures of merit are retained for further analysis.

In the second discriminator example, we extract the first few thousand out of the original millions of sample members and use the complete fault tree and the complete risk-compatible physical response models to evaluate how close each comes to producing the undesired outcome. From this subsample, we identify a few tens of members which appear to have the highest potential for the undesired outcome. These few members are used as "seeds" for paring down the original sample. The sample members whose sampling parameters lie closest to the seeds are retained for further analysis.

To test these processes and their ability to speed the convergence to a solution, we devised a sample problem. This problem simulates, qualitatively, the key characteristics of a nuclear weapon system's response to fires, including the thermal races between strong links and weak links that determine the outcome. At the same time, it incorporates a set of simplifications that enable us to obtain an "exact" analytical solution for the mean probability of the simulated undesired outcome and the statistical parameters that characterize its uncertainty. In addition to providing a means for benchmarking the intelligent searching algorithm, the sample problem provides insights into the challenges involved in searching for small regions of vulnerability in a large parameter space.

At Sandia, we have been exploring intelligent searching processes that are based on the principles of importance sampling. Figure 1 illustrates one of these processes which is applicable to the nuclear weapon system application. In this example, we first derive an accident scenario (event tree) model using the EVNTRE code.[4] We exercise this model to provide environment parameter distributions for a Latin hypercube sampling analysis performed by the code LHS.[5] We also input distributions for weapon system orientations, material physical properties, and component physical thresholds.

The presentation will describe the modeling concepts, searching algorithms, and code interfaces that we have discussed above. It will also present the derivation of the sample problem, identify how it relates to risk applications such as the nuclear weapon system example, describe how it was used as a test bed for searching algorithms, and discuss the insights gained from these analyses. Finally, it will identify some of the potential uses of intelligent searching algorithms for a variety of risk-related problems.

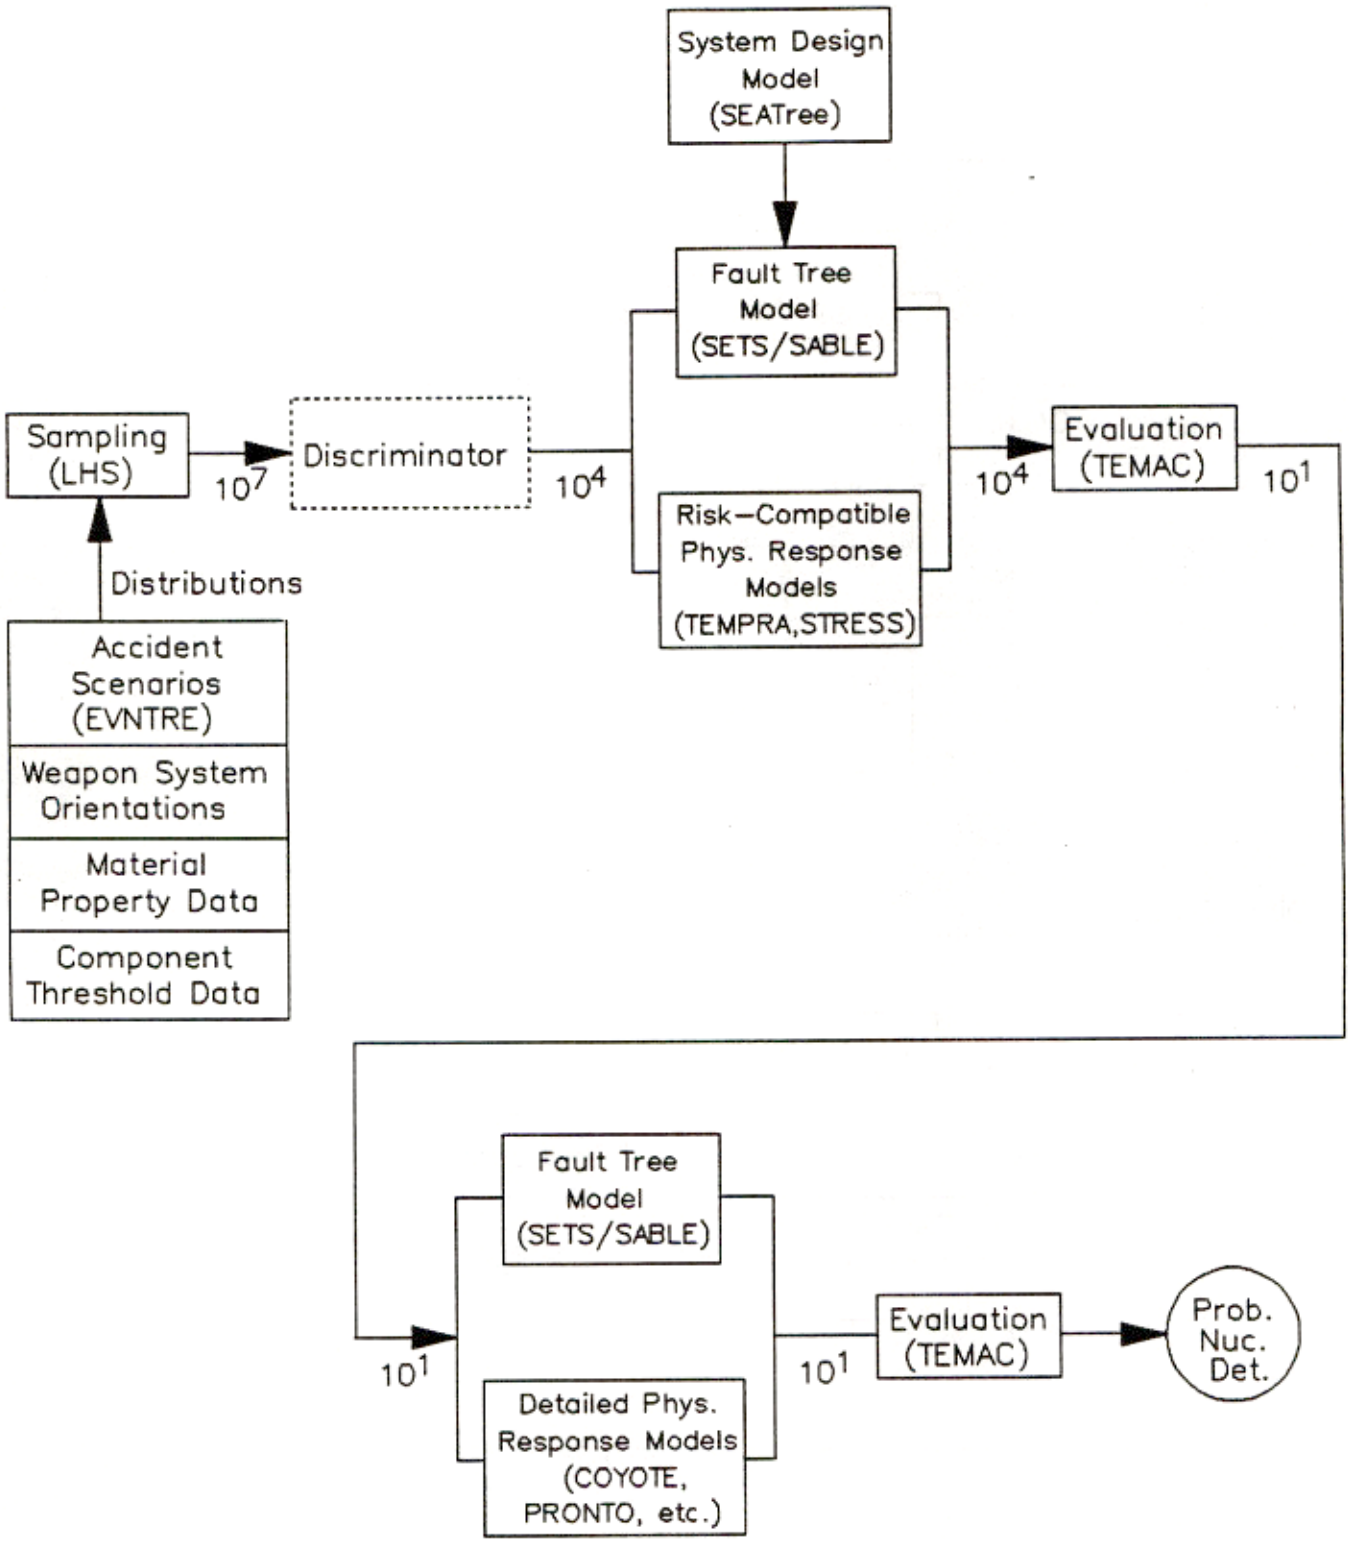
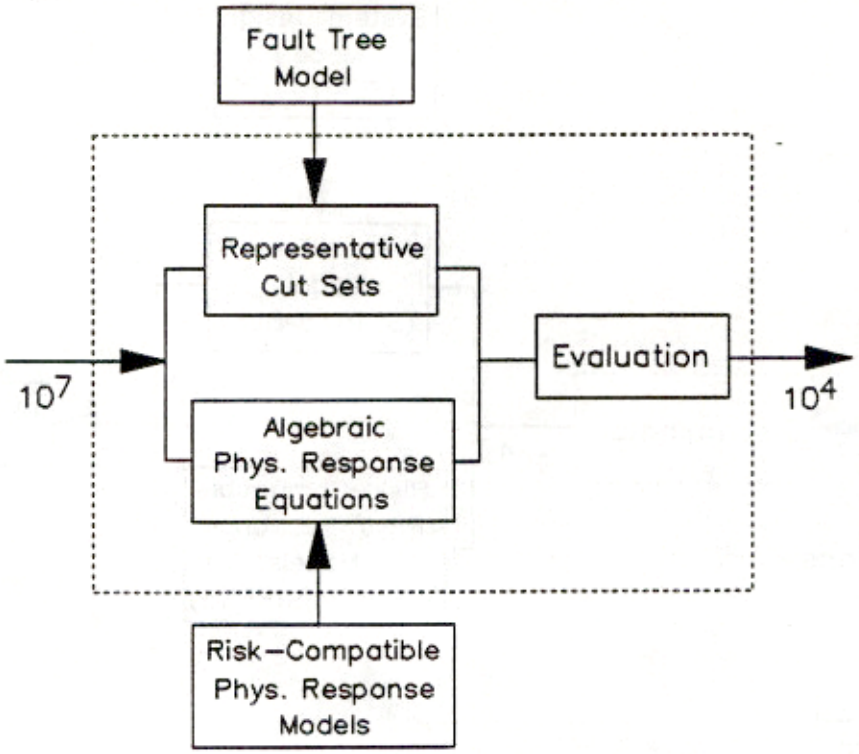


Figure 1. Illustration of a Process for Evaluating Nuclear Detonation Pathways.

Discriminator (A)



Discriminator (B)

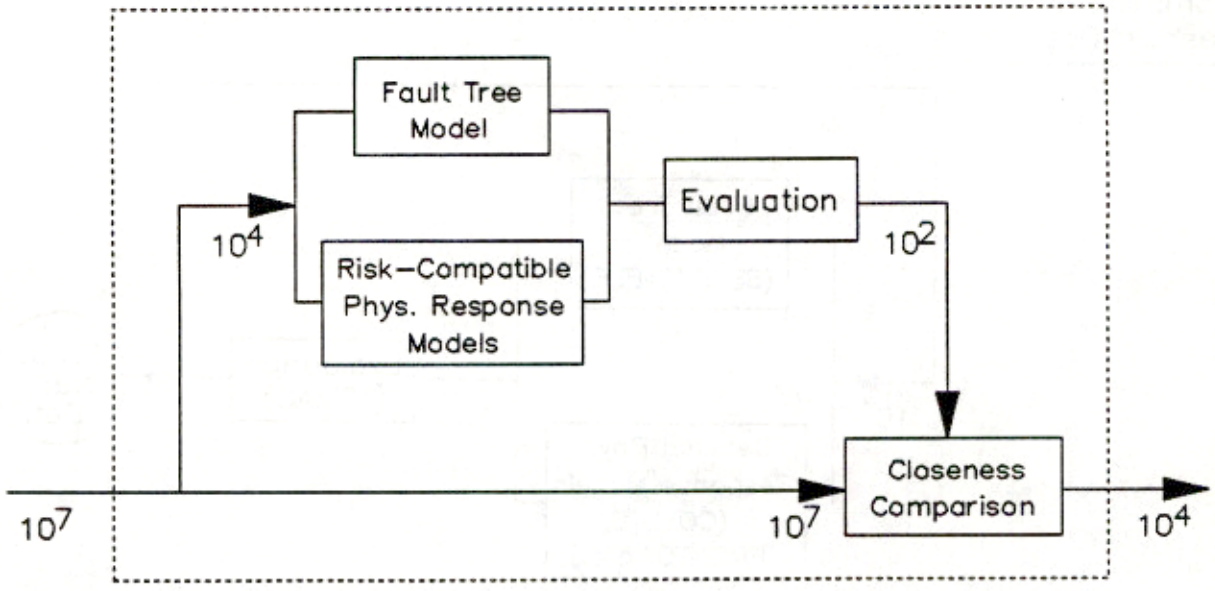


Figure 2. Illustration of Two Discriminator Subprocesses.

REFERENCES

1. "Severe Accident Risks: An Assessment for Five U.S. Nuclear Power Plants," NUREG-1150, 1990.

2. D. D. Carlson, A. S. Benjamin, R. J. Breeding, and D. M. Kunsman, "Probabilistic Risk Assessment Methods for Nuclear Weapon System Applications, Volume 2 - Methods," SAND91-0323/2, August 1991.

3. D. C. Aldrich, et al., "Technical Guidance for Siting Criteria Development," NUREG/CR-2239, SAND81-1549, 1982.

4. J. M. Griesmeyer and L. N. Smith, "A Reference Manual for the Event Progression Analysis Code (EVNTRE)," NUREG/CR-5174, SAND88-1607, 1989.

5. R. L. Iman and M. J. Shortencarrier, "A FORTRAN 77 Program and Users Guide for the Generation of Latin Hypercube and Random Samples for Use with Computer Models," NUREG/CR-3624, SAND83-2365, 1984.

6. A. S. Benjamin, "Development of Fast-Running Thermal and Structural Response Models for Probabilistic Analyses of Complex Systems," SAND93-1235C, This Conference, 1994.

7. B. Bingham and S. Walker, "SEATree, A SETS Fault Tree Drawing Tool," Science and Engineering Associates, Inc., 90-258-20-A:3, 1989.

8. R. Worrell, "SETS Reference Manual," NUREG/CR-4213, SAND83-2675, 1985.

9. R. L. Iman and M. J. Shortencarrier, "A Users Guide for the Top Event Matrix Analysis Code (TEMAC)," NUREG/CR-4598, SAND86-0960, 1986.

IMPORTANCE AND SENSITIVITY ANALYSIS IN ASSESSING THE SYSTEM RELIABILITY

Yun-Fu Wu[1] and J. D. Lewins[2]

1. Taiwan Power Company, (Room: 1706)
   242, Roosevelt Rd, Sec. 3, Taipei, Taiwan, R.O.C.
2. Cambridge Univ. Engineering Dept,
   Trumpington St, Cambridge CB2 1PZ, England, UK.

INTRODUCTION

The importance of a component in a system is a good index for the system engineers to plan activities within the system. For example, inspection, maintenance, and failure detection can be carried out in their order of importance for the components, and system upgraded by improving components with relatively large importance. As to the sensitivity coefficient of the components in a system, it is a good index used to upgrade or optimize the design of the system. Therefore, the importance and sensitivity of components in the system are the most important and valuable information desired by the system designer. In a conventional Monte Carlo simulation[1,2], to calculate the importance or sensitivity coefficients for the components in the system, each perturbed case (say a change in the failure rate or repair rate of one component) needs a separate Monte Carlo simulation. For example, a system with n components each having two parametric coefficients would call for $2n + 1$ separate simulations to obtain all the sensitivity coefficients. Not only is this expensive in CPU time, it is statistically inefficient for the calculation of the sensitivity coefficient because of the uncertainty of the difference between two Monte Carlo runs. To alleviate the consumption of CPU-time and improve the accuracy of the calculated results, we incorporate the correlated-sampling technique[1,3] into the conventional Monte Carlo simulation to do the importance and sensitivity analysis.

CORRELATED-SAMPLING SCHEME

To evaluate the perturbed system behaviour after system made a small change in its component characteristic (e.g., change in a component failure rate), the correlated sampling scheme was used in this study. Briefly, the way to apply correlated sampling in Monte Carlo simulation to evaluate the property changes is as follows: one set of histories is obtained for the unperturbed case (base case) and then the same histories are modified by adjustment of weighting values that represent the changes in parameters to obtain the modified results. In other words, the perturbed history is forced to follow the unperturbed one along the same

tracks in the phase space, but the effect of changes in parameters are taken into account by adjusting of the "weight" of each perturbed history. So that the perturbed behaviour of the system can be predicted by multiplying a weighting factor to the simulation result from a unperturbed case.

To see practically how we retrieve perturbed information from an unperturbed history, let us use one randomly generated history (Figure 1) of a three-component system, having one (say component 2) of its component failure rate changed from λ_2 to λ_2', to illustrate the details of this information retrieval. As a result of this change, there are two things which will change: one is the free flight time between two consecutive collisions; the other is the occurrence probabilities of events at a "collision" and the natural probability of having this random history. Figure 1 shows the picture of these two changes.

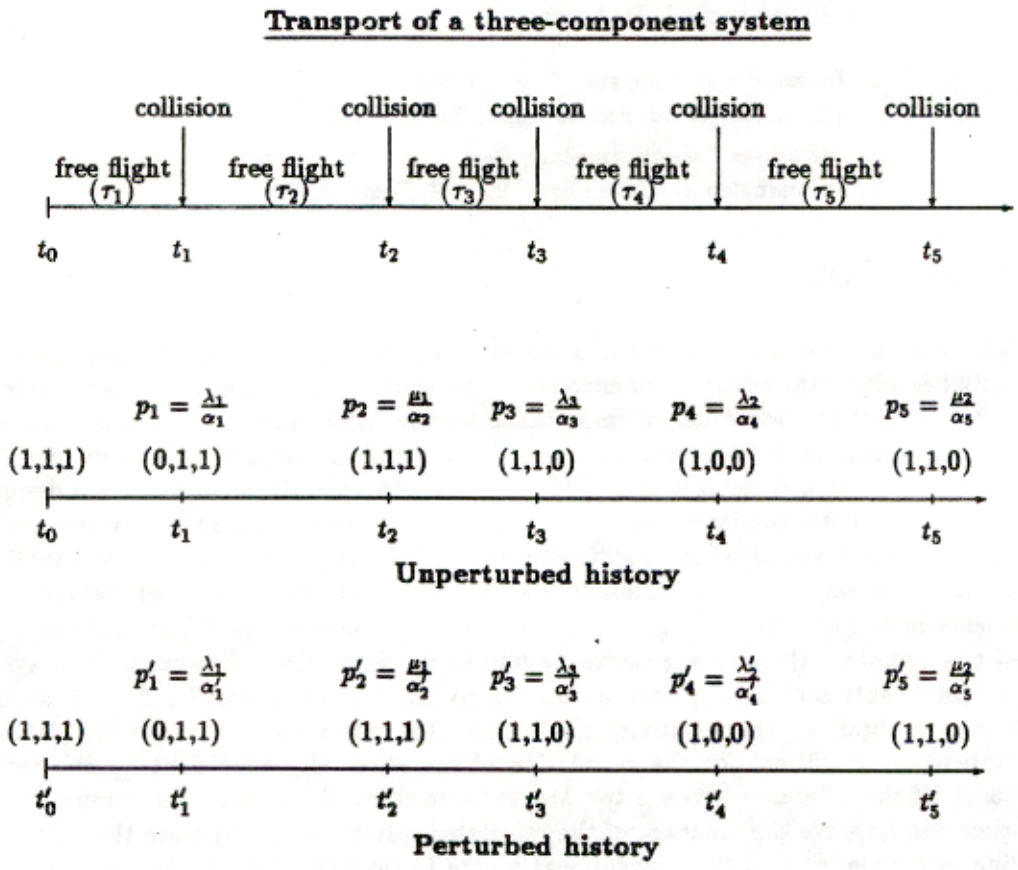


Figure 1: **Changes in one random history after λ_2 changes to λ_2'.**

First, let us look at the change of the free-flight time (τ_k) in one time segment, say k-th segment, caused by the change of the total system change rate from α_k to α_k'. The perturbed free flight time τ_k' is determined based on the same generated random number, if correlated-sampling scheme is applied,[1] as

$$\tau_k' = -\frac{1}{\alpha_k'}\ln(1-\eta) = \frac{\alpha_k}{\alpha_k'} - \frac{1}{\alpha_k}\ln(1-\eta) = \frac{\alpha_k}{\alpha_k'}\tau_k, \qquad (1)$$

where α_k is the total change rate of the system and is defined as

$$\alpha = \sum_{i=1}^{n} \beta_i, \tag{2}$$

and

$$\beta_i = \begin{cases} \lambda_i & \text{if component i is operational,} \\ \mu_i & \text{if component i has failed,} \end{cases}$$

and α'_k is the total change rate of the perturbed system before making its k-th collision.

Second, let us look at the change of event-probability at each collision and the natural probability of the history caused by the change of the component failure rate. Since the event-probability at each collision for the perturbed history will change, therefore, the natural probability of the history, which is the product of the event-probability at all collision points before the history is terminated, will change too. In a general Monte Carlo simulation, the weight of each generated history for the unperturbed case is set to be unity. However, for the perturbed case, since the event-probability at each collision point has changed compared to its correlated history in the unperturbed case, the weight of each history no longer equals unity. It has to be corrected based on the relative probability of its occurrence by comparison with its correlated history. In this example history, the weight (w_j) of its correlated history will be determined[1] as:

$$w_j = (\frac{p'_1}{p_1} \times \frac{p'_2}{p_2} \times \frac{p'_3}{p_3} \times \frac{p'_4}{p_4} \times \frac{p'_5}{p_5})_j,$$

where this example history is the j-th history in the Monte Carlo simulation.

Therefore, in such a Monte Carlo simulation study, not only is the effect on the free-flight time of perturbed condition corrected, but also the weight for each perturbed history is determined and recorded. Then the system reliability and availability before system made change (R and A) and after system made change (R' and A') can be determined as

$$R = \frac{1}{N} \sum_{j=1}^{N} r_j, \tag{3}$$

$$A(t) = \frac{1}{N} \sum_{j=1}^{N} a_j(t), \tag{4}$$

$$R' = \frac{1}{\sum_{j=1}^{N} w_j^r} \sum_{j=1}^{N} w_j^r r'_j, \tag{5}$$

$$A'(t) = \frac{1}{\sum_{j=1}^{N} w_j^a} \sum_{j=1}^{N} w_j^a a'_j(t), \tag{6}$$

where N is the total number of simulated histories and

- r_j and r'_j will be "one" if a simulated history has not failed before it passed the requested mission time; otherwise, "zero" would be assigned,
- a_j and a'_j will be "one" if a simulated history operated successfully at the time t of interest; otherwise, "zero" would be assigned,
- w_j^r is the weight of derived history used to estimate the reliability of perturbed case,
- w_j^a is the weight of derived history used to estimate the availability of perturbed case.

IMPORTANCE AND SENSITIVITY ANALYSIS

Different definitions of the importance have been given and used for different purposes. Two commonly used definitions are listed below for reference:

- **Birnbaum Importance**

$$I_i^B(t) = \frac{\partial \bar{A}(t)}{\partial q_i(t)} = \lim_{\Delta q \to 0} \frac{\Delta \bar{A}(t)}{\Delta q_i(t)},$$

where $\bar{A}(t)$ is the unavailability of the system and q_i is the unavailability of the i-th component.

- **Criticality Importance**

$$I_i^{CR}(t) = \frac{\partial \bar{A}(t)}{\partial q_i(t)} \cdot \frac{q_i(t)}{\bar{A}(t)} = \lim_{\Delta q \to 0} \frac{\Delta \bar{A}(t)}{\Delta q_i(t)} \cdot \frac{q_i(t)}{\bar{A}(t)},$$

As to the sensitivity coefficient in relation to the generic parameters (failure and repair rates, etc.) of the component, it is heuristically defined in the literatures as:

- **Classical Sensitivity**

$$S_i^C(t) = \frac{\partial \bar{A}(t)}{\partial \beta_i} = \lim_{\Delta \beta_i \to 0} \frac{\Delta \bar{A}(t)}{\Delta \beta_i},$$

where β_i is one of the i-th component's parameters. It can be either λ_i or μ_i in the study.

- **Fractional Sensitivity**

$$S_i^F = \frac{\partial \bar{A}}{\partial \beta_i} \cdot \frac{\beta_i}{\bar{A}} = \lim_{\Delta \beta \to 0} \frac{\Delta \bar{A}(t)}{\Delta \beta_i} \cdot \frac{\beta_i}{\bar{A}}.$$

To demonstrate the application of correlated sampling scheme to calculate the importance and sensitivity of the components in a real system, we study the standby-liquid-control (SBLC) system in a boiling water reactor unit. The SBLC system provides an alternate method to shutdown the reactor. The system is normally at standby as shown in Figure 2.

The failure mode of the system we concerned is "SBLC system fails to function on demand". The failure mode of the components are identified and tabulated in Table 1.

The importance and sensitivity of all the components in the SBLC system are determined by giving a 10% change of the component failure rate. Results are tabulated in Table 2.

CONCLUSIONS

It is known that correlated sampling is a powerful variance reduction technique, but this is not the only advantage of applying correlated sampling techniques to our system reliability studies. As we have already demonstrated in this study, after the correlated sampling technique has been incorporated into a conventional Monte Carlo simulation, it becomes possible that the calculation of the index of the importance and sensitivity of all components in a system can be obtained in only one Monte Caulo simulation.

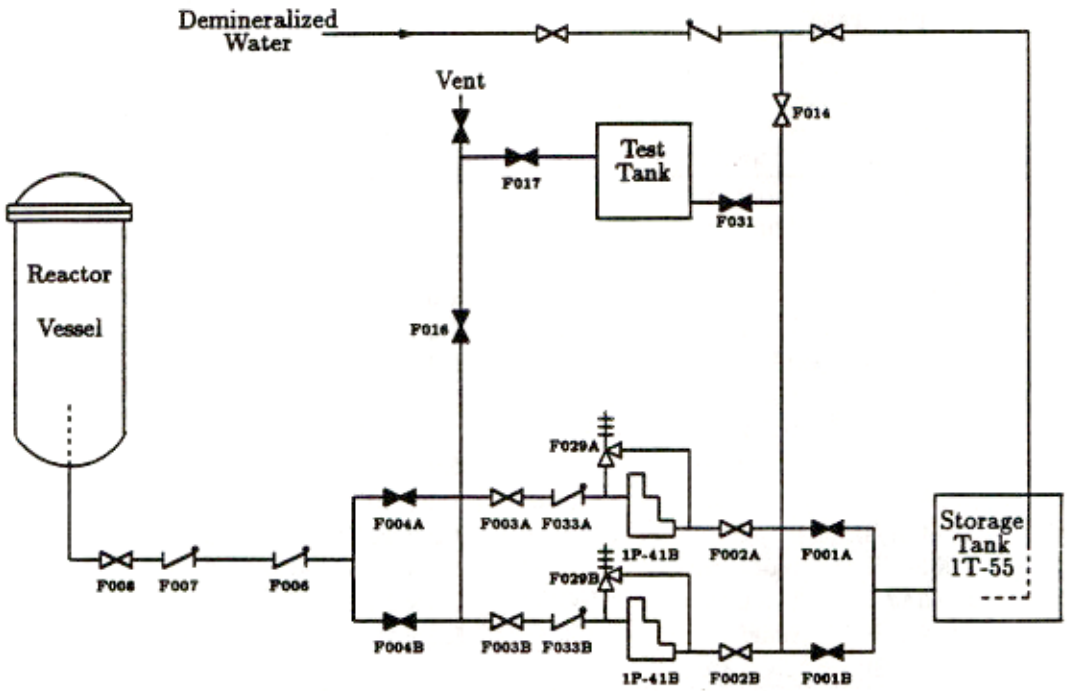


Figure 2: The flow diagram of SBLC system (standby mode).

REFERENCES

1. Wu, Y-F. and Lewins, J. D., "Monte Carlo Studies of Engineering System Reliability", Annals of Nuclear Energy, 19, p825-859, (1992).

2. Goldfeld A. and Dubi A., "Monte CArlo Methods in Reliability Engineering", Quality and Reliability Engineering Int. 3, p83-91, (1987).

3. Rief, H., "Generalized Monte Carlo Perturbation Algorithm for Correlated Sampling and A Second-order Taylor Series Approach", Annals of Nuclear Energy, vol 11, p455-476, (1984).

Table 1. Component's failure mode and the corresponding failure and repair rate.

I.D.in the Flow Diagram	Component's Function	Failure Mode	Failure rate (λ_o, day^{-1})	Repair rate (μ_o, day^{-1})
F014	block valve	NCFO	1.0×10^{-4}	0.25
F031	block valve	NCFO	1.0×10^{-4}	0.25
F006	check valve	NCFC	1.0×10^{-5}	0.25
F007	check valve	NCFC	1.0×10^{-5}	0.25
F008	block valve	NOFC	1.0×10^{-4}	0.25
F004A/B	explosive valve	NCFC	1.0×10^{-3}	0.25
F016	block valve	NCFO	1.0×10^{-4}	0.25
F017	block valve	NCFO	1.0×10^{-4}	0.25
F001A/B	block valve	NCFC	5.0×10^{-4}	0.25
IT-55	tank	rupture	1.0×10^{-7}	0.05
1P-41A/B	pump	fail to run	1.0×10^{-3}	0.25
F002A/B	block valve	NOFC	5.0×10^{-5}	0.25
F003A/B	block valve	NOFC	1.0×10^{-4}	0.25
F029A/B	relief valve	NCFO*	5.0×10^{-4}	0.25
F033A/B	check valve	NCFC	1.0×10^{-5}	0.25

NCFC = Normally Closed - Fails in the Closed position.
NCFO = Normally Closed - Fails Open.
NOFC = Normally Open - Fails Closed.
NCFO* = premature opening.

Table 2. Component's importance and sensitivity index.

I.D. in the	(I_i^B)		(I_i^{CR})		(S_i^C)		(S_i^F)	
flow diagram	mean	σ_{IB}	mean	σ_{ICR}	mean	σ_{SC}	mean	σ_{SF}
F014	0.6423	0.0181	0.0611	0.0017	222.64	6.28	0.0599	0.0017
F031	0.6375	0.0182	0.0607	0.0017	220.99	6.29	0.0595	0.0017
F006	0.6683	0.1078	0.0065	0.0010	239.67	38.67	0.0065	0.0010
F007	0.6744	0.1076	0.0065	0.0010	241.87	38.58	0.0065	0.0010
F008	0.6253	0.0183	0.0595	0.0017	216.77	6.34	0.0584	0.0017
F004A/B	0.1930	0.0051	0.1571	0.0042	47.60	1.26	0.1282	0.0034
F016	0.0364	0.0204	0.0035	0.0019	12.61	7.07	0.0034	0.0019
F017	0.0344	0.0203	0.0033	0.0019	11.91	7.03	0.0032	0.0019
F001A/B	0.1015	0.0070	0.0450	0.0031	30.23	2.09	0.0407	0.0028
IT-55	0.3500	0.3862	0.0000	0.0000	125.16	138.13	0.0000	0.0000
1P-41A/B	0.2630	0.0054	0.2141	0.0044	64.88	1.32	0.1747	0.0036
F002A/B	0.2027	0.0163	0.0097	0.0008	71.62	5.76	0.0096	0.0008
F003A/B	0.1944	0.0197	0.0185	0.0019	67.37	6.84	0.0181	0.0018
F029A/B	0.2273	0.0067	0.1008	0.0030	67.73	2.00	0.0912	0.0027
F033A/B	0.2008	0.1049	0.0019	0.0010	72.00	37.63	0.0019	0.0010

Note 1 : Calculation is based on 10% perturbation.
Note 2 : The number of simulated histories is 50000.

MULTIPLE WEIGHT STEPWISE REGRESSION

Joel Atkins[1,2] and James Campbell

[1]Statistics Department
University of California
Berkeley, CA 94720
[2]Sandia National Laboratories
Manufacturing Systems Reliability
Albuquerque, NM 87185

INTRODUCTION

Background

In many science and engineering applications, there is an interest in predicting the outputs of a process for given levels of inputs. To develop a model, one could run the process (or a simulation of the process) at a number of points (a point would be one run at one set of possible input values) and observe the values of the outputs at those points. These observations can be used to predict the values of the outputs for other values of the inputs. Since the outputs are a function of the inputs, we can generate a surface in the space of possible inputs and outputs. This surface is called a response surface. In some cases, collecting data needed to generate a response surface can be very expensive. Thus, in these cases, there is a powerful incentive to minimize the sample size while building better response surfaces. One such case is the semiconductor equipment manufacturing industry. Semiconductor manufacturing equipment is complex and expensive. Depending upon the type of equipment, the number of control parameters may range from 10 to 30 with perhaps 5 to 10 being important. Since a single run can cost hundreds or thousands of dollars, it is very important to have efficient methods for building response surfaces.

A current approach to this problem is to do the experiment in two stages. First, a traditional design (such a fractional factorial) is used to screen variables. After deciding which variables are significant, additional runs of the experiment are conducted. The original runs and the new runs are used to build a model with the significant variables. However, the original (screening) runs are not as helpful for building this model as some other points might have been. This paper presents a point selection scheme that is more efficient than traditional designs.

Latin Hypercube Sampling (Iman and Conover) is used to select design points. Stratification intervals are chosen in a way that will help us fit a better quadratic screening model. Once the selected input points have been run, a stepwise method is used to add the terms that give the most significant improvement in the model. Each time a term is added, terms that are no longer significant are dropped.

Once we have fit a quadratic model, that model is used to screen input variables and to decide which are important. Input variables identified as important are then used in a similar method to fit the final response surface.

Notation

Assume that the space of possible input values is rectangular, or in other words, that the possible values of any given input variable do not depend on the values the other input variables take. This is equivalent to assuming that $\bar{x}_m \in \prod_{i=1}^{n}\left[min(x_i), max(x_i)\right]$.

Throughout this discussion, we will use the following notation: d for the number of input variables, n for the number of points in the design, y for the output (which we will assume is one-dimensional), $\bar{x}_m$ for the values of the input variables at the mth point in our design, and $x_{i,m}$ for the value of the ith input variable at the mth point.

Initial Screening Model

Our model will include a mean, and a linear and square term for each independent variable. It will also include interactions between pairs of independent variables, i.e., $x_j x_k$. Thus our initial model will be of the form:

$$y = \mu + \sum \beta_i x_i + \sum \alpha_i x_i^2 + \sum_{j<k} \gamma_{j,k} x_j x_k .$$

CHOOSING THE DESIGN POINTS

We want our screening model to be "good" with respect to some importance function, $\lambda(\bar{x})$. (Assume that $\lambda(\bar{x}) = \prod_i^n \lambda_i(x_i)$).

For One Input Variable

Here, we choose the values of one input variable, x_i at the design points. We are interested in developing the best quadratic fit for y, i.e., $\alpha_i x_i^2 + \beta_i x_i + c_i$ which has the minimum mean squared error of all quadratic functions under weight $\lambda_i(x)$. For our purposes we can treat $\lambda_i(x_i)$ as a probability density. If we let $d_{i,j} = E\left(x_i^j | \lambda_i\right)$, then we will have:

$$\begin{bmatrix} \alpha_i \\ \beta_i \\ c_i \end{bmatrix} = \frac{\begin{bmatrix} d_{i,2} - d_{i,1}^2 & d_{i,2}d_{i,1} - d_{i,3} & d_{i,3}d_{i,1} - d_{i,2}^2 \\ d_{i,2}d_{i,1} - d_{i,3} & d_{i,4} - d_{i,2}^2 & d_{i,3}d_{i,2} - d_{i,4}d_{i,1} \\ d_{i,3}d_{i,1} - d_{i,2}^2 & d_{i,3}d_{i,2} - d_{i,4}d_{i,1} & d_{i,4}d_{i,2} - d_{i,3}^2 \end{bmatrix}}{d_{i,4}d_{i,2} + 2d_{i,3}d_{i,2}d_{i,1} - d_{i,4}d_{i,1}^2 - d_{i,2}^3 - d_{i,3}^2} \begin{bmatrix} E\left(x_i^2 y | \lambda_i\right) \\ E\left(x_i y | \lambda_i\right) \\ E\left(y | \lambda_i\right) \end{bmatrix}$$

The above equation is true because the weight measure is a product of one-dimensional weight measures and thus α_i, and β_i are orthogonal to all of the other terms in the model. Thus, α_i,β_i and c_i will be linear functions of $E(y|\lambda_i)$, $E(x_i y|\lambda_i)$, and $E(x_i^2 y|\lambda_i)$. We can divide the range of x into n strata of the form $[l_{i,m-1}, l_{i,m}]$ and choose a point from each strata using the distribution

$$\lambda_{i,m}(x) = \frac{\lambda_i(x)}{\Lambda_i(l_{i,m}) - \Lambda_i(l_{i,m-1})} 1_{(l_{i,m-1} \le x \le l_{i,m})}.$$

After we choose the points,

$$\sum_{m=1}^{n} \left[\Lambda_i(l_{i,m}) - \Lambda_i(l_{i,m-1})\right] y_m ,$$

$$\sum_{m=1}^{n} \left[\Lambda_i(l_{i,m}) - \Lambda_i(l_{i,m-1})\right] x_{i,m} y_m , \text{ and}$$

$$\sum_{m=1}^{n} \left[\Lambda_i(l_{i,m}) - \Lambda_i(l_{i,m-1})\right] x_{i,m}^2 y_m$$

will be unbiased estimates of $E(y)$, $E(x_i y)$, and $E(x_i^2 y)$. If we set $w_{i,m} = \Lambda_i(l_{i,m}) - \Lambda_i(l_{i,m-1})$, then:

$$\widehat{E(y)} = \sum_{m=1}^{n} w_{i,m} y_m ,$$

$$\widehat{E(x_i y)} = \sum_{m=1}^{n} w_{i,m} x_{i,m} y_m , \text{ and}$$

$$\widehat{E(x_i^2 y)} = \sum_{m=1}^{n} w_{i,m} x_{i,m}^2 y_m .$$

Thus,

$$\hat{\alpha}_i = \sum_{m=1}^{n} w_{i,m} \left[a_0 y_m + a_1 x_{i,m} y_m + a_2 x_{i,m}^2 y_m\right],$$

$$\hat{\beta}_i = \sum_{i=1}^{n} w_{i,m} \left[b_0 y_i + b_1 x_i y_i + b_2 x_i^2 y_i\right],$$

$$\text{and } \hat{c} = \sum_{i=1}^{n} w_{i,m} \left[c_0 y_i + c_1 x_i y_i + c_2 x_i^2 y_i\right],$$

are unbiased estimates of a, b, and c. We want to choose our intervals to minimize the expected value of the mean squared errors of our estimates, or:

$$MSE_\lambda = \int_{\hat{a},\hat{b},\hat{c} \in \Psi} \left[\int_{\min(x)}^{\max(x)} E\left(\left(y - \hat{a}x^2 - \hat{b}x - \hat{c}\right)^2\right) \lambda(x) dx \right] d\psi$$

We cannot minimize this function, or even evaluate it, if we do not know y. However, we can approximate this function for a given set of intervals, and then minimize our approximation. Towards this end, we will break y into $f(x) = E(y|x)$ and

$g(\bar{x},\varepsilon) = y - f(x)$. Thus, $E(g(\bar{x},\varepsilon)) = 0$, and $\text{var}(g(\bar{x},\varepsilon)|x) = \sigma^2(x)$. We can now argue that within each interval $g(\bar{x},\varepsilon)$ will have more variation than *f(x)* will have. If we assume that $\sigma^2(x)$ is fairly constant, we can minimize the part of the variance of $\hat{a}$, $\hat{b}$, and $\hat{c}$ that comes from $g(\bar{x},\varepsilon)$ by choosing shorter intervals near the ends of the range and longer intervals near the middle.

Since nothing we do in the one variable case depends on *y*, we will use the same methodology if there is more than one output variable of interest.

More than One Independent Variable

Once we have chosen x_i^m: m=1,.....,n, for each *i*,these marginal points are combined to create the design points. In creating the design points, we want to minimize the correlations between each pair of independent variables. One popular method for doing this is using a Cholesky decomposition variation of Latin Hypercube Sampling (Iman and Conover).

When the problem involves more than one input variable, we must consider interaction terms. We can show that $\hat{\gamma}_{j,k} = \sum_{m=1}^{n} w_{j,k,m}\left(d_{j,k,0} + d_{j,k,1}x_{j,m}x_{k,m}\right)y_m$. For $w_{j,k,m}$, we will use the average of the two relevant weights, giving us $w_{j,k,m} = \dfrac{w_{j,m} + w_{k,m}}{2}$.

To estimate the mean, we will use the average of the weights corresponding to significant terms. This, if x_1,x_2, and x_3 are all significant, then we let $\hat{\mu} = \sum_{m=1}^{n} \dfrac{w_{1,m} + w_{2,m} + w_{3,m}}{3} y_m$.

FITTING A MODEL AT EACH STEP IN THE STEPWISE ALGORITHM

Stepwise Technique

To build our model, we will cycle through two steps. In the first, we will consider whether adding any term will improve the fit. If any terms do improve the fit, then we add the term that gives the biggest improvement. If no term improves the fit, then we stop and keep the model from the last iteration. Each time that we add a new term, we will go through the second step, dropping any terms that are no longer significant. We will continue going through this cycle until both (a) there are no terms left which would significantly improve the model, and (b) all of the terms already in the model are significant.

Estimating Individual Terms

As we mentioned in the previous section,

$$\hat{\mu} = \sum_{m=1}^{n} \frac{\sum_{i \in model} w_{i,m}}{|model|} y_m$$

$$\hat{\alpha}_i = \sum_{m=1}^{n} w_{i,m}\left[a_0 y_m + a_1 x_{i,m} y_m + a_2 x_{i,m}^2 y_m\right],$$

$$\hat{\beta}_i = \sum_{i=1}^{n} w_{i,m}\left[b_0 y_i + b_1 x_i y_i + b_2 x_i^2 y_i\right],$$

$$\hat{\gamma}_{j,k} = \sum_{m=1}^{n} w_{j,k,m}\left(d_{j,k,0} + d_{j,k,1}x_{j,m}x_{k,m}\right)y_m.$$

Unfortunately, this leaves a great deal of confounding between the different effects. In the next subsection, dealing with this confounding will be discussed.

Estimating Terms Simultaneously

Here, we will discuss getting around the confounding and the different sets of weights. To do this, we will show an example where we are interested in fitting the model $\hat{y} = \mu + \beta_1 x_1 + \alpha_2 x_2^2 + \gamma_{3,4} x_3 x_4$. To do this, we can look for values of $\hat{\mu}$, $\hat{\beta}_1$, $\hat{\alpha}_2$, and $\hat{\gamma}_{3,4}$, so that evaluating $\hat{y}$ at each of the design points, and then using these values of $\hat{y}$ to estimate the individual terms will give us the same values of $\hat{\mu}$, $\hat{\beta}_1$, $\hat{\alpha}_2$, and $\hat{\gamma}_{3,4}$ as we used to construct $\hat{y}$. We can find these four values by solving the following system of linear equations.

$$\sum_{i=1}^{n} \frac{w_{1,m} + w_{2,m} + w_{3,m} + w_{4,m}}{4} y_m = \sum_{i=1}^{n} \frac{w_{1,m} + w_{2,m} + w_{3,m} + w_{4,m}}{4} \left(\hat{\mu} + \hat{\beta}_1 x_{1,i} + \hat{\alpha}_2 x_{2,i}^2 + \hat{\gamma}_{3,4} x_{3,i} x_{4,i} \right)$$

$$\sum_{i=1}^{n} w_{1,m} \left(b_{1,0} + b_{1,1} x_{1,i} + b_{1,2} x_{1,i}^2 + \right) y_i = \sum_{i=1}^{n} w_{1,m} \left(b_{1,0} + b_{1,1} x_{1,i} + b_{1,2} x_{1,i}^2 \right) \left(\hat{\mu} + \hat{\beta}_1 x_{1,i} + \hat{\alpha}_2 x_{2,i}^2 + \hat{\gamma}_{3,4} x_{3,i} x_{4,i} \right)$$

$$\sum_{i=1}^{n} w_{2,m} \left(a_{2,0} + a_{2,1} x_{2,i} + a_{2,2} x_{2,i}^2 \right) y_i = \sum_{i=1}^{n} w_{2,m} \left(a_{2,0} + a_{2,1} x_{2,i} + a_{2,2} x_{2,i}^2 \right) \left(\hat{\mu} + \hat{\beta}_1 x_{1,i} + \hat{\alpha}_2 x_{2,i}^2 + \hat{\gamma}_{3,4} x_{3,i} x_{4,i} \right)$$

$$\sum_{i=1}^{n} \frac{w_{3,m} + w_{4,m}}{2} \left(d_{3,4,0} + d_{3,4,1} x_{3,i} x_{4,i} \right) y_i = \sum_{i=1}^{n} \frac{w_{3,m} + w_{4,m}}{2} \left(d_{3,4,0} + d_{3,4,1} x_{3,i} x_{4,i} \right) \left(\hat{\mu} + \hat{\beta}_1 x_{1,i} + \hat{\alpha}_2 x_{2,i}^2 + \hat{\gamma}_{3,4} x_{3,i} x_{4,i} \right)$$

SIMULATION RESULTS

Here, three simulations were run to see how our method compares to stepwise regression. In all three cases, we used both uniform strata intervals and improved strata intervals with standard regression and MWSR (the method described in this paper). For each case, we used a uniform measure of the space $[-1,1]^d$ and found the mean squared error between the fitted models and the true model. Towards this end, quadratic screening models were fitted. Then, we used the quadratic model as a starting point and fit a cubic model (without the three term interactions). Starting with the quadratic model helps to minimize problems with collinearity.

The first model uses $1 = x_1 + x_2 + x_3 + x_4 + x_5 + \varepsilon$ and three nuisance input variables. For each case, 100 sets of 30 data points were run giving the summary statistics shown in the Table 1. For a benchmark, the MSE for the best constant fit is 1.6667 and the MSE for the best cubic fit is 0.

Table 1

sd(ε)		MWSR with improved intervals	stepwise regression with improved intervals	MWSR with uniform intervals	stepwise regression with uniform intervals
.05	ave(MSE)	0.0187	0.00516	0.0113	0.0102
	sd((MSE)	0.13617	0.00233	0.0091	0.0051
.2	ave(MSE)	0.1158	0.1152	0.1827	0.2047
	sd((MSE)	0.2332	0.1553	0.2108	0.1893
.7	ave(MSE)	0.9635	1.0301	1.7437	2.0444
	sd((MSE)	0.5890	0.5722	8.0330	1.3773

The second model uses $.5e^{2x_1+2x_2} + x_3^5 + \varepsilon$ with three additional nuisance input variables. Here 50 sets of 40 data points were run for both improved and uniform intervals. We then used both methods to analyze each data set, giving the summary statistics shown in Table 2. The MSE for the best constant fit is 11.7275, and the MSE for the best cubic fit is 2.9520.

Table 2

sd(ε)		best model with improved intervals	MWSR with improved intervals	stepwise regression with improved intervals	best model with uniform intervals	MWSR with uniform intervals	stepwise regression with uniform intervals
.3	ave(MSE)	4.6807	6.3097	4.6926	4.5370	6.3698	4.7347
	sd((MSE)	0.7669	1.2481	0.7853	1.0661	1.6630	1.5852
.9	ave(MSE)	5.4554	6.8793	5.7373	5.9334	7.3627	6.2314
	sd((MSE)	1.0510	1.3362	1.6547	1.5352	2.0608	1.8226
2.5	ave(MSE)	9.2659	10.9132	10.7266	10.1907	12.0027	11.3293
	sd((MSE)	3.5003	4.5304	3.9029	3.9314	4.3843	4.9634

The improved intervals cases did better than the uniform interval cases almost uniformly. We feel that this is due to the fact that the selection scheme is always a good compromise between a fractional factorial design and using the LHS with uniform intervals. We would advocate using this scheme in most situations.

On average, standard regression did better than MWSR. While this was a disappointment, there were a number of data sets for which MWSR did better. In Table 2, one notices that using the best model is significantly better than using the standard regression model. Thus, we would also advocate using both methods and then using some goodness of fit criteria to choose between the models.

REFERENCES

Iman, R.L., and Conover, W.J., 1982, A Distribution Free Approach to Inducing Rank Correlation among Input Variables, *Commun. Statist. - Simula. Computat.* 11, 3, pp. 311-334.

Li, K.C., 1992, On Principal Hessian Directions for Data Visualization and Dimension Reduction: Another Application of Stein's Lemma, *J. Amer. Statist. Assc.* **87** 1025-1039.

A SOFTWARE SYSTEM FOR PROBABILISTIC UNCERTAINTY AND SENSITIVITY ANALYSIS OF RESULTS FROM COMPUTER MODELS

B. Krzykacz, E. Hofer, M. Kloos

Gesellschaft für Anlagen- und Reaktorsicherheit (GRS) mbH
Forschungsgelände
D-85748 Garching b. München

INTRODUCTION

Along with the development and application of computer models, it has been increasingly recognized that the results of model runs are associated with uncertainty due to the lack of knowledge from various sources. The probabilistic uncertainty and sensitivity analysis provides quantitative information on the combined influence of many of these uncertainties and rankings of the individual contributors to model output uncertainty. The analysis can be summarized as follows:

1. Identification of all phenomena, modeling assumptions, and parameters that are potentially important contributors to model output uncertainty and representation of all uncertainty sources by uncertain parameters
2. Quantification of the state of knowledge of parameters by adequate probability distributions and measures of dependence interpreted in the sense of the subjectivistic probability concept
3. Generation of a random experimental design according to a joint probability distribution that satisfies the marginals and dependences specified under 2
4. Performance of respective model runs providing a random sample from the unknown probability distribution of the model output
5. Derivation of quantitative statements about the uncertainty in the model output
6. Derivation of a ranking of the parameters with respect to their contribution to model output uncertainty

The software System for Uncertainty and Sensitivity Analysis (SUSA) was developed to guide through the analysis and to perform the necessary mechanistic tasks (under 3 to 6). SUSA offers various types of probability distributions and dependence structures for modeling states of knowledge probabilistically, alternative parameter sampling procedures, and different types of quantitative uncertainty statements and sensitivity measures. SUSA supports the selection of input options, manages the data transfer within the software system, prints the input and output of the analysis in a form suitable for documentation, and presents

all relevant results graphically. The user interface of SUSA is a dialog program that guarantees a largely continuous analysis flow requiring as little user interference as possible.

The mainframe version of SUSA is appropriate for complex and time consuming applications, SUSA-PC for less demanding tasks.

HOW DOES SUSA SUPPORT THE ANALYSIS ?

Identification of Uncertainty Sources and Documentation of the Application

The concept of SUSA enables the user to fully concentrate his effort on the preparation of the analysis input that comprises the identification of all possibly important uncertainties, the appropriate representation by uncertain parameters, and the quantification of the respective states of knowledge including any relevant dependences.

For the purpose of documentation, SUSA requires some information on the current model application. The user is prompted to specify the name and the current version of the computer model as well as the corresponding assessment questions that are to be anwered by the model output.

State of Knowledge Quantification

After the model characteristics were specified, documentation and state of knowledge information on each uncertain parameter have to be entered. The documentation information comprises full name, physical unit, reference and best estimate value of the current parameter. SUSA-PC additionally prompts for the short program name and the symbol of the current parameter. Furthermore, the user is required to provide references to the literature and/or document his reasoning in writing.

The subsequent steps of the uncertainty and sensitivity analysis expect the state of knowledge of each uncertain parameter to be modelled in terms of subjective probability distributions. SUSA supports the transition from the expert's quantifications generally provided in terms of ranges and quantiles to the subjective probability distribution.

The following distributions are offered:

- (Log.) Histogram
- Polygonal Line
- Discrete Distribution
- (Log.) Uniform Distribution
- (Log.) Triangular Distribution
- Beta Distribution
- (Log.) Normal Distribution
- Weibull Distribution
- Gamma Distribution
- Extreme Value Distribution (Type I and Type II)

Suppose the state of knowledge does not allow the complete specification of the probability distribution, then SUSA derives the distribution parameters from quantifications like quantiles, truncations, median and factor, or expectation value and standard deviation.

After the distributional information on each uncertain parameter was given, SUSA prompts for the quantification of state of knowledge dependences. The information may be provided in terms of either population measures or sample rank correlations.

The population measures of association are:

- the Ordinary Correlation (Pearson)
- the Quadrant Measure
- Kendall's TAU

- Spearman's RHO

If the parameter sample to be generated should strictly reproduce the degree of the assumed population association, Spearman's Sample Rank Correlation is an appropriate empirical measure.

SUSA additionally offers the following alternatives for the quantification of state of knowledge dependences:

- complete dependence, in case the values of an uncertain parameter are suspected to satisfy an unknown monotonic relationship to another parameter
- conditional distributions, in case the state of knowledge of an uncertain parameter is different for each interval out of an exhaustive set of mutually exclusive intervals of another parameter
- design extension, in case the value of an uncertain parameter is a function of values selected for other parameters with fully specified marginals
- inequality, in case the value of an uncertain parameter is restricted by the value selected for another parameter

Generation of a Random Experimental Design

SUSA includes two alternative sampling procedures for Monte Carlo simulation: The Simple Random and the Latin Hypercube Sampling. The Simple Random Sampling selects values at random from a multivariate distribution satisfying the marginals and the dependence information. The Latin Hypercube Sampling provides sample values covering the range of each uncertain parameter evenly in probability. This is done by dividing the range of each parameter into eqiprobable intervals and selecting one value from each interval. This value may be the median of the interval, or it may be selected at random. It is impossible to derive uncertainty statements in terms of the commonly used tolerance limits (see paragraph: Uncertainty Statements) from a Latin Hypercube Sample.

Model Runs

Model runs are performed to obtain an estimate of the unknown probability distribution of the output quantities. Having access to the code, SUSA may initiate the runs automatically. If the program source including the model code is supplemented by adequate statements that read the sample of parameter values and finally produce a sample of model results, the corresponding JCL may be submitted from within SUSA. With another alternative offered by the mainframe version, SUSA runs a program that handles the data transfer from the generated experimental design to the model input file, then the JCL of the model code is submitted to derive the output. Details are described in (Kloos, 1993). If the model code is not too complicated, SUSA -PC may be used. The user has to insert the code into a Fortran skeleton, then SUSA-PC calls the compiler and performs the calculations. Even a code that is installed on a workstation may be initiated by SUSA-PC. In this case, SUSA-PC handles the data transfer from the PC to the workstation and vice versa and triggers the model runs on the workstation.

Finally, the results of the computation constitute a random sample from the respective unknown probability distribution of the model output.

Uncertainty Statements

SUSA offers the possibility to analyze scalar and index dependent model results. While in the first case, any output quantity of the model run is a single value, e.g. the peak cladding temperature in a loss of coolant calculation for a nuclear power plant, in the second case, it is a series of values over time, space, etc., like the cladding temperature at a specific location.

Details of the uncertainty and sensitivity analysis of the illustrative example are described in (Gläser et al., 1993).

Statistical (95%, 95%)-tolerance intervals serve as quantitative uncertainty statements saying that one can be at least 95% confident that at least 95% of the combined influence of all quantified uncertainties are between the interval endpoints.

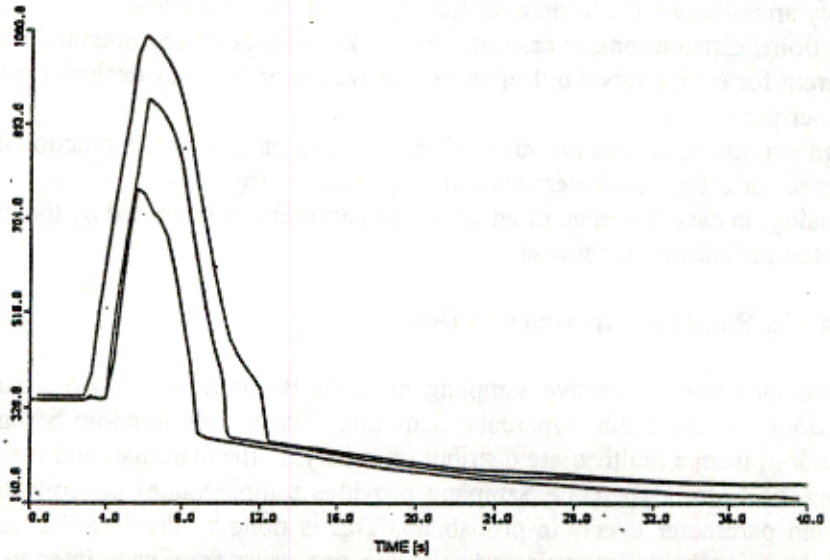


Figure 1. Statistical (95%, 95%)-tolerance intervals of the time-dependent cladding temperature at a specific location together with the solution obtained with the best estimate values of the uncertain parameters.

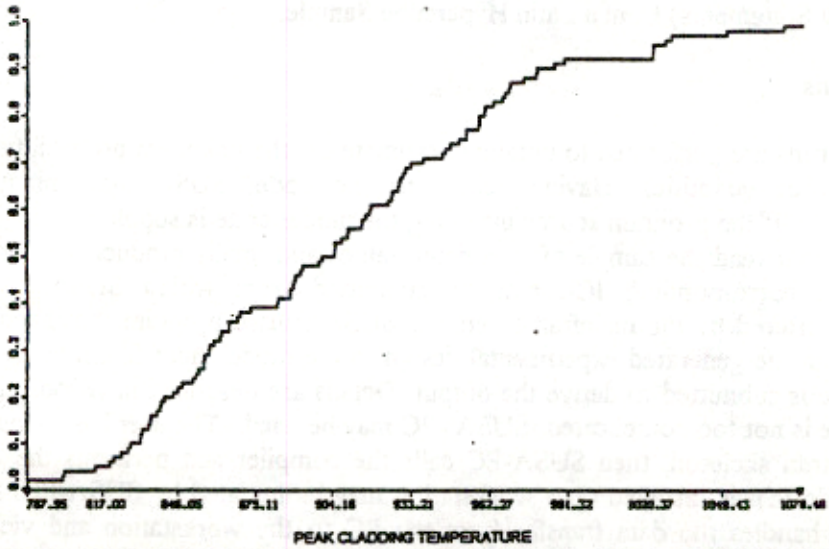


Figure 2. Empirical distribution function for the peak cladding temperature and the statistical (95%, 95%)-tolerance interval indicated on the temperature axis.

In addition to the tolerance limits, the user may require the computation and graphical representation of some statistics like minimum, maximum, mean, median, standard deviation, and the coefficient of variation of the model output as a function of the respective index set (time, space, etc). Upon request, the scalar uncertainty analysis with SUSA performs significance tests showing whether the selected parametric distribution provides a good fit to the empirical distribution of model results. If yes, the user may derive uncertainty statements from the fitted distribution analytically. If requested, SUSA establishes a regression function in the most important parameters (fitted response surface) selected by a

stepwise regression procedure and performs a Monte Carlo simulation using the response surface instead of the original code.

Sensitivity Statements

SUSA computes and graphically represents all important sensitivity measures of correlation type for raw and rank transformed output and parameter values. These are the simple and the partial correlation coefficient and the standardized regression coefficient derived from one of the classical types of correlation (Pearson's, Blomqvist's, Kendall's or Spearman's). In case some parameters represent modeling uncertainties in the form of indices of alternative model formulations, the correlation ratio is an appropriate sensitivity measure. It is less susceptible to the order in which the alternative models are indexed. In addition to the required sensitivity measures, SUSA calculates the corresponding coefficient of determination R^2 to indicate how much of the variation in the model output values is explained by the linear regression function. For scalar sensitivity analyses, SUSA additionally offers the quadrant measure from a 2x2 contingency table and the standardized (rank) regression coefficient from stepwise regression.

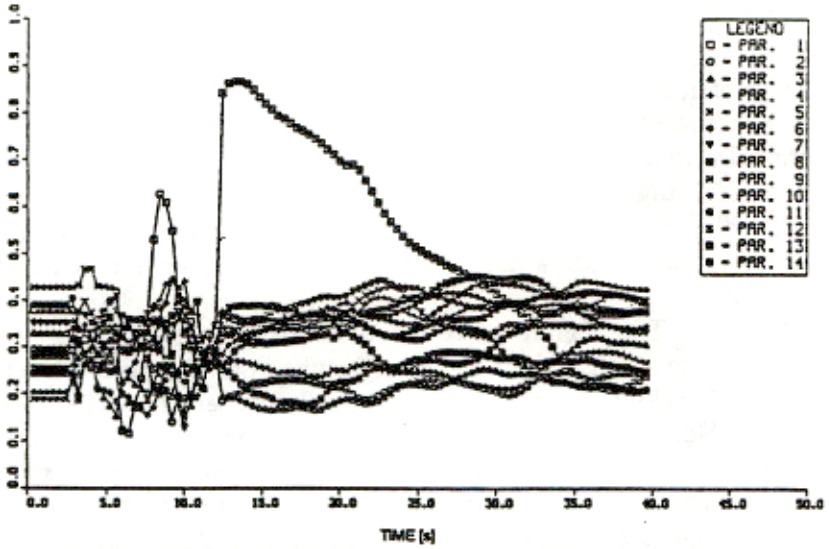


Figure 3. Correlation ratios for the first 14 uncertain parameters with respect to the computed cladding temperature at the specified location.

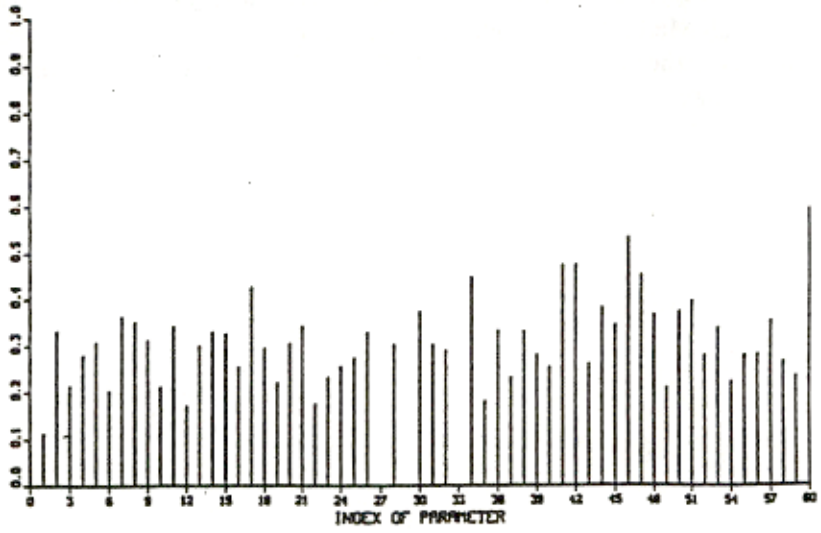


Figure 4. Correlation ratios with respect to the peak cladding temperature

SYSTEM REQUIREMENTS

SUSA-PC:
- IBM 386 or 486 PC or compatible with a math coprocessor
- at least 4 MB installed memory
- MS Windows 3.1 or higher
- MS Excel 4.0 or higher
- MS Fortran Compiler 5.1 or higher (necessary to handle non-correlation type dependence information or to perform model runs under the control of SUSA)
- Mouse driver
- PC-NFS 5.0 or higher for the workstation version

Mainframe version of SUSA:
- IBM compatible mainframe computer
- operating system: MVS/XA or MVS/ESA using ISPF/PDF
- Graphics: DISSPLA, SAS/GRAPH

REFERENCES

Gläser H., Hofer, E., Kloos, M., Skorek, T., 1993 "Uncertainty and sensitivity analysis of a post-experiment calculation for the OMEGA Rod Bundle experiment no. 9 with the thermal hydraulics code ATHLET", SMIRT-12, Post-Conference Seminar No. 15, Heidelberg, Aug. 23-24,1993

Hofer, E., 1993, "Probabilistische Unsicherheitsanalyse von Ergebnissen umfangreicher Rechenmodelle", GRS-A-2002, Gesellschaft für Anlagen- und Reaktorsicherheit (GRS), Garching

Kloos, M., 1993, "The Mainframe Version of the Driver of the Software System for Uncertainty and Sensitivity Analysis of Results from Computer Models", Version 2, GRS-A-, Gesellschaft für Anlagen- und Reaktorsicherheit (GRS), Garching

Krzykacz, B., Hofer, E., 1988, "The generation of experimental designs for uncertainty and sensitivity analysis of model predictions with emphasis on dependences between uncertain parameters", G. Desmet, ed., Reliability of Radioactive Transfer Models, Elsevier Applied Science Publishers Ltd., London and New York

Krzykacz, B., 1990, "SAMOS, EQUUS, TUSSIS: Programs for the derivation and graphical representation of uncertainty statements and sensitivity measures", Proceedings of the Seminar on Methods and Codes for Assessing the Off-site Consequences of Nuclear Accidents, Athens, 7.-11. May 1990, EUR13013, Commission of the European Communities

064 Risk Based Prioritization

Chair: G.E. Cummings, LLNL

Prioritizing Risk-Reducing Activities: Integrating Risk Management Into Business Planning

D. Brooks (Arizona St. U)

An Integrated Resource Management System and the Risk-Based Prioritization Model in Use at Martin Marietta Energy Systems

M.I. Sparks (Martin Marietta)

Risk-Based Prioritization: A Comparison of Approaches

R.S. Strait (LLNL)

PRIORITIZING RISK-REDUCING ACTIVITIES: INTEGRATING RISK MANAGEMENT INTO BUSINESS PLANNING

Daniel G. Brooks

Arizona State University
Tempe, AZ 85287-4206
(602) 838-8807

INTRODUCTION

The bases upon which priorities are set vary considerably from organization to organization, and often vary even within a single organization. The approach to prioritizing risk-reducing activities often depends on the source of the risk, rather than the risk itself or the value of averting the risk. For example, most organizations are highly formal in the way in which employee health-care risks are managed. Actuarial studies are the basis upon which employee health-care risks and benefits are determined; insurance and HMOs are some of the mechanisms that manage the risk. On the other hand, many organizations are self-insuring with regard to environmental risks and related impacts. This paper presents a sound methodology for managing risks: a systematic and logical means for prioritizing risk-reducing activities. The process is defensible, adaptable to a wide range of business activities, and provides an "audit trail" so that management decision makers can document the basis upon which they allocate resources to mitigate risks.

The sections that follow show how decision analysis using multiattribute preference theory provides a prescriptive methodology for aiding prioritization decisions when there are multiple risk objectives to be achieved, multiple stakeholders affected by the decisions, significant uncertainty about the outcomes of some of the activities, and the need for coordinating technical inputs from disparate disciplines.

DEVELOPING THE RISK-BASED PRIORITIZATION PROCESS

The overall process of risk management within which the prioritization of risk-reducing activities takes place is illustrated by the diagram in Figure 1. The process, incorporating both the definition of activities in the work breakdown structure and then the risk-based prioritization of these activities, is the basis for integrating risk management into the regular budget planning process: it combines the technical expertise needed to accurately assess risks with the policy judgments necessary to managing those risks. Developing a risk-based prioritization process that can be integrated with the long-term business planning of an organization means that the prioritization must

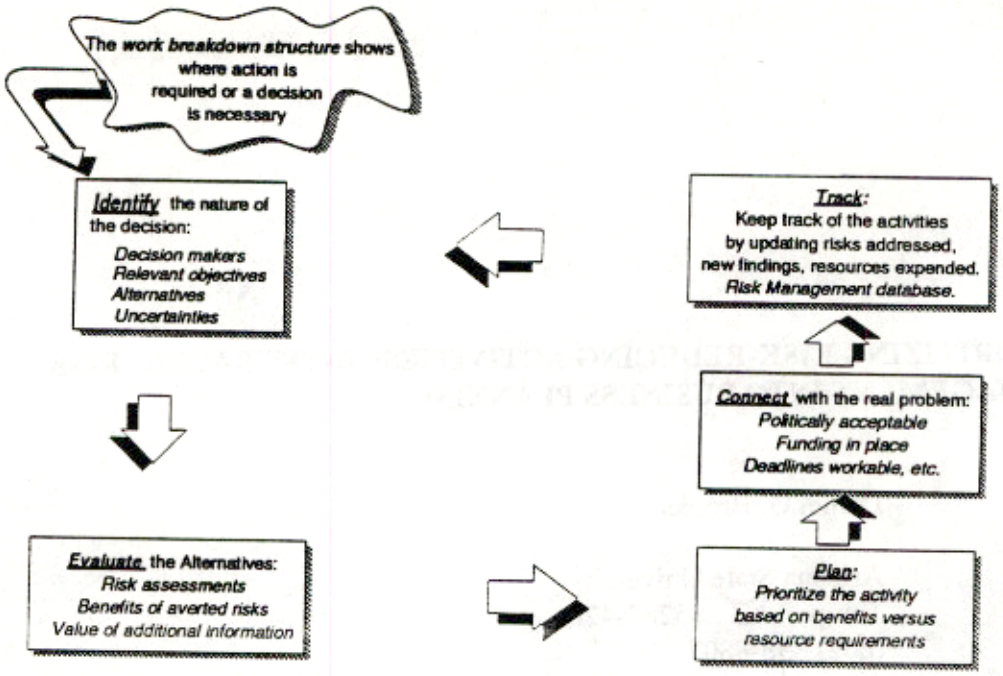


Figure 1. Components of the Risk Management Process

be consistent with the organization's strategic objectives.

Development of Risk Management Objectives

Risk is often thought of as the possible occurrence of an undesirable outcome. This definition includes the idea of *uncertainty*, since the outcome has not already occurred, and it includes the idea of *value*, since the outcome is identified as "undesirable." To usefully measure the uncertainty and the undesirability, quantitative methods are part of the risk management process. To measure "undesirability," the values of the organization must be defined. This is done by specifying the organization's objectives and then developing measures to show the impact on these objectives as a result of taking (or not taking) certain actions. The objectives that are identified by the organization for the risk management process define the dimensions of risk that will be used to evaluate proposed activities. A series of interviews and discussions are conducted with personnel involved in managing risk-reducing activities; the sets of objectives identified in these elicitations represent the basis for *evaluating* the effectiveness of a risk reducing activity: it is the degree to which that activity effectively enhances achievement of these risk-management objectives.

The prioritization process identifies the activities that are the most effective fixes by evaluating the amount of *averted risk* that can be attributed to those activities. Figure 2 illustrates a hierarchy of risk-management objectives. Note that this is prioritizing *solutions*, not *problems*. The most effective activity may not be addressing the most pressing problem.

Development of Risk Measures

To define what "undesirable" means for each of the evaluation measures requires development of *value functions*, where the value function specifies what is preferred as an outcome and what is not preferred over each measure. This is done in two steps: (1) develop a *measure* of impact for each category and then (2) develop a *value function*

over the range of possible measures, showing what is the best outcome and what is the worst. The development of measures and value functions for the measures requires more detailed definition of each category by subject matter experts. The level of detail in these definitions should match the needs of the risk management system applications.

The diagram in Figure 2 illustrates the conceptual process of developing a *measure* for one of the risk categories. Similar development takes place for each of the risk evaluation criteria defined by the set of risk management objectives. Few activities, if any, will require evaluation on every dimension of risk defined by the risk management objectives, but the comprehensive set of scales makes sure that no matter what dimension of risk is impacted by a project, that impact can be evaluated as part of the assessment of the value of the risk averted by the activities.

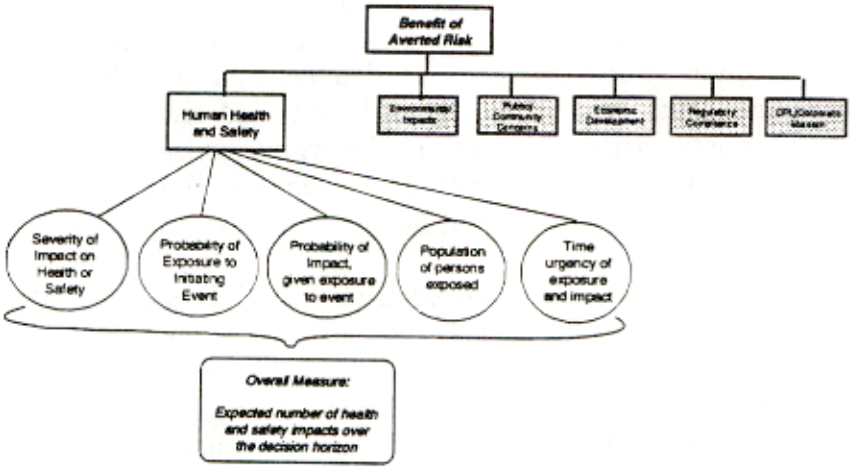


Figure 2. Developing a Risk Measure for Health and Safety

Assessment of the Overall Risk Impact of an Activity

Each evaluation criterion has a *measure* so it is possible to measure the degree to which an activity achieves the operational objective associated with that criterion. Each measure must then have a *value function* associated with it so that the *value* of various levels of achievement can be measured.

The value functions associated with the risk evaluation measures reflect the fundamental values of the risk managers and, thereby, the values of the organization. They are *not* a result of mass questionnaires, actuarial studies or summarizing data from Worker's Comp claims. These statistical averages show only what costs are associated with certain incidents, not the value of mitigating them.

To *combine* the impact values from each of the evaluation measures, there must be given "tradeoff" weights reflecting their importance relative to each other. It is important to remember that these tradeoffs are made routinely by the way in which funds are actively allocated. Addressing the tradeoffs explicitly makes managers aware of the tradeoffs they are using implicitly in making funding decisions. These weights represent an organization's policy in addressing competing risks with limited resources and time. These relative weights will be presented in the next section in economic terms.

Using the value function for the impacts within each category and the weights showing the relative importance's between categories, an overall risk impact for an activity can be computed by taking a weighted combination of the value of impacts within each of the categories. The combination rule depends on the degree of

independence of the risk categories. These independence properties are tested using a series of formally structured "independence" tests.

Estimation of Economic Value of Averted Risks

To give an economic value to the overall risk averted by an activity, some value must be attached to risks or impacts in each of the criteria areas. This value is determined using a series of tradeoffs between outcomes where the decision makers indicate their *willingness to pay* for certain types of outcomes (or willingness to pay to prevent these outcomes). This process is similar in concept to the bidding activity that determines the value of any product on the market. Although some of the outcomes may initially seem to be less tangible than others, a value can still be attached to these outcomes in the same way that the legal courts are able to attach some value to a person's good reputation after a slander trial. It is also important to note that the value of mitigating an activity is *not* a function of the cost of mitigating it. This confounds cost-benefit analysis: the whole point is to *separate* cost and benefit so that managers can identify over-priced fixes (cost exceeds benefit) and bargains (benefit exceeds cost). Making the benefit of fixing a problem a function of the cost of the fix *hardwires the benefit-to-cost ratio* and undermines the value of the entire process.

Instead, the monetary value of averted risk is based on management's fundamental values for averting competing types of risk. These values are determined through structured valuation of tradeoffs presented to managers. The monetary value of the overall risk averted by an activity is determined by estimating the averted risk for each evaluation measure and then using their weighted combination.

APPLICATION OF THE RISK-BASED PRIORITIZATION MODEL

An abbreviated example of an actual application follows; a fuller description of this is available in a paper by the author. In this application, the focus is on environmental management (an overhead operation) within a Federally-funded national laboratory. The risk-management objectives are shown in Figure 3. The risk measures associated with each of the operational objectives were developed and the set of tradeoff weights are illustrated by the values in Table 1. These values are illustrative only.

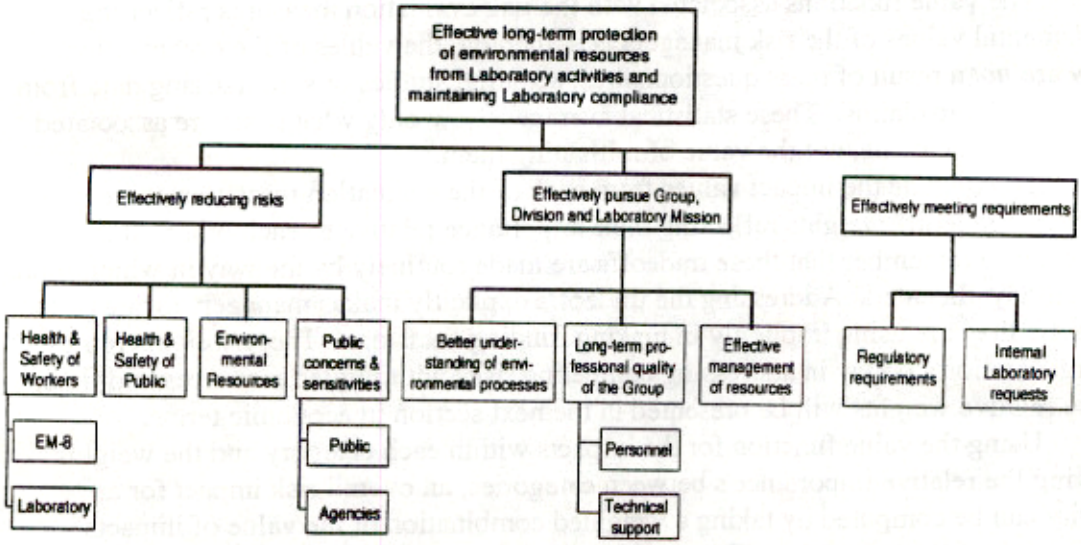


Figure 3. Risk Management Objectives for Environmental Management at a Federal Institution.

Risk Category	*Value of Coefficient*	*Highest-level outcome*
Health & safety of EM Workers	$2.5 million	Statistical fatality (weighted by population)
Health & Safety of Lab Workers	$2.5 million	Statistical fatality (weighted by population)
Health & Safety of the Public	$5.0 million	Statistical fatality (weighted by population)
Environmental Resources	$20 million	Elimination of a threatened and endangered species or a wetland equivalent
Public Concern	$5 million	Public demonstrations and negative national coverage
Regulatory Compliance	$25 million	Shut-down of operations
Professional Quality of Division	$2.5 million	Elimination of professional competence to judge environmental issues
Understanding of Environmental processes	$2.5 million	Elimination of ability to understand environmental risks
Management of Resources	$1 million	Inability to effectively operate the Division activities
Laboratory Requirements	$3 million	Inability to respond to a Lab need that effectively shuts down other programs at the Lab

TABLE 1: Risk Tradeoff Coefficients

Using this model, the economic benefit of averted risks compared to costs of their aversion can be used as the basis upon which activities will be prioritized. Table 2 summarizes costs and benefits for four (out of many) activities. The final column of the table shows the benefit-to-cost ratio. This information can be an aid to allocating resources so that *risk can be mitigated in the most effective way possible, given limited time, money and professional resources.*

Activity ID	*Activity Name*	*Benefit*	*1 Yr Cost*	*5 Yr Cost*	*Benefit/Cost*
WE1A-7	Air-net Monitoring: Pueblo Monitoring	$2,800,000	$51,000	$51,000	54.90
WE3E-7&8	NPDS Outfall Compliance: Additional Analyses	$2,062,000	$115,000	$575,000	3.59
WE6A-6	NEPA Compliance: Additional preparation of DECs	$306,000	$315,000	$1,575,000	0.19
WE1J-6	Ambient Air Quality Monitoring: Criteria Station	$4,000	$200,000	$600,000	0.01

TABLE 2: Summary Benefits and Costs of Four EM Activities

Figure 4 shows the four activities in terms of the cumulative contributions: the cumulative benefit shows the total value of the averted risks due to implementing the activities and the cumulative cost shows the total amount of support needed to implement the activities. Two specific capabilities of this decision-aiding process are very important if it is to be integrated into the true business planning process: (1) the results provide a basis for determining where to stop spending on activities. Only those activities with benefits at least as attractive as the cost should be supported in their current form. Where costs exceed benefits, there is incentive to re-examine the way in which those benefits are achieved. (2) Because resources come from different sources in different years with different limits, it isn't possible to just "follow a list until you run out of money." A *portfolio* of activities can be developed showing how to allocate budgets over the entire business planning cycle so that managers can consider how to manage risks over a multi-year planning process.

The portfolio is the result of solving a "packing problem" using linear-integer programming structure. Table 5 illustrates this capability. Selecting the set of action plans that yields the greatest benefit and can still be funded under the set of resource constraints (by year and source) is an optimization problem.

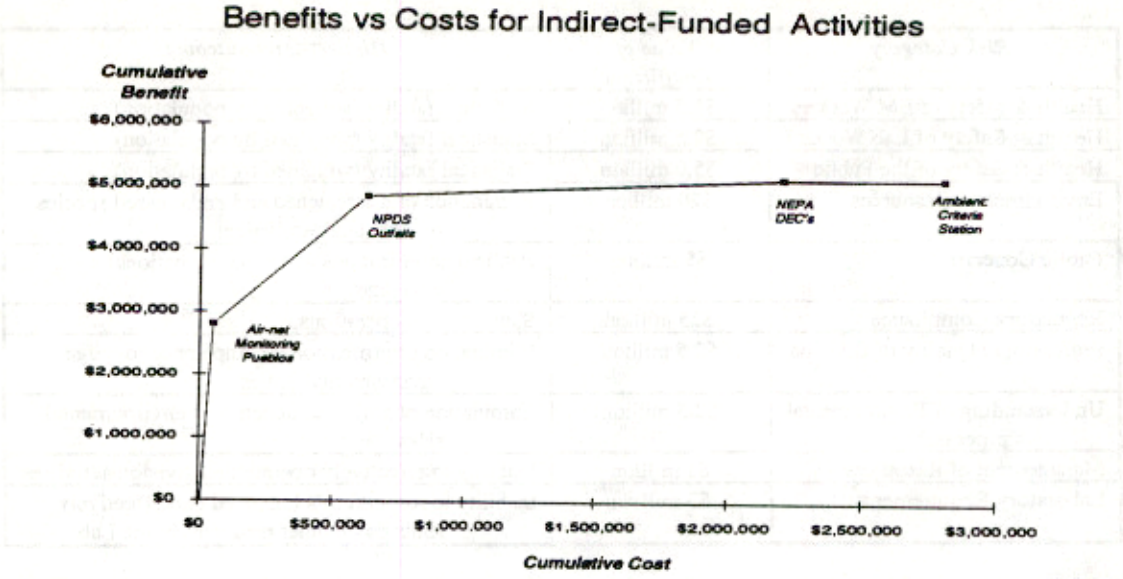


FIGURE 4: Cost versus benefit for four activities

Figure 5 shows the total portfolio of activities (beyond the four shown in Figure 4) to implement over the five-year planning horizon if the objective is to mitigate as much risk as possible over that time frame. The final link to actual project management, a crucial step of planning of risk management is to be converted to actual risk reduction, is a direct application of the results in this portfolio.

Start in FY92-FY93				Start FY94		Start FY95		Start FY96		Start FY97	
Start only these in FY92				Start FY93							
WS-12	42.2	EM-48	-.55	RP-12	-.99	PT-01	3.54	MGT-10	59.1	MGT-19	88.4
WS-02	36.4	EM-56	-.57			EM-50	3.54	FP-02	33.0	FP-14	26.2
PT-04	24.7	EA-02	-.76			EM-55	1.25	EM-52	17	EM-52	17.3
EM-18	13.3	MA-06	-.80			FP-01	.98	PT	4.9	EM-31	2.8
EM-35	11.5	MS-02	-.81			EP-07	.92			MGT-09	2.0
OA-01	8.2	PP-10	-.83			PP-09	-.01	EM-24		EM-23	.19
TS-02	7.4	EM-46	-.86			EM-14	-.28	EM-41	.03	EM-10	-.22
EM-25	4.6	ES-02	-.89			EM-34	-.67		.21	EM-43	-.23
NP-04	2.4	NP-06	-.91			ES	-.85	EM-29	-.33	EM-40	-.31
TC-04	2.2	MA-02	-.92			QV	.90	EM	-.35	MGT-11	-.63
QV-04	1.5	EM-15	-.93			NP-0		WS-0	-.85	MGT-12	-.79
OP-15	1.4	FP-03	-.94			FP-06	.9	CS-05	-.89	OP-04	-1
MGT-16	1.33	RP-20	-.95					PP-07	-.96	FP-10	-1
FP-13	.80	OP-19	-.99					EM-37	-.97	FP-15	-1
PT-03	.78	WS-13	-1	W	13 -1			FP-09	-.97		
RP-17	.31	AX-03						EM-28	-1		
OP-02	.06	EM	-1					FP-05	-1		
PP-03	-.17		-1			MGT-16	-1	MGT-16	-1		
NP-07	-.25		-08								
AX-07	-.25	FP	-1								
RP-21	-.26	MA	-1								
EM-32	-.36	MGT-1	-1								
MGT-02	-.41	RP-03	-1								
MGT-15	-.54										

EXAMPLE

Figure 5. Example portfolio of Risk Reducing activities for five-year Planning Horizon

SUMMARY

This approach to prioritizing risk-reducing activities is based on multiattribute preference theory. It provides a basis for documenting the definition of risk, the means by which its reduction is measured, and the effectiveness with which limited resources are used to mitigate risks. It has been reviewed by the National Academy of Sciences, outside technical review groups and peer reviewed in the risk-management literature. It provides a useful decision-aiding tool for managers responsible for managing risks.

ACKNOWLEDGEMENTS: This work supported in part by an LDRD Grant from Los Alamos National Laboratory for FY1993

AN INTEGRATED RESOURCE MANAGEMENT SYSTEM AND THE RISK-BASED PRIORITIZATION MODEL IN USE AT MARTIN MARIETTA ENERGY SYSTEMS, INC.

Mack I. Sparks

Policy and Management Systems
Martin Marietta Energy Systems, Inc.
Oak Ridge, TN 37831

This paper is intended to assist those who manage activities and tasks where there is more work to be accomplished than resources available.

The Integrated Resource Management System (IRMS) described in this paper was developed to provide Martin Marietta Energy Systems, Inc. (Energy Systems), the Department of Energy (DOE), and other DOE contractors with critical tools for more effectively balancing available resources with the needs competing for those resources. The IRMS process provides means to make better decisions and permits management to document and trace the basis for decisions relating to the expenditure of resources and the benefits expected from the resources expended. The IRMS was first employed within Energy Systems to help identify and focus resources on problems posing the highest risks to workers, the public, and the environment, and was intended to help ensure that relatively low-risk issues did not inappropriately detract from progress on high-risk issues. It also served as a way to integrate priorities of problems and solutions from many different sources.

A risk-based prioritization model is the "centerpiece" of the IRMS. Risk-based prioritization is not a new concept within Energy Systems and was developed from models used regularly within the commercial nuclear industry. Within Energy Systems it was first adapted for restart planning of the High Flux Isotope Reactor at Oak Ridge National Laboratory, and later evolved as an overall management tool to provide Energy Systems management with processes and decision-support tools to allocate resources, not just in the health and safety area, but in all aspects of business management .

During the 1980s, many commercial utilities invested in prioritization schemes that were consistent with a 1983 Atomic Industrial Forum white paper, "Living Schedule," that was an outgrowth of work (a risk-based prioritization system) initially developed by and implemented at the Duane Arnold facility.[1] The process of risk-based prioritization on which the IRMS model is based was developed from lessons learned from prioritization methods in use in commercial utilities to support upgrades required to meet Nuclear Regulatory Commission requirements.

Because of the extensive health and safety upgrades implemented throughout Energy Systems in 1989, a prioritization technique with credibility and universal application across multiple and diverse business units was needed. The IRMS has since evolved to become universally used within Energy Systems and has been adopted by DOE Headquarters for use in the DOE safety and health five-year planning process.

The functionality of the IRMS is based heavily on a prioritization matrix (Attachment I) composed of numerical values used in evaluating and prioritizing activities or projects in a detailed fashion. Energy Systems uses this risk-based prioritization methodology to ensure that valuable constrained resources are allocated in the most beneficial manner to achieve DOE objectives as expeditiously as possible. Overall, IRMS provides for the systematic collection of issues from all sources, the systematic analysis of those issues, and the proposal of solutions to those issues. Solutions are evaluated on the basis of the level of risk reduction that implementation will provide and the cost of that implementation. This process is then integrated into business management budgeting activities so that activities or solutions of greatest net benefit to the total organization are funded, and those of lesser benefit are unfunded or deferred until funding becomes available. This rational process enables both contractor and funding source to agree on a logical basis for funding or not funding initiatives, activities, or tasks. This agreement includes acceptance of the consequences associated with unfunded items.

A management review is performed which entails a project-by-project review of scope, cost, and risk reduction to ensure that all aspects have been properly considered and to make changes as deemed appropriate. Budget allocations are then made on the basis of a final prioritized list of projects (Attachment II).

The risk-based prioritization methodology has the ability to prioritize environmental, safety and health, safeguards and security, compliance, various legal requirements, mission objectives, support system effectiveness, efficiency activities, and upgrades in a manner consistent with Energy Systems objectives and DOE expectations. Prioritization is applied across multiple sites that respond to the different needs of several DOE customers and with different levels of maturity in affected programs. However, the application of prioritization for determining funding does not cross funding lines, but functions within funding categories and is typically used to distribute overhead or indirect resources.

The risk model employs risk analysis decision theory, public opinion, input from numerous technical areas, and management input, all of which are subjects of considerable discussion and differing expert opinion. In establishing this model, it was recognized that a perfect risk-based prioritization methodology is probably unattainable; however, a qualified system was sorely needed as a management tool to ensure that limited resources were effectively used. This qualified system was to be "reasonably accurate" and "easy to use." The model described in this paper is both. For purposes of this model, "risk" is defined as "the combined measure of consequences and probabilities of an undesired event." Therefore, the importance of a project is associated with a measure of the risk reduction it affords. This is based on a characterization of the risk in not performing an activity or completing a project, the risk inherent in the performance of a project, and the risk remaining on completion of the project or activity.

The risk prioritization model was developed after research of applicable literature and cost data, and after review and revision by Energy Systems management and, more recently, by management from various sections of DOE. An extensive external peer review/critique was also performed. The model selected employs multi-attribute utility theory to determine values associated with each of the consequences including, cost, impact and unwillingness.

Potential consequences that apply generically to all projects were identified and divided into seven general categories: public health and safety, environmental protection, site personnel safety, regulatory compliance, external confidence, mission and operational performance, and business efficiency. Consequences within the category of ***public health and safety*** address adverse effects on the health and safety of the surrounding population. Consequences associated with ***environmental protection*** include those that could result in physical degradation of the ecological system. ***Site personnel safety*** consequences address adverse effects on the health and safety of workers and other individuals who are voluntarily inside site boundaries. ***Regulatory compliance*** consequences consider the risk of fines, imprisonment, or loss of confidence by DOE stemming from violation of DOE requirements and other legal or contractual obligations. This category also recognizes the importance of adherence to proper procedures and good industry practices. Consequences associated with ***external confidence*** consider public perception or socioeconomic effects that arise from loss of public confidence in the ability of Energy Systems to operate in a sound manner, consistent with the public's expectations.

Mission and operational performance includes consequences negatively affecting the ability of Energy Systems to carry out its various missions, including shutdown of research reactors, significant failure to meet DOE production needs or mandated programmatic directives, and failure to meet DOE goals, objectives, and mission imperatives. Consequences in the category of ***business efficiency*** include those that result from capital losses or equipment damage from fires or other avoidable events that significantly increase production costs or prevent cost savings, such as lost opportunities to reduce production costs or lawsuits. The investment dollar limits within this category are generally based on limits established by DOE Order 5480.7, *Fire Protection.*[2]

The seven consequences categories are further subdivided into a total of 22 consequences. Combining the 22 identified consequences with their probability of occurrence leads to the risk matrix presented as Attachment I. To determine the risk reduction afforded by a particular activity or project, one must know the events for which the project was designed, the consequences that will be influenced by the project, and the probability of occurrence of each consequence. The scoring process calls for the selection of only one consequence/probability value from each of the seven groups. A project is evaluated by considering the difference in the sum of these consequence/probability scores for an issue before the project is implemented and the sum of the consequence/probability scores resulting from the application of the project.

The risk-based ranking is performed by teams where the individuals are very knowledgeable in the areas subject to evaluation, such as health or safety, environmental protection, mission, or regulatory compliance. These teams are usually called project evaluation groups. Their task is to evaluate each project based on information provided by the project owner and to score the risk associated with an issue or the risk reduction afforded by implementing a project. The procedure for evaluation of projects also calls for the consideration of benefit in terms of risk reduction relative to the cost of the project. The data base used to collect IRMS information will calculate scores and order project lists by score (total benefit) ranking or score/cost (benefit/cost) ratio. This facilitates management decisions on the best ranking for projects and the associated funding allocations.

The risk-based model provides a reasonable approach to prioritization, which is evidenced by the fact that management makes relatively few changes to the prioritized lists developed by project ranking teams. The risk-based prioritization model serves as a tool to assist management with their job of making informed, appropriate decisions concerning resource allocation; but, the model is not a substitute for common sense and management expertise. Conducting risk assessments on the basis of expert judgment has been practiced for years. The methodology utilized at Energy Systems provides a process for documenting technical decisions and for assisting management to make more effective decisions — and most importantly, it works.

---

[1] Atomic Industrial Forum (1983) "Living Schedule" Concept: A White Paper for Industry. Prepared by Committee on Reactor Licensing and Safety Subcommittee of Back Fit Requirements.

[2] C.B. Smith, "Basis Documentation for the Martin Marietta Energy Systems, Inc., Risk-Based Prioritization Methodology (1991).

ATTACHMENT I

Martin Marietta Energy Systems Issue/Project Evaluation Matrix

CONSEQUENCES	LIKELIHOOD OF OCCURRENCE A VERY HIGH	B HIGH	C MEDIUM	D LOW
PUBLIC HEALTH AND SAFETY				
1. Immediate or eventual loss of life/permanent disability	3,000	300.0	30.0	0.3
2. Excessive exposure and/or injury	300	30.0	3.0	0.03
3. Moderate to low-level exposure	30	3.0	0.3	0.003
ENVIRONMENTAL PROTECTION				
4. Extreme damage to the ecological system	1000	100.0	10.0	0.1
5. Extensive damage to the ecological system	200	20.0	2.0	0.02
6. Significant damage to the ecological system	50	5.0	0.5	0.005
SITE PERSONNEL SAFETY				
7. Immediate or eventual loss of life/permanent disability	1500	150.0	15.0	0.15
8. Significant injury requiring hospitalization with significant lost-time, or exposure > occupational limits	150	15.0	1.5	0.015
9. Exposure near limits (20 to 100%) or lost-time injury requiring medical treatment, or contaminated wound	15	1.5	0.15	0.0015
10. Minor injury requiring first aid, or exposure < 20% of limits, or reportable, removal skin contamination	1	0.1	0.01	0.0001
REGULATORY COMPLIANCE				
11. Noncompliance with Federal, State, or local law involving significant potential fines or penalties	100	10.0	1.0	0.01
12. Noncompliance with DOE Category A Orders	50	5.0	0.5	0.005

Martin Marietta Energy Systems Issue/Project Evaluation Matrix (continued)

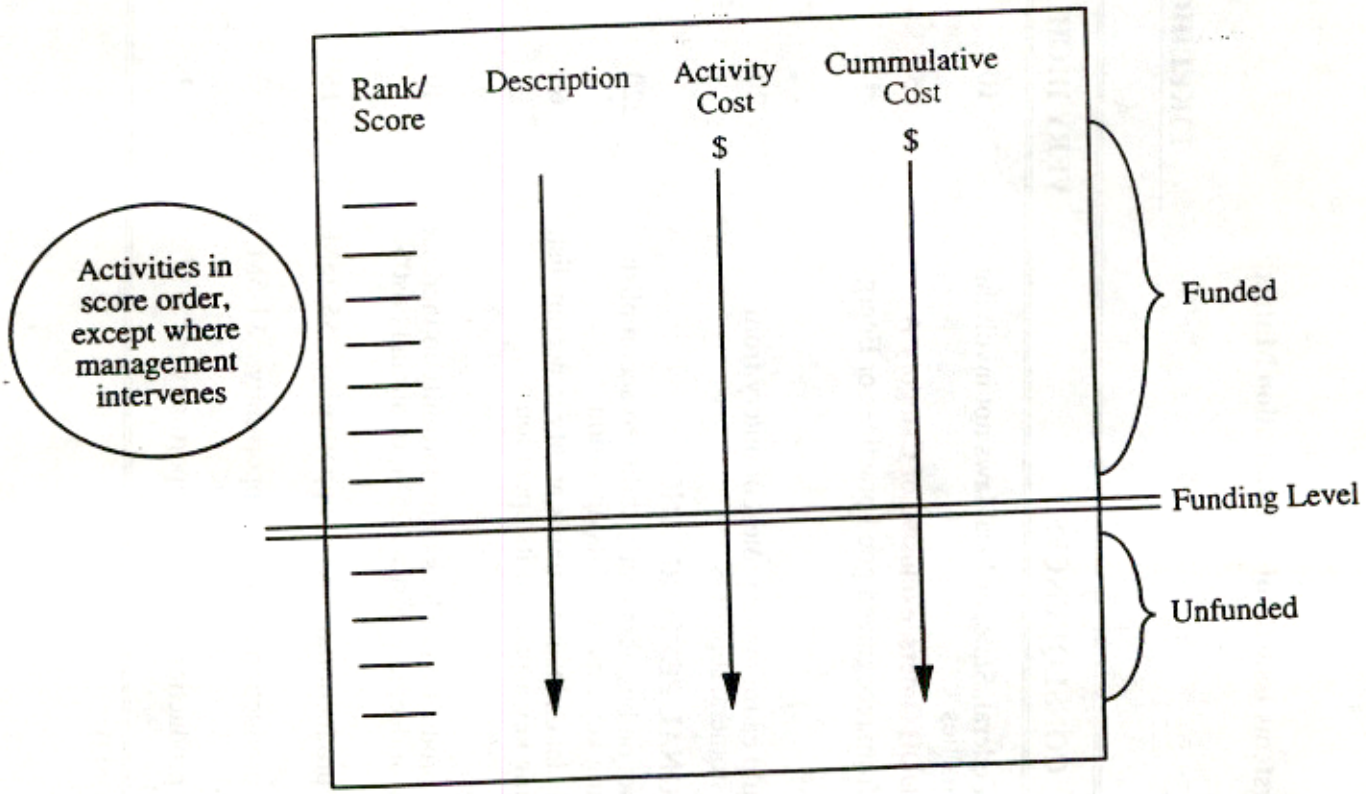
CONSEQUENCES	LIKELIHOOD OF OCCURRENCE A VERY HIGH	B HIGH	C MEDIUM	D LOW
14. Noncompliance with Federal, State, or local laws not involving significant fines or penalties	10	1.0	0.1	0.001
15. Noncompliance with DOE Orders, exclusive of Category A	5	0.5	0.05	0.0005
16. Significant deviation from recognized good practices or Energy Systems directives	1	0.1	0.01	0.0001
EXTERNAL CONFIDENCE				
13. Issues that have or could cause major protest or outcry from public, employees, customers, suppliers	20	2.0	0.2	0.002
MISSION AND OPERATIONAL PERFORMANCE				
17. Serious negative impact on business unit's ability to accomplish its mission by sustaining key aspects of its operation	220	22.0	2.2	0.022
18. Moderate negative impact on business unit's ability to accomplish its mission by sustaining key aspects of its operation	60	6.0	0.6	0.006
BUSINESS EFFICIENCY				
19. Loss of investment or production or loss of opportunity to increase quality or productivity or decrease cost > $25M, or annual cost > $5M	40	4.0	0.4	0.004
20. Loss of investment or production of loss of opportunity... $5-25M, or annual cost $1-5M	15	1.5	0.15	0.0015
21. Loss of investment or production or loss of opportunity... $1-5M, or annual cost $0.2-1M	3	0.3	0.03	0.0003
22. Loss of investment or production or loss of opportunity... <$1M, or annual cost <$0.2M	1	0.1	0.01	0.0001

(Energy Systems Matrix, 9/30/92)

MARTIN MARIETTA
MARTIN MARIETTA ENERGY SYSTEMS, INC.

ATTACHMENT II

"A simple process that supports and focuses management thought and decision making and does not replace them . . ."



RISK-BASED PRIORITIZATION: A COMPARISON OF APPROACHES[1]

R. Scott Strait

Lawrence Livermore National Laboratory
Livermore, California 94551

INTRODUCTION

The use of analytical techniques for risk-based prioritization, both quantitative and non-quantitative, is becoming much more wide spread and organizations are tailoring analytical techniques for risk management to their particular needs. This has led to the use of different techniques for different applications within the same organization, and correspondingly, the need for management to reconcile the process and results of different risk-based prioritization approaches. Different risk-based prioritization methods may be used to prioritize different types of activities, but eventually these prioritizations may need to be integrated to allocate organizational resources at a higher level or to develop an overall organizational budget.

This paper focuses on comparing and contrasting the different risk-based approaches. The objective is to help management understand the different approaches, choose the best approach for their needs, and reconcile different approaches. Most managers will use risk-based prioritization to assist in determining which risk reduction activities are most cost-effective or, even, which are cost-beneficial. While there are a large number of risk analysis approaches that are appropriate for choosing between a small number of risky alternatives, the number of approaches that are appropriate for risk-based prioritization is much smaller. Such approaches must be designed for evaluating large numbers of alternatives, while requiring only minor amounts of analysis. (All of the approaches could use the results of more extensive analysis, e.g., probabilistic risk analysis, if they were available.)

CHOICE OF A RISK-BASED PRIORITIZATION APPROACH

Choosing a risk-based prioritization process is a management decision, which should balance the pros of the approach against its cons. Generally this is a tradeoff between the approaches' ability to meet management needs and the effort required to implement it. The factors that determine an approach's ability to meet management needs include: its ability

[1] Work performed under the auspices of the U.S. Department of Energy by the Lawrence Livermore National Laboratory under Contract W-7405-Eng-48.

to differentiate among options; its defensibility and traceability; and its ability to accurately capture the risk reduction effect of an activity. In general, more complex approaches meet more of management's needs, but require more effort to implement. They require additional effort for development, scoring, and analysis with commensurate gains in other areas, e.g., theoretical soundness and the ability to differentiate among options.

The effort required to implement a risk-based prioritization approach is highly dependent on the problem, including the nature and number of the risk reduction activities being considered and the organizational constraints. Much of the effort required to implement a risk-based prioritization approach is independent of the approach chosen and arises from the need for data collection and documentation. To a large extent, it is this data collection and documentation that provides most of the defensibility and traceability of the prioritization results and that contributes to the overall quality of the process.

This paper focuses only on the abilities of different approaches to meet management's needs, and not on the effort to implement a prioritization approach. The effort required is much more problem-related than approach-related. The comparison will center on how each approach treats the components of the standard equation for risk:

$$P \times C = R$$

where: P is the probability of loss; C is the adverse consequences of the loss; and R is the expected loss, often used as a measure of risk. This paper examines each of four general types of approaches and how they deal with each of the components of this equation. It addressess how they measure probability; combine probability with consequences; measure consequences, including the population at risk; consider the tradeoffs among different types of risk; and prioritize the benefit of a risk reduction activity.

APPROACHES FOR COMPARISON

Although there are wholly non-quantitative approaches for risk-based prioritization, this paper will cover only those that involve at least a minimal amount of scoring and aggregation of scores. Approaches with less structure do not meet the needs of most managers seeking to prioritize activities. To represent the range of possible risk-based prioritization approaches, four of the more common ones will be considered. These four, in order from least structured to most structured, are: binning systems (e.g., the "hazard totem pole," Grose, 1987); quantitative aggregations of scores and weights, sometimes termed "rate 'em and weight 'em" (e.g., analytic hierarchy process, Saaty 1980); matrix approaches (e.g., the priority planning grid, Smith, 1991); and multi-attribute decision analysis approaches (e.g., Keeney and Raiffa, 1976).

Binning systems have separate sets of ordered bins for each factor considered important by the decision-maker in evaluating risk reduction activities. For example, an analysis may have a set of bins for worker safety (minor injury, major injury, disability, and death); a set for the probability of an accident (0.0 to 0.1, 0.01 to 0.1, 0.1 to 0.5, and 0.5 to 1.0); and a set for the cost of the risk reduction activity. Each possible combination of bins is then ordered according to the decision-maker's preferences, creating a "risk totem." Quantitative aggregation approaches typically have a relative 0 to 10 scale for each factor considered important by the decision-maker. Each factor is assigned a weight by the decision-maker (possibility using the analytic hierarchy process) and the benefit of a risk reduction activity is the weighted sum of the scores of an activity. Matrix approaches typically list levels of possible consequences on a vertical axis and ranges of probabilities of the consequences on the horizontal axis. Weights are assigned to each level on the vertical axis by assuming and assessing a linear additive multi-attribute utility function to reflect the decision-maker's preferences. Multi-attribute decision analysis approaches develop continuous or discrete scales, as appropriate, for every factor considered to be

important. Scores of an activity on the scales are combined using a multi-attribute utility function.

In comparing the different approaches, this paper will consider how the approaches are generally applied as suggested by their practitioners. All of the approaches have theoretical capabilities that are beyond the usual practice; however such theoretical possibilities will not be considered here. Because their practitioners do not usually apply these capabilities, this paper assumes that they are not in the "spirit" of the approach as applied to risk-based prioritization and should not be the basis for a management choice of a risk-based prioritization approach. Conversely, all of these approaches have, at one time or another, been implemented in an incorrect or inconsistent manner, but this comparison will assume that the approaches are properly implemented. For any particular application, all of these approaches can be modified or combined to meet specific management requirements, but only generic applications will be considered in this paper.

TREATMENT OF UNCERTAINTY

The different risk-based prioritization approaches vary in how they measure uncertainty. Binning systems and matrix approaches measure uncertainty similarly. They both imply discrete ranges or bins of probabilities. Generally these ranges are fairly broad with adjacent bins differing by an order of magnitude. Usually matrix approaches define these levels of probabilities as one axis on the matrix while binning systems use ordered bins. The number of discrete levels of probabilities is only limited by the size of the matrix or desired number of bins. Quantitative aggregation approaches use relative assessments of likelihoods such that an event is placed on a scale from least likely to most likely. These likelihoods are usually not defined as absolute probabilities. Decision analysis approaches use continuous scales corresponding to the actual probabilities. In general the different approaches are reconcilable and there is a natural migration from bins and levels on the matrix axis, to continuous probabilities. Reconciling the other three approaches with the quantitative aggregation approaches requires information about the likelihoods that is not normally assessed and requires some assumptions about the functional form of the relative scale.

Despite the similarities in how the different approaches measure uncertainty, they vary more in how they combine uncertainty with the levels of possible consequences. Binning systems use a qualitative and additive system that is very difficult to reconcile with the other approaches. The combination of consequences and probabilities in a "risk totem" is not consistent with the usual definition of risk. The quantitative aggregation approach is quantitative, as its name implies. However, it usually combines probability and consequence in an additive fashion, in a manner inconsistent with the usual definition of risk. There are no constraints that prohibit a multiplicative model, but such a model would complicate other aspects of the approach's implementation (see the discussion below about the tradeoffs among different types of objectives). The matrix and decision analysis approaches are both quantitative and follow the multiplicative form of the standard risk equation. Those two approaches are wholly consistent in this regard.

TREATMENT OF ADVERSE CONSEQUENCES

There are several aspects to the treatment of the impacts of an adverse event. One is the measurement of the severity of the impacts or consequences The binning system requires the user to assign an event to generally defined bins of adverse effects. There may be different series of bins for different types of consequences. The quantitative aggregation approach provides the user scales for different consequences and requires the user to assign a relative rating on each scale. Matrix approaches use well-defined discrete levels of consequences. These comprise one axis of the matrix. Decision analysis uses well-defined discrete or continuous scales, depending on the effect. In general the different approaches

are reconcilable and there is a natural migration from bins, to levels on the matrix axis, to decision analysis scales. Reconciling the other three approaches with the quantitative aggregation requires information about the consequences that is not normally assessed and requires some assumptions about the functional form of the relative scale.

Often an important factor in determining the consequences of an event is the population at risk. The binning, quantitative aggregation, and matrix approaches all usually assume a constant number of injuries or casualties for an analysis. The quantitative aggregation approach sometimes treats the number of casualties as another type of consequence, although this has the problem of leading to a linear, rather than multiplicative, combination of the number of injuries or casualties and the their severity. The decision analysis approach treats the population at risk explicitly. The binning and matrix approaches can be reconciled with the decision analysis approach, but reconciling the quantitative aggregation approach requires information about the population that is not normally assessed. The treatment of the population at risk by the binning, quantitative aggregation, and matrix approaches usually leads to a lexicographic ordering.

For many risk-based prioritization efforts, multiple types of consequences are of concern, e.g., public health and safety, environmental issues, and business objectives. This requires that prioritization approaches reflect management's tradeoffs among the different types of consequences. The binning system approach handles the tradeoffs among types of consequences qualitatively in the ordering of different combinations of bins, e.g., the risk totem. The quantitative aggregation approach weighs each consequence scale as part of a linear combination. Both the matrix and decision analysis approaches use utility theory to measure and model the relative tradeoffs among types of consequences. However matrix approaches are more restrictive and assume linear additive utility functions. If similar types of consequences are considered and if the utility functions of the matrix and decision analysis approaches are linear, then some reconciliation among all four approaches is possible. As in other areas, reconciling the other three approaches with the quantitative aggregation requires information about the consequences that is not normally assessed.

BASIS FOR PRIORITIZATION

The cost effectiveness of a risk reduction activity is usually determined by a benefit/cost ratio. Activities with higher benefit to cost ratios are assumed to be more cost effective. The different approaches vary in how they measure the benefit of a risk reduction activity. The binning system and quantitative aggregation approaches usually measure benefit as the level of risk. Therefore resources are allocated to risk reduction assuming that all risk reduction activities completely eliminate the related risk. The matrix and decision analysis approaches explicitly calculate averted risk as the measure of the benefit of a risk reduction activity. These differences among the approaches are difficult to reconcile, although in many applications it may be reasonable to approximate the averted risk as the total risk without skewing the results.

The ratio between benefits and costs of risk reduction activities under the binning system and quantitative aggregation approaches is only relative. That is, it is possible to compare the cost effectiveness of one activity relative to another (given the qualifications noted above), but it is not possible to ascertain whether the risk reduction is cost beneficial in an absolute sense. If monetary losses are one of the types of consequences, then both the matrix and decision analysis approaches provide an absolute measure of cost effectiveness. In this case, risk reduction activities whose benefit cost ratios exceed 1.0 are cost effective and those with ratios below 1.0 are not.

CONCLUSIONS

The preceding comparison of risk-based prioritization approaches is summarized by Table 1 on the following page. The different risk-based prioritization approaches vary in

terms of how they treat the components of risk and how they measure the benefits of risk reduction activities. The comparison should be useful to managers in: understanding the important assumptions inherent in the different approaches; choosing a prioritization approach; interpreting the results of prioritization models; and reconciling prioritizations determined by different approaches. This comparison of approaches may also be useful in developing two-step prioritization approaches. In a two-step approach, a simpler prioritization approach would be used to screen risk reduction activities and a more complex approach would be used to analyze a subset of activities (e.g., those near the budget limit).

Table 1. The different risk-based prioritization approaches vary in terms of how they treat the components of risk and how they measure the benefits of risk reduction activities.

	Risk -Based Prioritization Approaches			
	Binning System	Quantitative Aggregation	Matrix	Decision Analysis
Measure of probability	Bins of probabilities	Relative assessments	Discrete levels	Continuous
Combination of probability and consequences	Qualitative and additive	Quantitative	Probabilistic and quantitative	Probabilistic and quantitative
Measure of consequences	Bins	Undefined relative scale	Well-defined discrete levels	Well-defined continuous or discrete scales
Population at risk	Constant number of injuries or casualties	Treats number of injuries as if another consequence	Constant number of injuries or casualties	Addressed explicitly
Measure of benefit	Risk	Risk	Averted risk	Averted risk
Treatment of tradeoffs among types of consequences	Qualitative	Ad hoc quantitative	Utility theory	Utility theory
Benefit/cost ratio	Relative	Relative	Can be absolute	Can be absolute

Management's choice of a risk-based prioritization approach needs to balance the effort required to implement an approach against the approach's ability to differentiate among options; its defensibility and traceability; and its ability to accurately capture the risk reduction effect of an activity. More structured approaches provide more capabilities, but require more effort to implement. They require additional effort for development, scoring, and analysis with commensurate gains in other areas, e.g., theoretical soundness

and the ability to differentiate among options. The effort required to implement a risk-based prioritization approach is highly dependent on the problem, including the nature and number of the risk reduction activities being considered and the organizational constraints. Much of the effort required to implement a risk-based prioritization approach is independent of the approach chosen and arises from the need for data collection and documentation.

REFERENCES

Grose, V.L., 1987, *Managing Risk: Systematic Loss Prevention for Executives*, Prentice Hall, New York

Keeney, R. and H. Raiffa, 1976, *Decisions with Multiple Objectives*, Wiley, New York.

Saaty, T., 1980, *The Analytic Hierarchy Process*, McGraw-Hill, New York.

Smith, C., 1991, *Basis Documentation for the Risk-Based Prioritization Methodology*, Tenera, L.P., Knoxville, TN.

065 Meaning and Interpretation of Probability

Chair: A. Mosleh, U. Maryland

The Meaning and Use of Probability in Probabilistic Safety Assessment
G.W. Parry (NUS)

Is Probability of Frequency Too Narrow?
H.F. Martz (LANL)

On Bayesian Approaches of Parametric and Modeling Uncertainty
J.S. Wu (ASCA)

On The Objectivity of Subjective Theory of Probability for PSA Applications
A. Mosleh (U. Maryland)

THE MEANING AND USE OF PROBABILITY IN PROBABILISTIC SAFETY ASSESSMENT

Gareth W. Parry

NUS
910 Clopper Road
Gaithersburg, Maryland 20878

INTRODUCTION

In Probabilistic Safety Assessments (PSAs), probability is used in two very distinct ways. Firstly, and arguably this is the origin of the use of the word probability in the term PSA, probability is used to characterize outcomes of so-called random processes. Examples are to be found in the reliability modeling techniques used to quantify fault tree models to predict system unavailabilities, or in the dispersion models used to predict off-site consequences. The second use of probability is as a measure of an analyst's uncertainty about the validity of a proposition, whether that proposition concerns the value of a parameter, or concerns whether a particular phenomenon will or will not occur for a given set of boundary conditions. This latter case is an example of modeling uncertainty. To lessen the confusion, some analysts [1] have proposed to use (relative) frequency for the first usage, and reserve the word probability for the second. We will use the word probability for both usages, and rely on the context to provide the distinction.

The use of probability to characterize uncertainty in parameter values for reliability models has become generally accepted practice, though it has not been without its critics. However, considerable discussion and some controversy has been generated over the issue of how best to address questions related to modeling uncertainty of the type referred to above. This was first raised in relation to the development and quantification of containment event trees where the branches relate more to the representation of physical process consequences than engineered system behavior[2]. This issue has re-emerged with the interest in reactor systems with passive safety features which rely on physical processes rather than mechanical equipment for successful mitigation of accidents.[3] As in the case of containment event trees, it is to be expected that the branches on the event trees may represent differing outcomes of physical processes rather than traditional system availability issues. The issues that can be raised with respect to such branch points are: firstly, whether these branch points should be described with parameters that express randomness or variability, or with probabilities that represent the analyst's relative degree of belief about which path

through the branch is appropriate for the conditions addressed in the PSA model, or by a mixture of both, and secondly, whether indeed there is any need to maintain this distinction in the evaluation and presentation of the risk profile of the plant.

Some analysts have claimed that distinguishing between these two types of uncertainty, one representing the random or variable behavior of the system, the other the degree of belief of the analyst concerning possible outcomes of a process, is artificial, and furthermore is difficult to achieve. Others, on the other hand, believe that these state of knowledge uncertainties should be treated differently from the parameters characterizing the random or variable processes. They are concerned, for example, that failure to do so will lead to an underestimation of the uncertainty.

This paper presents a discussion of these issues by addressing the following questions:

- Is it necessary to have two types of uncertainty in the model, one characterizing outcomes of random or variable processes, the other the analyst's modeling uncertainty?
- If so, can they indeed be distinguished?
- Should they be separated?

WHY TWO TYPES OF UNCERTAINTY?

It is probably true to say that, as indicated below, from a purely subjectivist point of view, there is no need to have two measures of uncertainty. However, as the following discussion shows, the approach that has been historically adopted in PSAs does indeed have this feature.

It can be argued that the concept of randomness or variability is primarily invoked as a result of our inability to model the things we are interested in sufficient detail, given the level of definition of the boundary conditions adopted, that we can predict with certainty the outcome of an event. To overcome this inability we choose to describe our predictions in terms of probabilities. These probabilities essentially represent the relative fractions with which various outcomes would be expected given the situation being analyzed were hypothetically to be repeated a large number of times.[4] What these fractions represent is not "inherent" randomness in nature, but the fact that, at the level we have chosen to define events, they must be described that way. There are variabilities in underlying conditions that would govern the behavior of individuals in the population that are not being explicitly accounted for. Instead they are implicit in the form of the (probability) models we use for the basic events. Thus, these fractions, or whatever we choose to call them, are parameters that are used to describe our model of the world. This "World" encompasses the hidden variables by forcing us to characterize the system behavior as random or variable. For this to be a useful model, we need a mechanism for estimating the values of the parameters of the model and of updating them given new information. In most recent PSAs, this has been accomplished by adopting a Bayesian or subjectivist framework[5] which uses probability as an index of the analyst's assessment of possible values of the parameters, thus introducing a second level of "uncertainty".

Of course, it can be argued that, if the initial problems had been defined differently, only one measure of uncertainty would have been necessary. For example, if, instead of adopting a model of the world based on relative frequencies, the question had been phrased, "Will event X occur next year?", the outcome is binary, the event will occur or it will not, and all the hidden variability must be encompassed in the

However, in answer to the first question, it is clear that the way the PSA community has traditionally chosen to perform PSAs has led to a two-level description of uncertainty, one representing randomness as a means of capturing the variability that is not explicitly modeled, the other uncertainty in our knowledge about this variability.

CAN WE DISTINGUISH RANDOM AND STATE OF KNOWLEDGE COMPONENTS OF UNCERTAINTY?

This question arises when we start looking at "events" on event trees that address the outcomes of physical phenomena, the validity of models, or the opinions (not to be confused with estimates) of experts. An example of the first type of branch point is that which was introduced in Ref. 2, concerning the question of the potential for a steam explosion. Suppose the question appears on a containment event tree that is to be analyzed for a core damage sequence taken from a traditional event tree. Since what we are concerned with is the performance of a PSA, we need to provide an assessment of the "likelihood" of this containment event tree so that it can be combined with the probability or relative frequencies associated with the system unavailabilities.

It is true that, in a PSA, the conditions for which the containment event tree branch point is being evaluated are not defined precisely, but more generally. Because our underlying models are based on a description of the world as exhibiting random or variable behavior, the initial conditions should also appear to be random or variable. The core damage sequence is in reality representative of a class of possible real world scenarios which may vary in many aspects. Thus, the answer to the question we are asking should be evaluated taking this variability into account. If this were the case, the "likelihood" could be a measure of the relative fraction of the realizations of that sequence which lead to the occurrence of the steam explosion. This "likelihood" would then be a parameter of our model of the world and would address the variability implicit in that model. However, our reliability models generally do not address that variability in the degree of detail necessary (after all we earlier suppressed our knowledge about the origins of that variability), and the analyst may find assessing this fraction is extremely difficult. In order to assess this fraction he must construct, if not a mathematical, at least a mental model of the underlying variability and its consequences.

In some cases, the variability among different scenarios represented by the same core damage sequence may, however, be thought to be small relative to our understanding of the phenomenology related to steam explosions. In this case an analyst may choose to represent his understanding of the phenomenon in a deterministic way. In other words, as an example, he may choose to express his likelihood based on the assumption that, if a core damage sequence of a certain class were to occur, it would always lead to one, but only one, outcome. However, he may identify two (or more) possible outcomes. In this case, the model of the impact of the phenomenon is constructed with two (or more) possibilities. The expert expresses his degree of confidence about which possibility is correct through the probability, or degree of credence, he gives to the alternatives. In the case of two possibilities, he would construct a bi-modal probability distribution, with the "likelihood" being either 0 or 1 corresponding to the two possibilities. It should be recognized that this is simply an approximation to the true situation, since it ignores the variability inherent in the scenario definitions. However, if that variability is small relative to the state-of-knowledge uncertainty, then the degree of approximation may be acceptable; this approximation is also 'conservative' in the sense that it will tend to overstate rather

than understate the total uncertainty.

We can conclude that the two types of uncertainty can in principle be distinguished. That which analysts can representative of random behavior or variability is inherent in the way they set up the model of the world, and it is appropriate when they ask questions about (hypothetical) repetitions of events rather than a single event. That which analysts call uncertainty describe their knowledge about the parameters or their beliefs in which are the appropriate models to use to characterize the world.

However, while it is relatively straightforward to discuss the philosophical differences between variability and state-of-knowledge uncertainty, it may be difficult in practice for the subject matter expert to determine which approximation he should be using, particularly if he is not very conversant with these subtleties.

The interpretation of the "likelihood" associated with the branches representing the choice between two different model options is much clearer; it can only be a reflection of the degree of belief of the analyst. Again, the true model of the world is one and only one of two possibilities.

DO WE NEED TO TREAT RANDOMNESS AND UNCERTAINTY DIFFERENTLY?

As discussed in the previous section, because the parameters that are constructed to address randomness or variability, and the probabilities that are used to express uncertainty are dealing with different issues, the first being our model of the world, and the second our uncertainty about what that model should look like, there is no option but to treat them separately. The fact that both these sets of parameters may obey the mathematical laws of probability, and thus can be called probabilities in the mathematical sense, does not alter that. That this is so should be clear from considering how new evidence can alter the PSA. By collecting more information we can indeed decrease our uncertainty with respect to parameter values and modeling issues, within the context of the structure of the model, using Bayes theorem as a basis. However, to decrease the "random" content requires restructuring the model itself. This particular issue is not addressed in detail here. Instead, we illustrate the impact the distinction between the two types of characterizations of branch points has on the presentation of the results of a PSA, using a simple example.

Consider a branch point that is representing the choice between two hypotheses. Reference 2 illustrated this with a calculation related to the impact of steam explosion occurrence on consequence evaluation. As another example, suppose a core damage event tree has only two branches. The first (event A) represents the question of whether the required system is available or not. The probability associated with this branch is clearly a relative frequency, and let us assume it has a value of 10^{-3}. The second event (event B) asks whether the system A will be effective in preventing core damage. Suppose we have two experts. The first expert recognizes that there are many unspecified variables that characterize the plant status and has determined that, for a certain subset of the range of those variables, system A will be effective, and for the complementary subset, it will not. Furthermore, he attempts to estimate the fraction of time that the plant is in the subset in which A will indeed be effective and comes up with a mean value of .998. In this case he will determine that the frequency of core damage is given by f. ($A + \bar{A}.B$) where f is the frequency of the initiating event, assumed to be 10^{-1}/year. In this case, expert 1's assessment of core damage frequency is approximately 10^{-1}. (10^{-3} .9989 + 1.10^{-2}) or approximately $1.1x10^{-3}$.

The second expert believe that the best way for him to describe his knowledge about event B is that he feels very strongly that system A is effective, but wants to express the fact that he is not absolutely certain. In this case, he states that the

fraction of times either branch would be followed is 0 or 1. That this is so, is precisely because the PSA community has chosen to set up the PSA model of the world as one in which we will allow the experiment to be repeated (hypothetically) many times, and allow the underlying boundary conditions to vary. In this way the PSA thought experiment is a sampling from the world of possible boundary conditions for the plant in question. When an analyst does not carry the variability in boundary conditions through, and characterizes certain branches as essentially being deterministic, they have to be understood as implying that all sequences go one way or the other, rather than some go one way, and some go another. Suppose that the second expert makes a statement that he is 99% confident that A will work. His assessment of core damage frequency therefore is (f.A + f.$\bar{A}$.0), or 10^{-4} with 99% probability, or is (f.A + f.$\bar{A}$.1), with 1% probability. It should be noted that, for expert 2's assessment to be meaningfully different from expert 1's assessment, he must use a different approach to assessing his probability i.e., his must not be based on constructing a model of the underlying variability.

It is understood here that expert 1 will also provide a statement of his uncertainty on his estimate for the likelihood of event B by constructing a probability distribution on the value of the likelihood. It is usually claimed that expert 2 should not provide an uncertainty about his probability. However, Mosleh and Bier[6] have pointed out that there are conditions under which it makes sense to do so, that and their discussion parallels some of the thoughts in this paper.

CONCLUSIONS

When we are trying to represent the impact of a variability in initial or boundary conditions, that we cannot capture because of modeling or resource constraints, it has been customary to talk about our model of the world as being based on random processes, and the model will have parameters that characterize the system, or more accurately, the ensemble of "identical" systems. Even if there values are assessed subjectively they are parameters of the model of the world.

It is logical to represent uncertainty about the values of the parameters that characterize randomness or variability, using subjective probabilities.

There is a significant difference in impact on the results of the model between saying that a branch point represents uncertainty as to which of two possibilities is the correct (and only) one, and saying that both are possible because of underlying variability in the boundary conditions. It is up to the subject matter expert for the particular modeling issue to determine the most appropriate way for him to characterize the issue. It is, however, the responsibility of the person in charge of integrating the results of the PSA to make sure that the subject matter expert understands what question is being asked and what assumptions he will be making by choosing to characterize his branch points in different ways. It is clear that we must be very careful in formulating the problem, and defining the event(s) of interest so that we know how we are going to give our answer, and how we are going to characterize our uncertainty, if our assessment is to be meaningful.

If a PSA contains both types of approximations to the characterization of branch points, the distinction between the two cases is important and should be maintained because it is crucial for updating as briefly alluded to in this paper, but also, given the experts' assessments, it is the correct representation of their uncertainty

Acknowledgements

The author has benefitted from discussions with all participants at an NRC sponsored meeting held at NUS' Gaithersburg office, June 30th - July 1st, 1992, to discuss issues related to passive reactor system PRAs.

REFERENCES

1. S.Kaplan and B.J. Garrick, On the quantitative definition of risk, *Risk Analysis*, 1, 11-27,(1981).
2. G.W. Parry, On one type of modeling uncertainty in probabilistic risk assessment, *Nuclear Safety*, 24, 634-7, (1982).
3. H.F. Martz, L.L. Hamm, W.H. Reed, and P.Y, Pan, Combining mechanistic best-estimate analysis and level 1 probabilistic risk assessment, *Reliability Engineering and System Safety*, (to be published).
4. S. Kaplan, On the use of data and judgement in probabilistic risk and safety analysis, *Nuclear Engineering and Design*, 93, 123-34,(1986).
5. G. Apostolakis, Probability and risk assessment, the subjectivist viewpoint and some suggestions, *Nuclear Safety*, 19, 305-15, (1978).
6. A. Mosleh and V. M. Bier, Uncertainty about probability: a reconciliation with the subjectivist viewpoint, 1992 preprint.

IS PROBABILITY OF FREQUENCY TOO NARROW?

Harry F. Martz

Group A-1, MS F600
Los Alamos National Laboratory
Los Alamos, NM 87545

Modern methods of statistical data analysis, such as empirical and hierarchical Bayesian methods, should find increasing use in future Probabilistic Risk Assessment (PRA) applications. In addition, there will be a more formalized use of expert judgment in future PRAs. These methods require an extension of the probabilistic framework of PRA, in particular, the popular notion of probability of frequency, to consideration of frequency of frequency, frequency of probability, and probability of probability. The genesis, interpretation, and examples of these three extended notions are discussed.

FREQUENTIST AND SUBJECTIVIST PROBABILITY

In PRA theory and practice, there are two basic interpretations of probability: *frequentist* (or objectivist) and *subjectivist* (or Bayesian). The distinguishing notion of the frequentist interpretation is that probability is defined as the *limiting relative frequency* of occurrence of a random event in an infinite sequence of identical trials. On the other hand, the subjectivist view considers probability to be a subjective statement of an analyst's personal *state-of-knowledge or degree-of-belief* regarding the occurrence of an event based on the available evidence. It is generally believed that both interpretations are appropriate for use in PRA. Parry (1986) and others argue that it is important to separate the effects of both classes of uncertainty on results typically of interest in PRA.

Apostolakis (1990) refers to the model under consideration as the *model of the world*. For example, in Level 1 PRA, the model of the world consists of the event trees, fault trees, etc. used to represent the response of a facility to a particular accident.

For purposes here, consider two basic classes of uncertainties in the model of the world. The first class are those uncertainties which are virtually impossible (either technically or economically or both) to either reduce or remove. Collectively denote the set of such irreducible uncertainties as *stochastic or random uncertainties*. For example, variability in the quality of manufactured materials, uncertainties in equipment performance, and operator errors are typical examples of random uncertainties. Various well-known probability distributions (the formal name for the simplest type of probability models of the world), such as the normal, binomial, or Poisson distributions, are commonly used to model basic random uncertainties. These probability distributions involve various parameters (such as the failure rate or intensity in the Poisson distribution) which are usually unknown and must be estimated. The frequentist interpretation of probability is generally used to express random uncertainties.

The second basic class of uncertainties are *state-of-knowledge* (also known as *systematic*) *uncertainties*. These uncertainties represent our lack of precise knowledge about uncertain parameters in our model of the world (or even uncertainty regarding our overall choice of model of the world). For example, we may not know the precise value of the failure rate parameter in an exponential distribution used to model the time-to-failure of

an operating pump. Based on the available data and experience with such pumps, we represent our uncertainty about this failure rate in the form of form of a probability distribution. Because this distribution represents our current state-of-knowledge about the unknown failure rate parameter (without any regard for consideration of repeated trials), it embodies a subjectivist interpretation of probability.

In contrast to irreducible random uncertainties, it is possible to reduce state-of-knowledge uncertainties by improving or enhancing our state-of-knowledge. In the limiting case, we have perfect state-of-knowledge and the corresponding probability distribution would be a delta function. We can improve (update) our state-of-knowledge by obtaining additional pertinent data or information. Bayes' theorem is appropriate for updating our state-of-knowledge distribution given additional data or information, a process known as *Bayesian updating*. Because there are often alternative models for certain phenomena of interest in PRA, the uncertainty inherent in subjectively choosing a model (so-called *modeling uncertainty*) is a state-of-knowledge uncertainty.

Subjective probability is also an appropriate way to express uncertainty about events, the outcome of which cannot easily be interpreted as an imagined series of repetitions. Such events are often referred to as *propositions*. For example, we may be interested in the event "a manned mission to Mars occurs prior to the year 2050." This is certainly an uncertain event (as of 1993), and a repetitive series of trials of this event is difficult to conceive because it involves a specific (unique) period of time which will occur only once. Thus, it is appropriate to interpret our uncertainty about such events using subjective probability. Notice, however, that determining whether or not a particular event can be imagined as a repetitive series of trials is also often quite subjective, especially if we permit the use of "thought" experiments. While one person may be able (and willing) to entertain such an imagined series of trials, another person may find an imagined series of identical repetitions to be totally inconceivable. Thus, the most appropriate notion of probability to use to describe an uncertain event is also sometimes a subjective decision.

The bottom-line quantities of interest in a PRA are usually either uncertain exceedence probabilities or frequencies of occurrence per year (such as the probability with which damage level $\dot{x}$ or greater occurs per year or the frequency of occurrence of severe core damage per year). Clarotti (1988) and Howard (1988) argue that the frequency of occurrence of rare events considered in PRA are fictional parameters which should not be interpreted within the context of a repetitive series of trials. Thus, there is some disagreement regarding whether the frequency of occurrence of rare events per unit time should be interpreted using relative frequency or degree-of-belief.

To distinguish between the two notions of probability, I shall denote frequentist (relative frequency-based) probability as probability$_f$ (where the subscript "f" stands for "frequentist") and subjectivist (degree-of-belief-based) probability as probability$_s$ (where the subscript "s" denotes "subjectivist").

PROBABILITY OF FREQUENCY

Kaplan and Garrick (1981) introduced the notion and use of *probability of frequency* in PRA. This approach is further described in Kaplan et al (1981). They consider the relative frequency of occurrence of an event of interest (that is, the fraction of times that it occurs) in a series of identical repeated trials in which the event has a constant probability$_f$ of occurrence p on each trial. They denote this relative frequency as ϕ. Of course, by definition, ϕ approaches p in the limit as the number of repeated future trials approaches infinity. Thus, interest in ϕ (over an infinite series of trials) is equivalent to interest in p.

Because of inherent parameter and modeling uncertainties in our model of the world, we are usually uncertain about the precise value of p (especially so in the case of rare events). By performing an *uncertainty analysis*, we can express our state-of-knowledge about the unknown value of p in the form of a probability$_s$ distribution (or curve) over p$_f$. Kaplan and Garrick refer to these as *probability$_s$ of probability$_f$* (or *probability$_s$ of frequency*) curves. The outer probability$_s$ is interpreted as *Bayesian confidence* (*confidence$_s$* or *credibility*). For example, using the probability$_s$ of frequency approach, we can make statements such as "we are 95% confident$_s$ that the annual probability$_f$ of 10 or

more early fatalities resulting from internally initiated accidents at Plant X is not larger than 10^{-7}".

FREQUENCY OF FREQUENCY

Empirical Bayes (EB) methods are a class of modern statistical inference techniques which are potentially useful in PRA applications for estimating unknown parameters in the model of the world. To illustrate the use of EB in PRA, Martz and Abramson (1994) use EB methods to estimate the average (over time-and-station) probability$_F$ that emergency diesel generators (EDGs) at US commercial nuclear power plants successfully load-run on a given demand (the so-called *load-run demand reliability*). For simplicity, denote the time-and-station average of this load-run demand reliability as p. Martz and Abramson are interested in estimating p using data from a large number of stations. EB methods are appropriate for use in those situations where a parameter, such as p, is not constant over a corresponding population of interest (in this case, the population of stations). EB methods assume that p varies randomly over the population according to some unknown *EB prior distribution g(p)* [Maritz and Lwin (1989); Martz and Waller (1991)]. The EB prior g(p) expresses the probability$_F$ distribution of p over the EDG population and, for this reason, is sometimes referred to as a *population variability distribution* (or curve). In our example, g(p) expresses the fraction (or relative frequency) of stations whose average EDG load-run demand reliability falls in each reliability increment.

It is important to note that, in spite of the reference to Bayes in the name "EB", EB methods make no use of degree-of-belief (that is, probability$_S$). The reference to Bayes is due to two aspects of EB: (1) the notion and use of a "prior distribution" for p, and (2) the use of Bayes' theorem as an integral part of the EB procedure. The advantage of EB over Bayesian methods (when both can be used) is that EB methods involve less subjectivity, and thus are more "data driven", than pure Bayesian methods.

In a plant-specific PRA, we are not primarily interested in the overall population of stations; rather, we are interested in the EDG reliability of a particular station. Bayes' theorem is the tool for "specializing" (by updating) the EB prior on p to the particular value of p, say p_i, for a particular station (the i^{th} here). As in Bayesian estimation, we determine an EB posterior distribution for p_i. Unlike a Bayesian posterior distribution, this station-specific EB posterior distribution embodies probability$_F$ (as does the EB prior), and is a *frequency of frequency* distribution.

What interpretation do we give to this station-specific EB posterior distribution? Suppose that, for our particular station, we have observed $x_i = 2$ load-run failures in $n_i = 117$ total station EDG demands. The EB posterior distribution of p given $x_i = 2$ and $n_i = 149$ [that is, $g(p|x_i=2; n_i=117)$] expresses the fraction (or relative frequency) of stations in the population for which 2 failures are observed in 117 demands whose load-run demand reliability falls in each reliability increment. Because we have actually observed 2 failures in 117 demands for our particular station, this is the appropriate station-specific probability$_F$ distribution that we use to make inferences (such as point and interval estimates) about p_i.

FREQUENCY OF PROBABILITY

Consider a one-time event which either will or will not occur—except that we don't know which—and for which a repetitive series of trials is inappropriate or impossible to consider. For example, consider again the event that "a manned mission to Mars occurs prior to the year 2050." Now, this future event will certainly either occur or it won't occur—we just don't know which. In PRA parlance, this is known as a *zero-one* event.

Kaplan et al (1981) liken this situation to that in which either a two-headed or a two-tailed coin is to be flipped, and we just don't know which coin it is. Anyone can then express his or her state-of-knowledge (degree-of-belief) regarding the occurrence of the Mars event as the probability$_S$ p that the two-faced coin is two-headed. Further, suppose that we are interested in assessing the variability of p_i over some population of interest; for example, the population of Los Alamos National Laboratory employees. If we now ask each Los Alamos employee for his or her value of p (and everyone complies), then we can

construct a population variability curve that expresses the fraction (or relative frequency) of employees at Los Alamos whose value of p falls in each increment. Again, using the Kaplan and Garrick nomenclature, this population variability curve on p is a *frequency of probability$_S$* distribution.

Figure 1 shows the frequency distribution for the probability$_S$ p of the Mars event (in the form of a relative frequency histogram) based on a sample of size n = 56 Los Alamos employees. It is interesting that this distribution tends to be trimodal, having a delta function at 0.5 (which appears in Figure 1 as a rectangle) and a "bump" on either side of (but distinctly away from) 0.5. The two bumps correspond to strongly optimistic or pessimistic employees, while those who responded with p = 0.5 are indicating that they no nothing (and are thus "guessing"), regarding the occurrence of the Mars event. For these data, we find that the mean is 0.49 and that Probability$_F$ (0.01 < p < 0.92) = 0.80.

The formal use of expert opinion is an important part of modern PRA theory and application [Wheeler et al (1989)]. Suppose that we expand our PRA model of the world to include the representation of the variation in opinions of different experts. Within this expanded model, the notion of frequency of probability$_S$ is appropriate. Let me explain. In PRA analysis, many issues can often best be interpreted as zero-one events in which an expert is asked to quantify his or her degree-of-belief regarding the occurrence of the event. The uncertainty is usually elicited in the form of a probability$_S$ p. For example, in the reactor context, consider the event "in a proposed advanced reactor design, natural processes (such as passive heat removal) fail to successfully cool the reactor core." Suppose that we consider a panel of experts for this event to be a sample from a larger population of experts. In this case, frequency of probability$_S$ is an appropriate way to express the variation in degree-of-belief within the population of experts.

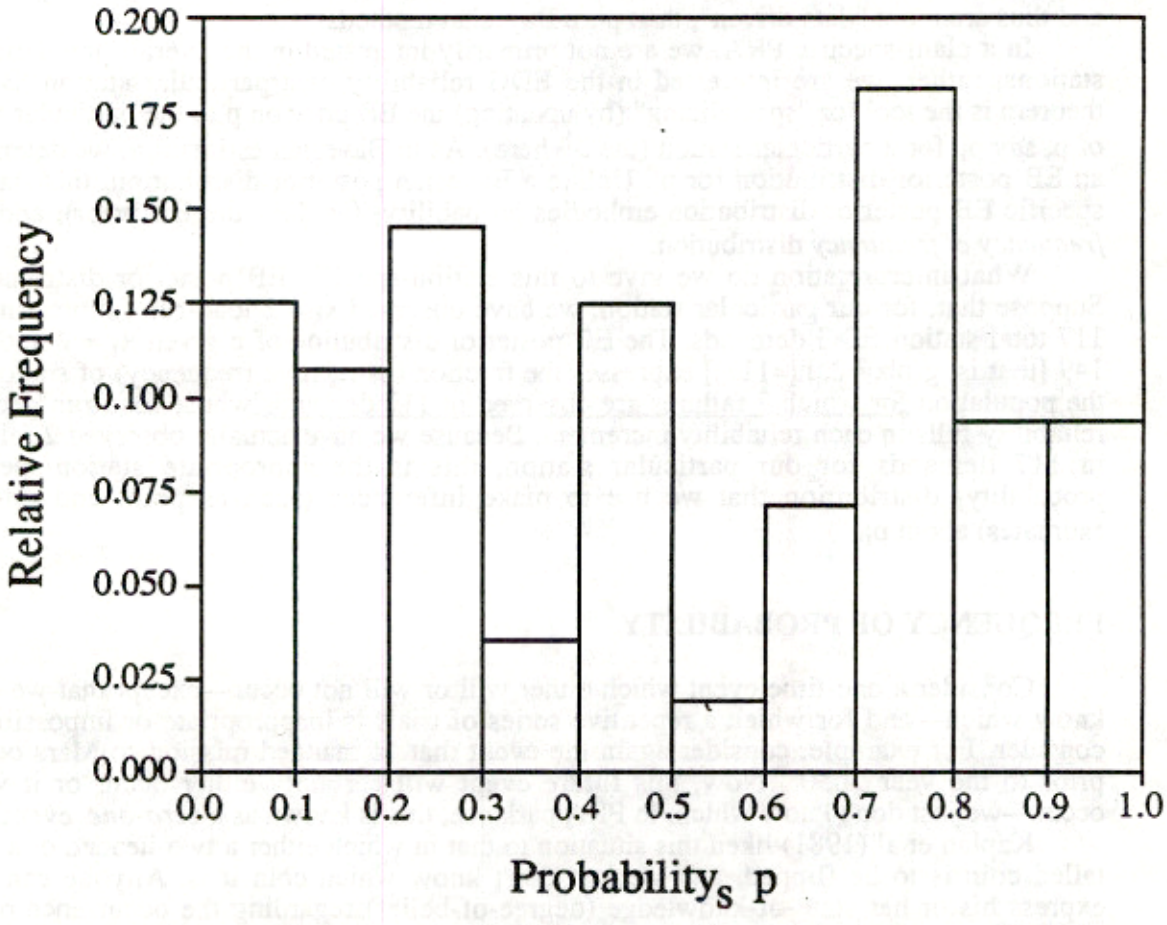


Figure 1. Frequency of probability distribution for the probability$_S$ p of a manned mission to Mars prior to the year 2050.

How would we update a frequency of probability$_S$ distribution as new information is obtained? We must be careful here to define what we mean by "new" information. Because the event of interest is zero-one, a single trial (if such would even be possible) is sufficient to completely remove all uncertainty about the event. This is not what I mean by new information. Rather, by "new information" I mean any change that an expert makes in his or her own subjective probability$_S$ p as a result of additional information or new knowledge. For example, in the Mars event, new knowledge might consist of an emerging global political and economic climate advocating that such a mission be undertaken. Each expert would then be asked to update his or her own value of p in light of this new information. Given the new value(s) of p, these estimates would similarly be used to determine a new (updated) frequency of probability$_S$ distribution in light of the new information. Note, however, that such updating is not performed using Bayes' theorem.

PROBABILITY OF PROBABILITY

Here we are basically concerned about what is believed about what is believed, and, at first glance, such a statement appears to be nonsense. If, however, I am basically concerned about expressing what an analyst believes about what an expert (or experts) believes (regarding the zero-one event), then this statement makes perfectly good sense. Further discussion regarding this statement is given by Pearl (1988), pp. 358-359.

As in the previous case, we now expand our PRA model of the world to include an expression of the variety and quality in the opinions of different experts as viewed by an analyst acting as a "super expert". Thus, within this expanded model, the notion of *probability$_S$ of probability$_S$* is defined. There are at least two possible interpretations of probability$_S$ of probability$_S$.

First, the analyst can represent his/her belief regarding the expert opinions in the form of a probability$_S$ distribution over the elicited probabilities. This situation is similar to the case of frequency of probability$_S$, except that now the outer distribution is a probability$_S$ distribution that must be subjectively assessed by the analyst. For example, the commonly used decision to equally weight each expert judgment is one such (outer) probability$_S$ distribution.

There is another interpretation which is also appropriate. Suppose that the analyst is interested in the elicited value from a particular expert, say Jane. Prior to eliciting Jane's value of p, p_{Jane}, the analyst may wish to express his/her personal belief about p_{Jane} in the form of a "prior" probability$_S$ distribution on p_{Jane}. This distribution expresses any substantive information that the analyst may have regarding Jane's ability in such tasks (including any known or perceived biases).

How do we update the analyst' prior once Jane has given her value? One possibility is to consider an appropriate subjective likelihood for Jane. This likelihood is a statistical model that relates the likelihood that Jane will report a value p_{Jane} relative to a true value of p. What do I mean by a "true value of p"? Subjective opinions being what they are, there clearly is no true value of p. However, one can define the true value of p to be the average value of p that would be elicited from a large population of experts. Examples of such conceptual likelihoods involving expert opinion are considered by Mosleh and Apostolakis (1984) and Martz and Bryson (1984). An important feature that is captured in such models is the consideration and inclusion of possible "biases" that are known to accompany expert opinions. Given the analyst' prior distribution and the assumed conceptual subjective likelihood model for Jane, Bayes' theorem is then used to obtain the desired updated posterior distribution for p_{Jane} given Jane's estimate.

CONCLUSIONS

Future PRA applications will likely use modern statistical methods, such as empirical and hierarchical Bayes methods, for estimating uncertain parameters in the PRA model of the world. In addition, the model of the world will be enlarged to include additional concerns involving the elicitation and combination of expert opinions. Because neither of these extensions are adequately considered using probability$_S$ of frequency, I believe that

this notion must be expanded to permit consideration of frequency of frequency, probability$_S$ of frequency, and probability$_S$ of probability$_S$.

Many authors have pointed out the importance of carefully distinguishing between random and state-of-knowledge uncertainties in PRA [Cornell (1984), Parry (1986), Apostolakis (1990)]. The primary concern in using these extended notions is how to separate the contributions from random and state-of-knowledge uncertainties in the results. It is not altogether clear that a clean separation is possible. If this task proves impossible (and well it might), perhaps the only reasonable alternative is to argue from a pure Bayesian perspective that all uncertainty, regardless of its origin, is really of the degree-of-belief type [Lindley (1985)].

Clearly there are other issues involved in using these extended notions which have not been considered here. For example, issues involving Monte Carlo simulation using these new notions are not discussed, and Bayesian updating has not been adequately addressed. However, I hope that this paper will stimulate additional interest in these and related issues.

ACKNOWLEDGMENT

This work was supported by the US Nuclear Regulatory Commission, Office of Nuclear Regulatory Research, under contract FIN L1836-3 at Los Alamos National Laboratory. I gratefully acknowledge this support.

REFERENCES

Apostolakis, G., 1990, The concept of probability in safety assessments of technological systems, *Science* 250:1359.

Clarotti, C. A., 1988, *in* : "Reliability Engineering," A. Amendola and A. Saiz de Bustaman, eds., Kluwer Academic, Dordrecht, Netherlands, pp. 49-66.

Cornell, C. A. ,1984, Seismic hazard analysis: some current directions ans some comments on its role in seismic PRA, presented at the *1984 Pressure Vessels and Piping Conference*, San Antonio, Texas, PVP-Vol. 79, p. 73, ASME.

Howard, R. A., 1988, Uncertainty about probability: a decision analysis perspective, *Risk Analysis* 8:91.

Kaplan, S. and Garrick, B. J., 1981, On the quantitative definition of risk, *Risk Analysis* 1:11.

Kaplan, S. et al, 1981, Methodology for probabilistic risk assessment of nuclear power plants, *Pickard, Lowe and Garrick, Inc. Report* PLG-0209, Irvine, California.

Lindley, D. V., 1985, "Making Decisions (*2nd* ed.)," John Wiley, New York.

Maritz, J. and Lwin, T., 1989, "Empirical Bayes Methods (*2nd* ed.)," Chapman and Hall, London.

Martz, H. F. and Abramson, L. R., 1994, Empirical Bayes estimation of the average reliability of nuclear power plant emergency diesel generators to load-run upon demand, in: "Proceedings of the PSAM-II Conference," G. Apostolakis, ed., Plenum, New York.

Martz, H. F. and Bryson, M. C., 1984, A statistical model for combining expert opinions, *IEEE Trans. on Rel.* R-33:227.

Martz, H. F. and Waller, R. A., 1991, "Bayesian Reliability Analysis," Krieger, Malabar, Florida.

Mosleh, A. and Apostolakis, G., 1984, Models for the use of expert opinions, in: "Low-Probability/High-Consequence Risk Analysis," R. A. Waller and V. T. Covello, eds., Plenum, New York, pp. 107-124.

Parry, G. W., 1986, A discussion on the use of judgment in representing uncertainty in PRAs, *Nucl. Eng. Des.* 93:135.

Pearl, J., 1988, "Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference," Morgan Kaufmann, San Mateo, California.

ON BAYESIAN APPROACHES OF PARAMETRIC AND MODELING UNCERTAINTY

Jya-Syin Wu

ASCA, Inc.
2250 East Imperial Highway, Suite 200
El Segundo, CA 90245-3547
Tel: (310) 648-6658
Fax: (310) 648-6659

INTRODUCTION

The growing demand for risk-related decision making from both public and private sectors has made the understanding of risk and the presentation of risk-related information fundamental to the successful management of risk. The increasing complexity of modern engineering systems, however, has complicated the understanding process and added to the difficulty of the presentation. In both understanding and presenting risk, the uncertainty is an important but intricate subject that often puzzles the risk analysts.

There are open issues related to the modeling of uncertainties in the probabilistic risk assessment (PRA). First, the terminology currently used is confusing. For example, in one case [1] the uncertainties are grouped into "stochastic uncertainty" and "state-of-knowledge uncertainty." In this setting, if the uncertainty is originated from an unknown outcome of a chance (random) event, it is termed as "stochastic uncertainty;" on the other hand, if it is a result of the lack of knowledge of the analyst's, the term "state-of-knowledge uncertainty" is used. In other occasions [2,3], uncertainties have been grouped into "parametric uncertainty" and "modeling uncertainty." For example, in reference [2], the term "parametric uncertainty" is used for the kind of uncertainties due to the lack of knowledge in the numerical estimates of parameters in a particular model, and "modeling uncertainty" for uncertainties due to the lack of knowledge in the model itself. In another attempt [3], the "parameter uncertainty" is also used for uncertainties due to the lack of knowledge in the numerical estimates, but limited to probabilistic parameters; similarly, "modeling uncertainty" is limited to uncertainties due to the lack of knowledge in physical phenomena. Still, in another application [4], the terms "data uncertainty," and "modeling uncertainty" are favored.

Then there is the modeling of the processes for assessing uncertainties and, under some circumstances, for reducing the uncertainties. On such issues, Bayesian and non-Bayesian practicians hold distinct difference on philosophical grounding and have developed significantly different approaches for real application [5-7].

This paper presents a Bayesian viewpoint of the modeling of uncertainty in PRA. A brief discussion on the non-Bayesian approaches will be presented at the end of the paper.

THE BAYESIAN APPROACH

The purpose of a PRA is to draw reasonable inferences on the risk associated with a complex system. The sources of information for a PRA include the knowledge of the analysts, the collective operating experience of the industry, and the judgment of the field experts. The final product of a PRA would be a set of risk curves that provides not only the range of risk, but also the analyst's degrees of belief in these curves.

The two types of information provided by a PRA are corresponding to the two types of uncertainty in the scientific reasoning process: stochastic uncertainty and knowledge uncertainty. Each point on a given risk curve represents the likelihood of the occurrence of some undesirable events of a given consequence and is referred to as the "stochastic uncertainty." The set of risk curves reflects the uncertainty of the analysts' in his assessments, and it is referred to as the "knowledge uncertainty." Both types of uncertainty are the inevitable byproducts of the processes of scientific inference, and should be discussed in the context of inferential processes of the analysts.

Two types of events are of major concern in the inferential processes: deterministic events and probabilistic events (hereto referred as chance events to avoid confusion). Deterministic events are logical consequences of a given set of verifiable characters, and therefore, there should be no "stochastic uncertainty," which concerns only with the likelihood of the occurrence of an event, associated with them. Since "knowledge uncertainty" concerns the lack of knowledge of the analysts', it could remain high for the case of deterministic events.

Many events in real world are, however, probabilistic in nature. Theories describing chance events do not follow logical consequences of verifiable characters. For chance events, both stochastic uncertainty and knowledge uncertainty are likely to exist. As pointed out in [8], while it is possible to reduce knowledge uncertainty by improving the understanding of the physical characters of the event, the stochastic uncertainty persists as part of the nature of the event.

Measures of both stochastic and knowledge uncertainties satisfy the axioms of calculus of probability. The terms "probability" and "frequency" are often used for the measure for the stochastic uncertainty, and "degree-of-belief" for the measure of knowledge uncertainty. In [7], the probability of a chance event is considered as one of the physical characters of the event. Thus, the statement "the probability of an event A is p" is taken semantically equal to the statement "the mass of A is x," rendering the modeling of knowledge uncertainties for deterministic events identical to that for chance events.

The sources of knowledge uncertainty could be multiple. First, for a given problem, there may be more than one model been proposed, each of which leads to a different prediction. In PRAs, the modeling of events after core damage and the modeling of external events are areas where more than one deterministic models would have been proposed. These are events occurred under operational conditions in which most of the physical phenomena were poorly understood and little amount of data has been collected. The knowledge uncertainty expressed in these events would be the degrees of belief of the analyst's on the physical models. The kind of knowledge uncertainty discussed here has will be referred to as the "modeling uncertainty" in this paper.

Modeling uncertainty does not need to be limited to physical phenomena or deterministic events. Depending on the event type, the models proposed in the analysis could be either deterministic or probabilistic in nature. In the case of chance events, if an

analyst believes that more than one statistical models--for example, two families of probability distributions--are "seriously possible" for the situation one models, the uncertainty against these statistical models can be expressed by the degrees of belief on these statistical models.

In real practice, the deterministic or statistical models are often presented in terms of parameters. The uncertainty associated with these parameters will be referred to as "parametric uncertainty" hereafter. Strictly speaking, parametric uncertainty is a subset of modeling uncertainty where the focus is on specific parameters. Mathematically, since parameters often take values over continuous intervals, operations on continuous functions are used more often in the parametric uncertainty than in the "modeling uncertainty."

THE BAYESIAN FRAMEWORK OF UNCERTAINTY

Figure 1 presents the generalized Bayesian framework of the modeling of uncertainty in scientific reasoning. Starting from the very left, the analyst first needs to identify all seriously possible hypotheses, h_i, of the event under consideration. The knowledge uncertainty of the analyst's on these hypotheses is expressed in terms of his degree of belief, $p(h_i)$, on each hypothesis h_i. Note that a tree structure can be constructed for presenting both modeling and parametric uncertainties. The modeling uncertainties, $p(m_i)$, and parametric uncertainty, $p(\mu_i|m)$, can be, then, represented by branches at different levels of the knowledge uncertainty tree, as shown in Fig. 1.

The next level of the Bayesian framework is to model the stochastic uncertainty (or likelihood) associated with the outcomes of the event. The outcomes and the conditional probabilities of these outcomes are determined by the model the parameters chosen in the setting. Given a model h_i and parameter μ_j, the stochastic uncertainty is measured, then, by the conditional probability, $p(e_k|m_i,\mu_j)$, for a specific outcome e_k.

For deterministic events, one and only one $p(e_k|m_i,\mu_j)$ will receive a probability 1, while all other branches receive 0's. This is because the predicted outcome is a logical consequence of the given hypotheses and would be a sure outcome. Events considered in physical models and events associated with physical parameters are often deterministic events. Knowledge uncertainties of these events are expressed on both physical models on physical parameters.

For a chance event, the stochastic uncertainty of the event is the likelihood function, $p(e_k|m_i,\mu_j)$ of an event e_k, conditioned on a given set of hypotheses. Since the events thus described are stochastic in nature, the models would be statistical. Similar to the case of deterministic events, knowledge uncertainty are be expressed on both statistical models and statistical parameters.

PRA APPLICATIONS

Discussions on modeling and parametric uncertainties in the PRA literature can be viewed as special cases of the generalized Bayesian framework shown in Fig. 1. The following paragraphs focus on how some of these discussions are consistent with the generalized framework of uncertainty. Several examples of PRA applications are also discussed within the proposed framework.

The concepts of "state-of-knowledge uncertainty" and "stochastic uncertainty" discussed in reference [1] matches the high level concepts of the Bayesian framework. In this setting, the state-of-knowledge uncertainty reflects the degrees of belief of the analysts; the stochastic uncertainty is associated with the random (chance) nature of the physical world.

In reference [2], both the "parametric uncertainty" and the "modeling uncertainty" are associated with the knowledge uncertainty of the analyst's, and both are measured by the degrees of belief. Chance (random) events are not of concern in this setting. While the "parametric uncertainty" is understood as due to the lack of knowledge in the numerical estimates of parameters in a particular model, the "modeling uncertainty" gives the measure of the lack of knowledge in the model itself. The models can be either deterministic or statistical, and so does the parameters.

Judging from the writing of reference [3], "parameter uncertainty" referred in [3] is specifically associated with the numerical estimates of statistical parameters in describing chance events, and "modeling uncertainty" is associated with physical phenomena, and, thus, the knowledge of deterministic events. The types of uncertainties discussed in [3] are special cases of those described by the generalized Bayesian framework-- one related to chance events and the other deterministic events.

To apply the generalized Bayesian framework of uncertainty modeling to specific PRA application, the following paragraphs present two extreme cases. As the first case, let us consider the modeling of component failure rates in PRA. For a given failure mode of a component, the failure probability of a component at time t is said to act according to an exponential distribution $F(t|\lambda)$, where λ is the failure rate of the component. The uncertainty over λ is represented by a lognormal distribution $\pi(\lambda)$.

The Bayesian interpretation of such uncertainty assignment is as the following: Under the specific situation of concern, the analyst judged that the component failure to be a chance (random) event. He also judged that one statistical model, the exponential distribution, would be all that is needed to model the random nature of the failure. In so doing, there is no "modeling uncertainty" in this specific assignment. The analyst, however, is uncertain regarding where λ would be exactly located, and he has assigned $\pi(\lambda)$ to express his knowledge uncertainty over parameter λ. $\pi(\lambda)$ is, thus, the "parametric uncertainty" of λ, given that the analyst has already chosen to go with exponential distribution for the component failure. The measure, $\pi(\lambda)$, is the degree of belief of the analyst's. The stochastic uncertainty in this application is presented by the conditional probability of the component failure at time t, $F(t|\lambda)$, given a failure rate λ.

A very different reasoning process is used in the prediction of the occurrence of steam explosion after core damage. As discussed in [3], while some theoretical models lead to the conclusion that steam explosion is a sure consequence of a damaged core, other theories lead to the opposite conclusion. Facing the two conflicting conclusions, the analyst would assign degrees of belief, p_1 and p_2, to these two models. The degrees of belief thus assigned would be the measure of the knowledge uncertainty of the analyst's on the issue. The event considered in this case is deterministic; therefore, there is no stochastic uncertainty considered in this application. The conditional probabilities take value only 1 or 0 in this case.

NON-BAYESIAN APPROACHES

A few words must be said regarding non-Bayesian approaches on the issue of the modeling uncertainty. One of the many shortcomings of non-Bayesian approaches is that there is, unlike the Bayesian framework described in this paper, no approaches in the non-Bayesian world that could model both deterministic events and chance events in a uniform fashion. For deterministic events, philosophers have spent significant amount of efforts studying the notion of confirmation and have tried to develop approaches in which deterministic theories could be confirmed by observations. Popper's attempt to solve the problem of induction [9], for example, is a deterministic theory which, in addition to its other weaknesses, fails to model chance events . For chance events, statisticians, including

Fisher [10], Neyman, and Pearson [11], have developed various forms of statistical inference in the "classical" sense. These models, besides other criticisms, cannot find a uniform approach in their counterpart deterministic theories. A detailed discussion on this issue can be found in [5].

CONCLUSIONS

A generalized Bayesian framework is proposed for the modeling of uncertainty in the PRA. Before assigning numerical values for uncertainties, the analysts should investigate first the nature of the event that one is to assessing and the kind of uncertainties associated with the reasoning process. Issues in the PRA literature--knowledge uncertainty, stochastic uncertainty, modeling uncertainty, and parametric uncertainty--are all important parts of the overall picture and are closely connected to each other. The Bayesian framework provides a uniform approach on these issues and on uncertainties associated with both deterministic and chance events. The author, therefore, willfully assigns a high degree of belief in choosing the Bayesian framework over its rival approaches.

REFERENCES

[1] G. Apostolakis, Uncertainty in probabilistic safety assessment, Nucl. Eng. Des., 115 (1989) 173-179.

[2] N.O. Siu, D. Karydas, and J. Temple, Bayesian assessment of modeling uncertainties: application to fire risk assessment, in: B. M. Ayyub, M. M. Gupta, and L.N. Kanal (Eds.), Analysis and Management of Uncertainty: Theory and Applications, Elsevier, Amsterdam, 1992, pp. 351-361.

[3] G.W. Parry, Technical note: on one type of modeling uncertainty in probabilistic risk assessment, Nucl. Saf., 24 (1983) 624-627.

[4] G.W. Parry and P.W. Winter, Characterization and evaluation of uncertainty in probabilistic risk analysis, Nucl. Saf. 22(1981) 28-42.

[5] C. Howson and P. Urbach, Scientific Reasoning: The Bayesian Approach, Open Court, La Salle, Illinois, 1989.

[6] I. Levy, Gambling with Truth, Knopf, New York, 1967.

[7] I. Levy, The Enterprise of Knowledge, MIT Press, Cambridge, Mass, 1980.

[8] G.W. Parry, Use of judgment in representing uncertainty in PRAs, Nucl. Eng. Des., 93 (1986) 135-144.

[9] K.R. Popper, The Logic of Scientific Discovery, Hutchinson, London, 1959.

[10] R.A. Fisher, Statistical Methods and Statistical Inference, Oliver and Boyd, Edinburgh, 1956.

[11] J. Neyman, Outline of a theory of statistical estimation based on the classical theory of probability, in A Selection of Early Statistical Papers of J. Neyman, Cambridge University Press, 1967, pp.333-380.

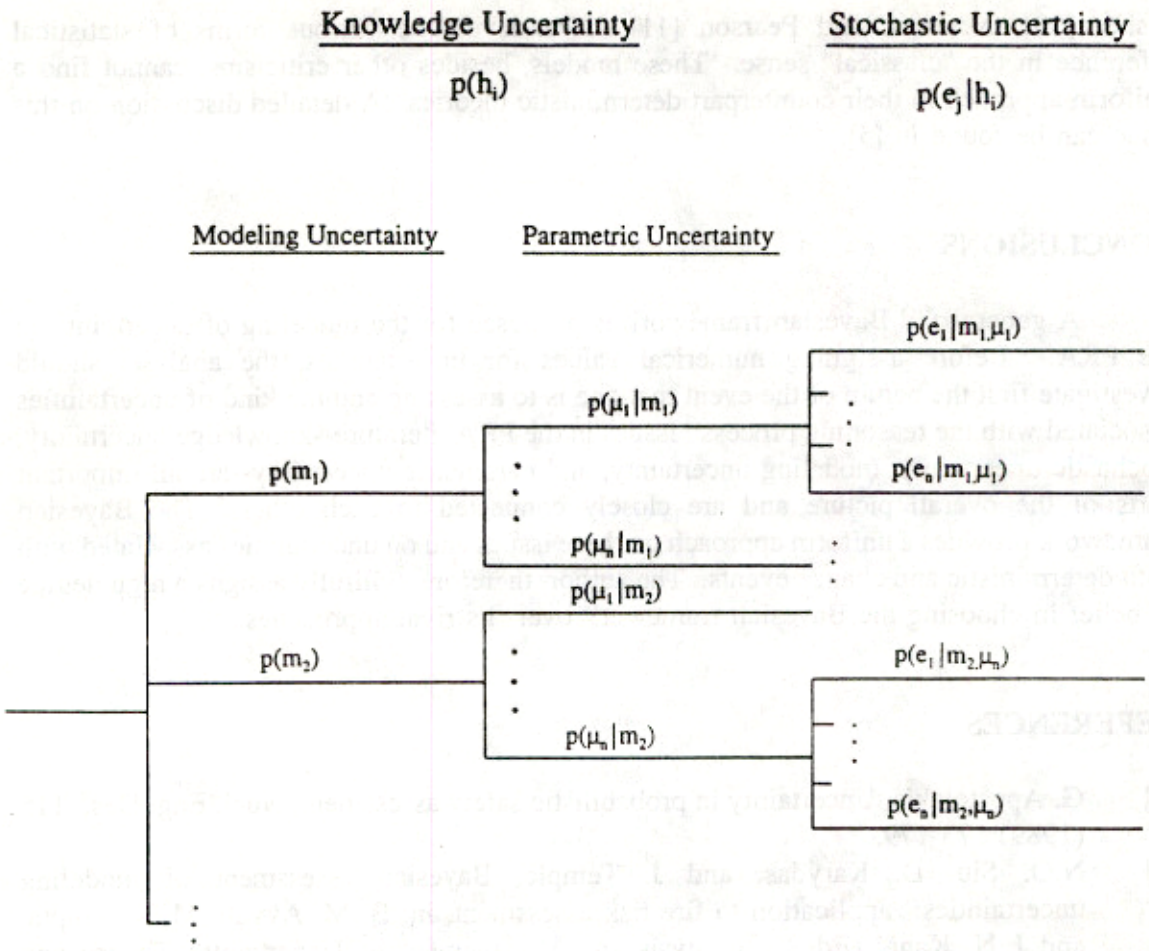


Fig. 1 The generalized Bayesian framework for the modeling of uncertainty

ON THE OBJECTIVITY OF SUBJECTIVE THEORY OF PROBABILITY FOR PSA APPLICATIONS

Ali Mosleh

Center for Reliability Engineering
University of Maryland
College Park, MD 20714-2115, USA

INTRODUCTION

Probabilistic Safety Assessment (PSA) in its relatively short history as a discipline has brought several issues regarding the interpretation and use of probability which have been traditionally only of interest only to probability theorists and philosophers, to a level of practical interest and subject of much debate. Some of the most frequently debated issues are:

1. Is probability an adequate measure for quantifying and communicating uncertainty? In the PSA context the more specific issues include the adequacy of probability as an effective measure for characterizing hardware performance (e.g., component unavailability) and phenomenological uncertainty.

2. What exactly is the probability of an event? Does it represent an inherent characteristic of the event or phenomenon, or is it just a measure of how much we know about it? What is frequency and what is its relation to probability? In the PSA context do we need both of these concepts or is there a unified concept that can be applied in all cases? A related question is concerned with the type of uncertainty, more specifically "stochastic" vs. "state-of-knowledge" uncertainty. Are they different, fundamentally? Practically? Do we need both in the PSA context?

3. How do we estimate or assess a probability? Can we be uncertain about our probability assessment?

These questions have been asked mostly in practical settings and in addressing specific issues in modeling, quantifying risk of technological systems and in making decisions suing risk assessment results. Solutions provided by theorists and practitioners have been debated, but often adopted without a final resolution of the objections and concerns. For example the use of subjective theory and particularly the Bayesian approach to statistical inference is widespread while there is disagreement on

some of the basic issues regarding the meaning of probability, its relation to frequency, and implications of different interpretations.

Similarly, many are still of the belief that both classical and Bayesian interpretations are needed and that in order to model events which are "stochastic in nature", one would have to use the frequentist or classical theory, whereas cases involving lack of knowledge about the event or phenomenon in question should be treated using "sate-of-knowledge" interpretation[1].

In this paper we will argue that the subjective theory of probability is a complete framework for the "P" in PSA. In particular we will explain how the subjective theory can be used in representing the various sources and layers of uncertainty including uncertainty in model parameters . It will be argued that there is no theoretical or practical reason to resort to the use of two different interpretation of probability in PSA applications.

UNCERTAINTY: ITS SOURCE AND REPRESENTATION

In PSA applications we are primarily concerned with physical events such as functional state of a component, occurrence of fire in a specific time interval, form of an operator action , etc. The questions that we have all relate to state of nature at a specified time or time interval. In an accident sequence for example, we ask whether a component will be in a failed or functional state when it is needed. In some cases the possible states of nature are not known due to limitations in the information base of the analyst, including limitation in the state of the art. In all these cases uncertainty arises from lack of knowledge, as opposed to something inherent in nature. This can be said with the highest degree of confidence as long as the objects of interest in risk analyses remain outside the domain of quantum mechanics. (As a side note this author believes that even in the quantum domain the uncertainty is in the observer's lack of knowledge about the sate of nature as opposed to an inherent indeterminism in nature).

As an example of the root of uncertainty in risk and reliability analysis consider the case of a component. Suppose the component can have two possible states of "Failed" and "Not Failed". For a continuously operating component the key question is whether the components will operate for a specified mission time. Ideally we want to know the exact time (T) at which the component fails. In principle given enough science and money we should be able to predict the failure time with enough precision. However , technical and economical constrains among others make such predictions difficult if not impossible. The result is that we are we are uncertain about T. This uncertainty is expressed (in standard reliability theory and application) by a probability distribution, f(t) over the possible values of T.

In this sense f(T) is nothing but a measure of the analyst's uncertainty about a physical phenomenon .i.e, the time at which the component fails. At this point there is no need to speak of uncertainty in nature or stochastic behavior on the part of the component. The probability f(t) could be assessed based on a model of the physical characteristics of the component and its environment ("Physics of Failure" approach) or by using the history of that and similar components as the basis for predicting the future behavior of the component in question (statistical approach). Either way lack of ability to know T is due to lack of knowledge rather than randomness in nature.

Randomness is nothing but a reflection of the level at which we choose to model the phenomenon of interest. It is a characterization we use for system behavior when we are unable or unwilling to model the system or phenomenon in enough detail such that for a given scope and set of boundary conditions the behavior can be predicted (deterministically). The aspects of the phenomenon and boundary conditions which cannot be determined are often lumped and represented by "parameters" in our models. Since these aspects and conditions represent classes of situations (each of which may result in a different outcome or realization of the phenomenon) which cannot be directly

traced back to specific causes, conditions, and characteristics of the physical entity or phenomenon, they are attributed to inherent "randomness" or "stochastic'' behavior.

Once we agree that there is no randomness in nature, all uncertainties in the outcomes of events can be characterized a state-of-knowledge uncertainty. This allows us to use one interpretation of probability and one method of inference consistently and in all cases, as has been argued by others previouly[2]. These are respectively the subjective interpretation of probability, and Bayes theorem.

In recent years, other measures of uncertainty have been proposed for use in PSA. These include fuzzy set theory, confidence factor, and theory of evidence, to name a few. While there seem to be some room for some of these measures in communication of uncertainty and risk results, there is no evidence that probabilistic representation of uncertainty within risk models has encountered any difficulty, conceptually or computationally. As long as no need is identified to resort to a different theory, probability should remain as the measure of uncertainty is PSA studies.

UNCERTAINTY ABOUT PROBABILITY ?

One of the reasons that some of the theorists and practitioners have felt that there is need for both stochastic and state-of-knowledge uncertainty notions in PSA is the controversy about uncertainty about probability. Many theorist such as de Finnetti, have concluded that it is inherently meaningless to be uncertain about a probability. This has been viewed as being too restrictive in practice since practitioners feel that many of the probability estimates cannot be made with the precision and unambiguity being demanded by the theoreticians. Alternatives (practical and theoretical) have been offered[3]. In the following lines it is argued that the theory of subjective probability is sufficient for handling practical and theoretical issues related to uncertain probabilities.

Sources of Uncertainty about Probability

The various theorists have generally failed to distinguish between different sources of uncertainty about probabilities. In particular, they have failed to distinguish between two very different reasons that judgments of subjective probability may be uncertain: uncertainty about the underlying events, propositions, or hypotheses on which those probabilities are conditioned; and imprecision resulting from inherent cognitive limitations. Distinguishing between these two sources of uncertainty is an important step in resolving the persistent theoretical controversy over uncertainty about probability.

As argued by Mosleh and Bier[4], uncertainty about probabilities frequently arises from uncertainty about the underlying events, propositions, or hypotheses on which those probabilities are conditioned. In other words, if P(E | X) is the probability of an event E given a set of conditions X, then the quantity P(E | X) can be viewed as a function of X, which may itself be uncertain (and would therefore be a random variable). Since every function of a random variable is itself a random variable, it follows that P(E | X) is a random variable and is therefore described by a probability distribution. In particular, one can construct the distribution for the quantity P(E | X) from the following pieces of information: (1) the form of the function P(E | X); and (2) the underlying probability distribution over X.

This formulation, agrees with de Finetti that it is inconsistent with the subjective theory of probability to be uncertain about the value of P(E | X) when the conditions X are fully known. Similarly, it is also inconsistent to be uncertain about the expected or unconditional probability of the event in question, P(E), since the unconditional probability can be found by simply taking the expectation of P(E | X) over all possible values of X.

As an illustration, let us consider an example in which the assessor is asked to give the probability of observing a head in the first toss of a fair coin. Coherence in

this case requires the assessor's response to be 0.5. Let us now replace the coin and inform our subject that the coin has either two heads or two tails (each with probability 0.5). Again, we would expect the probability of observing a head in the next trial to be 0.5. However, this value is the expected probability that the coin has two heads. The conditional probability of heads would be 1.0 if the coin had two heads, and 0.0 if the coin had two tails.

To formalize this, let X be the unknown condition of the coin; i.e.,

$$X = \begin{cases} \text{HH if two heads} \\ \text{TT if two tails} \end{cases}$$

Then the conditional probability of observing a head given X is

$$P(\text{head} \mid X=HH) = 1 \text{ (with probability 0.5)}$$

$$P(\text{head} \mid X=HH) = 0 \text{ (with probability 0.5)}$$

In other words, P(head | X) is now a random variable taking on values 0 and 1, each with probability 0.5. This constitutes a probability distribution over the conditional probability of a head on the next toss.

Naturally, the unconditional probability of observing a head in the next trial is simply equal to 0.5, since this is the expected value of the conditional probability, P(head | X), taken over all possible values of X. This is exactly the same as the unconditional probability of heads for the fair coin. However, there are clearly a great many decision problems (e.g., cases involving multiple tosses of the same coin) for which the two coins are far from equivalent.

The notion of uncertainty about underlying conditions is a simple but powerful explanation for the phenomenon of uncertainty about probability. Since this explanation depends only on the concept of conditional probability, such conditional uncertainties are clearly consistent with the axioms of subjective probability theory. Therefore, uncertainties of this type do not necessitate the development of new axioms, as is done in some of the approaches in the literature. It is worth nothing that this type of conditional uncertainty is neither a "higher" nor a "lower" order probability, but simply the probability of another event or verifiable proposition. Since the uncertainty in such cases stems not from some "failure of introspection" but rather from uncertainty about the conditioning event, the validity of a given conditional probability is therefore verifiable, based on whether the conditioning event actually occurs. Viewed from this perspective, conditional probabilities are verifiable propositions, about which uncertainty is permitted.

Imprecision

In the above discussion, we have considered uncertainty about probability arising from uncertainty about some underlying condition or event. However, there is another source of uncertainty or ambiguity, arising from cognitive limitations in our ability to assign unique numerical values to states of knowledge. As a result of these limitations, most probability judgments are subject to a certain degree of inherent imprecision, described "inexact measurement"[5] of beliefs. In other words, people are often unable to distinguish between different numbers intended to represent a given state of knowledge; the different probabilities simply don't "feel" sufficiently different. This can be viewed as analogous to an inability to distinguish between similar colors. In fact, a certain degree of imprecision is probably inherent in the capabilities of the human brain, just as any other physical system (e.g., a computer or a measuring device) is also subject to

certain limitations on its numerical accuracy.

This type of imprecision is inherently context-dependent, since our abilities to represent and process information will vary from one context to another. Thus, in attempting to estimate the probability of getting two heads in a row from a coin where P(head) is equal to .537, our range of imprecision may be relatively narrow, subject only to limitations in our ability to multiply three digit numbers in our heads. By contrast, in estimating the probability of an event that has never happened, the demands on our information-processing capabilities are much more complex, and the range of imprecision may therefore be much greater. Imprecision may also be motivation-dependent. In other words, we may be able to do a better job if we choose to; e.g., by performing a few simple calculations, or by introspecting more deeply. Finally, calibration training may be able to reduce imprecision, at least in some contexts. Risk analysts may be reasonably well-calibrated in estimating component failure probabilities on the order of 10^{-2} to 10^{-4}, but may nonetheless be poorly calibrated for more commonplace events such as tomorrow's weather.

In contrast to the case of conditional uncertainty, as discussed earlier, the axioms of subjective probability are not adequate to explain imprecision. In practice, we believe that most cases in which uncertainty about probabilities is large enough to be important are likely to reflect primarily the effects of conditional uncertainty rather than imprecision, and uncertainty due to imprecision will often be sufficiently small that it can simply be neglected.

In some contexts, however, the uncertainty due to imprecision may actually be as large as that due to lack of knowledge about underlying conditions, and imprecision might then need to be taken into account in decision making along with conditional uncertainty. However probability distributions can be used in practice to represent imprecision in such case.

REFERENCES

1. G. Parry , On the meaning of probability in probabilistic safety assessment, *Reliability Engineering and System Safety*, 23 (1988).
2. Apostolakis, G. Probability and risk assessment: the subjectivist viewpoint and some suggestion, *Nuclear Safety*, 19, (1988).
3. Kaplan, S. and B.J. Garrick, On the quantitative definition of risk, Risk Analysis, 1, (1981).
4. Mosleh A. and V.M. Bier , Uncertainty about probability: A reconciliation with the subjectivist viewpoint, submitted for publication , *Journal of Risk and Uncertainty*, (1993).
5. Suppes, P. , The measurement of belief, Journal of the Royal Statistical Association: Series B, 36, (1974)

066 Cognitive Modeling and Simulation in Human Reliability Analysis

Chair: Y. Fujita, Mitsubishi Atomic Pwr. Ind.

A Simulation Study of Errors of Commission in Nuclear Accidents
K-S. Hsueh (NUS); L. Soth (Commonwealth Ed.); A. Mosleh (U. Maryland)

Simulating Operator Cognition for Risk Analysis: Current Models and CREWSIM
V.N. Dang (MIT); N.O. Siu (INEL)

Quantification of Multiple Error Expressions in Following Emergency Operating Procedures in Nuclear Power Plant Control Room
A.P. Macwan, P.S. Wieringa (Delft U. Technol.); A. Mosleh (U. Maryland)

A SIMULATION STUDY OF ERRORS OF COMMISSION IN NUCLEAR ACCIDENTS

Kae-Sheng Hsueh[1], Larry Soth[2], Ali Mosleh[3]

[1]NUS, 1411 Opus Place Suite 103, Downers Grove, IL 60532
[2]Commonwealth Edision Company, Downers Grove, IL 60532
[3]University of Maryland, College Park, MD 20742

INTRODUCTION

Nuclear industrial experience and operating events show that human factors have a substantial influence on nuclear safety. Among the different types of operating crew, responses to abnormal conditions and emergency situations (post-accident phase) are of particular concern for human reliability analysis. Human errors modeled in current probabilistic risk assessments primarily include errors of omission or failure to perform the task within the specified time window. These models are developed primarily for categorizing and assessing the probabilities of different types of operator errors, primarily errors of omission. Some have relied on an analysis of factors influencing operator response and performance and have used subjective methods to estimate error probabilities.[1,2] Others have focused on collecting and analyzing simulator-based performance data in support of probability models.[3,4] In recent years, research efforts[5,6,7] are being directed towards modeling errors of commission and cognitive errors by elements of human information processing that leads to action (perception, intention, and action). These elements as well as their overall interactions are used to describe how errors are generated. However, these methods are in their early, conceptual stages of development.

The purpose of this study is to develop a modeling methodology for errors of commission focusing on abnormal situations in nuclear power plants, a very constrained environment in which interactions are primarily guided by written or memorized procedures. Specifically, the development of a realistic human model, which addresses on errors of commission in the post-trip operator actions, is the subject of this paper. To demonstrate the proposed methodology, steam generator tube rupture (SGTR) event and loss of heat sink (LOHS) event are analyzed by using a simulation based accident analyzer (ADS).[8]

THE MODELING APPROACH

In the area of decision making behavior, the human operator can be modeled by three basic cognitive activities: the monitoring, the interpretation, and the decision making. Human errors can arise from different activities in various formats. Some of the human error forms, such as errors due to inadequate training, errors due to misreading, and execution errors, have been extensively studied.[1,2,3] Therefore, the proposed model only focuses on modeling intention-related operator errors.

Transient events in NPPs involve very complex interactions among the reactor system, emergency safeguard systems, and operator actions. Our approach to modeling operator behavior is towards developing a model which generates sets of actions likely to be taken by control personnel of NPPs based on the characteristics of perceived symptoms and operator mental state. During accidents, control room operators, directed by procedures or the operator judgement, continuously monitor the plant physical parameters and safety systems.e.g., a high containment pressure or abnormal radiation in containment is a strong cue for diagnosing the event as a loss of coolant accident (LOCA) which will stimulate operator to focus on primary side pressure and inventory control. In other words, we assume that there are two kinds of operator monitoring activities: EOP-driven and intention driven.

In the interpretation activity, operators process observed symptoms and retrieve relevant knowledge to formulate the possible explanation. It is our assumption that the EOPs have been taken as a major influencing factor for operator behavior. Therefore, the EOP can be used as the basis of the human model knowledge base. The symptoms and the recovery action plans associated with different kinds of initiating events, and expected plant responses versus system states are also included.

Based on interpretation and the information perceived, the operator makes decisions to select action strategy. According to guidelines, control room personnel, as required by regulations, have to follow emergency operating procedures. However, there are several incidents which involve procedure violations. Therefore, it is conceivable that operator behavior can be described as the consequence of the interactions of operator thinking processes and the EOPs' underlying recovery schemes. Intentional mistakes and cognitive errors are often generated by the mismatch between operator interpretations and expectations and EOPs. Although guided by the EOPs, operators form the their own interpretations and predictions from their perception of the plant status and system response and of the evolution of the event. Based on their knowledge and experience, control room personnel may feel enough confidence so as to deviate from EOPs. Accordingly, operator commission errors are primarily the results of deviation.

Our approach to modeling operator behavior is towards developing a model which generates sets of actions likely to be taken by control personnel of NPPs based on the characteristics of EOPs, perceived symptoms, and operator mental state. According to our studies,[8,9] types of deviations from procedures that are relevant to NPP emergency operation are identified as described below:

- omission (intentional and unintentional)
- shortcut
- delayed action
- fallacy

The deviation action categories described here when implemented as a part of the human cognitive model can capture most important human errors in nuclear power plant control rooms. As a minimum the recognition of the causes and mechanisms of these errors will help human reliability analyst to develop a better model of operator behavior which will also facilitate the probabilistic assessment process by providing a more tangible basis for the error.

THE INFERENCE RULES

To generate spectrum of likely actions in a given context, the modeling approach is to identify factors, Performance Shaping Factors (PSFs), that influence selection of recovery strategy and affect operator decision making behavior. Then, the functional relationships between PSFs and operator actions are formulated as a set of generic inference rules. In each instant, context-dependent information and generic inference rules are used to specialize operator decisions and responses for specific situations.

Investigation of underlying human inference mechanisms of operator deviation actions shows that post-accident human actions are primarily affected by the following, but not limited to, PSFs:

- perceived symptoms
- perceived system states
- perceived stress or time pressure
- procedure context and intent
- consequence of actions
- operator confidence and prediction
- strengths of reasoning paths

The specific inference of these PSFs are based on our interviews with operators, review of operating events, and findings from literatures.

Perceived symptoms: This PSF has the major influence on the selection of recovery procedures. Under the current symptom based guidelines, operators will enter recovery procedures only if the EOP specific entry conditions are satisfied. Therefore, except for the reception of false signals or incomplete information, the chance of misdiagnosis is insignificant.

Perceived system states: During the accidents, operators are directed to verify the functionality of safety systems. Once a failed system has been identified, operators will face a high stress and less available resources. The chances for operators to take shortcuts or erroneous actions will increase. The underlying mechanism is that, under this circumstance, operators will try several ways to relieve the cognitive burdens.

Perceived stress or time pressure: This PSF causes a threat or time-stress to the operator. For example, if the operator finds high level in a steam generator in the event of SGTR, he would conclude that the steam generator is about to become solid, resulting in water going to the steam lines. This may lead him to take a shortcut type action. Our investigation shows that this PSF is associated with other PSFs such as the perceived symptoms. In the SGTR example, if operators fail to diagnose the event, then this PSF will have less influence on their decisions.

Procedure context and intent: Depending on the intent of the procedures, EOPs can be classified as verification, diagnosis, re-diagnosis, and action steps. There are several distinct attributes, such as importance to the incident, associated with procedures. Potential actions in these situations have been comprehensively discussed.[8]

Consequence of actions: Perceived Consequences of actions may cause the operator to delay actions such as feed and bleed type of cooling or containment spray which will jeopardize power operations. The "cost/benefit" considerations of action consequences would make the operators hesitate to take actions in time or wait for other alternatives as opposed to EOP specific actions.

Operator confidence and prediction: Expectations about plant behavior and confidences on the diagnosis are generated as a result of operator diagnosis and the actions performed on the plant. If the behavior of the plant does match the operator expectations, operator's confidence will increase which may lead to overlook other indications. When a mismatch happens, it would cause the operator to lose confidence in his own judgement and perform re-diagnosis.

Strengths of reasoning paths: This PSF depicts the effects of deep reasoning process. Human errors arise when the strength of erroneous reasoning paths is so strong such that wrong judgment is reached regardless of missing, contradictory or defective evidences.

In summary, this paper proposes a cognitive approach in characterizing operator actions as procedure following with various ways of deviation and in the classification of human decision making mechanisms and influencing factors.

THE CASE STUDY AND RESULTS

To demonstrate the feasibility of proposed approach and to test the role of various cognitive factors, the proposed model is implemented in a dynamic simulation environment.[8] Key human cognitive factors are delineated and traced in terms of operator mental state including:

- an operator diagnosis state which includes operator's perception of plant system functioning status and the perceived accident symptoms.
- an emergency procedure state which indicate which procedure is to be followed and what step to be executed.
- a stress factor represents a measurement of operator psychological stress and operator's confidence of diagnosis.
- an action index which denotes which erroneous action path (including no error) being followed.

The Westinghouse symptom based procedures and a Westinghouse plant have been chosen for the case study. All the major response procedures (E-0, E-1, E-2, and E-3) as well as functional restoration guidelines are thoroughly analyzed using the proposed model. SGTR and LOHS are selected as initiating events. Furthermore, various system configurations are also used to examine possible operator actions. Due to the lack of credible data source, this study uses the probability values adapted from the relative likelihoods.

Table 1 summaries the important operator actions found in this simulation study. Because of the significant uncertainty about the operator actions times and probabilities, only the qualitative results are discussed here.

- Shortcuts and Omissions:

The simulation results show that the short cuts and omission errors, although can occur relatively more frequently, have minor impacts on the nuclear safety. This is because that the structure of EOPs comprises of several redundant steps and the diagnosis strategy.

- Stress:

The effect of operator stress only have minor impact on the overall results. Our study shows the continuous monitoring of critical function tree dictated in the EOPs is a very effective way to prevent human errors. However, our study is limited to slow transients only. For the cases such as large LOCAs, human error rates can be much higher.

- Delayed actions:

In the case of LOHS, the delayed action due to perceived consequence has been found to be the dominant contributor to core deterioration. Even EOPs explicitly ask operator to conduct bleed and feed, operator tends to resist in taking the necessary actions and use the primary side conditions, instead of steam generator water level, as the decision reference.

- Instrument failures:

The current EOPs use the accident symptoms as the cue for diagnosis. However, throughout the whole procedures, there is no redundant steps or instructions to handle false signals and instrument failures. The failed SG radiation alarms in the case of SGTR events, for example, have been found to be the root cause of diagnosis difficulty.

- Fallacy:

This error mode which is embedded in the process of reasoning network has been found the dominant contributor to the core melt. It is not unusual that plant operators often misuse "rules of thumb" even with missing, contradictory or defective grounds.

CONCLUSIONS

A methodology has been presented which provides a systematical framework for analyzing and implementing operator errors during NPP accidents. EOPs and intentional deviations from EOPs form the basis for defining operator actions. Major performance shaping factors as well as the inference rules are also discussed in this paper. In order to support this new expanded realm, a simulation based model- ADS is used to implement the proposed operator model.

The case studies show that this approach is a realistic, yet implementable human model, which explicitly formulates errors of commission in the context of emergency operating procedures (EOPs). Integrated with the dynamic accident sequence analysis methodology (ADS), the simulation results show the proposed human model has the

capability to capture the cognitive elements of human interactions and can be used in large scale PRA applications.

However, further work needs to be carried out to develop cognitive human error rate data that are used in the methodology presented. Simulator exercises, exhaustive reviews of operating events, or expert judgements may provide useful information.

REFERENCES

1. Swain, A., and Guttmann, H., "Handbook of Human Reliability Analysis with Emphasis on Nuclear Power Applications," NUREG/CR-1278, 1983.
2. Embrey, D., Humphreys, P., Rosa, E., Kirwan, B., and Rea, K., "SIIM-MAUD: An Approach to Assessing Human Error Probabilities Using Structured Expert Judgment, " NUREG/CR-3518, 1984.
3. Hannaman, G.W., Spurgin, A.J., and Fragola, J.R., "Systematic Human Action Reliability Procedure (SHARP)," Interim Report, NP-3583, Electric Power Research Institute, June 1984.
4. Hall, R.E., Frogola, J., and Wreathall, J., "Post Event Human Decision Errors: Operator Action Tree/Time Reliability Correlation," NUREG/CR-3010,1982.
5. Amendola, A., Bersini, U., Cacciabue, P.C., and Mancini, G."Modelling Operators in Accident Conditions: Advances and Perspectives on a Cognitive Model," Int. J. Man-Machine Studies, 27, 599, 1987.
6. Reason, J. "Recurrent Errors in Process Environments: Some Implications for the Design of Intelligent Decision Support Systems," in Intelligent Decision Support in Process Environments, Hollnagel, E., Mancini, G., and Woods, D.D. (EDITORS), NATO ASI Series, 1985.
7. Woods, D.D., Roth, E.M., and Pole, H.Jr., "Modelling Human Intention Formation for Human Reliability Assessment," Reliability Engineering and System Safety, 22, 169-200, 1988.
8. Hsueh, K.S., "An integrated simulation model for plant/operator behavior in accident conditions", Ph.D. thesis, UMCP, 1992.
9. Hsueh, K.S. and Mosleh, A., "Categorization of Operator Errors in the Context of Symptom-Based Emergency Operating Procedures," SMiRT, Tokyo, 1991.

Table 1. The simulation results

Events	Damage state	Reason	Likelihood
SGTR	SG overfill	Failed SG rad. alarm	1E-1
SGTR	Extended leak	False LOCA signal	2E-2
LOHS	Core melt	Consequence consideration	4E-3
LOHS	Core melt	Failed PORV with shortcut	3E-3

SIMULATING OPERATOR COGNITION FOR RISK ANALYSIS: CURRENT MODELS AND CREWSIM

Vinh N. Dang[a] and Nathan O. Siu[b]

[a]Nuclear Engineering Department
Massachusetts Institute of Technology
Cambridge, MA, 02139, USA

[b]Idaho National Engineering Laboratory
P.O. Box 1625
Idaho Falls, ID, 83415-3855, USA

ABSTRACT

This paper compares three models of nuclear power plant operators developed to address operator cognitive errors in abnormal or emergency scenarios: the Cognitive Environment Simulation (CES), COSIMO (for Cognitive Simulation Model) and CREWSIM. The comparison highlights the designs of the models and the relationships of these designs with the behaviors that the models can exhibit. It also reviews the capabilities of these models with respect to the needs of probabilistic safety assessment.

INTRODUCTION

Cognitive errors* are significant as a potential cause of erroneous operator actions because subsequent actions may be associated with the same underlying cognitive processes. Through the situation assessment that motivates the operator response and the response planning that generates the intentions to act, the cognitive processes are a source of dependencies linking human actions.

Consideration of these internal, cognitive processes can be problematic in typical human reliability analysis (HRA) methodologies for a number of reasons. First, the processes of monitoring, situation assessment, and response generation are complex and interdependent. Second, analyzing operator cognition requires a "richer" description of context than is explicitly carried in methodologies such as event trees.[1] Finally, each erroneous action, especially the inclusion of actions not in the nominal response, leads to a subscenario, the evolution of which may impact the probability of erroneous actions. A non-failure/success oriented analysis is needed to treat these subscenarios.

This paper reviews three models that have been developed to address control room operator cognitive errors during accidents. The Cognitive Environment Simulation (CES), developed at Westinghouse with the sponsorship of the U.S. Nuclear Regulatory Commission, models the cognitive activities of nuclear power plant operators in dynamic fault management. It emphasizes capturing the demands imposed by the problem-solving

*"Cognitive error" refers to causes of erroneous actions associated with internal cognitive processes.

environment.[2-3] The Cognitive Simulation Model (COSIMO) is being developed at the Ispra Joint Research Centre of the Commission of the European Communities to model human behavior in simulated accidents in order to identify situations that may lead to erroneous behavior and, more generally, to assess and design man-machine systems.[4] The CREWSIM simulation model, being developed at the Massachusetts Institute of Technology (M.I.T.), is aimed at evaluating operator behavior for risk assessment purposes; it models both the procedure-guided and cognitively-based crew responses in accidents and treats the interactions between these. CREWSIM builds on an earlier crew simulation model also developed at M.I.T.[5] These models are potentially useful for supporting probabilistic risk assessment; this review examines model characteristics that might affect such an application.

MODELING APPROACH AND SCOPE

The dynamics of the plant evolution are central to the analysis of cognition; they determine the indications that can be observed by the operators. Both the dependence of this evolution on operator actions and, through intention formation processes, on cognition, as well as the allocation over time of operator cognitive resources involve dynamic interactions. Together, these issues lead to the simulation approach shared by CES, COSIMO, and CREWSIM. All three models are dynamic simulations which treat data monitoring, situation assessment, and intention formation. Table 1 summarizes each model's scope.

Table 1. Scope of models.

Characteristic/ Process	CES	COSIMO	CREWSIM
model of:	cognitive environment or system	individual	crew as interacting individuals
monitoring	√	√	√
situation assessment	√	√	√
intention formation	response planning	response generation	response generation
action execution		√	√
communication			√
procedures /procedure-following	implicit in knowledge base	no	explicit

The cognitive system modeled by CES is distributed across multiple people and may include machines like diagnostic aids, all of which perform cognitive work. The current version of CES models "the cognitive environment that reveals the cognitive functions that must be performed to successfully handle a domain task, rather than a performance model that describes the detailed processes by which these functions are achieved."[2]

The cognitive processes in COSIMO and CREWSIM are both based on the frame concept, which is assumed to dominate resource-limited processing in the control room/emergency response work environment. Although the models can incorporate reasoning processes for planning, e.g. goal decomposition, the frame-based implementations emphasize intention formation based on action scripts.

Both COSIMO and CREWSIM model the execution of operator actions, including erroneous actions not caused by cognitive errors (execution-related errors). COSIMO's treatment is general; CREWSIM currently only models execution-related errors in selected actions. CREWSIM explicitly models the plant's written procedures and the interaction with the operators' assessment, permitting consideration of potential conflicts. Note that procedure following represents a load on: a) the operators' attentional resources that can be devoted to diagnosis, and b) the time left to execute actions such as diagnosis-driven monitoring (as opposed to procedure-driven monitoring).[2]

Currently, all three cognitive models have been applied in scenario-specific analyses. CES has been used to treat steam generator tube rupture (SGTR), loss of heat sink, interfacing systems loss of coolant accident, and variants on these scenarios. COSIMO has been used to model a variety of auxiliary feedwater system (AFWS) transients. The initiating event for the reported COSIMO sequences is reactor scram; the AFWS then starts automatically. The operator's role is to regulate flow and maintain steam generator water levels to remove decay heat until other systems can be placed in operation when the plant reaches cold shutdown conditions. CREWSIM's predecessor models operator crews in SGTR scenarios; CREWSIM responds to loss of residual heat removal variants with the reactor in the midloop condition, a shutdown plant state.

BEHAVIORS AND DESIGN

As discussed above, the three models reviewed have similar basic structures. Where they differ is in the detailed implementation of these structures. Differences can be observed in the treatment of: the structure and contents of the knowledge base, the cognitive processes that manipulate the knowledge base, and the control of these processes. They also differ in their treatment of factors/parameters that shape operator performance.

Knowledge Base Structure

The structure of the knowledge base, i.e., how knowledge is represented, is important because it affects the implementation of cognitive processes with specific characteristics and thereby influences the ease of modeling various behaviors. The range of knowledge representations in the three models extends from inference rules to larger structures known as frames or schemas.

Inference rules represent knowledge in detail, at a low or fine-grained level. This detailed representation entails more processing (in the operator's mind as well as in the simulation) and more conscious processing (in the mind). The increased cognitive load, however, is counterbalanced by increased flexibility in the application of the knowledge.

Frames, on the other hand, are similar to compilations or linkages of rules and tend to imply recognition rather than detailed understanding. Conceptually, the frame links a set of cues (the triggers for a group of rules) with the final output of processing this group of rules. The linkage or compilation is only strictly applicable for the specific situation with these cue values. However, because cues for natural situations are often redundant, recognition processes may be triggered by observation of only part of the cue set. This is typical of recognition processes and modeled by frame-based algorithms. The frame concept also models another characteristic of recognition, the tendency to assume or at least expect the unobserved cue values. Finally, the frame also associates an appropriate response to the recognized entity.

CES is an implementation of a production rule system (EAGOL), a system for processing inference rules.[2] While production rules are typically if-then structures, CES's basic knowledge unit, the "coupler," can be used to make inferences about the effects of an occurrence (event) and, in the reverse direction, to infer possible causes for an observed effect. A strength parameter expresses the strength of an inference. Conditions can be placed on the coupler, representing the circumstances under which the rule applies (like an inhibit gate in a fault tree), and may also vary the nature of the coupler relationship. The CES knowledge base also includes "scenarios." These represent likely sequences of events that might unfold following an initial disturbance. Scenarios are similar to frames, although they emphasize expectations over time given a situation rather than the recognition of the situation. These expectations are used to narrow the range of possible diagnoses.

COSIMO uses "rule-based frames" (RBFs) to represent knowledge about initiating event evolutions and appropriate responses to these. A RBF lists the characteristic cues (both physical variable values and system indications), a subjective frequency, and an appropriate response sequence in a single data structure. RBFs are essentially static snapshots; dynamic evolutions are represented by sequences of frames.

CREWSIM uses schemas to structure knowledge. Schemas are conceptually related to frames. Like frames, schemas are related to recognition processes. However, they empha-

size cognition in action-oriented circumstances. Different types of schemas are used for different tasks and the associated knowledge. The generic diagnosis (situation) schema resembles COSIMO's RBFs; it consists of slots for physical cues, indications, and the appropriate operator response. The operator response itself is specified using an action schema, which describes the situation calling for this action, the goals of the action, and the plant response. The schema types form abstraction hierarchies[6] that cover the range of the operator's cognitive activities, from situation assessment to response generation to action monitoring, at different levels of detail, e.g. component- system-plant, action-sequence.

Knowledge Base Content

The contents of the knowledge base determine what a given cognitive model can respond to and what it can consider and include in this response. By changing the knowledge base content (with a corresponding change in the physical model), the model structure can be exercised on different scenarios. For a given scenario, performance can also be shaped by varying the knowledge base content to reflect individual differences in depth and accuracy of understanding, in how the knowledge base is organized, and in priority, availability, and even confidence.[3,5]

CES incorporates many kinds of knowledge in its knowledge base; this variety provides the model with considerable power and flexibility. In addition to abnormal events and their symptoms, CES knows about: automatic systems; critical parameters; scenarios; operator actions; goals, options, and payoffs; and plant layout. The detailed knowledge of automatic system actions enables CES to sort out the multiple influences on plant response. Given a symptom, CES can hypothesize its causes, generate expectations, and determine confirmatory and inconsistent symptoms. Its plant knowledge is used to locate faults. Using its goal-options knowledge, CES can determine the options available in a given context and choose among these.

The knowledge represented in COSIMO's rule-based frames is described in the preceding section; this knowledge is sufficient to model the operator's cognitive activities in a recognition-based mode. Heuristics and general engineering and physical principles will eventually be contained in "knowledge-based frames." These are being developed in conjunction with mechanisms for direct inference to model knowledge-based planning, plan adaptation, and cognitive processes in general.

CREWSIM's knowledge of situations and responses is represented in coupled schema hierarchies, with schemas for situation assessment and for actions or responses. While the situation schemas are similar to COSIMO's rule-based frames, CREWSIM explicitly models the operator's access to the operating procedures. Consequently, the linkage of diagnosis and an associated response (in the situation schema) may conflict with the procedure-directed response. Note that knowledge associated with written operating procedures is incorporated to some degree in the knowledge base for all three models since the procedures reflect the overall training of the operator. However, unlike CES and COSIMO, part of CREWSIM's knowledge base is derived directly from the procedures because the schemas are action- or response-oriented. Conceptually, schemas are acquired from experience; hence, they tend to emphasize procedural ("what-to-do") over declarative knowledge ("assertions"). The CREWSIM assumption is that this procedural knowledge is closely related to the written procedures.

Cognitive Processes

While the knowledge base constrains the behaviors of the model, the cognitive processes that manipulate the knowledge base and their control generate the behaviors. These processes can be described as model-driven or data-driven. For example, model-driven or diagnosis-driven monitoring is monitoring for cues confirming or disconfirming a diagnosis (the situation model). On the other hand, data-driven monitoring is based on indicator prominence and related to the operator's focus of attention.

Table 2 shows the major cognitive activities and types of processes included in each of the models. Although the situation assessment activity is aimed at explaining data, it is characterized in this table as model-driven because of its central role in a model-driven response. All three models incorporate a model of data-oriented monitoring. For instance, the operator models will detect the auditory and visual cues associated with plant alarms. If

task loading is modeled, as in CREWSIM, identification of the alarm has to be incorporated into an "action list"(a queue) and executed. Data-driven response planning refers to responding to symptoms rather than causes and implies that understanding the situation is assumed or not attempted. This process can be represented by a rule, "if-symptom, then-response," or a schema where cue recognition would trigger an associated response.

Table 2. Cognitive activities and processes.

Activity	Model/diagnosis -driven		Data-driven
	knowledge-based inference	"rule"-based (frame/schema)	
Monitoring	CES (e.g. behavior analyst)	- COSIMO, CREWSIM - CES (scenarios)	all 3 models
Situation Assessment	CES (e.g. situation analyst)	COSIMO, CREWSIM: recognition (matching)	N/A
Response Planning/Generation	CES knowledge-based planning (option comparison)	COSIMO, CREWSIM: associated with diagnosis frame	none

The production rule system of CES is processed by multiple independent processing agents or "analysts." The analysts are intended to model the multiple types of processing (e.g. monitoring, explanation building) and lines of reasoning that occur during problem-solving. The analysts primarily use inference on different parts of CES's knowledge base to accomplish the cognitive activities; the major exception is related to the "scenario" since this knowledge structure is essentially a frame storing expected symptoms, cue evolutions, for the scenario.

In COSIMO and CREWSIM, the frame or schema is the basis for situation assessment and response generation. Situation assessment uses the cue set of the frame to diagnose; the diagnosis process invokes the frames whose cue sets match the observed cues. Diagnosis-based monitoring (to confirm or revise a diagnosis or to resolve alternative hypotheses) checks for cues not observed or for expectations provided by the frame. Whereas CES generates by inference a response based on a diagnosis, COSIMO uses the predetermined response associated with the diagnosis in the frame. This tight coupling of response and diagnosis in COSIMO is not present in CREWSIM since the CREWSIM operators may continue to follow the procedures.

Regarding the selection of frames/schemas, COSIMO uses a "similarity matching"[7] mechanism that compares the degree of similarity (degree to which the observed cues match the frame's cue set) with an "evaluation threshold" for acceptance. If multiple frames satisfy the acceptance criterion (match the observed cues sufficiently), a "frequency gambling"[7] mechanism may be used and selects the frame with the highest subjective frequency (a frame attribute). CREWSIM's matching mechanism is a probabilistic rather than a deterministic function of the threshold-similarity relationship. Because procedure following is an option, the CREWSIM operator may accept multiple alternatives and monitor for further cues.

Controlling Cognition and Shaping Performance

In problem-solving, an intelligent system or agent has to identify problems to solve, decide what knowledge to use, and select methods and strategies to apply. It also has to decide how to evaluate alternative problem solutions (hypotheses), evaluate these alternative solutions, and select criteria to allocate attention to selected problems.[8] In modeling human operators, an important aspect of control is the allocation of cognitive and task resources. Of the cognitive resources, attention is especially critical, given the large number of control room indications.

CES control is currently knowledge-driven. Thus, analysts are generated when they are needed or can contribute to the solution and alternative hypotheses are pursued until a

knowledge-based resolution is obtained. Because it does not model the attention and processing resource limits of people, the CES performance is consistently faster and more thorough at following the implications of disturbances than actual operator crews in the same scenarios, although it uses the same line of reasoning as the operators.[2]

In the current version of COSIMO, control focuses on the cognitive activities underlying a rule-based behavioral mode, with an implementation based on the frame. The control problem then primarily concerns hypothesis (i.e., frame) acceptance and revision. COSIMO offers a rich set of operator characteristics that can be modified to influence performance. COSIMO considers all frames whose degree of similarity with the observed cues exceed an evaluation threshold. Because COSIMO's response is always based on the current hypothesis, multiple frame matches are resolved immediately either by choosing the best match or the most subjectively frequent frame, depending on the relative weights of the similarity-matching and frequency-gambling mechanisms. If none of the frames exceed the evaluation threshold, the operator continues to evaluate plant data in an "endless diagnostic stage," without selecting and carrying out a response. Hypothesis revision, on the other hand, depends on the degree of similarity of the current frame, which is recalculated as the scenario evolves.

In this model, resource limitations are not modeled explicitly. Instead, the frame-based behavioral mode and its underlying processes are assumed to reflect these limitations. In the monitoring and detection activity, for example, plant information is filtered on the basis of salience, a weighted function of external, physical factors and cognitive factors; the latter include the indications' relevance to the current hypothesis or expectations. A "filter threshold" models the extent of the field of attention. An "interpretation threshold" models the operator's sensitivity to indication trends or changes.

In CREWSIM, the control issue is broader because the operators have several possible modes of behavior. In procedure following, operators allocate resources both to the foreground task (procedure following) and to situation assessment. The actions associated with the schemas used for situation assessment can be inserted into or replace the procedure-guided response, providing the alternative response modes. Resource allocation to the dual tasks of procedure following and situation assessment is determined by an operator's "commitment to own assessment," which affects the level of processing resources committed as well as their interruptibility. Related control decisions concern the resolution of conflicts between the operator's assessment and the procedures. "Confidence in procedure" and "initiative" are combined into "commitment to procedures," which influences operator behavior in situations where the procedures appear inadequate.

As in COSIMO, the CREWSIM model also provides control parameters that influence diagnosis revision and direct monitoring attention. However, the frequency gambling mechanism is not emphasized to the same extent. In following procedures, for instance, the operator may delay resolution of alternative diagnoses until the data can support resolution. In a diagnosis-oriented mode, the "action/assessment orientation" may favor either responding prior to or "independently of" a global assessment or data gathering.

CONCLUSIONS

The comparison presented in this paper shows how each model's capabilities and behaviors relate to its design. The CES performance represents "an absolute lower limit on how quickly a disturbance can be diagnosed, i.e. the time at which all the evidence required ... becomes available". It characterizes the scenario more than the dynamics of the operator response and thus can be used in an analysis of the cognitive demands of problem-solving situations.[2] One application in risk assessment of such an analysis may be to provide criteria for successful human performance.

To model human performance, CES needs to include "psychologically plausible mechanisms to limit what data CES can 'see' on each time step, what knowledge it can access, and the thoroughness of its reasoning."[2] The COSIMO and CREWSIM models are two attempts to model operators with such mechanisms, i.e., to model resource-limited cognition. Work is continuing on the COSIMO project to develop "direct inference" cognitive processes; validation efforts have shown that this form of reasoning is important and that it is used for continued situation assessment while carrying out a selected response plant.[4]

This processing mechanism would add reasoning capabilities similar to those now exhibited by CES.

The CREWSIM model emphasizes a realistic treatment of the cognitive and task loads of the control room crew. Modeling the 'overhead' associated with the allocation of the resources available to a crew has led to a model of interacting individuals because the overhead includes, for instance, communication tasks for commands and reports, interruption of task performance and cognitive processes, as well as potential interference between the assessments of individual operators.

At the individual level, the CREWSIM model currently handles the interactions of multiple cognitive tasks. At least initially, these tasks are procedure execution and continued situation assessment. The situation assessment in CREWSIM relies on schema-based processes (rather than inference); CREWSIM's schema hierarchies provide sufficient combinatorial richness for a limited model of response in unfamiliar situations. While COSIMO is aimed at a realistic model of human performance that includes both human limitations as well as human capacities for a flexible, innovative response in unfamiliar situations, CREWSIM reflects a shorter-term goal to understand the dynamics of human interactions and their potential impacts on accident sequences. The CREWSIM model as a whole may be viewed as a dynamic model of how context, procedures, and other factors normally treated in human reliability analyses interact to result in the operator response. It is expected to be useful in supporting current risk assessment applications much in the same manner that detailed thermal hydraulic models are used; CREWSIM provides a tool to examine a range of crew responses to a nominal scenario. The model may also have applications in analyzing human performance data.

ACKNOWLEDGMENTS

This work has been supported in part through the Nuclear Engineering and Health Physics Fellowship program administered by ORISE for the U.S. DOE. Partial support has also been provided through the EG&G Idaho Laboratory Directed Research and Development Program under DOE Idaho Field Office Contract DE-AC07-76ID01570. The authors would like to thank Dr. Emilie Roth of Westinghouse Corporation for her helpful discussion of the CES project.

REFERENCES

1. N. Siu, Dynamic accident sequence analysis in PRA: a comment on "Human reliability analysis - where shouldst thou turn?", *Rel. Engr.Sys. Safety,* 29:359 (1990).
2. E.M. Roth, D.D. Woods, and H.E. Pople, Cognitive simulation as a tool for cognitive task analysis, *Ergonomics*, 35:1163 (1992).
3. D.D. Woods, H.E. Pople, and E.M. Roth. "The Cognitive Environment Simulation as a Tool for Modeling Human Performance and Reliability," NUREG/CR-5213, U.S. Nuclear Regulatory Commission, Washington, DC, 1989.
4. P.C. Cacciabue, et al., COSIMO: a cognitive simulation model of human decision making and behavior in accident management of complex plants, *IEEE Trans. Sys., Man, and Cybernetics*, 22:1058 (1992).
5. Y. Huang, V. Dang, and N. Siu, Modeling control room crews for accident sequence analysis, in *Proc. Probabilistic Safety Assessment Intl.Topical Mtg. (PSA '93)*, American Nuclear Society, La Grange Park, IL, 1993.
6. K.J. Vicente and J. Rasmussen, Ecological interface design: theoretical foundations, *IEEE Trans. Sys., Man, and Cybernetics*, 22:589 (1992).
7. J. Reason. "Human Error," Cambridge University Press, Cambridge (1990).
8. B. Hayes-Roth, Blackboard architecture for control, *J. Art. Intel.*, 26:251 (1985).

QUANTIFICATION OF MULTIPLE ERROR EXPRESSIONS IN FOLLOWING EMERGENCY OPERATING PROCEDURES IN NUCLEAR POWER PLANT CONTROL ROOM

A. P. Macwan[1], P. A. Wieringa[1], and A. Mosleh[2]

[1]Laboratory for Measurement and Control, Dept. of Mechanical Engineering, Delft University of Technology, Mekelweg 2, 2628 CD Delft, The Netherlands

[2]Dept. of Materials and Nuclear Engineering, University of Maryland College Park, MD 20742, U.S.A.

1. INTRODUCTION

Current HRA methodologies[1-4] are directed towards assigning probability for a single human error. These methods are based on expert estimation of the human error probability (HEP) using a decompositional approach[1], expert estimation of a success likelihood index which is then converted to HEP[2], or time-reliability correlation where HEP is calculated as a function of time and other appropriate variables[3,4]. These methods have been criticized[5,6] for the lack of a comprehensive treatment of errors of commission, in particular, and of intentional errors, in general. They do not explicitly address quantification of multiple error expressions (different ways the operators can err), but address operator interaction with the plant as binary in terms of performance or non-performance of required action. However, the methods provide a systematic, structured framework for assessing HEP, the influence of various factors on HEP, collecting expert assessments, combining judgements of different experts, etc.

Recent research in Human Reliability Analysis (HRA) and related fields include efforts in analyzing operator behavior, cognitive modelling; and includes error analysis, including categorizations, classifications, and taxonomies of error; and methodological work (primarily simulation-based) that incorporates various types of operator-plant interactions in risk and reliability analysis. Hsueh and Mosleh[7] pointed out possible categories of operator errors in following emergency operating procedures (EOPs). Macwan and Mosleh[5] expanded this categorization and provided systematic definitions of error expressions. Other research on operator behavior points to the need to include errors of commission[8] and other errors, besides the omission error. This implies that multiple errors will likely be incorporated in reliability and safety analyses. In turn, this would require that these errors be quantified.

The method presented in this paper employs use of quantification by expert estimation within the context of the multiple errors identified in the research mentioned above [5]. One of the reasons for using expert estimation is the lack of data about human error. Unless such data is available for the multiple errors, expert estimation in some form needs to be used to quantify these errors. The quantification method presented in this paper uses expert estimation. However, it can also be used for data collection. Quantification is important for applications to PSA as well as others such as simulation-based and analytical study of the causes and consequences of these errors.

The remainder of the paper is divided into four sections. The method of quantification is presented in the first section. Secondly, an application of the method to multiple errors in following EOPs is discussed. This section also includes definitions of specific errors and associated factors. In the third section, an example application of the method for a single operator-plant interaction is provided. Conclusions are presented in the last section.

2. QUANTIFICATION METHOD

The quantification method is based on a systematic classification of operator actions, especially errors. In order to assess relative likelihoods that are normalized to 1, appropriate factors are independently assessed as a function of performance influencing factors (PIFs). These factors are then combined so that the sum of the relative likelihoods is 1.

2.1 Classification of Operator Actions

A classification of operator actions, with focus on errors, is shown in Fig. 1. At the first level, an action is divided into success or an error. Error is broadly classified into intentional and unintentional. Intentional errors are further classified into procedural and non-procedural, where procedural errors are the ones where the error is an action within those prescribed in the procedures, while non-procedural errors refer to an action that is different from the ones prescribed in the procedures.

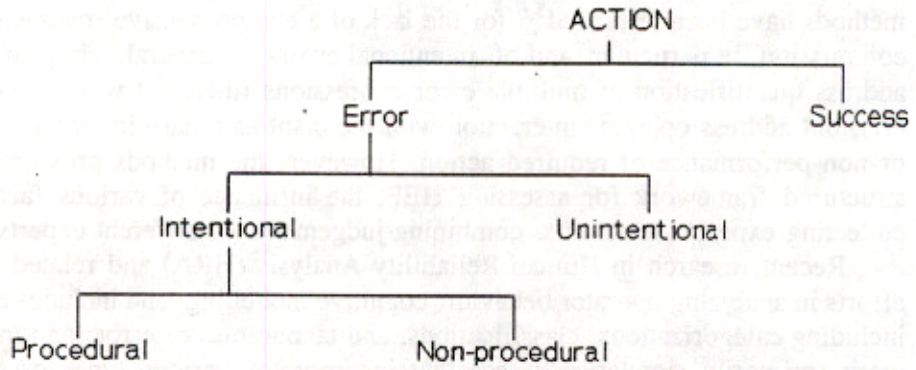


Fig. 1 Classification of Operator Action

This classification is general with respect to the error being intentional and within/outside procedures. Specific errors, procedural as well as non-procedural, are dependent on various factors such as the type of activity the procedure involves, time into the incident, etc. Examples of specific errors are given in Section 3.

2.2 Assessment of Relative Likelihoods

As mentioned earlier, the strategy for quantification is use of PIFs to assess relative likelihoods of different errors. Based on the recognition that some PIFs have effects that change during an accident scenario, while others have effects that do not change during the scenario, PIFs are classified into scenario-dependent and scenario-independent [Macwan and Mosleh[5]]. Using this distinction as a basis, the relative error likelihoods are assumed to depend mainly on the scenario-dependent factors, as discussed below.

The assessment is done in 3 steps, as follows.

1. Relative weights for different errors, called α factors, are assigned by experts, as a function of scenario-dependent PIFs. α's are assigned a value on a scale from 0 to 1.
2. Scenario-independent effects are accounted by assigning adjustors to the different α factors, designated by β's. Similar to α factors, β adjustors are assigned a value between 0 and 1, and are assessed by the experts.
3. The overall likelihood of each error is obtained by combining the appropriate α an β factors. Using a multiplicative approach, each β adjustor is multiplied with the corresponding α factor. Moreover, the likelihoods are calculated in such a way that they add up to 1.

3. APPLICATION

The method is applied to errors in following EOPs in response to accidents in nuclear power plants. Considering the EOPs as guidelines for operator response following an accident, deviations from the prescribed actions are broadly termed as errors. Also considered as errors are selection of incorrect procedures due to misdiagnosis. Major activities in following EOPs includes verification of automatic response of the plant protection systems, selection of procedure based on diagnosis of the root cause, and manual recovery actions. After identifyng applicable errors, different factors are defined to systematically quantify them.

3.1 Possible Errors

The following errors have been identified[5] and include errors in selection of procedure (first 3) and carrying out the instruction of the EOP (last 3).

* **Shortcut**, where a block of steps in the procedure is omitted with a specific end-objective in mind.
* **Incorrect Procedure Selection**
* **Delayed Procedure Selection**
* **Skip** refers to the operator omitting the action specified by the procedure.
* **Delayed Action**
* **Commission** refers to the operator taking an action that is different from the step prescribed in the procedure.

3.2 Scenario-Dependent Factors

The factors are based on the classification shown in Fig. 1. The applicable factors are α_s, α_{un}, α_{np}, α_{sk} and α_{md}. Each of them is defined below. For factors other than

α_s, equations similar to eq. 1 can be written. According to Macwan and Mosleh[5], likely errors are dependent on the type of activity the operator is performing according to the EOP. Thus, for each phase of EOP, a specific set of α's is applicable.

- α_s is the relative likelihood of the operator action being successful, and is defined as shown in equation 1. The overall likelihood of error is 1 - α_s since at this level, operator action is classified into success and error.

$$\alpha_s = \frac{w_{success}}{w_{success} + w_{errors}} \quad (1)$$

where w_i is the relative likelihood of the action being i

- α_{un} is the relative likelihood of the operator error being unintentional, given that the operator action is an error.
- α_{np} is the relative likelihood that the operator error is non-procedural, given that an intentional error has occurred. Commission is defined as a non-procedural error, thus the relative likelihood of procedural errors, given the same conditions, is 1 - α_{np}.
- α_{sk} is the relative likelihood of a procedural error being a skip, given that an intentional, procedural error has occurred. Thus, the relative likelihood of shortcut, given identical conditions, is 1 - α_{sk}.
- α_{md} is the relative likelihood of the operator having a misdiagnosis of the root cause of the incident. The likelihood of correct diagnosis, therefore, is 1 - α_{md}. It should be noted that α_{md} does not refer to the error expression but rather to misdiagnosis which is a cause of the error (e.g., in selection of procedure). Thus, it is not conditional like the other α factors (unintentional error given an erroneous action).

3.3 Scenario-Independent Adjustors

As mentioned earlier, these effects are assumed to modify the behavior suggested by the α factors. This is implemented in quantification by the so-called β factors, which are adjustors for the corresponding α factors (used to assign scenario-dependent weights). The values of β's are assigned by experts. β's are assigned on a scale from 0 to 1, where the value of 1 indicates an optimum crew with respect to the β being assigned. For example, a β_s of 1 means a crew behaviour is the same as that predicted by α_s, while a β_s of 0 means a crew with no likelihood of success.

It is proposed that not all α factors have corresponding β factors. The following assumptions have been made: (i) There is no scenario-independent effect on unintentional behaviour. Unintentional behavior is independent of the crew and plant specific characteristics. (ii) The misdiagnosis of an incident is primarily a function of the symptoms which are scenario-dependent. Thus, there is no effect of scenario-independent factors on the likelihood of misdiagnosis.

3.4 Relative Likelihoods of Errors

The overall likelihood of different actions including errors are obtained by combining α factors and β adjustors. An example of such combination is presented in Table 1. As can be seen in the third column of Table 1, β's are multiplicative factors. The weights are determined left to right within each row, and from top to bottom within

the table. The formulation is done with the constrain that all the weights add up to 1.

Table 1. Combination of α and β factors

α Factor	β Factor	WEIGHT (w_i)	ACTION
α_s	β_s	$w_s = \alpha_s \beta_s$	Success
α_{un}	---(1)	$w_{un} = \alpha_{un} [1 - w_s]$	Unintentional (Delay)
α_{np}	β_{np}	$w_{np} = \alpha_{np} \beta_{np} [1 - w_s - w_{un}]$	Non-procedural (Commission)
α_{sk}	β_{sk}	$w_{sk} = \alpha_{sk} \beta_{sk} [1 - w_s - w_{un} - w_{np}]$	Skip
α_{md}	---(1)	$w_{scm} = \alpha_{md} [1 - w_s - w_{un} - w_{np} - w_{sk}]$	Shortcut (Misdiagnosis)
---	---	$w_{scc} = [1 - w_s - w_{un} - w_{np} - w_{sk} - w_{scm}]$	Shortcut (Correct Diagnosis)

NOTE: (1) For detailed discussion, see section 3.3.

4. EXAMPLE

To demonstrate the method, the selected operator-plant interaction (as guided by the EOP) is of verification of containment isolation, phase A and B[9]. The set of SD PIFs includes operator diagnosis, time pressure, importance and memory. The set of SI PIFs includes operator experience, confidence, and recent experience with faulty signal[5]. The estimation of relative likelihoods of errors (w's) is presented below.

The first step consists of eliciting the α's from the experts as a function of the SD PIFs. Similarly, β's are elicited as a function of SI PIFs. Finally, they are combined using formulation of Table 1 to obtain w's, as shown in Table 2. The numerical values shown in the third column of Table 2 represent the relative likelihoods of the applicable errors.

Table 2. Example of Multiple Error Quantification

α Factor	β Factor	WEIGHT (w_i)	ACTION
$\alpha_s = 0.9$	$\beta_s = 0.9$	$w_s = 0.810$	Success
$\alpha_{un} = 0.1$	---	$w_{un} = 0.019$	Unintentional Delay
$\alpha_{np} = 0.1$	$\beta_{np} = 0.4$	$w_{np} = 0.007$	Commission
$\alpha_{sk} = 0.8$	$\beta_{sk} = 0.9$	$w_{sk} = 0.118$	Skip
$\alpha_{md} = 0.0$	---	$w_{scm} = 0.00$	Shortcut (Misdiagnosis)
---	---	$w_{scc} = 0.046$	Shortcut (Correct Diagnosis)

5. CONCLUSIONS

In this paper, a method for quantification of multiple errors is presented with focus on errors in following EOPs in NPPs. Recent attention on errors on commission, as

well as recognition of multiple ways that operators can err, point to possible applications of the method. An important characteristic of the method is recognition and incorporation of generic classification of possible operator actions. This renders the method to be applicable to non-nuclear industries. The relative likelihoods of possible errors that are normalized to 1 and thus can be incorporated into probabilistic analysis rather easily. Moreover, the error likelihoods are assessed conditional to the characteristics of the scenario as well as plant and crew characteristics. Thus, the method described in this paper can form a basis of data collection for validation, and use of this data for assessment of error likelihoods in the long run.

REFERENCES

1. A.D. Swain and H.E. Guttmann, "Handbook of Human Reliability Analysis with Emphasis on Nuclear Power Plant Applications", NUREG/CR-1278, U.S. NRC (1983).
2. D.E. Embrey et al, "SLIM-MAUD: An Approach to Assessing Human Error Probabilities Using Structured Expert Judgement", NUREG/CR-3518, U.S. NRC (1984).
3. G.W. Hannaman, A.J. Spurgin, and Y.D. Lukic, "Human Cognitive Reliability Model for PRA Analysis", NUS-4531, Electric Power Research Institure (1984).
4. E.M. Dougherty, Jr. and J.R. Fragola, "Foundations for a Time Reliability Correlation System to Quantify Human Reliability", IEEE Fourth Conference on Human Factors and Power Plants, Monterey, CA (1988).
5. A.P. Macwan and A. Mosleh, "Methodology for Analysis of Operator Errors of Commission during Nuclear Power Plant Accidents with Application to Probabilistic Risk Assessment", University of Maryland Report MDNE 93-001 (1993).
6. E.M. Dougherty, Jr., "Human Reliability Analysis: Where Shouldst Thou Turn?", Guest Editorial in *Reliab. Engr. and System Safety* 29, Pp. 283-299 (1990).
7. K.S. Hsueh, and A. Mosleh, "Categorization of Operator Errors in the Context of Symptom-Based Emergency Operating Procedures", SMiRT, Tokyo (1991).
8. S. Hirschberg, "Experience from International Peer Reviews of Probabilistic Safety Assessments", Proc. of International Symposium on Use of Probabilistic Safety Assessment for Operational Safety, Vienna, Pp. 35-46 (1991).
9. Westinghouse Electric Corporation, "Westinghouse Owners Group Emergency Response Guidelines", (1983).

067 Comparative Risk Assessment of Complex Technological Systems (I)

Chair: S. Haddad, New South Wales Govt.

Severe Accidents in Comparative Assessment of Energy Sources: Current Issues and Actual Experience Data

S. Hirschberg, C. Parlavantzas (P. Scherrer Inst., Switzerland)

An Analysis of Electricity Generation Health Risks: A United Kingdom Perspective

D.J. Ball, L.R.J. Roberts, A.C.D. Simpson (U. East Anglia, England)

External Costs as an Indicator for Comparative Risk Assessment of Different Energy Systems - A Case Study for Comparing the Nuclear and the Coal Fuel Cycle

W. Krewitt, P. Mayerhofer, R. Friedrich, A. Greβmann (U. Stuttgart, Germany); M. Dreicer (CEPN, France)

The Use of Multi-Attribute Analysis for Comparing the Health and Environmental Impacts of Different Fuel Cycles - A Feasibility Study for the Comparison Between Coal and Nuclear Fuel Cycles

M. Dreicer, P. Kuentzmann, T. Schneider (CEPN, France); W. Krewitt (U. Stuttgart, Germany)

SEVERE ACCIDENTS IN COMPARATIVE ASSESSMENT OF ENERGY SOURCES: CURRENT ISSUES AND ACTUAL EXPERIENCE DATA

Stefan Hirschberg and Constantin Parlavantzas

Paul Scherrer Institute
Würenlingen and Villigen
CH-5232 Villigen PSI
Switzerland

INTRODUCTION

Paul Scherrer Institute (PSI) and the Swiss Federal Institute of Technology in Zurich (ETHZ) established in 1993 a long-term project on "Comprehensive Assessment of Energy Systems" (Hirschberg et al., 1993). The ultimate goal of this project is an integrated evaluation, covering the risk-related, environmental and economic aspects, associated with different energy systems of interest for the Swiss conditions.

According to the latest major international review of the current status of comparative studies of health and environmental effects associated with the various energy sources, one of their weak points is the treatment of risk contributions from severe accidents (Chadwick, ed., 1991). Furthermore, in a recent editorial in Risk Analysis (Fritzsche, 1992) the international risk community was encouraged to direct efforts towards assembling and evaluating data on energy-related severe accidents, since the present status of such information is unsatisfactory.

The field of energy production is according to the existing (very incomplete) data evaluations a major contributor to man made disasters. The evidence based on actuarial data provides only a partial picture of risks associated with the different energy systems. First, the conditions (for example with respect to technology, safety principles and culture, physical and operational environment) characteristic for a specific event may be such that its applicability to other conditions will be questionable, possibly precluded. The quality, relevance and uncertainties associated with such assessment vary. Second, the actuarial experience, if available, in most cases only reflects some examples from a wide spectrum of potential accident scenarios. Third, for some energy sources the actuarial accident-related data are so scarce that in the best case only predictive estimates are available. Fourth, there is a need, if possible, to cover the risk contributions associated with all parts of fuel cycles (including up- and down-stream processes) and with stages in the plant life-cycle other than operation.

The present work at PSI aims at improvements of the state of knowledge on severe accidents, for use in the context of comparative studies. This paper presents some of the emerging results concerning both actuarial data and predictive assessments. In addition, a

number of issues related to this work is briefly discussed. At this stage there is no intention to provide a complete list of issues but rather focus on some important ones.

PSI SEVERE ACCIDENT DATABASE FOR THE ENERGY SECTOR

The database is in the process of development. Further extensions and updates will be made during 1994; this will be accompanied by detailed evaluations. For this reason, the numerical results presented in this paper should be regarded as preliminary. The focus of present evaluations is on electricity production.

Table 1 shows the principal sources of information used for the generation of the PSI database. The contents (with emphasis on the energy-related events) of these databases have been merged. This is augmented with events from literature. With the exception of WOAD, none of the databases used has been specifically intended for addressing energy-related events.

Table 1. Databases of actuarial accidents used to date as the input for the PSI database.

Organi-sation	Country of Origin	Full Name	Code Name
OFDA	USA	The Office of Foreign Disaster Assistance Database	OFDA
UK SRD	UK	The Major Hazards Incidence Data Service	MHIDAS
DNV	NORWAY	The Woad Offshore Databank	WOAD
	USA	The Minerals Management Service Accident Database	MMS
CEC-ISPRA	EEC	The Major Accidents Report System	MARS
EPA	USA	The Acute Hazardous Events Database	AHE
SR	CH	Disaster Data & the Oplis Database	OPLIS-SIGMA
UBA	BRD	The Accidents Handbook	

Figures 1 and 2 reflect the status of the PSI database. In Figure 1 man made accidents resulting in fatalities are accounted for and compared with the content in one of the external databases (Resources for the Future; RfF). The PSI database contains at present about 5500 events in spite of the fact that some of the originally present events (such as aircraft crashes) were screened out. Figure 2 shows the status of the PSI database with respect to the energy-related events (about 2300 occurrences). It needs to be pointed out that allocation of the events to specific energy activities is partially subjective. When compared to commercial and other data sources utilised by us, the PSI database exhibits a more balanced distribution of contributing countries.

The energy-related events in the PSI database were further allocated to the different fuel cycles and the normalised fatality rates were calculated. To our knowledge the best available, previously published corresponding data are those in (Chadwick, ed., 1991; after Fritzsche, 1989); see Table 2. Table 3 shows our preliminary results, using the same period of time (1969-1986) and the same normalisation. Delayed fatalities, particularly relevant for the Chernobyl accident, are not included in the tables. Further checks are needed of the hydropower accidents in the PSI database to verify that all of them occurred at facilities for energy production. Due to the much larger number of events for oil and gas, the estimated corresponding fatality rates based on the PSI database are significantly higher.

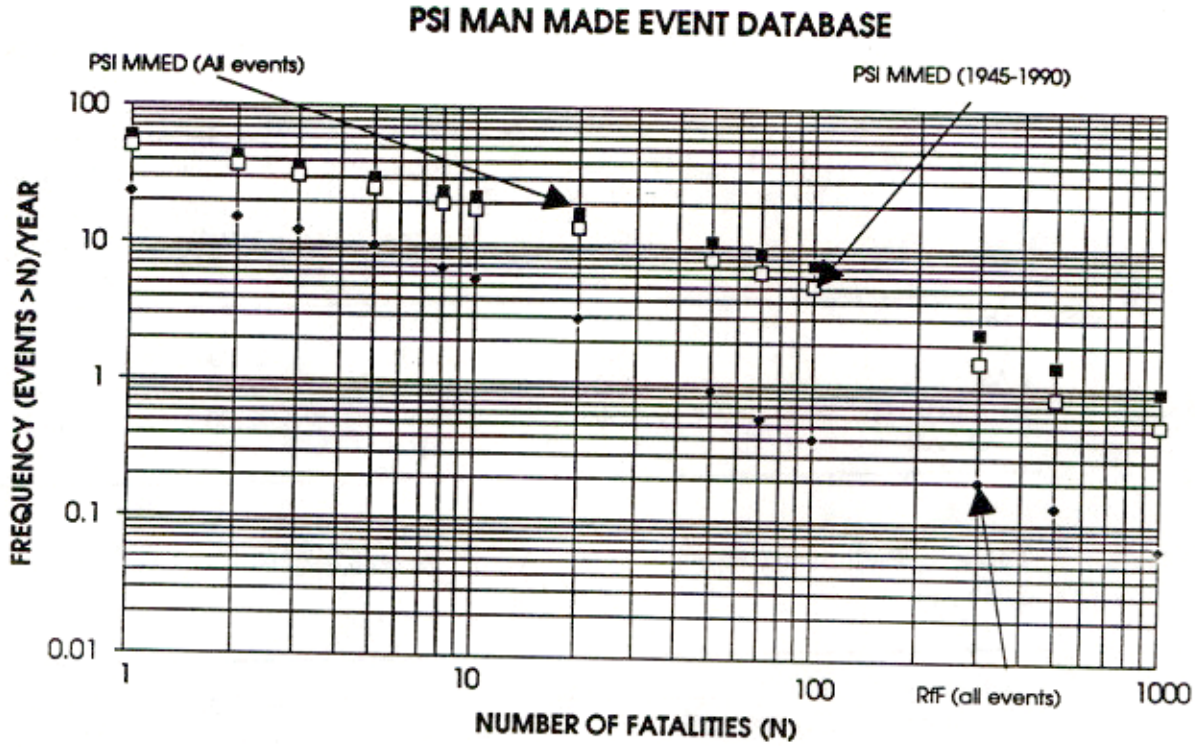


Figure 1. Current status of the PSI man made accident database, compared to the RfF database.

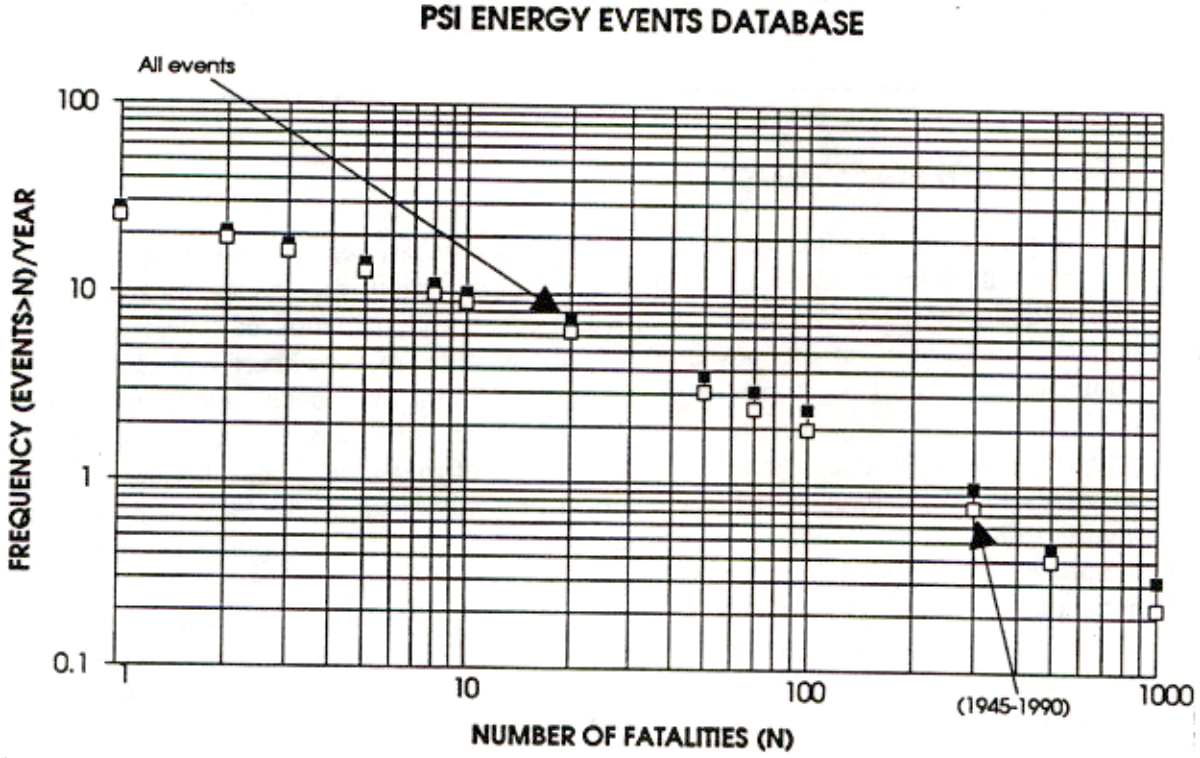


Figure 2. Current status of the PSI energy-related accident database.

Table 2. Normalised fatality rates for severe accidents (1969-1986) (Chadwick, ed., 1991; after Fritzsche, 1989).

Energy Option	No of Events	Immediate Fatalities/ Event	Total Immediate Fatalities	Energy Produced ($GW_e \cdot a$)	Immediate Fatalities/ Energy (#/$GW_e \cdot a$)
COAL	62	10 - 434	3600	10'000	0.36
OIL	63	5 - 500	2070	21'000	0.10
GAS	24	6 - 452	1440	8600	0.17
HYDROPOWER	8	11 - 2500	3839	2700	1.42
NUCLEAR	1	31	31	1100	0.03

Table 3. Preliminary normalised fatality rates for severe accidents (1969-1986) (PSI database).

Energy Option	No of Events	Immediate Fatalities/ Event	Total Immediate Fatalities	Energy Produced ($GW_e \cdot a$)	Immediate Fatalities/ Energy (#/$GW_e \cdot a$)
COAL	64	10 - 434	3944	10'000	0.39
OIL	170	5 - 2700	6423	21'000	0.31
GAS	83	5 - 550	3282	8600	0.38
HYDROPOWER	17	9 - 2500	5405	2700	2.00
NUCLEAR	1	31	31	1100	0.03

PREDICTIVE ASSESSMENT

In parallel to the work on the severe accident database we are addressing the issue of hypothetical severe accidents in the different steps of various fuel cycles, by application of probabilistic approaches. The initial activities concern contribution of severe accidents to the external costs of nuclear power. For details we refer to Hirschberg and Cazzoli (to be published 1994). This controversial issue has attracted much attention in recent publications. The discrepancies in the estimated external costs are enormous and are primarily caused by the varying (by several orders of magnitude) contributions of nuclear accidents to the external costs. These are in turn governed by the estimated number of fatalities that would result from extreme nuclear accidents. The published literature is dominated by simplistic approaches which use the Chernobyl specific total dose to the population as the basis for the calculations. Use of **one** specific accident at a plant with a design philosophy and an operating environment so different from "western" conditions is not an acceptable approach for the estimate of the external costs of nuclear power. Possible acceptable approaches could either be based on a careful extrapolation of results from a Probabilistic Safety Assessment (PSA) performed for a suitable reference case or on a plant-specific PSA for the case of interest. The second approach is naturally preferable but may not always be feasible.

The numerical consequences of the different approaches are illustrated by few examples. Table 4 shows some of the mean risk measures accounted for in (Hirschberg and Cazzoli, to be published 1994), and based on plant-specific Level III PSAs. Mühleberg is a

Swiss nuclear power plant (GE BWR4 with MARK-I containment); the results for Peach Bottom and Zion where obtained using the information in NUREG-1150 (USNRC, 1990).

Table 4. Calculated mean risk measures directly based on Level III PSAs.

Risk Measure (per $GW_e \cdot a$)	Mühleberg	Peach Bottom[1]	Zion[1]
Early fatalities	$<10^{-12}$	3.5×10^{-8}	1.0×10^{-5}
Latent cancer fatalities[2]	6.7×10^{-3}	6.2×10^{-3}	2.4×10^{-2}
Individual risk of cancer death	1.6×10^{-9}	6.7×10^{-10}	9.4×10^{-9}

[1] Only internal events

[2] Calculation to 800 km for Mühleberg and 1600 km for Peach Bottom & Zion

For comparison, the widely published results by Hohmeyer (1988) based on an invalid and arbitrary extrapolation of the Chernobyl specific experience, suggest a cancer rate associated with nuclear power accidents of about 140-1400 cancers per $GW_e \cdot a$ for German plants (implicitly supposed to apply to other western plants). Recent work by Masuhr and Oczipka (1993) leads to an estimate for Mühleberg of 0.048-1.75 fatalities per $GW_e \cdot a$. Also in this case Chernobyl specific total population dose was used as the reference; however, an attempt was made to partially account for the plant-specific frequencies for the different source terms and for few site-specific characteristics.

SOME INSIGHTS AND CURRENT ISSUES

Below follows a brief summary of insights and issues identified in the course of the work on the severe accident database and on predictive assessments, being in progress at PSI. Efforts to address some of the present limitations will continue within the current long-term project.

Completeness. In spite of a relatively modest effort the PSI database appears to have a significantly higher degree of coverage of energy-related accidents when compared with other existing databases known to the authors. It also provides a better balance between reporting levels for individual countries. The completeness of the PSI database in the absolute sense can only be addressed qualitatively due to lack of an objective criterion for such a comparison. In spite of the progress made, the authors have reason to believe that a significant number of events with considerable consequences may still be missing. Radical improvements of the confidence with respect to the completeness and accuracy would require a systematic evaluation of information in the vast amount of sources beyond those used to date. This implies commitment of large resources to the task.

Scope/Severe Accident Definition. Apart from hydropower no attempt has been made to evaluate past severe accidents related to renewable fuel cycles (e.g. manufacturing of solar cells). This would require use of data sources other than those employed until now. It would also be desirable to cover a wide range of consequences (such as acute and delayed fatalities, serious injuries, evacuation, extensive land, water or food contamination), when evaluating the events in the severe accident database. However, due to the scarcity of information, poor statistical evidence and lack of accuracy, at present only evaluation of acute fatalities and possibly serious injuries seems to be meaningful.

Uncertainties and Practical Limitations. Uncertainties are intrinsic to all steps in the evaluation of actuarial accident data. Here only some specific sources of uncertainty are commented on. Thus, categorisation of events as energy-related is frequently based on subjective judgement. In many cases it may not be clear whether a particular event should be allocated to oil or gas fuel cycles. In order to make the comparison possible the energy source-specific risk coefficients for coal, oil and gas calculated in Tables 2 and 3 were obtained by converting the total energy production to an equivalent electrical output. In reality there are alternative (other than electricity generation) uses of these energy options.

Validity and Applicability. As reflected by the nuclear example given in the present paper, the fact that relatively large uncertainties are associated with the evaluations of severe accidents does not justify use of invalid approaches. Use of a plant-specific PSA, if available, is the most rational basis for the evaluation of risks associated with hypothetical severe accidents. In case this is not feasible any extrapolations must be done with due attention paid to the differences that may influence the results. It is acknowledged, however, that in the context of some other energy sources compromises are necessary due to the non-uniform level of knowledge and limited scope of applications of risk analysis. When evaluating the actuarial severe accidents questions arise with respect to the relevant time span to be covered (technological and regulatory changes over time) and with respect to inclusion/exclusion of specific countries (differences in the level of technological development and in the safety culture).

Limitations of Quantifiable Results. Quantified risks are subject to intrinsic and practical limitations related to the data and methodological approaches used. In addition, the quantifiable results lack another important dimension. It is acknowledged that socio-political factors, whose evaluation is dependent on subjective judgement, need to be considered in the context of comparative risk assessment of energy sources. This limitation of quantifiable results is important to keep in mind, but their merit as one essential indicator among others is indisputable.

ACKNOWLEDGEMENT

This work is supported by the Swiss Federal Office of Energy (BEW).

REFERENCES

Chadwick, M. J. ed., 1991, Key issue paper no. 3 by International Expert Group 3: Comparative environmental and health effects of different energy systems for electricity generation, in: "Proceedings of Senior Expert Symposium on Electricity and the Environment", held in Helsinki, Finland, May 13-17, 1991.

Fritzsche, A. F., 1989, The health risks of energy production, Risk Analysis, 9:565.

Fritzsche, A. F., 1992, Severe accidents: Can they occur only in the nuclear production of electricity?, Risk Analysis, 12:327.

Hirschberg, S. et al., 1993, Assessment of energy systems (Ganzheitliche Betrachtung von Energiesystemen - GaBE), Detailed outline of the project, 2nd Version, Paul Scherrer Institute, Würenlingen and Villigen.

Hirschberg, S. and Cazzoli, E., to be published 1994, Contribution of severe accidents to external costs of nuclear power, Topical Meeting on PSA/PRA and Severe Accidents, Ljubljana, Slovenia, April 17-20, 1994.

Hohmeyer, O. (1988), "Social Costs of Energy Consumption. External Effects of Electricity Generation in the Federal Republic of Germany," Springer-Verlag, Berlin, Heidelberg, New York, London, Paris, Tokyo.

Mashur, K. P. and Oczipka, T., 1993, The external costs of nuclear electricity production, Prognos, Basel.

USNRC, 1990, Severe accident risks: An assessment for five U.S. nuclear power plants, Final summary report, NUREG-1150, U.S. Nuclear Regulatory Commission, Washington DC.

AN ANALYSIS OF ELECTRICITY GENERATION HEALTH RISKS: A UNITED KINGDOM PERSPECTIVE

D. J. Ball, L. E. J. Roberts and A. C. D. Simpson

Environmental Risk Assessment Unit
School of Environmental Sciences
University of East Anglia
Norwich NR4 7TJ, England

INTRODUCTION

There is renewed interest in the application of comparative risk analysis as a decision aid in the planning of electricity generation systems. New studies have been undertaken on behalf of DG XII of the Commission of the European Communities[1], and the International Atomic Energy Agency has hosted a series of meetings in Europe[2,3,4] on methodologies and procedures, aimed at achieving greater consistency in comparative risk analysis, in accord with recommendations of the 1991 Helsinki Symposium on Electricity and the Environment[5].

Likewise, Nuclear Electric plc has sponsored the Environmental Risk Assessment Unit to produce a study of the human health risks posed by alternative large scale (1,000 MW) electricity generation schemes which were envisaged as feasible in Britain during the next two or three decades(*). Earlier studies had been conducted in Britain by Ferguson[6] in 1981, Cohen and Pritchard[7] in 1980, and Fremlin[8] in 1985. However, significant changes had since occurred in some risks and renewables had not previously been assessed.

What follows is a brief account of results and conclusions which have emerged from the study. Interested parties are encouraged to obtain the full report[9] which provides details of the methodology, limitations and uncertainties. The full report also addresses major accidents.

SCOPE AND METHODOLOGY

The study examined the human health risks associated with seven representative technologies. These were tidal power (Severn Barrage); on-shore wind and off-shore wind

(*) Responsibility for the contract (CON/90080/VN) passed from the then CEGB to Nuclear Electric plc following the latter's investiture as a limited company on 31 March 1990.

power; nuclear power (PWR); and coal, oil and gas-fired power. A systems approach was adopted along the lines recommended by the IAEA working groups[2,3,4], in which risks associated with each stage of selected fuel cycles are examined. Thus, for the coal cycle, this would include risks of mining, transportation and so on, through to public risks associated with effluent emissions. However, continental and global scale risks of the type associated with acid deposition and greenhouse gas emissions were not included, these being covered by other studies. For the purpose of risk comparison, all risks are normalised per gigawatt-year of electrical energy (GWa), with attention having focused on those risks which are greater than or approaching 0.1 fatalities or serious diseases per GWa.

OVERVIEW OF ROUTINE OCCUPATIONAL RISKS

Table 1 summarises our estimates of the acute occupational fatality risks per GWa for stages of the seven electricity generation cycles examined. Risks in this Table have been calculated from either actuarial data, or from materials and man-power requirements and associated risk factors[8]. From Table 1 it can be seen that the main acute occupational risks are likely to be encountered during the construction phase and, to a lesser extent, during material and component fabrication, and operation. The analysis suggests that the largest of these risks is of the order of 0.5 to 1 per GWa and may be associated with the construction of off-shore wind farms. This conclusion stems from the higher risk factor which has been ascribed to off-shore construction activities in the North Sea. For the nuclear and fossil cycles, Table 1 shows acute occupational fatality risks to be associated with some or all of the fuel extraction, fuel preparation and reprocessing, plant construction and plant operation stages. None of these risks is very large. Even that of coal mining, traditionally a hazardous occupation, has been estimated to have fallen below one acute fatality per GWa for mines operated by British Coal.

The estimated occupational risks of serious disease including cancer are shown in Table 2. The largest risk appears to be that of pneumoconiosis in sub-surface coal mines though, as with prompt fatalities, this risk is far smaller than in earlier decades, for mines operated by British Coal. Routine occupational radiological risks of fatal cancer have been calculated for the nuclear and fossil cycles on the basis of collective doses and current ICRP risk factors[10].

OVERVIEW OF ROUTINE RISKS TO THE PUBLIC

The projected main public health risks arising are listed in Table 3. Practically all of the risks reported in this Table are of disease rather than acute fatality. The estimates are mainly hypothetical, especially those in parentheses, and most entail massive speculation. For example, in the case of the nuclear cycle, there are major philosophical difficulties associated with the calculation and interpretation of public collective radiation doses. Individual doses, even to the most exposed individuals, are much less than the doses from the natural background. Yet integration over vast time scales, to infinite time in some instances, or over global populations, produces collective doses which may appear significant. An inherent presumption in this kind of calculation is that the dose-response curve for radiation carcinogenesis is linear with no threshold down to minute doses orders of magnitude below those received every year from natural background radiation to which we are unavoidably exposed. Even were these risks realised, the relevant comparison should be with estimated fatalities over the same time scale as a result of exposure to natural background radiation.

Major uncertainties also apply when assessing the risks posed to public health by effluents from fossil-fired power stations. It has taken many years of research to demonstrate the link between fossil power plant emissions and acidification of vulnerable environments. To

determine the impact upon public health of the low doses associated with effluent emission from modern power plant presents a far more formidable task. The ranges shown in Table 3 are based on a review of estimates by other authors, all of whom stress the uncertainties involved, and are understandably very wide.

SOME CONCLUDING REMARKS

This paper has summarised normalised health risk estimates associated with the construction, routine operation and decommissioning of seven feasible large-scale UK electricity generation options. Occupational risk estimates of prompt and delayed fatalities are the most robust of the statistics, though still subject to many sources of unavoidable uncertainty. The indications are that the occupational risks for each of the seven cycles, when combined for each cycle, currently lie in the range of about 0.1 to 0.2 to a maximum of about 1 to 2 fatalities per GWa, with tidal power and gas lying at the lower end of this range, and coal and off-shore wind at the higher end.

Significant public health risks, at least of the type included within the remit of this project, were not identified for the renewable cycles. Comparisons between the nuclear and fossil cycles are fraught with difficulty. Individual radiation doses to the public from the nuclear (and fossil) cycles are insignificant compared to those from natural background radiation, but collective doses may appear significant if integrated over thousands of years and continental or global populations. The meaning of such estimates in terms of health cannot be assessed, and in any event the estimates pale into insignificance when compared with those resulting from exposure to natural radiation were this to be calculated on the same basis. This problem is further compounded when comparison between nuclear and fossil cycles is attempted, since the latter produce copious quantities of either solid, liquid or gaseous effluents, which may contain heavy metals, carcinogens and known respiratory irritants, but for which neither short nor long-term public health implications have been quantified with confidence.

REFERENCES

1. CEC/US, Joint Study on Fuel Cycle Costs, in preparation.
2. International Atomic Energy Agency. "Methods for Comparative Risk Assessment of Different Energy Sources." Proceedings of a Specialists Meeting, Studsvik, Sweden, 27-30 August. Organised and published by IAEA, Vienna (1991).
3. International Atomic Energy Agency. "Format and Structure of a Database on Health and Environmental Impacts of Different Energy Systems for Electricity Generation." IAEA TECDOC-645, Vienna, Austria. ISSN 1011-4289 (1992).
4. International Atomic Energy Agency. "Methods for Comparative Risk Assessment of Different Energy Sources. IAEA-TECDOC-671, Vienna, Austria. ISSN 1011-4289 (1992).
5. International Atomic Energy Agency. "Comparative environmental and health effects of different energy systems for electricity generation. Key Issues Paper No.3 chaired by Chadwick, M.J. In: *Senior expert symposium on electricity and the environment*, Helsinki, 13-17 May. Published by IAEA, Vienna.
6. R.A.D. Ferguson. "Comparative Risks of Electricity Generating Fuel Systems in the UK." Peter Peregrinus Ltd. ISBN 0 906048 66 4 (1981).
7. A.V. Cohen and D.K. Pritchard. "Comparative Risks of Electricity Production Systems: A Critical Survey of the Literature." Health and Safety Executive, H.M.S.O. (1980).
8. J.H. Fremlin. "Power Production: What Are The Risks?" Adam Hilger. ISBN 0 85274 133 2 (1985).
9. D.J. Ball, L.E.J. Roberts and A.C.D. Simpson. "An Analysis of Electricity Generation Health Risks - A United Kingdom Perspective, in preparation.
10. International Commission on Radiological Protection. "1990 Recommendations of the International Commission on Radiological Protection". Publication 60, Annals of the ICRP (1991).

Table 1.Summary of estimated acute occupational fatality risks per GWa by stage of electricity generation cycle. Data sources, methods of calculation, limitations and uncertainties are described elsewhere[8]. Note: transport risks include an element of public risk.

	Risk Magnitude per GWa						
	Tidal[1]	On-Shore[2] Wind	Off-shore[3] Wind	Nuclear[4]	Coal[8]	Oil[10]	Gas[13]
Fuel Extraction	na	na	na	0.005-0.2[5]	0.25-0.7[9]	0.07-0.11[11]	0.07-0.11[11]
Fuel preparation/ reprocessing	na	na	na	0-0.2[6]	-	0.008-0.08[12]	-
Materials/component fabrication	0.007-0.03	0.08-0.14	0.09-0.15	0.002-0.005	0.005-0.02	0.001-0.004	0.0005-0.003
Plant construction	0.1-0.2	0.2-0.3	0.5-1	0.1-0.2[7]	0.1-0.2[7]	0.14-0.24[7]	0.07-0.11[7]
Power plant operation	~0.03	~ 0.1	~ 0.1	0.02-0.2[7]	0.01-0.2[7]	0.01-0.2[7]	0.01-0.2[7]
Transport	~0.01	~0.03	~0.05	~0.01	0.002-0.06	-	-
Decommissioning	nk	0.001-0.002	0.001-0.002	-	-	-	-
Waste Disposal	-	-	-	-	-	-	-

[1] Based on a hypothetical 16km 8,600MW scheme for the River Severn with a projected 233 GWa output over 120 years.
[2] Based on 450 hypothetical 300KW turbines with a load factor of 30% over an assumed 25 year lifetime.
[3] Based on 22 hypothetical 22MW turbines with a load factor of 26% over an assumed 30 year lifetime.
[4] PWR requirements are assumed.
[5] Uranium ore mining risks are dependent on ore grade and mine type.
[6] To date, no acute fatalities have been recorded at Springfields or Capenhurst.
[7] Based on actuarial data for the construction and operation of nuclear and non-nuclear stations.
[8] Assumes a coal requirement of 3.54M tonnes per GWa.
[9] Based on a study of accident rates in mines operated by British Coal.
[10] Assumes an off-shore oil requirement of 2.04M tonnes per GWa.
[11] Based on UK continental shelf actuarial data. Oil and gas related fatalities are not easily separated and have been taken as equal.
[12] For refineries.
[13] Assumes a gas requirement of 1.7 Bm^3 per GWa.
(na = not applicable; nk = not known; - signifies not calculated but generally of minor or no significance on a per GWa basis)

Table 2. Summary of estimated chronic risks of occupational disease per GWa by stage of electricity generation cycle. Data sources, limitations and uncertainties are described elsewhere.

	Risk Magnitude per GWa						
	Tidal	On-Shore Wind	Off-shore Wind	Nuclear	Coal	Oil	Gas
Fuel Extraction	na	na	na	0.01-0.08[1]	0.2-0.9[4]	-	-
Fuel preparation/ reprocessing	na	na	na	~ 0.04[2]	-	-	-
Power plant operation	-	-	-	~ 0.09[3]	-	-	-
Decommissioning	-	-	-	~ 0.01	-	-	-
Waste Disposal	-	-	-	~0.006	-	-	-

[1] These figures are based on radiation exposures at surface and sub-surface uranium mines in several countries. The figures indicate the associated fatal cancer risk which has been derived using the current ICRP risk factors.

[2] Based on radiation exposure data from Springfields and Capenhurst, but with some speculation about the new NOFC and THORP complexes, and the total collective occupational doses which might be apportioned to the PWR cycle.

[3] Based largely on the CEGB's collective dose design target of 2 man Sv per annum for the Hinkley Point 'C' reactor.

[4] Calculations indicate 90% attributable to pneumoconiosis and 10% to cancer.

(na = not applicable; - signifies not calculated but generally of minor or no significance on a per GWa basis)

Table 3. Summary of estimated hypothetical chronic public risks per GWa by stage of electricity generation cycle. It should be noted that major philosophical problems and technical uncertainties surround many of the data in this Table and care is essential in their use. Data sources, methods of calculation, limitations and uncertainties are described elsewhere[8].

	Risk Magnitude per GWa						
	Tidal	On-Shore Wind	Off-shore Wind	Nuclear	Coal	Oil	Gas
Fuel Extraction	na	na	na	0.09 (7.5)[1]	-	-	-
Fuel preparation/ reprocessing	na	na	na	0.02 (1.6)[1]	-	-	-
Power plant operation	-	-	-	0.02 (1.9)[1]	0-6[2]	0-6[2]	0-0.2[2]
Transport	-	-	~ 0.05[3]	-	0.002-0.06[3]	-	-
Decommissioning	-	-	-	0.1	-	-	-
Waste Disposal	-	-	-	0.001 (1.6)[1]	4, -2[4]	-	-

[1] Two cancer risk figures are given. The first figure applies to a collective dose integration period of 100 years (and/or current UK population), whereas the figures in parentheses are based on the accumulated radiation dose over the entire global population over 10,000 years or more. Clearly, the figures in parentheses are based upon the summation over billions of persons of minute increases in radiation levels which are tiny compared with the magnitude of each individuals radiation exposure from natural sources. The relevance of such estimates to resource allocative decisions is therefore open to question.

[2] Risks associated with gaseous and particulate discharges from fossil stations are notoriously difficult to assess and apportion, whether arising from chemical carcinogens, emissions of radioactivity, or of flue gases with other implications for the human respiratory system. This uncertainty is reflected by the broad range of risk factors assigned to this category. Even so, these estimates take account only of relatively short-term effects. Possible long-term effects, arising from, say, increased availability of heavy metals are not included.

[3] Includes an element of risk to the workforce. It should be noted that the risks in this category are of acute, i.e. prompt fatalities, unlike the others in the Table which are mainly delayed i.e. shorten life but do not terminate it immediately.

[4] Related to the use of fuel ash from coal burning in building materials. Again, this risk factor is enormously uncertain and may even be positive or negative depending upon assumptions.

(na = not applicable; - signifies not calculated but generally of minor or no significance on a per GWa basis)

EXTERNAL COSTS AS AN INDICATOR FOR COMPARATIVE RISK ASSESSMENT OF DIFFERENT ENERGY SYSTEMS - A CASE STUDY FOR COMPARING THE NUCLEAR AND THE COAL FUEL CYCLE[1]

W. Krewitt, P. Mayerhofer, R. Friedrich, A. Greßmann[2] and M. Dreicer[3]

[2]Institut für Energiewirtschaft und Rationelle Energieanwendung
Universität Stuttgart
Heßbrühlstr. 49a
D-70565 Stuttgart GERMANY

[3]Centre d'Etude sur l'Evaluation de la Protection dans le domaine Nucleaire
Route du Panorama, B.P. 48
F-92263 Fontenay-aux-Roses FRANCE

INTRODUCTION

The health and environmental impacts of different energy systems are emerging as significant issues in today's energy planning. In order to integrate the complex aspect of environmental impacts into the decision making process, scientists are expected to provide comparative information as a result of an interdisciplinary research effort. An appropriate framework is required that allows a consistent quantification of health and environmental risks of various energy systems as well as a direct comparison of the quantified physical impacts.

In contrast to former comparative risk assessment studies, we are facing a multitude of attributes describing health and environmental impacts that require new concepts of comparison. A promising one seems to be the concept of external costs which is well known in welfare economics. Following the external costs approach, monetary values are assigned to the physical health and environmental impacts that describe the welfare losses due to these damages. Well established methodologies exist that allow an economic valuation even of non-market commodities, so that the monetary unit can be used as an overall indicator for direct comparison of various effects. Besides of comparison, the external costs approach allows the direct integration of environmental issues into the

[1]This work was performed with the funding from CEC DGXII JOULE Programme, CEC DGXII Radiation Protection Programme, and Stiftung Energieforschung Baden-Württemberg

decision making process which certainly supports future decision towards an optimal allocation of scarce resources.

AN EXCURSION INTO WELFARE ECONOMICS - WHAT ARE EXTERNAL COSTS?

The Concept of Externality

Before using external costs as an indicator in the context of comparative risk assessment, a brief excursion into welfare economics seems to be helpful to understand the underlying concept of externality.

For any given energy service there will be certain externalities which are not reflected in the market price of that energy service. External costs arise when the production of energy causes welfare and utility losses to some third party and that losses are not compensated by market processes (Pearce et al., 1992). When externalities are present the allocation of resources in the economy will be inefficient. An optimal allocation of resources is achieved by internalizing externalities, i. e. third party losses of wellbeing are compensated for by the creator of the externality. Society's optimum level of environmental damage is derived from the analysis of marginal abatement costs and marginal damage costs. At the optimum level, further avoidance of damage is inefficient because the abatement costs exceed the benefits. Once external costs are internalized, decisions referring to the choice of different energy systems are automatically taken that are optimal not only for the utilities, but also for the whole society.

Measuring Externalities

Problems in assigning a monetary value to an environmental damage arise as soon as no market exists for the respective good. Using the example of human health impacts, some methodologies are briefly described that derive monetary values based on individual's preferences.

The Willingness-to-Pay Approach. Founded on the economic theory of social welfare , the Willingness-to-Pay (WTP) approach attempts to identify the monetary tradeoffs individuals are willing to make for a reduction in the risk of death and injury (Soby and Ball, 1991). The WTP approach operates by averaging individual marginal rates of substitution of money for survival probability to obtain an estimate of the value of a statistical life (VSL). An individual's WTP to reduce the risk of injury and death might be obtained by individuals' real choice as revealed by their behaviour in the marketplace, or by individuals' expressed preferences when directly asked about their WTP. Various WTP-methodologies are illustrated in Figure 1.

Revealed preferences. By examining the money-risk tradeoffs individuals make, the amounts they are willing to pay for safety are revealed. The analysis of wage differentials (acceptance of a job with greater risk of death in exchange for a wage premium), consumer behaviour (how much do people spend on equipment that reduces the risk of injury and death) and regulatory decisions (decisions made by public institutions related to safety measures) points to an individual's money-risk tradeoff.

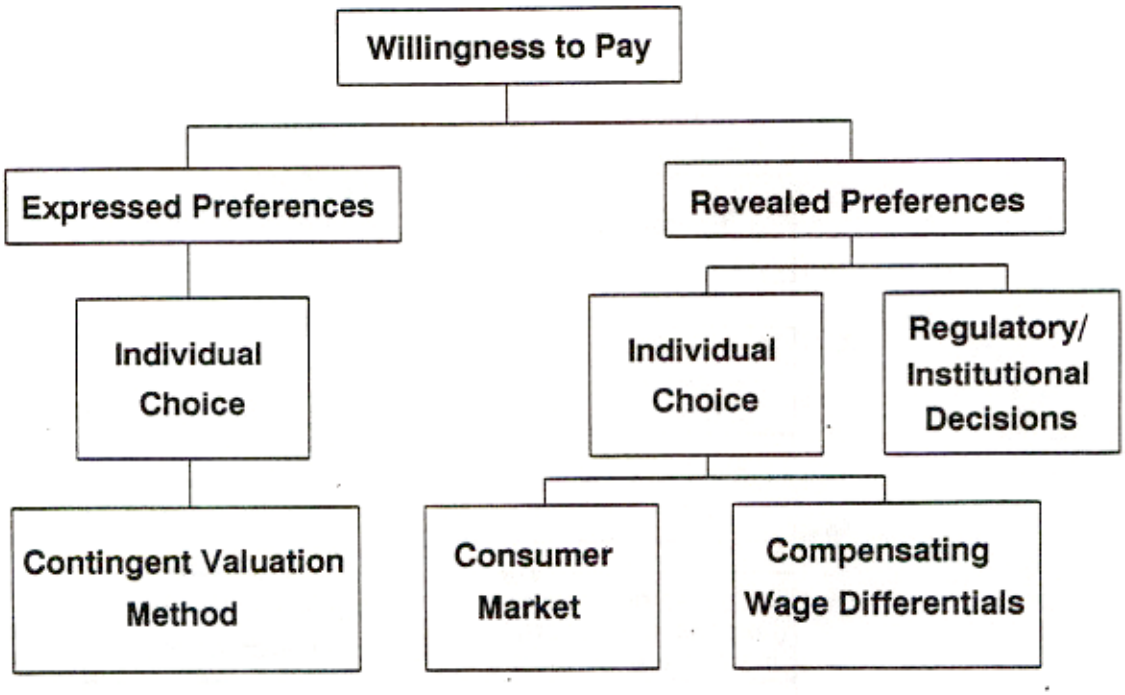


Figure 1. Willingness-to-Pay Methodologies (Soby and Ball, 1991)

Expressed Preferences. The Contingent Valuation Method (CVM) approach involves asking a sample of individuals what they would be willing to pay for varying levels of risk reduction. This approach requires people to explicitly express their preferences. In spite of some well known disadvantages (strategic behaviour and free rider problem, answers to hypothetical questions differ from actual behaviour, socially desirable answers etc.), the CVM approach is appropriate to reflect individual preferences in particular circumstances. CVM is a well established and accepted methodology in environmental economy. Monetary values used within the comparative case study described below are based on CVM studies.

As individuals in general attach less weight to a benefit or cost in the future than they do to a benefit or cost today, a discount rate might be used to take into account an individual's time preference. However, discounting human life is a heavily discussed issue. A social time preference rate of around 2-4% would be justified taking into account a sustainable rate of per capita growth and a rate of time preference, whereas rates of discount based on the opportunity cost of capital would lie at 5-7% for European countries (CEC, 1992). A 0% discount rate would give the same weight to all future generations' lives. Within the present case study, a range of discount rates including 0%, 3% and 10% is used to illustrate the dependence of different types of damage to this parameter.

THE CASE STUDY - RISK ASSESSMENT OF THE COAL AND NUCLEAR FUEL CYCLE

The impact pathway approach

To estimate marginal damages caused by an additional power plant, a producer related damage function approach must be applied, in which the pathway of a burden/pollutant is described in detail from the release to the impact on a receptor. Such a chain of causal relationships is called an "impact pathway". A comprehensive assessment of environmental

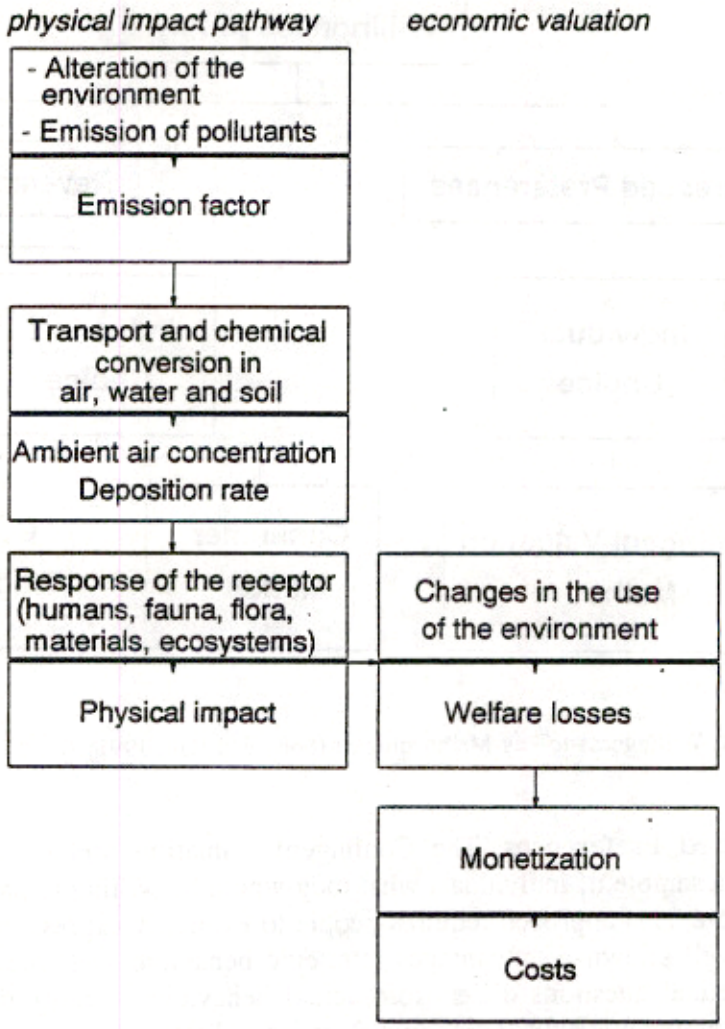


Figure 2. General structure of an impact pathway

impacts from an energy system requires the analysis of a multitude of different impact pathways, including various burdens, environmental media and receptors.

The impact pathway describes all processes and mechanisms from the origin of a pollutant or a burden to the physical impacts on a receptor. Based on the quantification of physical impacts, the resulting socio-economic impacts are quantified and described as monetary values (Figure 2).

In a first step, the environmental burdens resulting from the activity under analysis are quantified taking into account site specific data of the relevant facility. Emission data are an input to the second stage of the impact pathway that describes transport and chemical conversion processes of pollutants in air, water and soil. The transport calculation stage provides deposition rates or concentrations of a pollutant that cause an impact on a receptor. In general, humans, animals, plants, and materials are regarded as receptors, but in some cases it may be reasonable to look on a whole ecosystem as one receptor.

The quantification of physical impacts caused by a pollutant or a combination of pollutants is based on exposure-response functions describing the causal relationships between type and concentration of a pollutant and the resulting effect. It is worth mentioning that - although based on extensive scientific research efforts - the application of exposure-response functions is subject of major uncertainty within the pathway analysis. In the case that no reliable exposure-response functions are available, threshold values

related to effects that are expected as a consequence of having exceeded a threshold value might allow a first estimation of potential effects (concept of critical loads/levels).

The physical impacts derived by using biological and physical cause-effect models lead to socio-economic impacts on society. The objective of the pathway analysis is the monetization of the resulting socioeconomic impacts, so that the various environmental impacts can be quantified and compared, using the overall indicator of monetary values.

Taking into account the complexity of an energy system, it is obvious that not all of the possible impact pathways should be analyzed in detail. Hence, some priority pathways are identified that are assumed to cover major damages due to activities within the respective energy system. As impacts are expected to be site specific, a hypothetical 700 MW coal fired power plant with FGD and DENOX facilities located at Lauffen, north of Stuttgart, Germany, and a 900 MW pressurized water reactor at Tricastin, France, were selected as reference power plants. Site specific data on population distribution, land use, meteorology etc. were compiled and used as an input into the various stages of the pathway analysis.

EXTERNAL COSTS AS AN INDICATOR FOR COMPARATIVE RISK ASSESSMENT

The analysis of the priority pathways of the coal and the nuclear fuel cycle results in a considerable amount of health and environmental indicators that allow a reasonable quantification of many of the priority impacts of concern, although in fields like global warming the uncertainties are overwhelming. Following the concept of externality outlined above, the example of public health effects is used to demonstrate the use of external costs as an indicator for comparative risk assessment.

Adverse health effects within the general public due to electricity generation from coal or nuclear energy originate from an increase in ambient air pollution and an increment in radiation exposure respectively. The application of available exposure-response functions determine the set of health endpoints that is quantified. Table 1 shows the endpoints used for risk assessment in the present case study. The Value of a Statistical Life (VSL) used in this study is derived as the average value from a set of European CVM studies (CEC, 1993), it amounts to $2.6 \cdot 10^6$ ECU. As there is a lack of European literature on valuing non-fatal effects, monetary values on morbidity compiled in the US fuel cycle study (ORNL, 1992) are applied, although the transferability of monetary values from one context to another is known to be a critical issue.

Table 1 shows preliminary results from the analysis of adverse health effects within the general population due to an increment in air pollution level and radiation exposure at the two reference sites (not including severe accidents from nuclear power plants). Results are in no way comprehensive and therefore do not allow a comparison of the coal and nuclear fuel cycles at the current work status. While the relevant health impacts from the coal fuel system are acute effects which do not have to be discounted, impacts from the nuclear cycle include long term effects occurring years after the shutdown of the plant. It is obvious that the discount rate plays a dominant role in this case.

Results show that a monetary unit, expressing individuals' preferences even towards non-market commodities, is an appropriate indicator to normalize a set of very different physical endpoints to a common scale.

Table 1. Public health impact valuation (dr = discount rate)

Impact	Valuation [ECU]	Costs per kWh [mECU]		
		Coal[1]	Nuclear[2]	
Mortality	$2.6 \cdot 10^6$ (VSL)	0.85 - 1.6	dr=0%	$2.20 \cdot 10^{-3}$
			dr=3%	$0.39 \cdot 10^{-3}$
Non-fatal cancer	250 000		dr=0%	$0.51 \cdot 10^{-3}$
			dr=3%	$0.09 \cdot 10^{-3}$
Severe hereditary effects	$2.6 \cdot 10^6$		dr=0%	$0.44 \cdot 10^{-3}$
			dr=3%	$0.03 \cdot 10^{-3}$
Respiratory hospital admissions	6 600	$31 \cdot 10^{-3}$		
Emergency room visits	186	$2 \cdot 10^{-3}$		
Case of bronchitis	138.1	$2 \cdot 10^{-3}$		
Asthma attack	31.3	$340 \cdot 10^{-3}$		
Restricted activity day	62.4	$170 \cdot 10^{-3}$		
Symptom day	6.3	$580 \cdot 10^{-3}$		
Cough episode	5.0	$9 \cdot 10^{-3}$		
Chest discomfort day	6.3	$10 \cdot 10^{-3}$		

[1]) Electricity generation only, impacts within total FRG, discount rate = 0%
[2]) Electricity generation only, impacts within a circle of 100 km radius around the French reference site, not including severe accidents

REFERENCES

Pearce, D., Bann, C., Georgion, S., 1992: The Social Costs of Fuel Cycles, University College London, London

Soby, B. A., Ball, D. J., 1991: Consumer Safety and the Valuation of Life and Injury, School of Environmental Sciences, University of East Anglia, Norwich

Inhaber, H. 1979: Risk of Energy Production, 4th Edition. Ottawa: Atomic Energy Control Board

CEC 1993: Assessment of the External Costs of the Coal Fuel Cycle, ETSU, IER, Draft report prepared for DG XII of the Commission of the European Community

ORNL 1992: Damages and Benefits of the Coal Fuel Cycle: Estimation Methods, Impacts and Values, Draft report prepared for the U.S. Department of Energy, Oak Ridge National Laboratory, Tennessee

THE USE OF MULTI-ATTRIBUTE ANALYSIS FOR COMPARING THE HEALTH AND ENVIRONMENTAL IMPACTS OF DIFFERENT FUEL CYCLES - A FEASIBILITY STUDY FOR THE COMPARISON BETWEEN COAL AND NUCLEAR FUEL CYCLES[1]

M. Dreicer, P. Kuentzmann, T.Schneider,[2] and W. Krewitt[3]

[2]Centre d'étude sur l'Evaluation de la Protection dans le domaine Nucleaire
Route du Panorama, B.P.48
F-92263 Fontenay-aux-Roses FRANCE

[3]Institut für Energiewirtschaft und Rationelle Energieanwendung
Universitat Stuttgart
Hessbrühlstrasse 49
D-70565 Stuttgart GERMANY

INTRODUCTION

With the rising demand for energy and increased awareness of the importance of environmental protection, it has become necessary to focus the health and environmental impacts of energy systems. To make choices between systems is difficult. The indicators of risk for each fuel cycles can vary in type, magnitude and time scale. The public perception of acceptability varies greatly between different technologies for production of energy. By designing a consistent framework for the environmental and health risks associated with all fuel cycles and developing techniques for the comparison of complex and varied indicators, it is hoped that a link can be made between risk assessment and decision-making.

The differing scales of time and space that the impacts can potentially span, create difficulties in making direct comparisons. For example, the impacts from radioactive waste disposal are not likely to be seen for tens of thousands of years and will be generally local effects. The emission of greenhouse gases from fossil fuel cycles may have potential global impacts in the not too distant future.

[1] This work was performed with the funding from CEC DGXII Radiation Protection Programme, CEC DGXII Environment Programme, and EDF Mission Environment.

Occupational and public risks are generally considered differently and have different levels of importance in the decision making process. The integration of low probability of occurrence but high consequence events is also a difficult task, especially when the issues of the public's aversion to large impacts (no matter how low the probability) is considered.

In many comparative risk assessments risk indicators are transformed into monetary terms. By using a willingness-to-pay approach to quantification of the risks, some of these factors may be taken into account, however, some indicators are still very difficult to value in monetary terms. An additional drawback is that site-specific information is needed a monetary valuation method and therefore it is difficult to work interactively with these type of results to explore the outcome of different choices in a more global sense.

The risks resulting from the production of energy can be classified into environment and human health impacts. Past work has shown (Rostron et al. 1992) that the risks must be considered in terms of a time and space matrix. The time categories of short, medium and long can be designated depending on the impact being considered. In the case of nuclear power, the short-term can be considered to be 1 year, the medium-term would be an average persons lifetime and long-term would stretch into future generations. The context of space can be generally divided into local, regional and global. For the different populations (worker, public) and situations (normal, accidental) the impacts would fall into different categories of the time-space matrix.

Within the context of the comparison of energy systems, there has been a lot of work on the definition and quantification of common risk indicators that will allow for direct comparison between the fuel cycles being considered. For the purpose of testing the feasibility of using Multi-Attribute Utility Analysis (MAUA) to conduct a comparative analysis between coal and nuclear fuel cycles, we will consider only the indicators for impacts to human health.

MULTI-CRITERIA UTILITY ANALYSIS

The underlying principle behind a multi-criteria utility analysis is that a complex problem can be divided into a number of smaller decisions. The dimensions and attributes underlying each decision can be taken into account (Otway et al., 1991). The results of the smaller decisions are re-assembled to provide an overall decision on the larger picture. The attributes are defined as the factors that the decision-makers find relevant to the problem to be considered. MAUA provides a structure and process to be used by the decision-makers in considering these decisions. Its basis is the classic expected utility theory (Keeney and Raiffa, 1976). During this process, criteria are developed and structured according to priorities. Weights are assigned to the criteria and indicators are developed so that each criteria can be scored in a qualitative or quantitative way.

One of the largest difficulties in using this methodology will be the determination of the appropriate weights for the criteria. It is expected that these will be influenced greatly by the background and interests of the group making the decisions. In this feasibility study, nuclear and coal fuel cycle options will be compared using simple assumptions for the weighting of the criteria. The results were not intended to provide an actual comparison between nuclear and coal fuel cycles but to demonstrate the advantages and limitations of using MAUA. The comparison of two options was too simplistic an example, therefore, two other "options" were created by including additional information on the minimum risk for each fuel cycle. The impacts from Coal and Nuclear options were taken from the EC-US Project on the External Costs of Fuel Cycles (ETSU and IER, 1993, and Dreicer and Tort, 1993) and the minimum values were found in a literature review.

According to the main dimensions considered in this project, a set of criteria was defined (time/space, public/workers). Figure 1 illustrates the criteria and the risk indicators considered for the two fuel cycles and includes the simple assumptions for the corresponding weights that were assigned. It must be stressed that the weights were simply apportioned between the criteria and are not intended to represent reality.

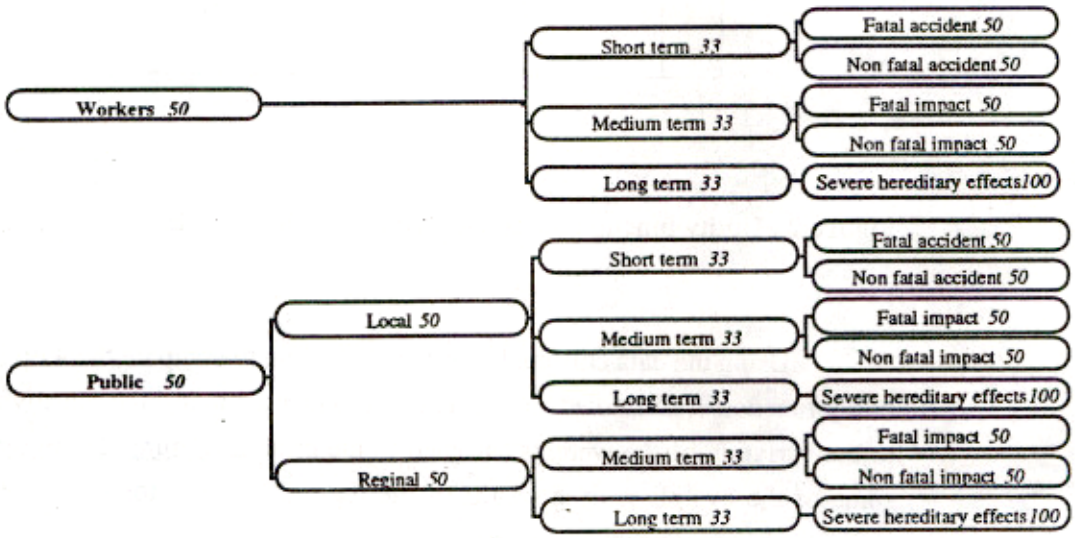


Figure 1. Weights assigned to the criteria used in the feasibility study.

For the four options designated (coal, nuclear, minimum coal, and minimum nuclear) it was not possible to directly compare the impact of a severe nuclear accident with the coal fuel cycle because an equivalent analysis does not exist for the coal fuel cycle. To integrate the impacts of a severe nuclear accident into the feasibility study, the results were weighted and added to the normal operation impacts for the nuclear options. It was decided that the publics' aversion to risk would be used in weighting the impacts of an accident. At this time, the author are not aware of any studies that have quantified a "risk aversion" factor but in general, risk aversion has been estimated to increase the impact by an

exponential factor of 1.2 to 2 for nuclear issues. For the feasibility study, a constant value of 1.2 was applied to the impacts from a severe nulcear accident and the impacts were added to the normal operation impacts.

In choosing a utility function for scoring the estimated risks (impacts) of the fuel cycle options, a different consideration for risk aversion was also required. It was considered that the populations' aversion to risk increases with the increasing magnitude of risk, regardless of the origin of the risk. With this in mind, an evolving risk aversion utility function was chosen. If the risk aversion factor is noted as a, then the shape of the utility function is defined by $0 < a(\text{minimum risk}) < 1$ and $a(\text{maximum risk}) > 1$. This is illustrated in the Figure 2 below.

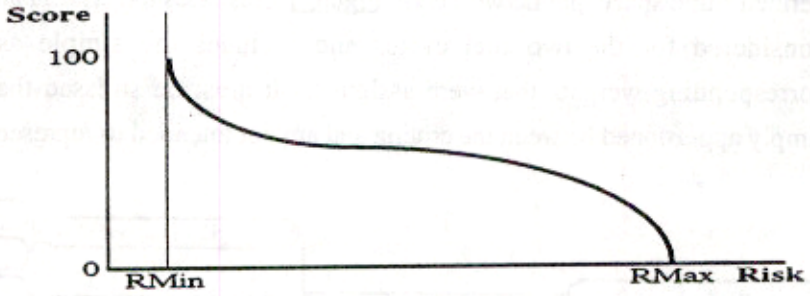
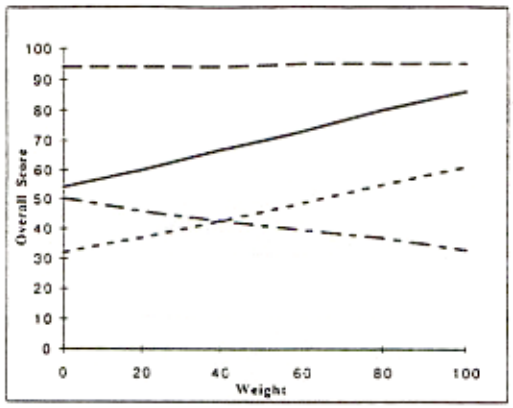


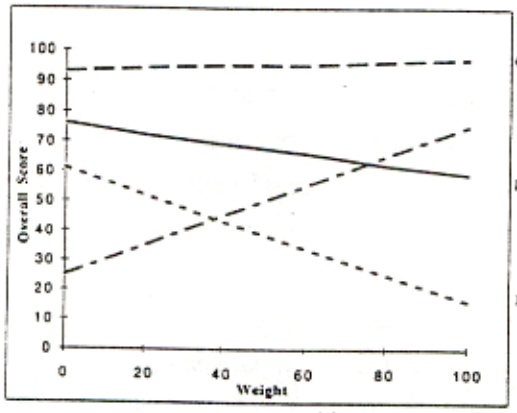
Figure 2 . Utility function for an increasing aversion to risk.

Using the data collected, a test case was run to compare the 4 options assuming a utility function as illustrated above. This was then used to test the sensitivity of the system to the variation of the weights assigned to the criteria (Figure 1). Figures 3 illustrates some of the overall scoring results of the feasibility study for the 4 options, given the simplistic assumptions that were taken.

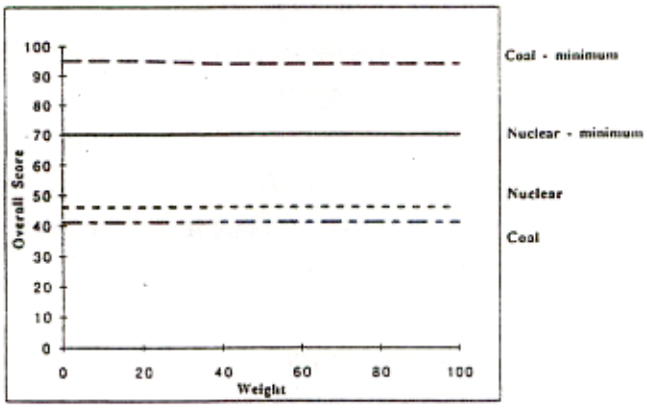
In Figure 3a at the place on the x-axis where the weights of the worker and public impacts are both 50%, the ranking of the options is Minimum Coal, Minimum Nuclear, Nuclear and Coal (from best to worst). If the weight is increased on the impacts of workers and therefore the weight of the impacts on the public is decreased, the order of the ranking does not change significantly, even thought the utility of Coal decreases. If the weight of worker impacts is decreased, the ranking between Coal and Nuclear change. Figure 3b illustrates the sensitivity of varying the weight given to the long-term health effects of workers in comparison to the weight given to the short- and medium-term effects. In this case, increasing the weight slightly changes the ranks of Coal and Nuclear. By contrast, Figure 3c shows how the weight given to the short-term fatal accidents for workers in comparison to short-term non-fatal accidents for workers has no impact on the final ranking of the options.



(a) workers/public



(b) workers:
long-term/low,medium-term



(c) short-term
worker accidents: fatal\non-fatal

Figure 3. Preliminary results of the sensitivity analysis of weights assigned to criteria.

The sensitivity of the resulting ranking of the 4 options to the risk aversion factor, for the impacts of a severe nuclear accident, was tested by varying the aversion factor, over the 1.2 to 2 range found in the literature review, while keeping rest of weighting factors unchanged. The Figure 4 illustrates how important the choice of values for this factor can have in the final outcome of the ranking.

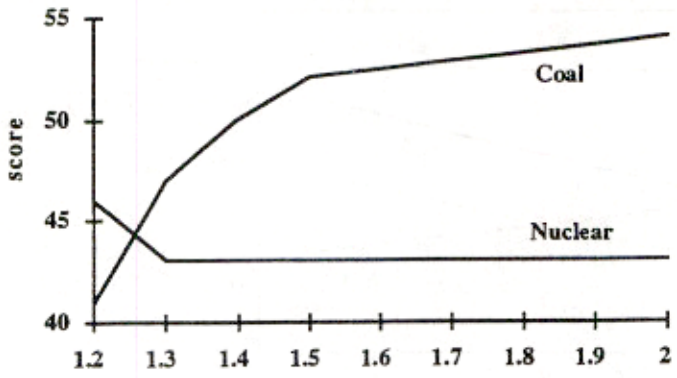


Figure 4. Impact of risk aversion factor on the score of the Coal and Nuclear Fuel Cycles.

From these preliminary results, it can be seen that the MAUA technique is useful. The sensitivity analysis has demonstrated the importance of continuing to work on determining reasonable weights for the criteria and reasonable values for the risk aversion factors. These values will be very specific to the population concerned with the ranking of the options.

The next step in the development a methodology for the comparative analysis of fuel cycles using MAUA, is to conduct a more realistic comparison of the coal and nuclear fuel cycle. It is not surprising that the feasibility study for MAUA techniques has shown the importance of the weights given to the different risk criteria. It is clear that these weights are very specific to the population involved in the choices to be made. Further work is planned for gathering information on weighting of the criteria, through further literature reviews, and the organisation of decision conferences for the investigation of the different weights that different population groups would assign to the criteria. The advantage of using MAUA, is that it is possible to work interactively with the software and the people that are making the choices.

REFERENCES

Dreicer, M. and Tort V., 1993 "Assessment of Physical Impacts and Monetary Valuation of the Priority Pathways - Nuclear Fuel Cycle," EC-US Joint Project, CEC DGXII, Brussels, Belgium.

ETSU and IER, 1993, "Assessment of the External Costs of the Coal Fuel Cycle," EC-US Joint Project, CEC DGXII, Brussels, Belgium.

Keeney, R.L. and Raiffa, H., 1976, "Decisions with Multiple Objectives: Preferences and Value Trade-offs", John Wiley, New York, USA.

Otway, H., Puckett, J.M., and von Winterfeldt, D., 1991, " The Prioritization of Environment, Safety, and Health Activities," LA-12175-MS/UC-907, Los Alamos National Laboratory, Los Alamos, USA.

Rostron,W., Schneider T., Thieme, M., and Dreicer M:, 1992, "Comparative Assessment and Management of Health and Environmental Impacts of Energy Systems," CEPN Report 203, Centre d'étude sur l'Evaluation de la Protection dans le domaine Nucleaire, Fontenay-aux-Roses, France.

068 Risk Assessment Failure Data Estimation

Chair: C. Fleming, PLG

Advanced Test Reactor Large LOCA Frequency Estimation
S.T. Khericha (INEL); S.A. Eide (LATA)

Component Failure Rate Comparison and Recommendations for Various Processing Fluids
L.C. Cadwallader, J.N. Wilkinson (INEL); S.A. Eide (LATA)

Application of Reliability-Centered-Maintenance to PRA Fault Tree Analysis for a BWR HPCI System
M.A. Feltus, Y.A. Choi (Penn. St. U)

ADVANCED TEST REACTOR LARGE LOCA FREQUENCY ESTIMATION

S. T. Khericha[1] and S. A. Eide[2]

[1]Idaho National Engineering Laboratory
P. O. Box 1625
Idaho Falls, ID 83415-3850

[2]Los Alamos Technical Associates, Incorporated
P. O. Box 51688
Idaho Falls, ID 83405-1688

INTRODUCTION

The Advanced Test Reactor (ATR) located at the Idaho National Engineering Laboratory (INEL) is one of several Department of Energy (DOE) test or production reactors. As part of the ongoing, comprehensive probabilistic risk assessment (PRA) for this reactor, a significant effort was expended to estimate the frequency of a large loss-of-coolant accident (LOCA). This effort was necessary because the ATR primary coolant system (PCS) differs significantly from commercial nuclear reactor PCSs. ATR differs in both its operating parameters and piping thickness. The PCS water temperature for ATR ranges from 125 to 175°F, compared with 500 to 600°F for a commercial plant. Also, the water pressure is approximately 350 psig, compared with 1000 or 2000 psig for commercial plants. Finally, the ATR PCS piping wall thicknesses range from 0.25 to 0.625 in., compared with 2.0 to 3.0 in. for commercial plants.

Estimating frequencies for large LOCAs is a difficult task. No such events have occurred in U. S. commercial nuclear power plants or DOE Class A (greater than 20 $MW_{thermal}$) reactors. Most PRAs of commercial nuclear power plants have used large LOCA mean frequencies in the range of 1.0E-3 to 1.0E-4/yr. The 1975 WASH-1400 estimate of 2.7E-4/yr (mean value) was based on a review of piping ruptures in nuclear and other industries and on judgement.[1] More recent estimates have been based on nuclear power plant history, resulting in a mean frequency of approximately 2.0E-4 to 5.0E-4/yr, depending on the estimation technique used. Several PRAs of commercial nuclear power plants attempted to refine the large LOCA estimate. For example, the German Risk Study of Biblis B estimated the large LOCA (break size of 30 $in.^2$ or larger) frequency to be less than 1.0E-7/yr.[2] This frequency was obtained based on fracture mechanics analyses, leak-before-break arguments, and use of piping materials and water conditions that minimize intergranular stress corrosion cracking (IGSCC) concerns. Also, probabilistic fracture mechanics analyses of U. S. nuclear power plants have resulted in large LOCA frequencies of less than 1.0E-8/yr.[3, 4, and 5]

METHODOLOGY

Because of the wide range of large LOCA frequencies used in PRAs and because of the unique ATR PCS features, a variety of methodologies was used to estimate the ATR large LOCA frequency. The results were then combined to obtain a suggested frequency and distribution. Three philosophically different approaches were included:

- statistical estimation based in historical data from other plants
- expert judgement elicitation
- phenomenological analysis of dominant failure mechanisms using the probabilistic fracture mechanics code PC-PRAISE.[6]

The statistical approach was further subdivided into three methods: integral (using data on entire PCSs for other plants), basic elemental (using leakage and rupture data for basic components from other plants), and Thomas elemental (using the Thomas method[7] for pipe breaks). All three approaches incorporated some ATR-specific information, with the integral statistical approach incorporating the least, and the PC-PRAISE analysis, the most.

Integral Statistical Approach

The integral statistical estimation methodology involved modeling the ATR PCS as a "black box", without identifying the various types and numbers of components (valves, pumps, feet of piping, etc.) within the system. The operating experience (no PCS ruptures) of commercial and DOE Class A reactors was used to estimate the large LOCA frequency for this approach. A Bayesian update of a noninformative prior was used to obtain a frequency from this experience.[8] The equation for the mean frequency, Mn, is the following:

$$Mn = (2n + 1)/(2T) \quad (1)$$

where n is the number of events in T system years of experience.

The PCS components for a commercial nuclear power plant typically are visually inspected for leakages before each start-up. This occurs several times each year. However, for the ATR this occurs approximately 25 times per year. This frequent inspection for leakages should reduce the potential for leakages to progress to ruptures before being detected and repaired. Therefore, the PCS rupture frequency obtained from Equation (1) was multiplied by 0.1 to account for the approximately factor-of-ten increase in the inspection frequency at ATR. Assuming that all U. S. commercial and DOE reactor experience is applicable, the large LOCA mean frequency from the integral statistical approach is 3.3E-5/yr.

Basic and Thomas Elemental Approaches

The two elemental statistical methods involved identifying the numbers and types of components in the ATR PCS, quantifying rupture frequencies for these components, and adding the individual component contributions to obtain the system (PCS) result. Two variations of this approach were used: one involving quantifying piping rupture frequencies using a piping rupture failure rate (expressed on a per foot basis) and the number of feet of piping, and the other using the Thomas method for quantifying piping

rupture frequencies. In both cases, leakage frequencies for components other than piping were obtained from data from commercial nuclear reactors and other sources, and then a rupture probability was applied.[9] Ruptures were defined as complete severance of piping or an external leakage rate of 50 gpm or greater. The large LOCA frequency estimate from the basic elemental method, with the 0.1 factor for increased inspections at ATR, is 7.4E-6/yr.

The Thomas elemental methodology differs from the basic elemental approach only in its calculation of the piping rupture frequency. Other components such as valves and flanges are treated identically. For piping ruptures, the Thomas method estimates a rupture frequency based on piping characteristics such as length, diameter, wall thickness, and number of welds. The piping rupture frequency, P_R, is calculated using the following equation:

$$P_R = X*C*(Q_P + A*Q_W)*F*Z*Y.$$

The parameters piping material quantifier, Q_P, and weld zone quantifier Q_W, are estimated using the equations as presented below:

$$Q_P = (D*L)/(t^2) \quad \text{and} \quad Q_W = (1.75N*D)/t.$$

The fraction of leakages (and ruptures) that are ruptures, X, was assigned the value 0.01. (In his paper, Thomas suggests a value of 0.05 to 0.1. However, based on the literature survey, 0.01 is more appropriate for PCS piping.) The global scaling parameter, C, was assigned a value of 1.0E-8/yr, as recommended by Thomas. For the weld zone penalty factor, A, the value 43.0 was used, as indicated by Thomas for U. S. nuclear reactor experience. The plant age factor, F, was assigned a value of 1.4, again from Thomas's recommendations. As discussed earlier, Z was assigned a value of 0.1 to account for an approximate order of magnitude increase in the ATR PCS walkdown inspection frequency compared with commercial nuclear power plants. The plant availability factor, Y, of 0.7 was used. The parameters D, L, t, and N represent pipe diameter, length, thickness, and total number of welds, respectively. The large LOCA frequency estimate using the Thomas methodology is 1.0E-5/yr.

Phenomenological Approach

Unlike the statistical methods that do not distinguish failure mechanisms, phenomenological analysis involves detailed modeling of one or more dominant failure mechanisms. Probabilistic fracture mechanics (PFM) is often used to model failure mechanisms such as crack growth in piping and vessels. The PFM computer codes simulate the growth of piping flaws by anticipated or accidental changes in stresses during the operation of the plant. When these flaws grow to some critical depth, they cause a pipe to leak or result in a catastrophic break.

The PC-PRAISE code was used for the ATR large LOCA analysis. PC-PRAISE is the PC version of the mainframe code PRAISE (Piping Reliability Analysis Including Seismic Event). The mainframe PRAISE code was initially developed in 1981 by Lawrence Livermore National Laboratory (LLNL) to assess the probability of pipe fracture due to fatigue crack growth in primary coolant loops of a nuclear power plant. The piping reliability model in PC-PRAISE covers two major failure mechanisms, fatigue failure and stress corrosion cracking failure. Various failure mechanisms such as vibratory stresses, residual stresses, seismic stresses, assembly stresses and operating stresses attributed to these two failure mechanisms are considered in the model. Initial interior surface flaws are assumed to exist along either the pipe circumferential direction or the longitudinal

direction. The PC-PRAISE code, however, is restricted to an exponential or a lognormal initial crack size distribution. Moreover, these distributions are restricted to the welds, since a failure is considered to be much more likely in a weld than in the piping base metal.

For this analysis, a representative highly-stressed weld was selected. It was conservatively assumed that the predicted failure probability for this weld was applicable for all other welds in the PCS. Fatigue crack growth in welds was modeled, using ATR-specific piping materials and dimensions, and thermal-hydraulic conditions and stress cycles. A large LOCA frequency of 1.0E-10/yr was estimated using the PC-PRAISE code for the phenomenological method.

Expert Judgement Approach

In the absence of historical or experimental data, reliance on defensible or informed opinion by experts for estimating the likelihood of rare events of potential risk significance is becoming a more widely acceptable practice in the nuclear industry. Methods for obtaining the expert opinions to estimate frequency of a rare event vary significantly. NUREG-1150 reports on the elicitation of expert opinion to obtain defensible estimates for likelihoods of rare events, such as a confinement failure, in the absence of historical or experimental data.[10] The nominal group technique was used to estimate a frequency of a large LOCA in the ATR PCS. This technique was selected based on considerations of available resources (i.e. availability and time of facilitator/analyst and experts), expert estimation theory, and past experience with other methods.

The ATR is a unique facility compared to a commercial power plant; therefore, it was thought prudent to select experts employed at the INEL with the academic and work experience relevant to the event to be judged. A list of potential candidates was developed based on their knowledge of the systems, maintenance, operations, materials and stress analysis of the ATR, or of PRA. From this list, ten experts were selected. Eight were selected on the basis that the person had some combination of expertise in the following fields: stress analysis of piping, operations, maintenance, plant aging, piping system integrity, fracture mechanics, or metallurgy specific to the ATR, in conjunction with data analysis, reliability engineering, human reliability analysis, reactor safety analysis, or PRA. The two other experts had the combination of expertise in the fields above but specific to a commercial power plant; however, they were also familiar with the ATR facility. The experts' relevant experience ranged from 5 to 33 years, with a mean of 20.5 years.

Due to scheduling constraints and cost considerations two sessions, approximately 2 hours each, were held on the same day. In the first session, as part of the problem definition phase of the elicitation session, presentations and discussions were conducted aimed at calibrating the experts regarding knowledge from the scientific literature about judgment bias, and the nature of the smallness of the frequency ranges upon which the judgments must be made. At the end of first session, the experts were asked to mark their judgement on the response scale sheet provided to them. The responses were presented on the board in the random order so that the graphical presentation on the board would not be in the same order as the experts' sitting order. This approach eliminated any initial association of specific estimate with the most vocal or aggressive experts, thus reducing the chances for initial personality or interpersonal style related bias on the judgement of the other experts. During the second session, the reasoning behind the first estimates were discussed. Following the discussion, the experts were instructed to make an independent, final judgment.

Expert judgments were analyzed and aggregated according to the guidelines presented in NUREG/CR-3688[11]. The following equation was used to aggregate the judgments:

$$\text{Aggregated Estimate} = \frac{\left(\prod_{i=1}^{m} EST_i\right)^{1/m}}{\left(\prod_{i=1}^{m} (1-EST_i)\right)^{1/m} + \left(\prod_{i=1}^{m} EST_i\right)^{1/m}}$$

where m is the number of experts and EST_i is the estimate of expert i.

The expert judgement elicitation process resulted in a large LOCA estimate of 7.4E-5/yr. The mean final-estimate went up 10% compared to the mean pre-estimate, suggesting an influence on the experts final estimates by the discussion of each expert pre-estimate rationales.

Initial data inspection suggested a trend of greater conservatism in the judgments of normative experts as compared to those of substantive experts. Also normative experts seemed to reduce in conservatism from their initial to their final judgment, and substantive experts seemed to increase in conservatism from their initial to their final judgment. Although these trends were not found to be statistically significant, they may be of interest as a topic for further investigation into the nature of expert judgment.

RESULTS

Each of the five large LOCA frequency estimation approaches resulted in a large LOCA frequency distribution, as discussed previously. Several different types of aggregation routines were considered to combine the five estimates. The five results were finally combined using a weighted geometric averaging scheme to obtain an overall large LOCA frequency distribution. The resultant mean is 5.0E-6/yr, with a lower bound of 2.0E-8/yr and an upper bound of 2.0E-5/yr.

CONCLUSIONS

The use of a variety of estimation techniques is appropriate for estimating the frequency of a large LOCA in the ATR PCS. Estimation of such a frequency is difficult and uncertain, and each estimation technique has its own strengths and weaknesses. The methodology employed was able to incorporate useful information from each technique.

ACKNOWLEDGEMENT

This work was supported by the U.S. Department of Energy, Assistant Secretary for Nuclear Energy, under DOE Idaho Field Office Contract No. DE-AC07-76ID01570.

REFERENCES

1. *Reactor Safety Study: An Assessment of Accident Risks in U. S. Commercial Nuclear Power Plants*, WASH-1400 (NUREG 75/014), U.S. Nuclear Regulatory Commission, Washington, D.C. (1975).

2. Gesellschaft fur Reaktorsicherheit (GRS) mbH, *Deutsche Risikostudie Kernkraftwerke, Phase B*, Gesellschaft fur Reaktorsicherheit (GRS) mbH, Karlsruhe, Germany (1989).

3. D. O. Harris et al., *Fracture Mechanics Models Developed for Piping Reliability Assessment in Light Water Reactors*, NUREG/CR-2301, U.S. Nuclear Regulatory Commission, Washington, D.C. (1982).

4. G. S. Holman et al., *Probability of Pipe Failure in the Reactor Coolant Loops of Combustion Engineering PWR Plants, Vol. 1: Summary Report*, NUREG/CR-3660, U.S. Nuclear Regulatory Commission, Washington, D.C. (1985).

5. G. S. Holman and C. K. Chou, *Probability of Pipe Failure in the Reactor Coolant Loops of Westinghouse PWR Plants, Vol. 1: Summary Report*, NUREG/CR-3660, U.S. Nuclear Regulatory Commission, Washington, D.C. (1985).

6. *PC-PRAISE Version 2.1, A Piping Reliability Software for Personal Computers*, Lawrence Livermore National Laboratory, Livermore, CA. National Laboratory.

7. H. M. Thomas, Pipe and Vessel Failure Probability, *Reliability Engineering*. 2:83-124 (1981).

8. *PRA Procedures Guide*, NUREG/CR-2300, U.S. Nuclear Regulatory Commission, Washington, D.C. (1983).

9. S. A. Eide et al., *Component External Leakage and Rupture Frequency Estimates*, EGG-SSRE-9639, Idaho National Engineering Laboratory, Idaho Falls, ID (1991).

10. T. A. Wheeler et al., *Analysis of Core Damage Frequency from Internal Events: Expert Judgement Elicitation*, NUREG/CR-4550, Vol. 2, U.S. Nuclear Regulatory Commission, Washington, D.C. (1989).

11. M. K. Comer et al., *Generating Human Reliability Estimates Using Expert Judgment*, NUREG/CR-3688, U. S. Nuclear Regulatory Commission, Washington, D. C. (1984)..

COMPONENT FAILURE RATE COMPARISON AND RECOMMENDATIONS FOR VARIOUS PROCESSING FLUIDS[a]

L. C. Cadwallader,[1] S. A. Eide,[2] and J. N. Wilkinson[3]

[1,3]Idaho National Engineering Laboratory
[1] P. O. Box 1625
[3] P. O. Box 4000
Idaho Falls, ID 83415

[2]Los Alamos Technical Associates, Inc.
P. O. Box 51688
Idaho Falls, ID 83405

INTRODUCTION

Risk assessments of facilities depend greatly upon accurate failure rates for the various components and failure modes included in the risk models. Failure rates often include both generic estimates and limited facility-specific failure data. Because of extensive data collection efforts for water systems in the commercial nuclear power plant field, a comprehensive set of generic failure rates can be assembled for risk studies of such plants. Recent examples include extensive reliance upon actual failure data from such plants.[1,2] For systems dealing with other types of fluids - liquid metals, chemicals, and compressed gases - much less actual failure data exist. In such cases, the risk analyst must decide whether existing generic failure rate estimates are appropriate for the system and fluid of concern.

In this paper, component failure data were compared for four processing fluids: water, liquid sodium, mild chemical, and compressed gas. This comparison was based on a limited review of existing data sources for such fluids. Because the most data exist for water systems, a base case set of component failure rates was generated from such data. The base case included typical components such as valves, pumps (or compressors), piping, heat exchangers, filters, and tanks. Then the more limited data for the other fluids were compared with the water base case. From this comparison, recommendations were made concerning the use of the base case water results for the other three types of fluids.

a. This work was supported by the U.S. Department of Energy, Director of Energy Research, Office of Fusion Energy, under DOE contract number DE-AC07-76ID01570.

DATA COLLECTION

A significant amount of actual component failure data exists for water systems, mostly from commercial nuclear power plants. References 1 and 2 contain a summary of many of the data that were available up through 1990. Aggregation of all of the data in those reports resulted in the component failure rates listed in Table 1 for water systems. Reference 2 was used for the external leakage and rupture failure rates, and Reference 1 provided the other entries. The Reference 1 data were obtained from the Nuclear Computerized Library for Assessing Reactor Reliability (NUCLARR) data base.[3] Many of the failure rates in the table are backed by tens or hundreds of failures, so most estimates have strong bases.

The failure data for components processing liquid sodium were obtained from the Centralized Reliability Data Organization (CREDO) data base[4], as interpreted in Reference 1. The component failure rates are presented in Table 1. Although there are as many entries for the liquid sodium components as for the water components, the numbers of failures are more limited for the liquid sodium entries.

Failure rates for components processing mild chemicals are also presented in Table 1. Failure data for these estimates were obtained from the *Offshore Reliability Data Handbook*[5] and limited data collected at the Idaho Chemical Processing Plant at the Idaho National Engineering Laboratory.[6] Failure rate entries in Table 1 for chemical components are much more limited than those for liquid sodium or water.

Finally, failure data for compressed gas systems were obtained mainly from the *Nonelectronic Parts Reliability Data* data base[7], and various data collection efforts associated with tritium systems.[8, 9, and 10] Failure rate estimates for compressed gas system components are also shown in Table 1. Again, the number of entries is limited compared with components associated with water or liquid sodium. Error factors for individual failure rates were considered, but are not listed here. The error factors can be found in the source documents.

FAILURE RATE COMPARISON

Using the fission plant water system failure rate estimates as the base case, the other three processing fluids were compared with the base case. The comparisons were made by dividing the estimate in question with the base case value for each component failure mode where both cases had failure rate estimates. The resultant quotients were then combined using a geometric average. A geometric average near 1.0 indicates that both sets of failure rates are similar, in general.

The component failure modes were divided into three categories:

Category 1. Fail to open/close, spurious operation, fail open/closed, fail to respond, fail to start, fail to run, overspeed, and fail to stop

Category 2. Plugged and internal leakage

Category 3. External leakage.

Internal rupture and external rupture were not included in the comparison because almost no failures existed for these failure modes. This division of failure modes was made based on the assumption that each category might show different trends with the four fluids reviewed.

Table 1. Component failure rate comparison for four fluids.

Component/Failure Mode	Component Failure Rates[a] Processing Fluids: Water	Sodium	Chemical	Compressed Gas
Standby or Safety Valve				
Manual				
Fail to open/close	5.0E-4/d	3.0E-4/d	-	3.0E-3/d
Plugged	3.0E-9/h	3.0E-8/h	-	-
Internal leakage	-	5.0E-8/h	-	-
Internal rupture	5.0E-8/h	1.0E-8/h	-	-
External leakage	1.0E-8/h	3.0E-7/h	-	1.0E-7/h
External rupture	4.0E-10/h	1.0E-8/h	-	-
Motor-Operated				
Fail to open/close	3.0E-3/d	1.0E-3/d	-	5.0E-4/d
Spurious operation	5.0E-8/h	5.0E-7/h	-	-
Plugged	5.0E-9/h	3.0E-8/h	-	-
Internal leakage	-	5.0E-7/h	-	-
Internal rupture	1.0E-7/h	1.0E-8/h	-	-
External leakage	1.0E-8/h	5.0E-7/h	-	-
External rupture	4.0E-10/h	1.0E-8/h	-	-
Air-Operated				
Fail to open/close	1.0E-3/d	3.0E-3/d	1.0E-2/d	3.0E-3/d
Spurious operation	3.0E-6/h	3.0E-7/h	3.0E-8/h	1.0E-6/h
Plugged	3.0E-8/h	3.0E-8/h	3.0E-8/h	5.0E-7/h
Internal leakage	1.0E-6/h	1.0E-7/h	5.0E-7/h	3.0E-6/h
Internal rupture	-	1.0E-8/h	3.0E-8/h	5.0E-7/h
External leakage	1.0E-8/h	1.0E-6/h	5.0E-7/h	1.0E-6/h
External rupture	4.0E-10/h	1.0E-8/h	1.0E-8/h	1.0E-8/h
Solenoid-Operated				
Fail to open/close	5.0E-4/d	-	-	5.0E-4/d
Spurious operation	5.0E-7/h	-	-	5.0E-7/h
Plugged	3.0E-8/h	-	-	-
Internal leakage	1.0E-6/h	-	-	-
Internal rupture	-	-	-	-
External leakage	1.0E-8/h	-	-	5.0E-6/h
External rupture	4.0E-10/h	-	-	3.0E-7/h
Check				
Fail to open	5.0E-5/d	1.0E-4/d	-	1.0E-3/d
Fail to close	1.0E-3/d	1.0E-4/d	-	3.0E-3/d
Plugged	5.0E-9/h	3.0E-8/h	-	-
Internal leakage	3.0E-6/h	5.0E-7/h	-	1.0E-5/h
Internal rupture	-	1.0E-8/h	-	5.0E-7/h
External leakage	1.0E-8/h	5.0E-7/h	-	1.0E-5/h
External rupture	4.0E-10/h	1.0E-8/h	-	3.0E-7/h
Control Valve				
Air-Operated				
Fail open/closed	-	-	5.0E-6/h	3.0E-6/h
Fail to respond	-	-	-	3.0E-6/h
Plugged	-	-	-	-
External leakage	-	-	-	5.0E-6/h
External rupture	-	-	-	1.0E-7/h

Table 1. (Continued)

Component/Failure Mode	Component Failure Rates Processing Fluids: Water	Sodium	Chemical	Compressed Gas
Pump or Compressor				
Motor-Driven				
Fail to start	3.0E-3/d	5.0E-3/d	1.0E-2/d	3.0E-5/d
Fail to run	3.0E-5/h	5.0E-5/h	5.0E-4/h	1.0E-5/h
Overspeed	-	-	-	3.0E-6/h
Fail to stop	-	-	-	-
External leakage	3.0E-8/h	3.0E-6/h	1.0E-5/h	5.0E-7/h
External rupture	1.2E-9/h	1.0E-7/h	5.0E-8/h	-
Turbine-Driven				
Fail to start	3.0E-2/d	-	-	-
Fail to run	1.0E-4/h	-	-	-
Overspeed	-	-	-	-
Fail to stop	-	-	-	-
External leakage	3.0E-8/h	-	-	-
External rupture	1.2E-9/h	-	-	-
Diesel-Driven				
Fail to start	1.0E-2/d	-	-	-
Fail to run	5.0E-3/h	-	-	-
Overspeed	-	-	-	-
Fail to stop	-	-	-	-
External leakage	3.0E-8/h	-	-	-
External rupture	1.2E-9/h	-	-	-
Strainer or Filter				
Plug	5.0E-6/h	3.0E-6/h	3.0E-6/h	1.0E-5/h
Fail open	-	3.0E-6/h	5.0E-6/h	1.0E-5/h
Tank				
External leakage	1.0E-8/h	1.0E-6/h	3.0E-6/h	3.0E-7/h
External rupture	4.0E-10/h	1.0E-7/h	3.0E-7/h	5.0E-8/h
Heat Exchanger				
Tube				
Plugged	-	1.0E-6/h	-	-
Internal leakage	1.0E-6/h	1.0E-6/h	5.0E-6/h	5.0E-5/h
Internal rupture	-	-	-	-
Shell				
External leakage	1.0E-8/h	1.0E-6/h	-	-
External rupture	4.0E-10/h	-	-	-
Piping (> 0.5 inch diameter)				
External leakage	3.0E-9/ h-ft	3.0E-9/ h-ft	-	-
External rupture	1.2E-10/ h-ft	3.0E-10/ h-ft	-	-

a. All failure rates were obtained from actual failure data for the processing fluid indicated.

For the liquid sodium comparison, the category 1 geometric average was 0.9, the category 2 result was 1.6, and the category 3 average was 65. This indicates that all of the component failure modes except external leakage (and rupture) appear to have similar failure rates for both water and liquid sodium components. For the chemical system comparison, the geometric averages were 1.5, 1.1, and 119, respectively. Again, except for external leakage and rupture, the failure rates are similar for both chemical system components and water system components. Finally, the compressed gas system comparison results were 0.9, 4.3, and 79 for categories 1, 2, and 3, respectively.

A special case for external leakage and rupture involves piping. No data were identified for mild chemical and compressed gas system piping external leakage and rupture events. We can draw no conclusions about those failure modes for those processing fluids. However, the liquid sodium piping external leakage failure rate is similar to that for water piping. Therefore, as a first approximation, the water system piping results for the rupture failure mode may be applied to sodium piping.

FAILURE RATE RECOMMENDATIONS

Based on the failure rate comparison results presented in the previous section, the following recommendations can be made for component failure rates for liquid sodium, chemical, and compressed gas systems:

1. For valves, pumps, tanks, heat exchangers, and filters, failure rates for water system components can be used for liquid sodium, mild chemical, and compressed gas systems, except for external leakage and rupture failure modes.

2. For external leakage and ruptures (except for piping), the water system results should be multiplied by approximately 100, based on our category 3 findings.

3. For piping, the water system results may be used directly for liquid sodium, mild chemical, and compressed gas systems.

These suggestions are only approximate; if at all possible, actual data for the system in question should be collected.

It is not known why the external leakage and rupture failure rates for liquid sodium, mild chemical, and compressed gas system components (except piping) appear to be much higher than those for water systems. It is possible that part of this difference is due to reporting differences. For example, a very small liquid sodium leak may be considered safety-significant, while a comparable water leak may only be an operational inconvenience.

REFERENCES

1. S. A. Eide, S. V. Chmielewski, and T. D. Swantz, Generic Component Failure Rate Data Base for Light Water and Liquid Sodium Reactor PRAs, EGG-SSRE-8875, EG&G Idaho, Inc., February 1990.

2. S. A. Eide, et al., Component External Leakage and Rupture Frequency Estimates, EGG-SSRE-9639, EG&G Idaho, Inc., November 1991.

3. D. I. Gertman et al., Nuclear Computerized Library for Assessing Reactor Reliability (NUCLARR), NUREG/CR-4639, June 1989.

4. J. J. Manning et al., A Guide for Completing Input Data Forms for CREDO, A Centralized Reliability, Availability and Maintainability Analysis Center for Liquid-Metal-Cooled Reactor and Test Facility Components, ORNL/TM-9892, Oak Ridge National Laboratory, November 1986.

5. Offshore Reliability Data Handbook (OREDA), OREDA-84, first edition, published by BP Petroleum Development Ltd., Norway, 1984.

6. J. N. Wilkinson, J. A. Vail, J. R. Wilson, and B. J. Harwood, Idaho Chemical Processing Plant Failure Rate Data Base, WIN-330, Westinghouse Idaho Nuclear Company, Inc., October 1991.

7. M. J. Rossi, Nonelectronic Parts Reliability Data, NPRD-3, Reliability Analysis Center, Griffiss Air Force Base, New York, 1985.

8. L. C. Cadwallader and M. A. Stolpe Gavett, Tritium Waste Treatment System Component Failure Data Analysis From June 18, 1984 to December 31, 1989, EGG-FSP-8973, revision 1, EG&G Idaho, Inc., November 1990.

9. L. C. Cadwallader and D. P. Sanchez, Secondary Containment System Component Failure Data Analysis from 1984 to 1991, EGG-FSP-10323, EG&G Idaho, Inc., August 1992.

10. L. C. Cadwallader and G. L. Taylor, Experimental Tritium Cleanup System Availability Analysis from 1984 to 1992, EGG-FSP-10603, EG&G Idaho, Inc., May 1993.

APPLICATION OF RELIABILITY-CENTERED-MAINTENANCE TO PRA FAULT TREE ANALYSIS FOR A BWR HPCI SYSTEM

Madeline Anne Feltus and Young A. Choi

The Pennsylvania State University
Nuclear Engineering, 231 Sackett Building
University Park, PA 16802
(814) 865-1341

ABSTRACT

This paper describes the application of Reliability Centered Maintenance (RCM) methods into plant PRA and safety for a BWR HPCI system. RCM is a technique, which is system function-based, for improving a preventive maintenance (PM) program, which is applied on a component basis. PM programs are based on time-directed maintenance tasks, while RCM methods focus on component condition-directed maintenance tasks. Stroke time test data for motor-operated valves (MOV) from four Emergency Core Cooling Systems (ECCS) was gathered from Philadelphia Electric Company's Peach Bottom Atomic Power Station to: (a) determine if MOV stroke time testing (as required by Section XI of the ASME code) was useful as condition-monitoring data; (b) determine and compare the plant-specific MOV failure data from a broad RCM philosophy time period vs. PM period, and also vs. generic industry MOV failure data; and (c) determine the effects and impact of the plant-specific MOV failure data in the PRA on core damage frequency.

The analyses showed that MOV stroke time testing was not a predictor for eminent failure, and should be considered as a go/no-go test. The preliminary failure data from the broad RCM philosophy showed an improvement vs. the PM-period failure rates in the ECCS MOVs. Also, the preliminary plant-specific MOV failure rates for both maintenance philosophies were shown to be lower than the generic industry estimates.

A significant decrease in core damage frequency (CDF) was found when comparing the plant-specific data vs. generic data. Only a small difference in CDF was found between plant-specific RCM data vs. PM data, which can be attributed to the few number of MOVs modeled in the HPCI PRA fault tree.

INTRODUCTION

Reliability centered maintenance (RCM) is a technique, which is system function-based, for improving a preventive maintenance (PM) program, which is applied on a component basis. Those components that truly impact plant function are identified and

maintenance tasks are focused on preventing their failures. The RCM evaluation establishes the relevant criteria that preserves system function so that a RCM-focused approach can be flexible and dynamic.[1] PM tasks are essential to nuclear plants, but are also very costly. Many past decisions on the type and frequency of PM tasks were based simply on vendor recommendations or engineering judgement with little regard to the component's history or function within the system. RCM is a systematic method of identifying those components that truly impact system functioning. Component failure histories and maintenance personnel interviews are conducted to determine proper RCM-based component tests and frequencies. Proper diagnostic/surveillance tests should be sensitive to component degradations, as RCM emphasizes condition-directed maintenance tasks vs. time-directed maintenance tasks. The RCM approach may provide benefits of greater system reliability with equal or reduced maintenance costs.

This study demonstrates how motor-operated valve (MOV) stroke time data information can be applied to a RCM approach, plant safety analyses, and PRA techniques.[2,3,4] The impact of using a RCM vs. PM approach can be demonstrated in the cascading effects in the PRA fault tree, and a Bayesian update[5] of the performance data.

This research effort investigated the impact of motor-operated valves (MOVs) in the Emergency Core Cooling Systems (ECCS) using a RCM approach. MOV stroke time testing has been conducted for many years as a PM task, as required by ASME Code Section XI. Quarterly MOV stroke time testing ensures operability of those MOVs which must function during an emergency event or cold shutdown.[6] This relatively frequent testing has provided a multitude of MOV stroke time data. This study analyzed MOV stroke time data for two purposes: (1) to determine if stroke time testing could be used as condition-monitoring data to detect valve degradations and eminent valve failure; and (2) to generate plant-specific failure data for the PRA by providing a conservative estimate for the number valve demands.

METHOD OF ANALYSIS

This study first determined if MOV stroke time testing could be used as condition-monitoring data. Condition monitoring data can be analyzed using a variety techniques. The most common and often appropriate method is trending, which uses a graphical plot of the data vs. time.[7] Stroke time data was gathered for MOVs in four BWR ECCS: (1) High Pressure Coolant Injection (HPCI), (2) Reactor Core Isolation Cooling (RCIC), (3) Residual Heat Removal (RHR), and (4) Core Spray (CS) systems at the Philadelphia Electric Company Peach Bottom Atomic Power Station (PBAPS)*. Corresponding corrective and preventive maintenance histories were gathered from 1980-1992. The most common MOV failure mode was packing leaks. The amount of packing friction can influence the motor speed, and may possibly be seen in the MOV's stroke time. Packing failures would be only considered as a degraded function, since packing leaks do not affect the MOV's primary function of opening and closing. Only HPCI valves were sampled to determine any correlations of stroke time results with packing leaks.

Motor-operator failures were also sampled from HPCI, RCIC, RHR, and CS systems since motor-operator failures were fairly uncommon. The motor-operator failures were divided into three categories: (1) motor-electrical failures, (2) mechanical failures, and (3) control switch failures. The failure categories were then further classified by their specific failure-related causes.

---

*This effort was funded by the Philadelphia Electric Company (PECo) and the Electric Power Research Institute (EPRI), with the assistance of NUS Halliburton. The authors would like to specifically thank Steve Hess (PECo), Steve Nornhold (NUS), and David Worledge (EPRI) for their technical contributions.

The correlation between stroke time and failure data was analyzed by plotting MOV stroke time data versus time, and then overlaying the MOV failure and PM data on the same time scale. Data points between the last MOV repair and subsequent valve failure of the identical MOV were evaluated using linear regression analysis to establish if there were any trends. The regression analysis used the last 8 data points or approximately six months from the identification of the MOV failure, whichever was shortest. Additionally, regressions using only the last four data points were completed to identify any immediate trends.

The second purpose of the stroke time analysis was to generate plant-specific MOV reliability data. The PBAPS HPCI fault tree[8] includes MOV failures as primary events with their corresponding failure probabilities; however, these failure probabilities were developed using generic industry data. Corrective maintenance and stroke time data was gathered for the HPCI system from 1980-1992. MOV failures were determined from the corrective maintenance history. Failures to open or close, or conditions which would prevent the valve from operating (e.g. motor burnout) were considered as MOV demand failures. MOV stroke times which exceeded their maximum stroke time limit would not have been considered as a demand failure under PRA concerns. This situation, however, did not arise in the analysis. The number of demands was estimated conservatively by the number of stroke time tests recorded. The estimate was deemed conservative since the MOV may have been required to stroke outside of stroke time testing.

The PBAPS underwent a general philosophical change from the PM philosophy toward a broad RCM approach during their 1987-1989 outage. As a result, MOV diagnostic testing was introduced to provide thrust and motor current analyses. The stroke time data from 1980-1986 and 1989-1992 were categorized as PM and broad RCM time periods, respectively. The demands were also categorized as open and close demands. The plant-specific MOV failure rates were generated by a ratio of failures/number of demands.

Ideally, a Bayesian update procedure should be used to incorporate the straight plant-specific MOV failure data into the generic data, and into the reliability (i.e., availability) values for the PRA. The generic industry failure data for MOVs uses a lognormal distribution as the prior distribution. The lognormal distribution should be converted into an equivalent Beta or Gamma distribution.[4,5] The parameters of the Beta or Gamma distribution are updated with the addition of the "new" data to form a posterior distribution.[5] The updated failure rate is then computed from the posterior distribution. A Bayesian update is currently in process for the plant-specific MOV data generated in this study. As the amount of the plant-specific data increases, the posterior distribution would gradually shift toward the plant-specific distribution. The use of plant-specific only failure rates, however, can show the largest extent of the potential difference between the plant-specific-only and the generic failure rates.

RESULTS OF ANALYSIS

Results of the regression analysis showed very little or no correlation between MOV stroke times and degraded function of the valve. Similar non-correlation results were also found for motor-operator failures using the same technique as described earlier. Thirty-one packing failures and twenty-nine motor-operator failures were identified and analyzed. The HPCI MOVs, however, did not experience a motor-operator failure. Figure 1 shows a plot of non-correlated stroke time data for a RHR MOV, which was typical for all the MOV failures sampled. Figure 2 shows the regression results for the identical RHR MOV. These results were also typical of the other MOV failures. The R^2 value, coefficient of determination, measures the percentage of variability that can be accounted for by the linear regression model. A high R^2 value shows a high correlation between the two variables, and ranges from 0 to 1. Accepted or desirable R^2 values are usually greater than 0.85. Typical

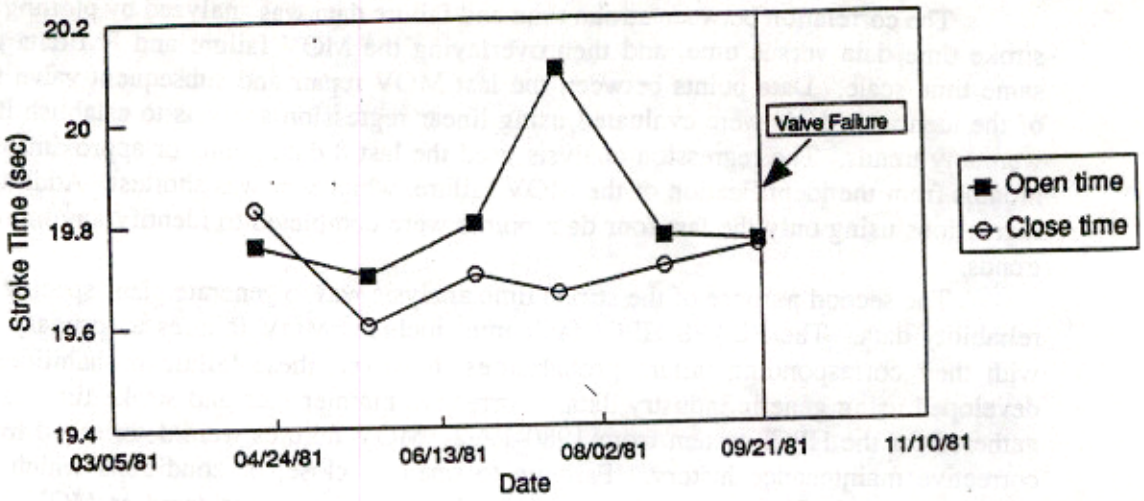


Figure 1. Typical Non-Correlated Stroke Time Data for Degraded and Motor-Operator MOV Failures.

Regression Output:

Constant	6.544264
Std Err of Y Est	0.157682
R Squared	**0.032079**
No. of Observations	6
Degrees of Freedom	4
X Coefficient(s)	0.000446
Std Err of Coef.	0.001224

Regression Output:

Constant	23.42007
Std Err of Y Est	0.087215
R Squared	**0.008404**
No. of Observations	6
Degrees of Freedom	4
X Coefficient(s)	-0.00012
Std Err of Coef.	0.000677

Regression Output:

Constant	70.02307
Std Err of Y Est	0.182697
R Squared	**0.157913**
No. of Observations	4
Degrees of Freedom	2
X Coefficient(s)	-0.00168
Std Err of Coef.	0.002748

Regression Output:

Constant	-1.10317
Std Err of Y Est	0.031194
R Squared	**0.525324**
No. of Observations	4
Degrees of Freedom	2
X Coefficient(s)	0.000698
Std Err of Coef.	0.000469

Figure 2. Motor-Operator Failure Regression Results.

Table 1. Plant-Specific Failure Rate and PRA CDF Effects.

	Failures/demands (open)	Failure rate	Decrease in CDF from Generic Data	Increase in HPCI Availability
PM era	7/2119	3.29 E-03	5.0 %	11.2 %
Broad RCM era	1/993	1.01 E-03	6.1 %	14.2 %

non-correlated regression results had very low R^2 values. Of the few regression analyses which produced an acceptable R^2 value, the actual stroke time deviation range was within stroke time measurement error. MOV stroke time data can only be used as condition-monitoring data if a definite correlation of stroke time changes to valve degradation exists. This analysis shows that stroke time data cannot be used for such purposes, and may be considered only as a "go/no-go" test. The plant-specific failure data is shown in Table 1. The generic industry MOV demand failure rate is 0.0118. Although a Bayesian update procedure has not been completed, preliminary results show an improvement in MOV failure rates since using a RCM approach. Also, the use of overall plant-specific data was an improvement to using generic data.

The effect of the plant-specific failure data on the core damage frequency (CDF) and HPCI subsystem unavailability is also shown in Table 1. The results from using the plant-specific data show a decrease in CDF compared to using generic data. Both plant-specific PM and RCM data yield higher availability. The RCM data yields even a slightly lower CDF than the PM plant-specific results.

The PRA, however, only considers 3 out of 13 HPCI MOVs (open demands only) as critical for HPCI to provide flow to the reactor core during severe accident scenarios. This observation may explain the relatively low difference in CDF observed during the PM vs. RCM failure data update. The Level 1 PRA is only concerned with providing adequate core cooling. Although the PRA may not model isolation and throttling purposes for maintaining HPCI functioning, the low number of valve demands modeled implies that certain MOVs have varying degrees of importance to preserve system functions. Although the plant-specific only failure rates are an order of magnitude different from the generic data (i.e., PM, 0.00329, RCM, 0.00101, and generic, 0.0118) this is not directly translated into an order of magnitude change in HPCI availability or CDF. This observation is because the MOV demands are only a fraction of the basic events in the HPCI fault tree; moreover, HPCI is only one system that is required during severe accidents which lead to core damage scenarios.

CONCLUSIONS AND SUMMARY

This study shows that stroke time testing cannot be used as condition-monitoring data, and has essentially no predictive failure capability. Stroke time testing should be considered as a go/no-go test. Although MOV stroke time testing may not be adequate as a condition-monitoring test for diagnostic purposes, PM stroke time testing is still required by the ASME Code.

The plant-specific failure rates were lower than the estimated generic industry data. Also, preliminary results for MOV reliability show that the RCM period data was an improvement over the PM period data. This result may imply that intensive RCM programs could greatly improve plant system reliability.

The CDF was decreased when using plant-specific data vs. generic data in the PRA. The PRA should be used as a guide to determine the valves which are truly critical. A RCM approach would include those valves found critical in the PRA in addition to valves which preserve system functioning. The maintenance programs should reflect those different levels of importance for those MOVs. MOVs which are not found to be critical should be tested less often.

Specifically, this study shows that stroke timing may not provide much benefit to measure the condition of MOVs. A simple operability test may provide almost the same information as the currently performed stroke time PM test. It is highly unlikely, however, that stroke time testing for MOVs will be eliminated as a PM task. Stroke time testing for "non-critical" MOVs should be replaced by operability tests. It may also make sense to lengthen the duration between MOV operability tests.

REFERENCES

1. Cox, S. J., N.R.S. Tait, Reliability, Safety, and Risk Management, an Integrated Approach, New York: Butterworth-Heinemann, 1991.
2. Frankel, Ernst G., Systems Reliability and Risk Analysis, Boston: Kluwer Academic Publ., 1988.
3. Lakner, Armand A., Ronald T. Anderson, Reliability Engineering for Nuclear and Other High Technology Systems: A Practical Guide, New York: Elsevier Applied Science Publ., 1985.
4. Lewis, Elmer E., Introduction to Reliability Engineering, New York: Wiley and Sons, 1987.
5. Sander, P., and R. Badoux, Bayesian Methods in Reliability, Boston: Kluwer Academic Publ., 1991.
6. American Society of Mechanical Engineers, Boiler and Pressure Vessel Code, Section XI, 1986.
7. Nuclear Maintenance Applications Center (NMAC), Flude, J.W., Nicholas, J. R., Predictive Maintenance Primer NP-7205, Palo Alto: Electric Power Research Institute, 1991.
8. Kolaczkoski, A.M., et al., Analysis of Core Damage Frequency: Peach Bottom, Unit 2 Internal Events, Sandia National Laboratories, U.S. Nuclear Regulatory Commission, NUREG/CR-4550, SAND86-2084, Vol. 4, Rev. 1 Parts 1 and 2, August 1989.

069 Waste Risk Management

Chair: B. Thompson, H.M. Insp. Pollution

The HMIP Research and Development Programme for Post-Closure Risk Assessment
B. Thompson (HMIP, UK)

Risk Analysis In Regulation and Risk Communication
S.M. Stearn (HMIP, UK)

Validation: What Should a Regulator Look for in a Safety Case?
R.D. Wilmot, D.A. Galson (Galson Sciences, England)

Aspects of Review of a Proponent's Post-Closure Safety Assessment on Behalf of the Regulator
T.J. Sumerling (Safety Assessment Mgt.); D. Read (WS Atkins Sci. & Technol.)

THE HMIP RESEARCH AND DEVELOPMENT PROGRAMME FOR POST-CLOSURE RISK ASSESSMENT

B G J Thompson

Her Majesty's Inspectorate of Pollution (HMIP)
London
United Kingdom.

ABSTRACT

The programme of work, undertaken by HMIP since about 1982, to develop, test and apply probabilistic system assessment (PSA) methods for the assessment of risks from underground disposal of radioactive waste in the United Kingdom is outlined. Attention is drawn to some key results from trial assessments and to the supporting research. Issues confronting client managers of such multidisciplinary contractor based projects are also highlighted. Current interests of regulators are raised and will be followed up in detail by subsequent papers in this conference.

INTRODUCTION

Her Majesty's Inspectorate of Pollution (HMIP) of the United Kingdom Government Department of the Environment (UK DoE) together with the Ministry of Agriculture, Fisheries and Food (MAFF) are the primary bodies responsible for authorising disposals of solid radioactive waste on land in the United Kingdom. In order to support this role, HMIP/UK DoE has developed since the early 1980's, through a team of contractors, the capability to perform independent assessments of the radiological performance of such disposal systems during the post-closure period. Large scale uncertainties are an inherent feature of such long-term forecasting and so a probabilistic systems assessment (PSA) procedure centred around Monte Carlo simulation has been developed and tested extensively by means of a series of trial assessments, Dry Runs 1, 2 and 3, as summarised in Thompson and Sagar[1], for instance.

The proposals and the safety documentation from a developer will be examined against a primary regulatory target, applicable to a single repository at any time, expressed as: "a risk to an individual in a year equivalent to that associated with a dqse of 0.1 mSv: about 1 chance in a million"[2]. (HMIP are currently reviewing this reference to provide revised guidance bearing in mind recent publications by the International Commission for Radiological Protection (ICRP) and others).

Currently, the principal concern of HMIP is with the deep repository, for low and intermediate level waste, proposed by UK Nirex for Sellafield. A planning application is likely to lead to a Public Inquiry in about 1997/98 at which HMIP may be expected to express a provisional view as to the likelihood of granting an eventual authorisation to dispose of active materials in the light of more extensive information about the site following a further period of

site exploration including the operation of a rock characterisation facility (RCF) acting as an underground research laboratory[3].

THE HMIP RESEARCH PROGRAMME

The overall research programme reported, for instance, in HMIP annual reports 1989/1990 through to 1992/93, see references[4, 5, 6 and 7], consists of four main areas:

(i) site investigation and evaluation,
(ii) geotechnical and repository engineering,
(iii) radionuclide release and migration processes in both the subsurface geology and in the surface environment, and
(iv) overall-site specific systems analysis and risk assessment.

The experience gained by HMIP of managing such multidisciplinary projects is summarised in reference[8] and results in the organisation shown here in Figure 1, where it is seen that a small in-house assessment management group is supported by two main lead contractors to enable the examination of Nirex safety documentation from two principal standpoints:

(a) against the overall scientific and technical knowledge that exists using a panel of acknowledged experts and covering both 'hard' and 'soft' scientific disciplines ranging from mathematics, physics and chemistry, operational research and statistics, through environmental and the earth sciences eventually to sociological issues and risk perception, and
(b) using site-specific understanding of the application of information from Nirex site studies and their revolving repository design to examine critically their approach to assessing risks and accounting for uncertainties.

RISK ASSESSMENT

The growth and capability of the HMIP programme during the past decade is illustrated here by:

- the extent of the repository/environmental setting formally modelled, and
- the coverage of an integrated, iterative performance assessment procedure including the treatment of all types of uncertainty and bias that are inherent in long-term forecasting.

The differences between in-house trials and a real assessment of a proponent's programme and safety documentation are outlined and the issues of interest to regulators raised for more detailed discussion in companion papers at this conference.

System Representation

The overall system affecting repository performance is outlined in Figure 2 taken from Thompson and Sagar[1]. During the early years 1982-1986 of the UK DoE/HMIP programme, effort was concentrated upon the central model of repository release and subsequent radionuclide migration through the groundwater system to the surface (biosphere) environment and thence to man (see Box 5).

Following results reported for Dry Run 1 and 2, it was decided to develop a new generation of Monte Carlo simulation codes called VANDAL, see reference[9], in which the groundwater flow in three-dimensions could be integrated (Box 4) and which could evaluate the influence of time-varying boundary conditions (of recharge, topographic changes, sea level etc) due to largely climate-driven environmental changes. These boundary conditions were supplied by a second family of Monte Carlo programmes TIME2, TIME4[10], encompassing

Boxes 1 and 2. However, models accounting for subsurface geological processes (fault development (say)) as in Box 3, have not yet been formulated and integrated.

Those processes that cannot be represented by mechanistic models may be included on an empirical stochastic basis, as for instance in Box 8. Local human actions, traditionally described under uncertainty by Poisson probability models, can also be integrated as suggested here by Box 6, and indeed the importance of incorporating an account of pathway interactions is also shown by Thompson and Sagar[1]. Woo[11] proposes Markov processes as a superior representation, however, and this is currently under consideration within the HMIP programme.

"Scenarios" of the kind discussed in the business literature[12] provide further external constraints, such as the influence of national energy policy on the inventory to be disposed (see Box 7).

Between 1989 and 1991, the third trail called Dry Run 3, Sumerling[13], demonstrated the much wider HMIP capability and showed for the first time, that climate change effects may have a significant influence on risk, as Figure 3 indicates. This result is of course specific to the particular site location and the simplified initial treatment of environmental changes on the groundwater system. However, it remains true that for any site climate change must be considered.

Assessment Procedure

An outline of an integrated, iterative risk assessment procedure was established by 1985, and matured by means of experience gained during the three trials. Figure 4 also taken from Thompson and Sagar[1], summarises current thinking and indicates the coverage achieved during the trials.

Dry Run 3 and to the lesser extent the previous trials considered Boxes 4, 5, 6 and 8 in some detail, with rather less emphasis upon Boxes 1, 2 and 9. The formal treatment of some parameter probability density functions (see Box 3), appropriate to the simplified models within the central PSA was also attempted. With the growing availability of site data from the extensive geological exploration at Sellafield, it has become possible to carry out the earlier stages of assessment in much greater detail with the intention of establishing a three-dimensional representation of site geology and from which one or more hydrogeological and geochemical transport models can be constructed to explore the implications of prevailing uncertainty. The formal exercise in Dry Run 3 to elicit a conceptual model structure may now be repeated in greater detail and the implications of different approximations explored more quantitatively. Indeed as a companion paper by Grindrod and Watson[14] in this conference describes, a procedure is being developed to complete the audit of perceived bias and residual uncertainty throughout the assessment process (Box 7).

SOME ISSUES OF CURRENT INTEREST FOR REGULATORS

Sumerling and Read[15], elsewhere in this conference, discuss the possible practical difficulties of carrying out assessments on behalf of regulators.

Such assessments are considerably different from the trials as undertaken previously:

(a) because of the much greater emphasis on exploring the implications of the detailed information about the site, and
(b) because of the lesser degree of control over the programme, of work that can be undertaken by the regulator when 'tracking' the proponents activities.

A key question arises concerning the extent to which a site can or should be investigated (for example, by surface boreholes or underground laboratories) to provide adequate confidence in any radiological risk assessment over the long term. Gallegos and Bonano[16] and others have argued that this requires iteration between site studies and performance assessments and Mackay[17] has conducted on HMIP's behalf a theoretical study for the Dry Run 3

site of the possible influence of gathering more data upon the estimation of risk. Figure 5 shows that after three attempts the results showed no trend towards the pseudo-exact answer from a geostatistical representation of the geology and the hydrogeology for this particular site.

The entire question of what a regulator should expect from a proponent in order to achieve a satisfactory level of confidence in safety has recently been examined in a seminar on Validation issues and this will be described by Wilmot and Galson in a companion paper[18].

Finally, credibility in the people, process and products of any decision making activity of this nature requires an understanding of the different viewpoints of the various stakeholders involved. Stearn[19] also at this conference, discusses this fundamental problem in terms of the United Kingdom regulatory background and the likely public inquiry process to be used for the examination of Nirex proposals.

ACKNOWLEDGEMENTS

The author thanks HMIP for permission to publish this paper. The work may be used for the formulation of UK Government policy but does not at this stage constitute UK Government policy.

REFERENCES

1. B.G.J. Thompson and B. Sagar. The Development and application of integrated procedures for post-closure assessment, based upon Monte Carlo simulations: the probabilistic systems assessment (PSA) approach. Submitted for publication in the Journal of Reliability Engineering and System Safety (1993).
2. Department of the Environment. Disposal Facilities on Land for Low and Intermediate Level Radioactive Wastes: Principles for Protection of the Human Environment. HMSO, (ISBN 0-11-751775-5) (1984).
3. UK Nirex Ltd, Nirex Report No. 327. Sellafield Repository Project. A Rock Characterisation Facility Consultative Document (October 1992).
4. HMIP Third Annual Report 1989/1990. HMSO. ISBN 0-11-752435-2.
5. HMIP Fourth Annual Report 1990/1991. HMSO. ISBN 0-11-752519-7.
6. HMIP Annual Report 1991/1992. HMSO. ISBN 0-11-752671-1.
7. HMIP Annual Report 1992/1993. HMSO. ISBN 0-11-752832-3.
8. B.G.J. Thompson, M.W. Wakerley and T.J. Sumerling. Recent management experience of UK performance assessments of radioactive waste disposal, Proceedings of International High Level Radiocative Waste Management Conference, Las Vegas (May 1993).
9. P. Kane. VANDAL Version 1.3 Technical Overview, UK DoE Report, HMIP/RR/92.095 (1992).
10. Dames and Moore. Technical Reference Manual for TIME4 Version 1.0 (2 volumes), UK DoE Reports, HMIP/RR/91.049 and 91.050 (1991).
11. G. Woo. A Statistical model of possible future human actions. Submitted for publication in the Journal of Reliability Engineering and System Safety (1993).
12. P. Ayton and G. Wright. Judgemental Forecasting. Publisher Wiley (1987).
13. T.J. Sumerling (Ed). Dry Run 3: A Trial Assessment of Underground Disposal of Radioactive Wastes based upon Probabilistic Risk Analysis: (in 11 volumes) Overview, UK DoE Report, HMIP/RR/92.039 (1992).
14. P. Grindrod and S. Watson. How uncertain we should be about the outputs of a PSA: the HMIP procedure (ibid).
15. T.J. Sumerling and D. Read. Aspects of review of a proponent's post-closure safety assessment on behalf of a regulator (ibid).
16. D.P. Gallegos and E.J. Bonano. Consideration of uncertainty in the performance assessment of radioactive waste disposal from an international regulatory perspective. Submitted for publication in the Journal of Reliability Engineering and System Safety (1993).
17. R. Mackay. A Study of the Effect of the Extent of Site Investigation on the Estimation of Radiological Performance: Overview, UK DoE Report, HMIP/RR/93.053 (March 1993).
18. R.D. Wilmot and D.A. Galson. Validation, What should a regulator look for in the Safety Case? (ibid).
19. S. Stearn. Risk analysis in regulation and risk communication (ibid).

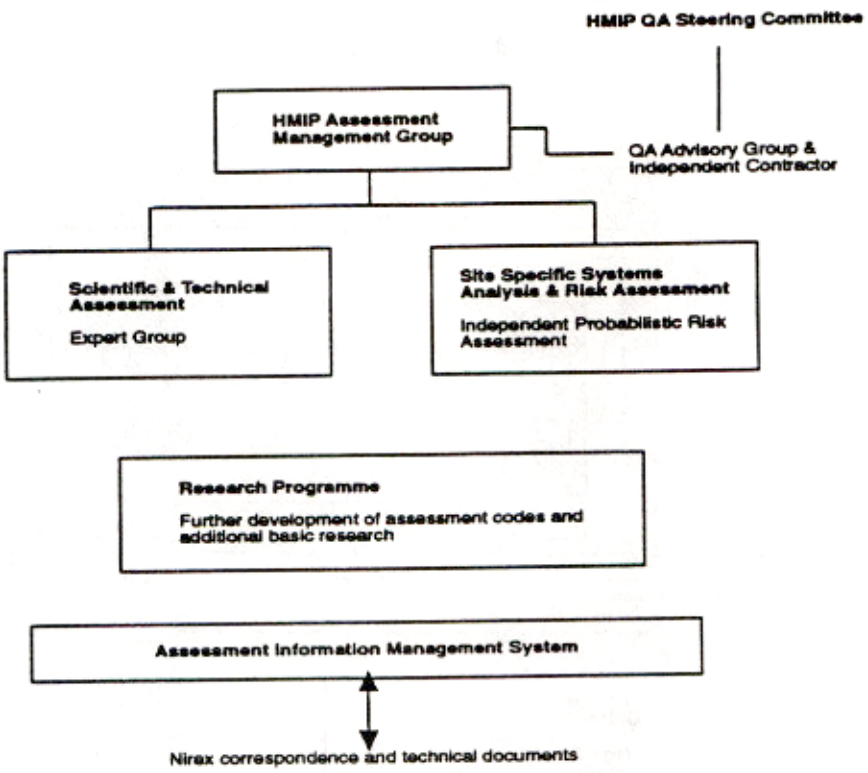


Figure 1. HMIP assessment programme structure.

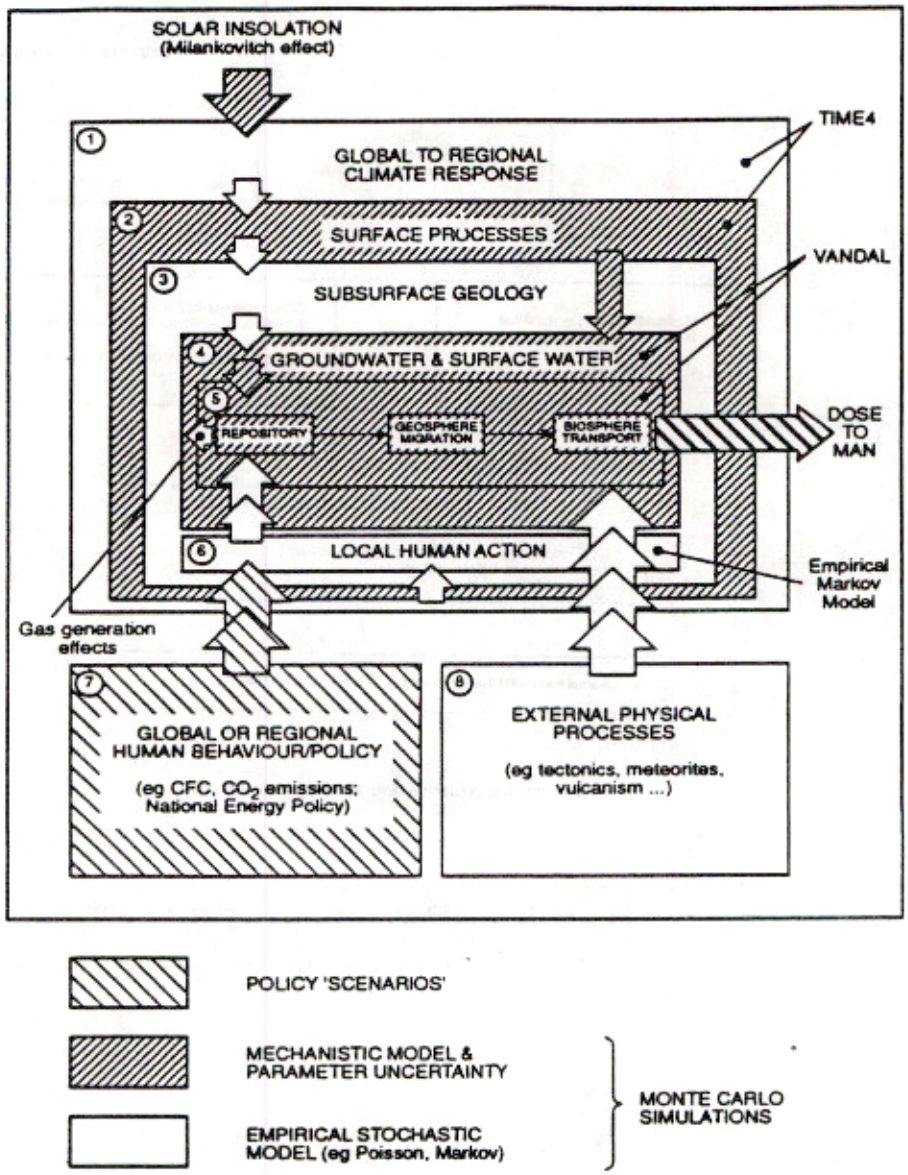


Figure 2. The overall system affecting repository performance.

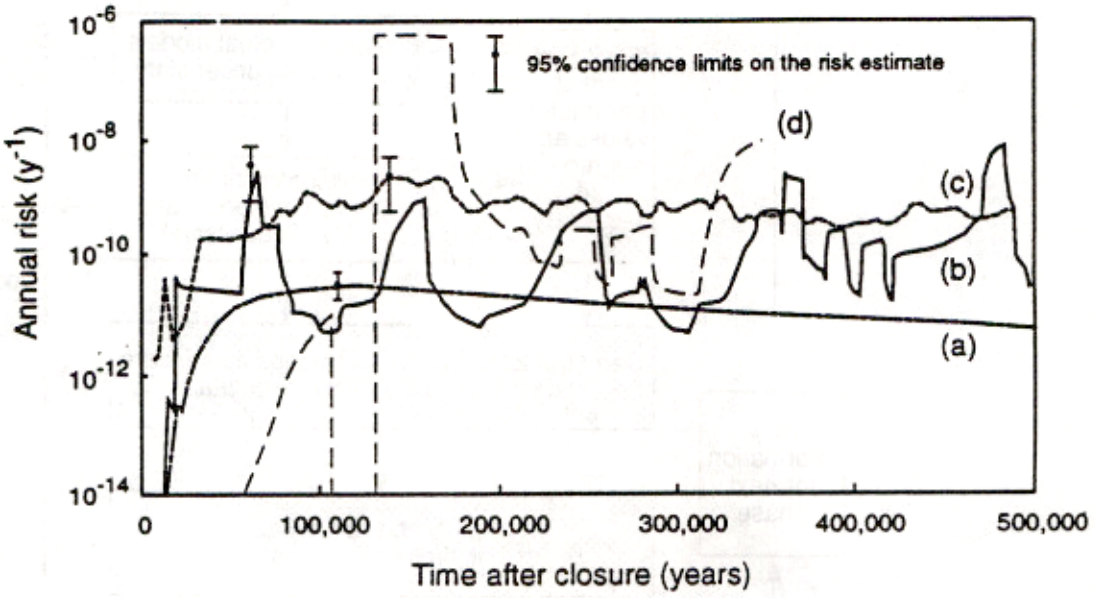


Figure 3. Comparison of risk estimates from Dry Run 3: (a) The VANDAL simulation for constant temperate conditions, (b) the TIME4-VANDAL simulation for the historical climate sequence, (c) the fully time dependent TIME4-VANDAL simulation and (d) the conditional risk from the highest dose run of (c).

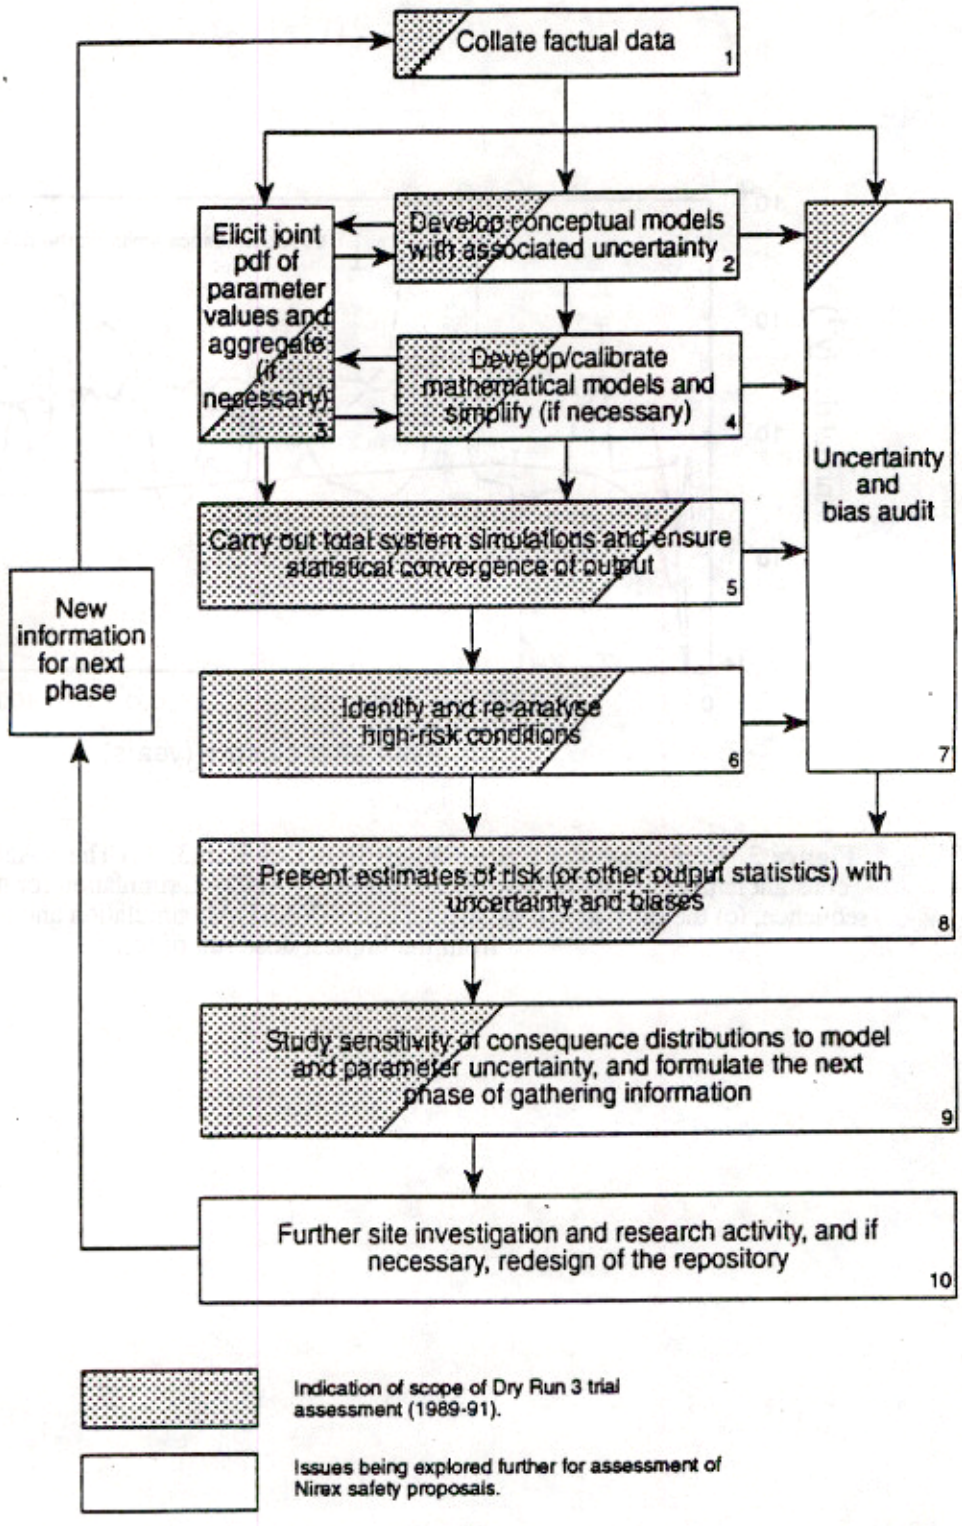
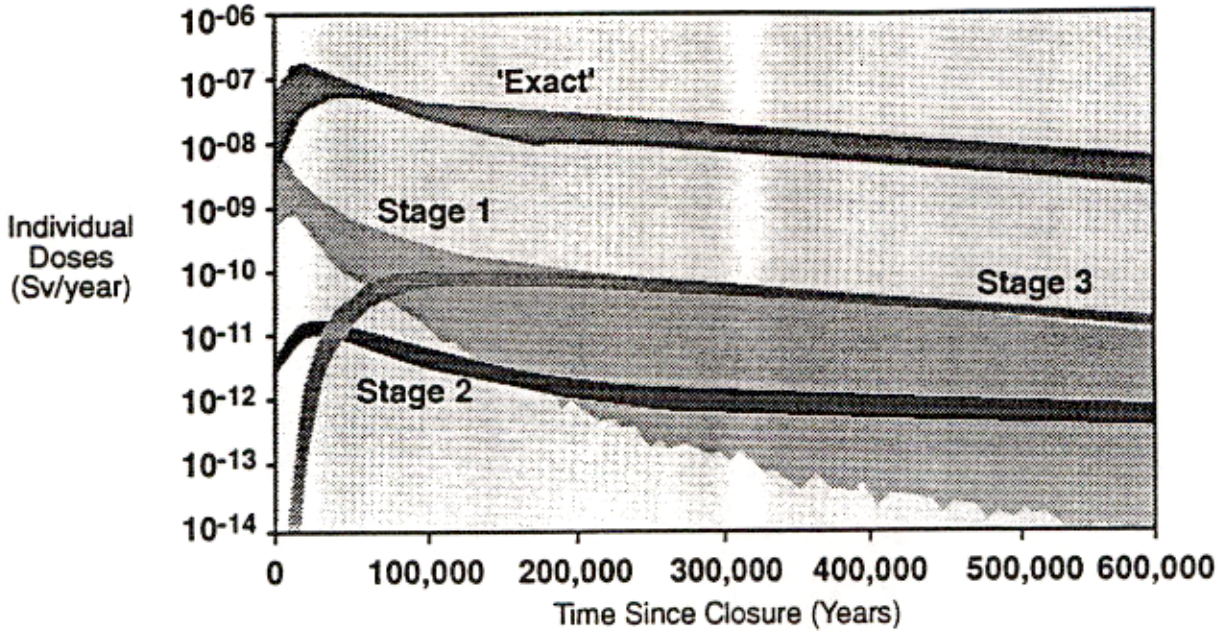


Figure 4. A general procedure for phased assessment.



Stage 1: 2 boreholes (as Dry Run 3); Stage 2: 5 additional boreholes over the repository foot print; Stage 3: 2 further regional boreholes.

Figure 5. The simulated effect of different amounts of borehole information on dose estimation (Mackay[17]).

RISK ANALYSIS IN REGULATION AND RISK COMMUNICATION

Dr S M Stearn

Her Majesty's Inspectorate of Pollution (HMIP)
London
United Kingdom

REGULATORY OBJECTIVES FOR THE COMMUNICATION OF RISK

Regulatory Requirements in the UK

Anyone who disposes of radioactive waste in the United Kingdom must hold an authorisation from the appropriate Authorising Department or Departments under the terms of the Radioactive Substances Act 1993 unless the disposal is trivial when it may be covered by an Exemption Order issued under the Act.

The Act itself sets out procedures but not prescriptive limits. These are matters to be determined by Government policy.

In brief the Act provides that an application is made, that it is publicly registered and that before issuing an authorisation the Authorising Departments, in the case of nuclear sites the Ministry of Agriculture, Fisheries and Food and Her Majesty's Inspectorate of Pollution (HMNII), will consult local and other public authorities in accordance with the requirement of the Act.

The stage of formal authorisation will not be reached until a facility has been constructed and is ready for operation. Such construction is regulated under the terms of the Nuclear Installations Act 1965 (as amended) and administrative arrangements between the Authorising Departments and the Nuclear Installations Inspectorate (NII) ensure that waste management considerations are taken fully into account at all stages in the design and construction of the facility.

The developer will also have to obtain permission under the Town and Country Planning Act 1971 and may, to obtain that permission, have to satisfy a Public Inquiry. Before a Public Inquiry is held into whether planning permission should be granted to construct such a facility the Authorising Departments will give their provisional view of whether the proposed facility would be suitable for authorisation. Consultation with local planning authorities will take place before that provisional view is given.

Once a disposal facility has been constructed, authorised waste producers consigning waste to that facility will have to apply to the Authorising Departments for separate authorisations for the disposal of waste from the site from which it originated. The conditions attached to these later authorisations will include the type of waste, the radionuclides contained in the waste, their radioactivity and their volumes and packaging. These conditions ensure that the movement of radioactive waste is fully regulated at all stages.

BASIC RADIOLOGICAL REQUIREMENTS

In deciding whether to give authorisations, and if so what conditions should apply, the Authorising Departments will have regard to the basic radiological requirements set out below. The National Radiological Protection Board is the statutory adviser to the Government on radiological protection, in particular by interpreting the recommendations of the International Commission on Radiological Protection (ICRP).

The Authorising Departments will require that proposals for a facility include quantitative estimates of the activity likely to be released from the radioactive waste and describe how the various components of the facility act to limit its radiological impact. They will also require that the information should be submitted in a form which permits them to undertake their own assessment of these factors should they wish. The safe disposal of solid wastes requires a system consisting of a number of barriers whereby site, waste-form and methods of emplacement all contribute to the effectiveness of containment. The Authorising Departments will need to be satisfied that adequate containment can be provided, either by the nature of the geological host rocks or by engineered safeguards or by a combination of the two.

In considering the radiological protection objectives, a distinction must be drawn between the successive stages in disposal. First there is the institutional management period which can be divided into the pre-closure period, likely to last several decades, in which disposal of waste is made into the facility and the post-closure period of institutional management, which is envisaged as lasting at most a few hundred years. In its later phase the facility has been closed but there is adequate assurance that specified forms of supervision will remain available. Finally there is the post-institutional management phase for which it must be established that on-site supervision is no longer required.

Broadly speaking existing standards for radiological protection of the public will be applied to land disposal facilities but two features require special consideration.

The first arises because radioactivity may be released from such a facility into the environment as a result of a range of situations with varying probabilities and leading to various radiation doses. The existing standards apply to situations in which it is reasonable to assume that doses will actually be received, but the application of these standards should be modified if the probability of a particular dose being received is low.

The other special feature in the application of existing standards arises from the long time period which may have to be

considered to assess the potential risk from disposal. There are, for example, uncertainties in predicting the exposures which could occur in the far future because population distributions and human habits may change in ways which cannot readily be foreseen. Recommendations to deal with this situation have been received from the IAEA and the Nuclear Energy Agency (NEA) of the OECD.

In the institutional management period steps must be taken to protect site workers and members of the public. During this period the responsibility for regulating the health and safety of workers will rest with the HMNII. The responsibility for assuring that the public is protected rests jointly with the HMNII and the Authorising Departments and in this connection they will wish to satisfy themselves that the objectives set out in the Annex to the 1982 White Paper (1) have been, and will be, achieved. These objectives are:

(a) All practices giving rise to radioactive waste must be justified.

(b) Radiation exposure of individuals and the collective dose to the population arising from radioactive waste shall be reduced to levels which are as low as reasonable achievable, economic and social factors being taken into account.

(c) The average effective dose equivalent from all sources excluding the actual background radiation and medical procedures to representative members of a critical group shall not exceed 5 millisieverts in any one year.

This latter limit is under review at the moment. As a result of the publication of ICRP 60 a dose constraint has been advised of 0.3 millisieverts for any member of the public from all exposure from any one source with an overall limit of 1 millisieverts in any one year from all sources. Authorising Departments are working to these limits now.

During post-institutional management the proper objective to apply is one based not on dose but on risk to the individual. That risk may be defined as the probability that a given dose will be received multiplied by the probability that such a dose will result in a fatal cancer.

Taking all the above considerations into account the Authorising Departments will be guided by the following objectives:

> "A risk to any member of the public in any one year from exposure to radiation from all sources other than background and medical exposure, should not be greater than that associated with a dose of 1 millisievert."

In order to take account of the exposure pathways and health effects not at present recognised and to avoid prejudicing any future decision that might lead to other activities which cause radiation exposure to the same members of the public, an additional safety margin is required. The appropriate target applicable to a single repository at any time is, therefore, a risk to an individual in a year equivalent to that associated

with a dose of 0.1 millisieverts, about one chance in a million. The risk to this dose will be substantially less than those arising from variations in natural background radiation within the United Kingdom.

During the operation of a facility, and thereafter, the developer will be required to demonstrate that he has done all he can to ensure that exposures from radiation are as low as reasonably achievable, taking into account economic and social factors. This should result in bringing doses down below target levels. During the post-institutional management periods, the Authorising Departments will require the operators to use the best practicable means to ensure that any radioactivity coming from a facility is at as low a level as reasonably practicable. The effect of this will be to ensure that exposures are as low as reasonably achievable. The Authorising Departments will also in conjunction with the waste producers apply the ALARA principles to decisions between overall management procedures relating to specific wastes and covering treatment of the wastes, method of immobilisation, and method of disposal. A fuller exposition of this issue is contained in (2).

USE OF PROBABILISTIC MEASURES

It is very tempting to produce a single number which hopefully demonstrates that dose limits will not be exceeded, thus ensuring that the proposal is acceptable to the public. But it is not realistic to proceed in such a way without exploring explicitly and honestly the uncertainties that are inherently present in any such analysis.

There are essentially four stages involved. In the first stage, hazard assessment, the nature of the hazards are explored and as much as possible must be found out about the disposal site in question.

The second stage is the estimation of risk which involves assessing exposures as carefully as possible, establishing probable inventories and the toxicity of the material involved. The next stage is the evaluation of risk in terms of the confidence that we have in our answers. The last stage is to manage and control the risk by the application of measures to reduce it to as low a level as is reasonable.

It is not possible to know everything about a disposal site but the operator needs to establish enough to be able to satisfy regulatory bodies that the proposal meets the conditions that are set out. In particular there needs to be a clear view about how individual uncertainties might affect the output of the system.

PUBLIC PERCEPTION

Any proposal for a nuclear waste repository is likely to cause considerable public interest. The public are aware that it is a complex issue that requires a good deal of care in its assessment. Therefore bland assurances are unlikely to be satisfactory.

Many members of the public are extremely well informed about these issues, and for this reason will want examine the data upon which the conclusions of regulators are based. They will perhaps

carry out their own analyses and will seek, certainly, both at the Public Inquiry stage and at other times to question the conclusions of regulating authorities.

In that context it is important that the regulatory system calls upon operators and developers to:

- to set targets;
- to set out a plan to assess the performance of their proposals against these targets;
- to explore critically areas of uncertainty;
- to provide clear quality assurance both of their data and any analysis they perform on this.

Members of the public may be concerned about the concept of risk. They may well be uneasy about being exposed to risk for, apparently, someone else's benefit. Arguments about cost benefit analysis, which had been explored, are not always attractive in this area. The public will rightly expect the risk to which they are exposed to be as low as possible. They will expect the regulators to require a demonstration that everything possible is being done to ensure that the risk to which they are exposed has been correctly assessed, is below, any published limit and is in any case as low as possible.

It is important in this context that the operator, the regulator and above all the public are satisfied that a genuine and vigorous attempt has been made to look for all areas of uncertainty to ensure that all the relevant data has been acquired and that it has all been properly considered.

Under the Radioactive Substances Act 1960, as originally enacted, much of this information could have remained essentially confidential. Under the terms of the Environmental Protection Act 1989 important provisions were incorporated in the Act to make information about applications under the Act public. This was carried over into the new Radioactive Substances Act 1993. This provides the framework for ensuring that the public are kept properly informed about the progress towards authorisation and have full and proper access to the information that enables them to judge that the Regulators are indeed doing the job that they expect of them.

ACKNOWLEDGEMENTS

The author thanks HMIP for permission to publish this paper. The work may be used for the formulation of UK Government policy but not at this stage constitute UK Government policy.

REFERENCES

1. "Radioactive Waste Management", White Paper. Cmnd 8607, HMSO 1982.

2. "Principles for the Protection of the Human Environment", Department of the Environment, December 1984.

VALIDATION: WHAT SHOULD A REGULATOR LOOK FOR IN A SAFETY CASE?

Roger D Wilmot & Daniel A Galson

Galson Sciences Ltd
35 Market Place
Oakham, Leicestershire, LE15 6DT
England

INTRODUCTION

Any proposal for the disposal or long-term storage of radioactive waste must include an assessment of the behaviour of the waste and any barrier systems, and hence any potential for transfer of radionuclides to the accessible environment. Such a long-term safety case will be prepared by a proponent with responsibilities for waste disposal and assessed in some way by a regulator with responsibilities for environmental and public safety.

The form of a safety case will depend upon the safety criteria and the degree of prescription imposed by legislation. A common element in all proposals and assessments, however, will be the use of models to demonstrate an understanding of the system and its possible future behaviour. It is upon such models that conclusions as to long-term safety will be drawn and it is therefore important that they are scrutinised by the regulator in determining whether the conclusions are justified. The large temporal and spatial scales inherent in these models pose problems for their validation, and this paper touches upon some of the measures which might be adopted to establish confidence in their use.

GENERAL CONCEPTS OF VALIDATION

Validity and the process of achieving or demonstrating it, validation, have a wide range of meanings, depending upon the context or subject area in which they are used. In mathematics and logic, for example, a theorem or argument is valid if the conclusion must be true given the particular set of premises. This type of validity, based on deduction, is not applicable in the natural sciences in which theories, hypotheses and models are developed to explain the world as it is observed to behave rather than as it **has** to behave through a set of formal rules. A theory developed to explain a particular set of observations is valid for that set and, if it is to be at all useful, will also explain other observations made under similar conditions. Since not all of the possible sets of observations can be made, however, such a theory cannot be regarded as universally valid; it can only be regarded as **not invalid**

until a set of observations shows it to be false. Successful tests of a theory do not, therefore, impart or enhance its validity, which can never be proven, but, provided the tests are not trivial, they will corroborate the theory[1].

In a simple system, a single set of observations at variance with a theory may be sufficient to invalidate the theory. In a more complex system, in which a number of processes are active, there are additional factors to be considered before a theory or model is rejected[2]. These factors can be categorised under three headings: data errors; initial and boundary condition errors; and appropriateness of data.

A complex set of observations may contain measurement errors in the original data or, since the parameters used in the theory are frequently derived parameters, there may be errors in the interpretation or analysis of such data.

Only in simple experimental set-ups can the initial and boundary conditions be controlled or even measured with any degree of precision; in a large-scale or complex system incorrect assumptions regarding these conditions may easily lead to erroneous conclusions. It is important that the uncertainties in boundary conditions are properly considered before the validity of the underlying theory is determined.

The way in which data are collected and interpreted may in itself be dependent upon the theory under test. This so-called theory-dependence of data may lead to problems. Data collected under one experimental or observational programme designed to test a particular theory may not be appropriate for tests of a different theory. It would, for example, be inappropriate to use permeabilities derived on the basis of flow through a porous medium to test a model or theory of flow through fractures. Such data are not incorrect, they are simply appropriate to a different conceptual framework. This distinction between conceptual frameworks can lead to alternative theories which are supported by appropriate observations but cannot be directly compared due to the lack of a set of observations valid for both[3].

If observational errors and boundary condition uncertainties are acknowledged, it is feasible for two or more theories to explain observational data (Figure 1), whether or not they are derived under different conceptual frameworks. The quest of science is generally to resolve this apparent ambiguity and retain only a single theory.

Although much of the philosophy of science stresses the fact that theories can never be proven, there is a need in practice to have theories and models which can be used to make predictions and on which decisions can be based. Some philosophers would contend that science progresses through the development of tests or experiments which seek to disprove

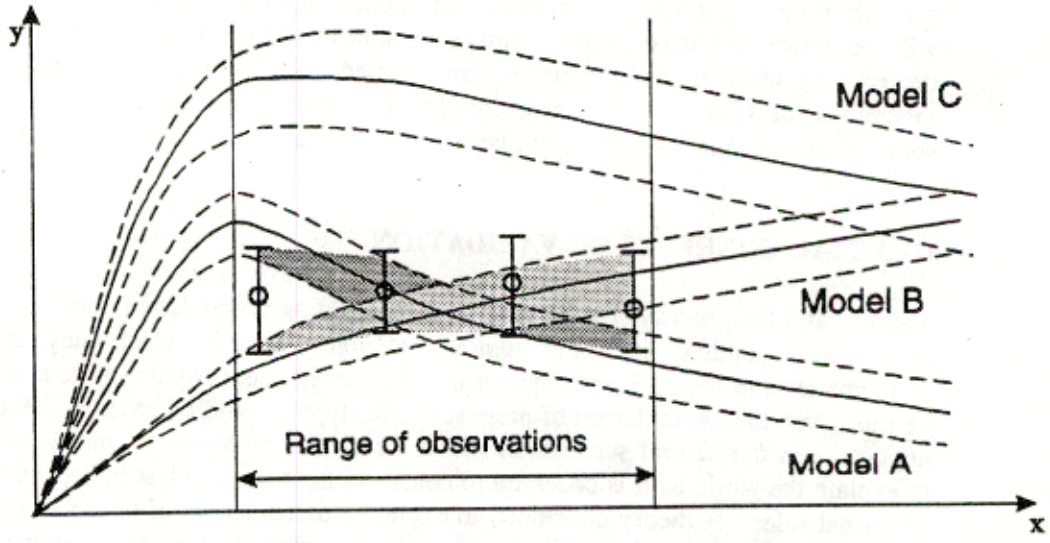


Figure 1. Model output, with a representation of uncertainties, compared with observations with representative error bars. Models A and B show equal levels of fit between predictions and observations.

or invalidate theories and hence lead to revised, more generally applicable, theories. A more pragmatic approach would be to seek tests that provide support for a theory and thus increase the level of confidence in predictions made using that theory[3,4]. In this sense, the IAEA definition of validation could be regarded as appropriate:

> *Validation is a process carried out by comparison of model predictions with independent field observations and experimental measurements. A model cannot be considered validated until sufficient testing has been performed to ensure an acceptable level of predictive accuracy.*[5]

Because of the potential for confusion between such a definition and the stricter use of validation in the theory and philosophy of science, it has been suggested that this process of comparing predictions with observations be termed *model evaluation*.

VALIDATION IN PERFORMANCE ASSESSMENTS

The difference between performance assessments of radioactive waste repositories and other uses of models lies essentially in the large spatial and temporal scales which must be considered in the former, and particularly the differences in scale between the system of concern and that of the observations which can be made. Large differences do arise in other subject areas between the scale at which observations are readily made (eg wind tunnel tests) and that of the systems of concern (eg atmospheric dispersion of gas and smoke). In such cases, however, observations of the real systems can be made in a few instances to provide confirmation of the upscaling techniques used[6]. Without such a basis, models which are equally acceptable over the scale accessible to observation (Figure 1) may lead to radically different solutions when extrapolated (Figure 2).

Depending on the requirements of the appropriate legislation, the period for which the behaviour of a radioactive waste disposal system must be considered ranges from ten thousand to a million years or more. In contrast, even a ten year period of observations on an experimental system is currently exceptional[7] and periods of a few days or weeks are commonly used as a basis for modelling studies. Depending on the repository design and the geological host rock, the volume of rock which may potentially be affected by migrating radionuclides will be in the order of several cubic kilometres. Detailed exploration of such a volume is impractical, notwithstanding the fact that too great a level of intrusive exploration may adversely affect the integrity of the site.

Validation in the sense of model evaluation discussed above cannot, therefore, be achieved for a total system performance assessment model intended to consider the long-term behaviour of a disposal system although it can be achieved for component models. Other

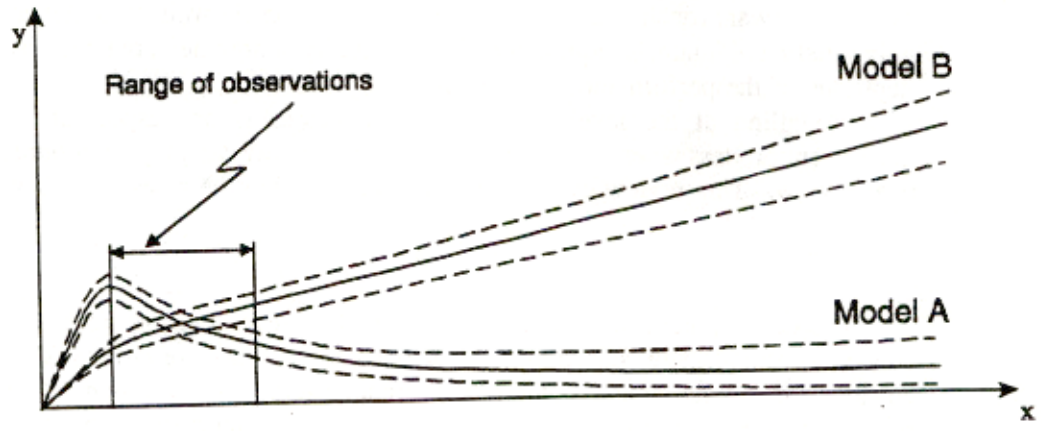


Figure 2. Possible divergence of model predictions when extrapolated beyond the range of observations used in model evaluation.

methods of establishing or building confidence in the use of total system models must be sought[8]. It should also be stressed that, since the aim of a performance assessment is a demonstration of safety, a bounding analysis rather than a representation of reality may be sufficient. The extent of validation or evaluation which is sought for different elements of the overall system should therefore be constrained by the intended purpose of such elements within the overall assessment.

CONFIDENCE BUILDING

The performance assessment of a disposal system will entail the use of a wide range of theories and models relating to particular processes and aspects of the overall system - waste form, repository design, near-field processes, geosphere processes, the biosphere and surface environment etc. The way in which such models are combined varies from programme to programme but much of the underlying conceptual model development work will be similar in all assessments.

Conceptual models consist of three principal elements: processes, geometry, and initial and boundary conditions. Extensive use of judgement is involved in each of these elements, and a full description of the role of such judgements and of the associated uncertainties is a critical aspect of a safety case[9]. A regulator would expect not only the judgements and assumptions to be documented but also the means of derivation, in order to ensure that all judgements and assumptions were traceable.

During development of a conceptual model, it is important that possible alternatives are properly considered. The discussion in previous sections has highlighted that there is no *a priori* method of determining which alternative is more applicable when extrapolated beyond the range of observations. The impact of alternative models and assumptions on safety should therefore be evaluated, and a structured and documented elicitation of such alternatives would build confidence that critical issues had been identified.

Processes

The processes to be assessed in a particular performance assessment must be determined. A direct development of such a list in an ad hoc manner could lead to the introduction of bias. A bottom-up approach, in which a full list of possible processes is first developed regardless of significance or modelling feasibility, is less likely to introduce bias. The reduction of such a comprehensive list can be structured in a variety of ways. Certain processes might be excluded based on policy decisions (eg human intrusion), probability (eg meteorite impact) or consequence (eg density-driven groundwater flow). The important aspect of any approach is that the reasons for exclusion are well-documented. This will not only satisfy a regulator's requirement for traceability of argument, but will also facilitate later iterations of the performance assessment.

Modelling at the level of individual processes is one aspect of the performance assessment process where the wider scientific community can play an important role through peer review. This review can help reduce biases within the assessment community.

Iteration

The development of a safety case will require several stages from the initial disposal concept, through site selection, site characterisation and repository design, repository construction, operation and monitoring, to final closure. This development will be paralleled by an increase in the information about the site and an increase in the level of confidence required from the performance assessment. There is, however, no guarantee that an increase in the amount of information derived from a site will necessarily lead to an increase in the

level of confidence in the models[10]. Particularly during the site characterisation programme, when there is a possibility for design changes to be made, it is important that there is adequate feedback from the performance assessments to subsequent phases of site characterisation. Such an iterative approach will ensure that processes or parameters shown to have a major impact on safety by preliminary assessments will be studied in more detail than those with little impact on safety[11].

A phased approach to site characterisation and repository construction will also give an opportunity for model evaluation. It is important that predictions which are to be compared with observations are made independently of the observations or experimental programme. A lack of such independence can easily result in a model calibration exercise rather than meaningful comparisons. An increase in understanding of the site will lead to a greater level of confidence, and this understanding can best be demonstrated by making predictions on the basis of initial site data and comparing these with subsequently determined data,

An iterative approach and clear documentation of all stages of development will help a regulator assess a safety case. A demonstration that the conceptual model of the site can be used to predict results from the later stages of site characterisation or construction may increase confidence in long-term safety assessments based on use of the conceptual model.

Initial and boundary conditions

Although there are limitations in the number of boreholes and other exploration methods which can be administered at a site, the site characterisation programme provides the best opportunity to establish boundary conditions. Even if relevant parameters must be inferred from other data, the ability to perform experiments and have direct access to groundwater and the rock matrix will reduce the number of assumptions otherwise required and hence reduce the overall level of uncertainty in the assessment. The predictions made by models using such boundary conditions cannot, however, be evaluated against any long-term behaviour of the system. In contrast, predictions made using models of past behaviour can be evaluated against present-day conditions. The initial and boundary conditions applicable to such historical models are no longer accessible to measurement, however, and will thus need to be inferred from present-day analogues.

Uncertainties over the initial and boundary conditions can limit the use of natural analogues or palaeohydrogeological models in performance assessment studies. Nevertheless, achieving an understanding of past behaviour and how present-day conditions have evolved is a major step towards understanding future behaviour. A well-documented argument demonstrating how this understanding has been developed and used as a basis for the safety case would increase confidence in an assessment.

Geometry

The geometry of the repository structure at closure will be known and can be anticipated in performance assessments. The geometry of the geosphere, including the orientation and properties of faults as well as the structure and spatial heterogeneities of different rock units, cannot be known in the same level of detail. Unless a mechanistic model of geological evolution is developed, which could be evaluated against the data available from site characterisation, recourse must be made to extrapolation or to a geostatistical model. Calibration or conditioning of such models with data from early phases of site characterisation and evaluation against data from later phases may be possible. The extent of heterogeneities in the geological environment, particularly in fractured media, may mean however that the available site characterisation data only poorly constrain a model. Predictions and evaluations should nevertheless be made and documented to ensure that the uncertainties in this element of the conceptual model are assessed.

CONCLUSIONS

Strict validation of any model or theory is an impossible task, although evaluation of component models does have a role to play in developing confidence in the total system models used in the performance assessment of radioactive waste repositories. In the absence of observations at appropriate temporal and spatial scales against which total system model predictions can be compared, confidence is best established through a full description of the judgements and assumptions made in model development and model application. A regulator will assess whether conclusions drawn by the proponent are justified by the analysis. This will require submission of a well-documented, traceable and defensible safety case.

The measures which might be adopted to establish confidence in the use of models within the safety case include:

- Systematic approaches to model building and consideration of alternative conceptual models;
- Iteration between model building, performance assessments and site characterisation;
- Use of natural analogues and palaeohydrogeological models to treat uncertainties arising from the temporal scales of concern;
- Use of geostatistics and conditional simulations to help treat uncertainties arising from the spatial scales of concern;
- Ongoing critical peer review and regular interaction between regulator and proponents prior to actual licensing.

ACKNOWLEDGEMENTS

This paper is funded by the UK Department of the Environment, but the views expressed do not necessarily represent Government policy. The authors would like to acknowledge the support of Dr Brian Thompson of Her Majesty's Inspectorate of Pollution in developing some of the ideas presented here.

REFERENCES

1. K.R. Popper. "The Logic of Scientific Discovery," Hutchinson, London (1959).
2. I. Lakatos. Falsification and the methodology of scientific research programmes, *in:* "Criticism and the Growth of Knowledge," I. Lakatos & A. Musgrave, eds., Cambridge University Press, Cambridge (1974).
3. T.S. Kuhn. "The Structure of Scientific Revolutions," University of Chicago Press. Chicago (1970).
4. H. Putnam. The 'corroboration' of theories, *in:* "Scientific Revolutions," I. Hacking, ed., Oxford University Press, Oxford, (1981).
5. "Radioactive Waste Management Glossary," IAEA-TECDOC-447, International Atomic Energy Agency, Vienna (1988).
6. J. McQuaid (ed). "Heavy Gas Dispersion Trials at Thorney Island - 2," *J. Haz. Matl*, 16, (1987).
7. "Long-Term Observation of the Geological Environment: Needs and Techniques," OECD, Paris, (1993).
8. U. Niederer. In search of truth: the regulatory necessity of validation, *in:* "Validation of Geosphere Flow and Transport Models (GEOVAL)," OECD, Paris (1991)
9. P. Grindrod and S. Watson. How uncertain we should be about the outputs of a PSA: the HMIP procedure, *in:* "Proceedings: PSAM-II", ed. G. Apostolakis, Plenum, New York (in preparation).
10. R. Mackay. "A Study of the Effect of the Extent of Site Investigation on the Estimation of Radiological Performance: Overview," UK Department of the Environment Report, HMIP/RR/93.053 (1993).
11. B. G. J. Thompson. The HMIP research and development programme for post-closure risk assessment, *in:* "Proceedings: PSAM-II", ed. G. Apostolakis, Plenum, New York (in preparation).

ASPECTS OF REVIEW OF A PROPONENTS'S POST-CLOSURE SAFETY ASSESSMENT ON BEHALF OF A REGULATOR

Trevor J. Sumerling[1] and David Read[2]

[1] Safety Assessment Management, Beech Tree House,
Hardwick Road, Whitchurch-on-Thames, Reading, RG8 7HW, UK.
[2] WS Atkins Environment, Woodcote Grove,
Epsom, Surrey, KT18 5BW, UK

INTRODUCTION

In the United Kingdom, disposal of radioactive wastes requires an authorization from the appropriate authorising Departments, or Departments, under the Radioactive Substances Act 1993, and, for an on-land repository in England, the Chief Inspector of Her Majesty's Inspectorate of Pollution (HMIP) together with the Minister for Agriculture Fisheries and Food will be responsible for granting any such authorization. At an earlier stage, both may be required to give a preliminary view as to the potential long-term safety and suitability of a proposed site at any Public Inquiry related to development of the site brought under the Town and Country Planning Act. Further discussion of the regulatory background and role of the regulator has been given by Steam (ibidem).

UK Nirex Limited (Nirex) is the company with the responsibility to provide and manage a national disposal facility for solid intermediate-level and low-level radioactive wastes. Nirex is currently carrying out a programme of investigation into the suitability of a deep underground site, close to the British Nuclear Fuels plc. complex at Sellafield in Cumbria, which is the site currently preferred by Nirex on grounds of transport cost and local acceptability (Nirex 1994).

Table 1 summarises the roles of the regulator and implementor during the planning phase leading to the granting of a construction license for a repository, based on discussion in McCombie (1992). It is generally accepted that performance assessment (PA) methods based on calculations of possible performance and radionuclide releases, coupled with sufficient information from proposed disposal sites, can provide the technical basis to decide whether specific disposal systems would offer a satisfactory level of safety to society (NEA 1991).

Since 1982, the UK Department of the Environment (UK DoE), the parent ministry for HMIP, has funded the development of a procedure for the post-closure assessment of the safety of underground disposal of radioactive wastes which is independent of that used within the nuclear industry. The methodology is based on the probabilistic system assessment (PSA), see Thompson (1987, 1993 and ibid.). The methodology has been applied to actual and proposed disposal facilities in the UK, eg. see Nicholls (1990), and in a series of "Dry Run" assessments of a hypothetical deep underground facility, the latest of which, Dry Run 3, is summarised in Sumerling (1992). HMIP consider that the possession of a risk assessment capability and an ability to carry out independent assessment are important in fulfilling their role as a regulatory agency; this view has been endorsed by Sandia National Laboratories during their review of the Dry Run 3 application (Zimmerman et al. 1992).

Table 1. Roles of the implementor and the regulator in the planning phase

1. Regulator	To define what is "safe", ie. to develop reasonable global safety standards (based on knowledge of radiation risks and comparison of other risks that must be regulated) and to consider how compliance with these requirements should be demonstrated.
2. Implementor	To develop safe solutions for disposal, ie. to propose a sound technical concept and suitable site, and to design with flexibility a disposal system that can meet the standards defined.
3. Implementor	To carry out and present a sufficiently comprehensive characterization and analysis of the system in order to demonstrate that the safety standards will be met.
4. Regulator	To judge whether the overall safety standards have been met, or shown to be met, based on examination of the assessments and safety case of the implementor and possibly on their own independent assessments.

Performance assessment is available to the implementor and to the regulator to tackle their appointed tasks, as summarized in Table 1, and to communicate with each other during this process. However, possible aims of the implementor and the regulator in their use of PA will differ, and hence the methodologies adopted may also differ quite widely, this is discussed in Sumerling and Read (1992).

In the UK, it is solely the responsibility of the proponent of a repository for radioactive wastes to produce a satisfactory safety case. The regulator will determine the adequacy of the case from the regulatory standpoint by examination of the quality of, and traceability to, the scientific basis, and by examination of the assessment procedure, in particular, the method for accounting for uncertainty and the traceability of data used. It is not within the remit of the regulator to offer advice on repository siting or design, or to lay down recommended procedures for a PA or safety case. It is within their remit to criticize and point to aspects of the implementor's PA and safety case that considered to be are poorly supported or unacceptable.

HMIP APPROACH TO EXAMINATION OF A PROPONENT'S CASE

It is assumed that, in general, a UK proponent of a major disposal facility will argue the long-term safety of their proposals upon:

(i) the overall scope and quality of their own, and related, research programmes;
(ii) the overall scope and quality of their site investigation programme;
(iii) their repository design and its compatibility with the natural geological environment;
(iv) their approach to quantitative assessment of impacts in the post-closure period.

HMIP will need to be convinced from their submissions that:

(a) the estimated risk and other end points to be considered in any authorization are soundly based, and adequately account for the uncertainties and biases due limited site information, waste characterization and scientific understanding of the relevant processes;
(b) the estimated risk and other endpoints are satisfactory in relation to the regulatory criteria;
(c) the information gathered as a result of (i), (ii) and (iii) above reduce the overall level of bias and uncertainty in the risk estimates sufficiently to give overall confidence in such estimates.

HMIP reserves the right, in the case of any particular proposal for radioactive waste disposal, to examine the proposal, and review the proponent's safety case by whatever method it deems most appropriate to the particular nature of the disposal and circumstances of the application. In the case, of the expected proposal for the construction of a national deep underground disposal facility for low- and intermediate-level radioactive wastes currently in preparation by Nirex, HMIP have opted for a review procedure based on **two complementary elements.**

HMIP will **firstly** review Nirex's case against the overall scientific and technical knowledge that exists using a panel of acknowledged experts. This will allow the scope and quality of points (i) and (ii) above to be assessed, but will not necessarily provide the detailed insights into the significance of various phenomena and processes to the long-term safety and risk estimates, for the particular site- and design-specific conditions. Hence, HMIP will **secondly** carry out independent analyses. These may include: examination of the growing body of site information in order to develop an independent model of the proposed site and prognoses for its future evolution; and risk analyses based on the PSA-centred approach developed by HMIP, see Thompson and Sagar (1993). These activities will provide the site-specific understanding of the use of the information from points (i), (ii) and (iii) above, and a critical basis for judging (iv).

It is not envisaged that the independent analyses will require the exhaustive effort that a full safety case involves, nor is it the intention of HMIP to prepare such a case; that is entirely the responsibility of the proponent. However, it is believed that a the parallel use of the same factual information to carry out significant parts of a full analysis will lead to a critical awareness of the strengths and weaknesses of the proponents case.

REVIEW OF THE SCIENTIFIC BASIS OF A PROPONENT'S CASE

To support HMIP in examining and preparing a response to Nirex's initial post-closure safety submissions, HMIP have appointed a panel of acknowledged experts in various technical and scientific areas of relevance. These experts will review the Nirex submissions in the context of the overall scientific and technical knowledge that exists, with special emphasis on their own particular areas of expertise, according to a coordinated plan. Their independent reviews will be discussed and the need for more detailed reviews of selected topics identified. The output will be individual review documents and a coordinated summary review document. The latter will act as advice to HMIP and may also be provided to Nirex.

Establishing the scientific merit of Nirex submissions unavoidably entails subjective judgements and places great reliance on the independence and expertise of the individual consultants. Therefore a set of guidelines needs to be followed in order to obtain a complete and consistent set of high quality review documents. The general topics on which guidance has been provided to reviewers in the current round of review are noted below.

Objective of the HMIP reviews. The objective of each review is to determine the scope, status and relevance of the work cited by Nirex in support of their case. Once complete, the reviewer should be in a position to assess: compliance to stated technical and quality requirements, traceability of factual information used to support arguments and the extent to which interpretations and conclusions presented are justified by the data.

Compliance of Nirex documents with their stated objectives. The objectives of the work should be clearly stated. The reviewer should assess to what extent the objectives have been met and highlight the areas where further work is necessary. The reviewer should give careful consideration to the conclusions and any wider interpretations and inferences drawn. The conclusions should be fully sustained by the information presented.

Scope of Nirex documents. Each Nirex document should be examined to ensure that its scope is compatible with the issue investigated, and that all material relevant to the case is accessible. The reviewer should identify important relevant areas of research or knowledge that have been overlooked. Where these, and other aspects which the Nirex authors have justified omitting, may influence the overall conclusions this should be commented on in the review. If the document describes work that is ongoing, the status of the work must be readily apparent to the reader.

Materials and methods. The reviewer should assess the suitability of analytical methods and codes employed with respect to accuracy and sensitivity for the task. The reviewer should comment on sampling protocols, data handling procedures and statistical treatment of results. The methods chosen by Nirex should be appropriate to the type and size of the data set under analysis. Results should include evaluation of the confidence limits associated with them. It should be evident which method or code has been used to generate each set of data and the extent of any subsequent data processing in the derivation of results.

Theoretical modelling. The use of modelling methods to represent changes which occur on spatial and temporal scales beyond those that can be investigated in experiments, will be fundamental to the Nirex safety case. These models will be of two principal types: conceptual and mathematical. The conceptual models selected by Nirex must be satisfactory representations of the system or sub-system under consideration, and decisions leading to their adoption made clear. The capabilities/limitations of the mathematical implementations must be also be made clear. The reviewer should comment on known deficiencies or omissions of the conceptual model(s) and mathematical implementation(s). The review of models employed by Nirex should consider acceptance tests reported and degree to which the models have been shown to be verified or validated, see Wilmot and Galson (ibidem).

Traceability. The component parts of the document should be referenced in such a way that the primary sources of information are ultimately traceable. Primary sources may include work published in the open literature and unpublished material. However, if an author cites unpublished material then it must be possible to access this information. The reviewer should satisfy him/herself that all the references cited in the text are represented in the reference list and are appropriate to the statements they support.

Use of expert judgement. The Nirex documents must make clear the methods by which expert judgements have been incorporated into the analysis. Where formal elicitation methods have been employed these should be documented, together with sufficient information to demonstrate the quality of the enactment and, for example, whether dissenting or alternative views have been properly recorded. The reviewer should assess whether expert judgement was only applied where more objective approaches were not possible, for example, due to lack of directly appropriate and statistically meaningful data sets. The reviewer should be satisfied that: the form and type of data used as input to elicitation was appropriate and complete; the reasoning of the experts leading to the stated conclusions or elicited data is described.

Recommendations. Each review document should conclude with a section advising HMIP of the implications of the review and a set of suggested actions, prioritised according to their likely impact on the ultimate acceptability of the Nirex safety case and documentation. These actions will generally be of the form of suggestions for additional work or requests for clarification, but may also call for independent analysis by the HMIP assessment team.

INDEPENDENT ASSESSMENT OF SITE AND DISPOSAL SYSTEM

Relation of HMIP Assessment to a Proponents Analyses

Before embarking on an independent assessment it is important to have a clear view on how the independent analyses will relate to the analyses performed by the proponent. Given the large uncertainties in underground geological environments and possible long-term evolution of a waste disposal system, and also given significant differences that may exist between the methodologies employed by proponent and regulator, it is clear that the "results" of the proponent's and regulator's analyses may not be immediately compatible; although, if both analyses honestly acknowledge the full range of uncertainties, then the envelope of results, incorporating uncertainty, should be consistent.

However, quantitative agreement is not being sought. The aim of the regulator's analyses are to selectively explore the uncertainty with a view to identifying key uncertainties and potential biases; the proponents case can then be examined to see how these have been dealt with there. The uncertainties referred to are not the simple parametric uncertainties within a defined model, but the uncertainties and biases related to interpretation of site data, development of site models, decisions on the inclusion and exclusion of processes from the models of site and radionuclide release/transport, and decisions on choice of mathematical formulation and statistical techniques.

The simplest form of analysis will take a single path of decisions based on expert judgements and conditioned by the availability of models and data. However, in principle, a multitude of parallel or divergent assessment calculations can be envisaged employing alternative basic data, conceptual models and mathematical formulations (Thorne and Laurens, 1989). These cannot be combined probabilistically since they are not independent entities. Rather the difference between results, with uncertainty, for each analysis path, is a measure of the bias or uncertainty associated with the decisions made.

Neither, the proponent nor the regulator can be expected to explore all possible analysis paths. Rather, the regulator will expect the proponent to produce reasoned arguments for the significant decisions taken in their assessment, and also, perhaps by limited examination of alternatives, to quantify the bias associated with their results and adequately caveat their conclusions. The regulator should be able to examine alternatives in a freer way and hence seek to identify conditions that might lead to unsatisfactory disposal system performance. If the regulator's analyses reveal such conditions, and these are not adequately covered by the proponents analyses, then the proponent would be asked to examine this more fully or advance arguments as to why such conditions could not occur.

In addition, within the HMIP programme, effort is being directed to developing procedures for *uncertainty and bias audit* (Thorne, 1992; Grindrod and Watson, ibid.) and *uncertainty and bias evaluation*. The latter is the subject of work in the current phase of HMIP assessments of Nirex proposals. It is hoped, this will lead to a satisfactory method of accounting quantitatively for uncertainties and biases due to the necessarily limited scope of assessment at a given stage of site data and model development, and also provide a means by which the assessments of the proponent and regulator may be compared.

Geological Understanding of the Proposed Site

Post-closure radiological assessments often begin by assuming a single geological and hydrogeological model of the disposal site for which the degree of uncertainty or confidence is not considered. Given the indirect nature of geophysical investigation techniques, very sparse direct sampling of the geological media by the means of boreholes, and the subjective and inferential nature of geological and hydrogeological interpretation techniques, this is a wholly unwarranted assumption. Experience shows that geological interpretations are notoriously unreliable and subject to change both due to additional data acquisition and re-interpretation of the same data (Muir-Wood et al., 1992). In addition, a trial study sponsored by HMIP based on synthetic site data, has shown that radiological risk estimates based on progressively increasing site data can vary significantly beyond the bounds of uncertainty accounted for within the models derived by the analysts, see MacKay (1993) (also Thompson, ibid.).

HMIP will expect Nirex to provide realistic estimates of the degree of uncertainty associated with their geological and hydrogeological interpretations and models. It is possible that alternative geological and hydrogeological models, representative in some justified way of the range of uncertainty, will have to be analyzed in a final assessment.

Understanding the site and data limitations and uncertainties will enable HMIP to judge the adequacy of the ongoing Nirex site investigation, and is also an essential precursor to erecting sensible independent assessment models. Within HMIP's programme of work related to assessment of Nirex proposals, two groups of geological studies have been put in hand. The **first** allows for QA audit of the Nirex site investigation programme; this includes, for example, site visits, examination of rock cores, geophysical and hydraulic measurement records and selective tracing of data from measurements through the Nirex contractor records and reports; selective audits of Nirex contractor interpretations, eg. interpretation of geological structure from geophysical data, will also be undertaken. The **second** allows for access to Nirex geological data catalogue, examination of selected data and independent detailed evaluation/interpretation of key aspects. Aspects which may be considered include structural interpretation, tectonic evolution, groundwater geochemistry, heterogeneity of the host rock, hydrogeological interpretation and geotechnical aspects. The possibilities for developing an independent HMIP geological database on which alternative interpretations can be tested and visualised are being investigated.

The first group of studies above mainly relate to testing the quality of Nirex's programme. The second group are essential to HMIP as a regulator: firstly, to provide an understanding of the geological uncertainties and the degree to which Nirex's planned investigations may be able to reduce these uncertainties; secondly, to allow a view on the interpretations proposed by Nirex and the degree to which they represent the uncertainty; thirdly, to provide the basis for alternative interpretations as candidates for examination in HMIP's radiological performance assessments

Post-closure Radiological Assessment

The objective of HMIP in commissioning an independent PSA of the Sellafield site is to provide a quantitative understanding of the site and key uncertainties affecting the estimates of risk; this in order to support the review of the Nirex safety case and post-closure analysis.

The core of the HMIP analysis is based on the use of the system codes and supporting codes mainly as available and already demonstrated in Dry Run 3 (Sumerling, 1992). Previous assessments have indicated the difficulties of undertaking site assessments with versions of code that have not undergone full testing. Where developments are necessary to represent site-specific processes not currently modelled, these will be made in parallel with the assessment activities and, only after testing, will the revised codes be released for assessment use. The key decisions leading to the definition of models and data employed in the assessment are recorded as part of the uncertainty and bias audit (Grindrod and Watson, ibid.). The bias due to the omission of particular processes, or due to a particular conceptual model choice, is investigated, preferably, through scoping calculations or supplementary analyses.

The overall outputs from the assessments are expected to be:

- estimates of risk and probability density function of dose as a function of time;
- estimates of bias due to omitted processes and/or model choices;
- sensitivity of risk to parameters and model choices and, hence, identification of key uncertainties;
- definition of requirements from Nirex site investigations, design studies or research to reduce or improve the quantification of these uncertainties;
- recommendation for longer-term model development and/or research by HMIP/UK DoE to improve the evaluation of risk from a repository at the Sellafield site.

The HMIP PSA activities have been organised in distinct phases, to coincide with the expected phases of Nirex provision of data and preliminary safety submissions. A **first phase**, completed in April 1993, was aimed at producing a pilot demonstration model based, as far as possible, on Nirex data and hydrogeological interpretation. This was implemented directly with the VANDAL radiological release, flow and transport code (Kane, 1992) and considered time-independent conditions of the site. A **second phase**, scheduled for completion in December 1994, will be based initially on Nirex site interpretations but incorporating a greater range of geological and hydrogeological uncertainty, and will be supported by independent detailed hydrogeological modelling of the site. The VANDAL implementation will be linked to a site-specific implementation of the TIME4 environmental simulation code (Dames and Moore, 1991), which will provide time-dependent boundary conditions to the VANDAL flow and biosphere models. The scope of a possible **third or subsequent phases** will be decided by HMIP in the light of Nirex's developing programme of submissions. The process of iterative assessments will build up an experience which eventually will be used as part of the basis for considering a possible authorization for the disposal of radioactive wastes.

SUMMARY

In the UK, it is solely the responsibility of the proponent of a repository for radioactive wastes to produce a satisfactory safety case. The regulator will determine the adequacy of the case from the regulatory standpoint by examination of the quality of and traceability to the scientific basis, and by examination of the assessment procedure, in particular, the method for accounting for uncertainty and the traceability of data used. In the case, of the expected proposal for the construction of a national deep underground disposal facility for low- and intermediate-level radioactive wastes currently in preparation by Nirex, HMIP have opted for a review procedure based on two complementary elements.

HMIP will firstly review Nirex's case against the overall scientific and technical knowledge that exists using a panel of acknowledged experts and, secondly, carry out independent analyses. These may include: examination of the growing body of site information in order to develop an independent model of the proposed site and risk analyses based on the PSA approach developed by HMIP since 1982.

It is not envisaged that the independent analyses will require the exhaustive effort that a full safety case involves. However, the parallel use of the same factual information to carry out significant parts of a full analysis will lead to a critical awareness of the strengths and weaknesses of the proponents case.

Acknowledgement

The authors thank HMIP for permission to publish this paper. The work may be used for the formulation of UK Government policy but does not at this stage constitute UK Government policy.

REFERENCES

Dames and Moore, 1991, "Technical Reference Manual for TIME4 Version 1.0," (2 volumes) UK DoE Reports DoE/HMIP/RR/91.049 and 91.050.

Grindrod P. and Watson S., ibid., How uncertain should we be about the outputs of a PSA: the HMIP procedure.

Kane P., 1992, "VANDAL Version 1.2 Technical Overview," UK DoE Report DoE/HMIP/RR/92.095.

Mackay R., 1993, A Study of the Effect of the Extent of Site Investigation on the Estimation of Radiological Performance: Overview, UK DoE Report DoE/HMIP/RR/92.053.

McCombie C., 1992, The role of the repository implementer in providing and demonstrating safe disposal of radioactive wastes, in "High Level Radioactive Waste Management", Proc. of the 3rd Ann. Conf., Las Vegas, Vol.2, pp.102-105.

Muir-Wood R., Woo G. and Macmilan, 1992, "The Nirex Sellafield Site Investigation: The Role of Geophysical Interpretation," UK DoE Report DoE/HMIP/RR/92.111.

NEA, 1991, "Disposal of Radioactive Waste: Can Long-term Safety be Evaluated?" published NEA(OECD) Paris.

Nicholls D.B., 1990, "A Preliminary Radiological Risk Assessment of Management Options for the Dounreay Waste Shaft," UK DoE Report DoE/HMIP/RR/90.082.

Nirex, 1994, *Nuclear Energy, Journal of the British Nuclear Energy Society*, February 1994.

Steam S., ibid., Risk analysis in regulation and risk communication.

Sumerling T.J. (Editor), 1992, "Dry Run 3: A Trial Assessment of Underground Disposal of Radioactive Wastes Based on Probabilistic Risk Analysis: Overview" UK DoE Report DoE/HMIP/RR/92.039.

Sumerling T.J. and Read D., 1992, Approaches to long-term performance assessment of deep underground disposal of radioactive wastes: A european perspective, in "Scientific Basis for Nuclear Waste Management XVI" Proc. Symp. Boston, Mass., December 1992, published Materials Research Society, Pittsburgh, pp951-961.

Thompson B.G.J., ibid., The HMIP research and development programme for post-closure risk assessment.

Thompson B.G.J. and Sagar B., 1993, The development and application of integrated procedures for post-closure assessment, based upon Monte Carlo simulations: the probabilistic systems assessment (PSA) approach. Submitted for publication in J. Reliability Engineering and System Safety.

Thompson B.G.J., 1987, The development of procedures for the risk assessment of underground disposal of radioactive wastes: research funded by the Department of the Environment 1982-1987, *Rad. Waste Management*, 9(1-3), pp231-262.

Thorne M.C., 1992, "Dry Run 3: Volume 8: Uncertainty and Bias Audit," UK DoE Report DoE/HMIP/RR/92.040.

Thorne M.C. and Laurens J-M., 1989, The development of an overall assessment procedure incorporating an uncertainty and bias audit. in Proc. IAEA/NEA/CEC Symposium, Paris, published NEA 1990.

Wilmot R.D. and Galson D.A., ibid., Validation, what should a regulator look for in the safety case?

Zimmerman D.A. et al., 1992, "Peer Review of the UK DoE Dry Run 3 Exercise," Sandia National Laboratories Report SAND92-1945 available with Addendum from HMIP, London.

070 Reactor PSA Studies

Chair: M. Hitchler, Westinghouse Savannah Rvr.

Probabilistic Safety Analysis of Ignalina NPP, RBMK-1500

G. Johansson (IPS); E. Shiverskly (IPE); E. Söderman (ES-Konsult)

Results and Insights from the Low Power and Shutdown Level 1 PRA of the Grand Gulf Nuclear Power Station

D.W. Whitehead, B.D. Staple, S.L.Daniel (SNL); J. Darby (Sci. & Eng. Assoc.); J. Yakle, S. Miller, J. Forester (SAIC)

Comparative Evaluation of the Performance of Advanced Semi-Passive Pressurized Water Reactors in the Loss of Offsite Power Event

M. Ouyang, M. Golay (MIT)

PROBABILISTIC SAFETY ANALYSIS OF IGNALINA NPP, RBMK-1500

Gunnar Johanson[1], Eugene Shiverskiy[2] and Erik Söderman[3]

[1] Industrial Process Safety AB
Svartviksslingan 11
161 29 Bromma, Sweden

[2] Research and Development
Institute of Power Engineering
P.O .Box 788 Glavpochtampt
Moscow 101000, Russia

[3] ES-Konsult AB
Box 3096
161 03 Bromma, Sweden

INTRODUCTION

The Barselina Project was initiated in the summer of 1991. It is a multilateral cooperation between Lithuania, Russia and Sweden. The long range objective is to establish common perspectives and unified bases for assessment of severe accident risks and needs for remedial measures for the Large Power Boiling Reactor design, (Reaktor Bolshoy Moshchnosty Kipyashchiy - RBMK). The Swedish BWR Barsebäck is used as reference plant and the Lithuanian RBMK Ignalina as application plant.

A general concept for analyzing RBMK reactors has been developed, which can be applied and used for all RBMK plants. The results are preliminary and the model and data will be further justified during the ongoing third phase of the project, March 1993 to June 1994. Therefore the results can, at this stage, not be used as a basis for an absolute overall risk assessment of the RBMK reactors. The results, however, clearly indicate both the need for and possibilities to improve the safety of the Ignalina NPP. The risk topography established can now be used to identify plant specific improvements in system design, operating or maintenance procedures. A special group is formed to perform this detailed analyses and to propose safety improvement changes.

RBMK CHARACTERISTICS IMPORTANT WHEN CONDUCTING PSA

To give some aspects on the Ignalina PSA model will we shortly present some interesting RBMK characteristics and features. These items create significant modelling differences in comparison with BWR and PWR models of Swedish plants.

Blockage events

After screening and grouping of initiating evens one important difference to vessel type reactors is the primary circuit circulation blockage events. Since the reactor actually constitutes of 1600 channels one can say each channel represent a small vessel reactor. The blockage event can be a blockage of one channel but also groups of channels.

Rupture location zones

The rupture zone definitions are illustrated in Figure 1. For the RBMK the number of zones is larger than for western reactors with primary system containment. The upper part of the circulation system and the steam separators are not connected to the confinement system, called Accident Localization System. Different system function requirements makes it necessary to treat ruptures in zones 1, 2 and 3 separately. For the Ignalina plant, the Design Basis for the

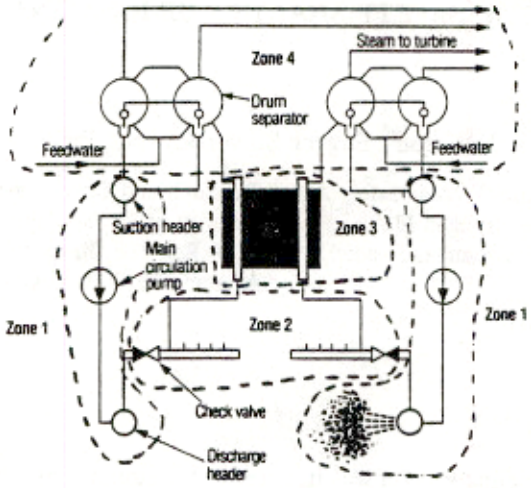


Figure 1. Rupture location zones for the Ignalina plant.

confinement system is a rupture of the main circulation pump discharge header (900 mm). In zone 4, the largest rupture is defined as main steam line (400 mm)

Hazard states

A Level 1 PSA usually implies the assessment of plant failures leading to core damage without distinguishing different levels of core damage. In the RBMK case, four hazard states have been defined. The final result greatly depends on the definition of what constitutes a core damage. The four core damage categories used for probabilistic safety indices evaluation are the following: Safe conditions (category S) when the safe operating limits are not exceeded. Accident sequences within safe condition limits can however lead to significant economical losses. Violation of the reactor safe conditions (category V). This category includes accidents with fuel elements cladding damage in a few channels. This category can be looked upon as a "mild consequence". Reactor core damage (category D). This category is characterized by significant deviation from the design scenario with fuel elements damage. This category can be looked upon as a "medium consequence". However such accidents do not lead to loss of core structural integrity and other most severe consequences. Severe accident (category A). Core melt or fuel damage at high pressure leading to loss of core structural integrity and other most severe consequences such as: 1- pressure increase in the reactor cavity leading to raising of reactor upper plate at high pressure 2-damage (melting) of fuel of a part or the core as a whole 3-fuel dispergation 4-core graphite burning.

Success criteria

The core cooling safety functions in the RBMK-1500 design are based on two six (6) train systems, ECCS and AFWS. For design basis LOCA scenarios the basic redundancy level for

each of these is ~6x25% and totally 12x25%. However, there are scenarios raising higher requirements than the design basis LOCA which yields a redundancy level of 12x15% in specific cases, i.e. 7 out of 12 trains are required (failure criteria, fc, 6/12). To complicate the criteria further the reactor core is divided in two halves, the core injection must arrive where its needed, so that in the 7 out of 12 case the success criteria for one half is 3 out of 12 (fc 10/12) and the other, ruptured, half is 4 out of 12 (fc 9/12). The final outcome, the functional failure expression, will in this case be an or gate with three inputs one for each half, fc1 & fc2, and one for the core as a hole, fc 3. The problem is further complicated by the support systems that have different success criteria dependent on what combination of ECCS pumps and AFWS pump that eventually are used. The fault tree model that represent this functional event becomes extremely complex, the number of boolean indicated failure combinations exceeds 10^{1000}. This can be compared with Swedish PSAs that in the most complex situations can reach to the order of 10^{120} combinations. This aspect raise very high requirement on the code used for evaluating the model, in our case the Risk Spectrum code.

MIDWAY RESULTS AND CONCLUSIONS

During phase 2, from April 1992 to February 1993, a qualitative PSA model of Ignalina unit 2 (INPP-2) has been developed for purposes of model testing and demonstration. A first set of quantitative results has been obtained, using screening data. Concluding remarks on the absolute risk level are, due to screening data and conservative assumptions, not relevant at this stage. The risk topography, however, gives a firm basis for further work. It shows that it is important to continue to improve the PSA models and to extend the study towards more initiating events.

An illustration of the risk topography is shown in figure 2. It shows the barrier relation between initiating event and emergency system barriers to core damage for different initiating

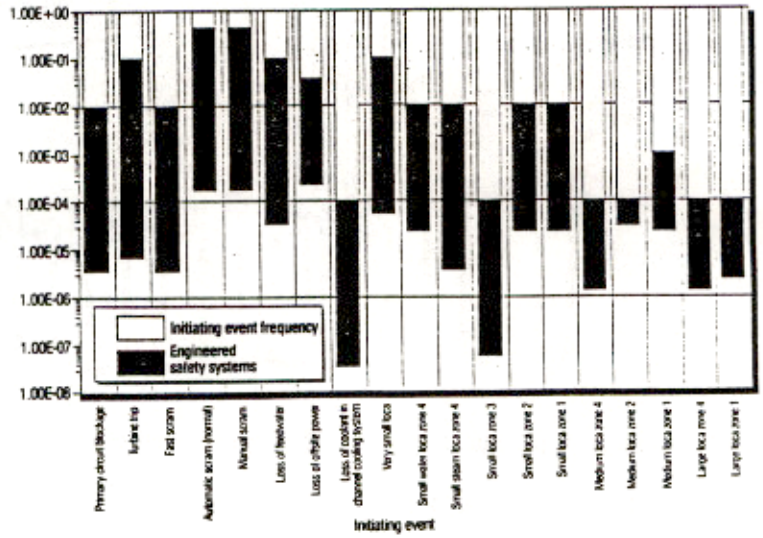


Figure 2. Results of a preliminary risk evaluation of Ignalina unit 2.

events, figure 2 demonstrate the principles of Risk Spectrum application with unconfirmed event frequencies and component data. The total barrier should be as large as possible and is illustrated in two parts, the upper one representing the assumed initiating event frequency and the lower (darker) one the barrier representing the engineered safety systems.

On the basis of the screening results, proposals for change of input data have been risen. Both data and other findings in the result have been used for a sensitivity analysis, showing examples of potential for improving the models, the data and the plant design and hence to change the risk topography and also the risk level of the plant.

An example of the summary of this improvement potential is shown in Figure 3. It is important to note that these improvements represent a total potential, which will have to be justified. This will change the risk topography, as will the extension of the study to new initiating events etc.

Consequences and functional requirements in the accident sequence model, success block diagrams and event trees, are often based on conservative assumptions. This is because of lack of reference deterministic analyses in many cases. Core hazard states assigned in the phase 2 model are preliminary and are not supported by accident process analysis. References to the accident process analysis must be formalized and used for establishing the successful sequence of events, as well as references for the assignment of core hazard states.

Core damage sequences are, using the screening initiating event frequencies, dominated by transients, such as loss of off site power, loss of feed water and manual or automatic reactor scram. Behind multiple common cause failure on all diesel generator systems or all emergency core cooling pumps, interesting findings and possibilities for safety improvements are identified. Still the findings are preliminary and may be due to mistakes or incompleteness in the models.

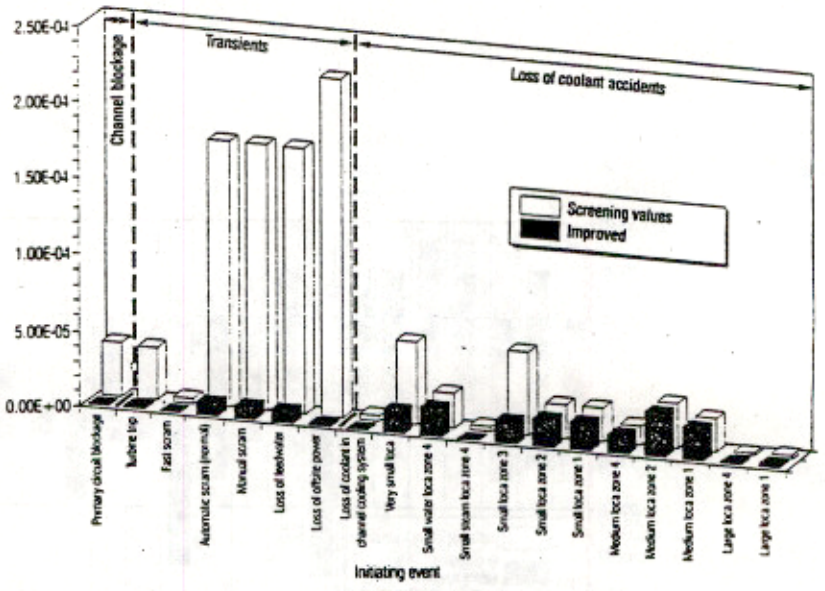


Figure 3. Summary of potential risk topography improvements at Ignalina unit 2.

ORGANIZATION FOR SAFETY IMPROVEMENTS

The organization and responsibility for safety improvements will be moved from the designer to the plant owner. The plant organization is changing their attitude and work formats. In this context the PSA play an important role. Currently the safety division at the plant has been complemented with a PSA office. INPP personnel contribute significant to the development of the PSA model, 40% of the total effort. The PSA has contributed in the process

of safety review and have introduced the organization to start to think and analyze in worst case scenarios in a systematic way. Taylor´s law "if anything can go wrong it will" must be implemented in safety management.

PSA APPLICATIONS IN SAFETY MANAGEMENT

The plant management has produced a large number of issues for safety improvements based on the operational experience. The PSA has contributed to this list in this first stage. Later on the most important use of the PSA will be to prioritized among the issues and to use the PSA as one element in the long term safety management and operation of the plant. A special PSA Implementation Group has been formed. This group will have as its special task:

- to analyze the results aiming at proposals for safety improvement changes in design or procedures and
- to analyze the list of safety improvements proposed by the designer or plant staff in view of the PSA knowledge obtained.

The Implementation group is constituted by PSA and plant specialists from the Ignalina and Barsebäck plants. They will work in close connection also to the human factors analysis group, which work on improved modelling of human factors and in analyzing both operational and maintenance procedures.

In the detailed analyses of dominating accident sequences, several possibilities for risk reduction have been identified. Some of the findings are:

ECCS logic: Following a large or medium loss of coolant accident, the auxiliary feed water pumps shall inject water to the emergency core cooling headers. The operator must close valves that could just as well be closed automatically. This event substantially contributes to the damage risk, since only two minutes are available to carry out this action, and have formed the basis for a complete review of the ECCS logic performed at the plant.

Support systems: The high redundancy - *e.g.* six diesel generator trains and six emergency core cooling pumps and six auxiliary feedwater pumps in the front line systems - tend to move the vulnerability to support systems, which are necessary for the long term emergency cooling. Several weaknesses has been identified and removed in these systems.

Removal of main circulation pump by pass: In the case of a rupture in the primary circuit that cannot be isolate backflow must be established through the affected channels. To provide backflow the MCP´s must be isolated to not loose water to the ruptured piping. This operator action could not be carried out successfully due to a by pass line without remote controlled isolation valves, to enter the compartment where the valves are located is impossible at LOCA conditions. The by pass was originally installed to ensure natural recirculation in loss of offsite power situation, but had not been removed when the by pass was shown to be unnecessary in full scale experiments.

During the result evaluation a long candidate list for detailed plant improvements and further study during phase 3 were established. This list will form one of the starting points for the work. The prospects for improvement is illustrated in figure 3.

PSA MODEL DEVELOPMENT

The manpower effort spent hitherto is in the order of 15 manyears, October 1993. The Swedish project organization started the project with a manpower rate of 3 manyears/year, or 80% of the total work in the initial stages of the project. At the end of phase 3 the Swedish effort will be only 20% of the total effort. The Russian side now works with a manpower rate of 4-5 manyears/year. The Lithuanian side, INPP staff and the Lithuanian Energy Institute, are putting in a manpower rate of 4-5 manyears/year. When phase 3 is completed, June 1994, the total manpower effort will be approximately 25 manyears.

The objective with phase 3 is to produce a full scope Probabilistic Safety Assessment, Level 1+, of Ignalina NPP, unit 2. This study will be reviewed internationally and published by the INPP staff and management.

Initiating events

The initiating event selection has been based both on Western (EPRI, Canadian) and Russian (RDIPE = the main designer) experience. The complete initiating event list defines 5 transient events, 3 blockage events, 4 Common Cause Initiator, CCI, events and 13 LOCA events. The existing phase 2 model is not complete. Hitherto the range of initiating events included is limited to internal events such as transients and loss of coolant accidents. Common cause initiators (CCI) such as fire, support system failures (power supply, service water, ventilation), malfunction of level indication or flow measurement are not considered. Such incidents may give considerable contributions to the risk profile. The continued work will extend the study to include CCI analyses and reactivity transients

Accident Sequence Modelling

The safety functions needed for preventing core damage of the Ignalina NPP have been identified. The required safety functions have been illustrated by functional block diagrams for successful accident protection. A preliminary set of plant function block diagrams for successful event response have been generated. Due to the lack of a western standard FSAR a large work by the designer and plant staff has been focused on the design of these success diagrams, which form the basis for the event trees. This work is currently supplemented with deterministic analysis to remove assumptions regarding success criteria.

System analysis

The overall system analysis, using fault tree technique, at present covers 18 systems for the Ignalina 1500 MWe NPP. Human Performance Analysis has been, so far, limited to screening analysis. The PSA model has concentrated on identifying the critical elements and types of human interaction. Several crucial human interactions have been identified in the accident sequences. These will be further analyzed during the next phase of the project. All functional, physical and human dependencies are specificly modelled in the plant model event trees or fault trees. Common Cause Failures are included in the models to represent the residual dependencies with unidentified causes not covered by the qualitative dependency analysis. The continued work include as the most important tasks to develop the plant models especially in the fields of functional dependencies and human factors.

Initiating Events Frequency data and Component Reliability Data

Only screening values for initiating event frequencies and for component reliability data have been used. This data is based on Swedish experience. A comprehensive work are carried out by the plant staff to establish a plant specific data base for component reliability, which will be implemented during next phase of the study. The continued work include to justify component reliability by review of operating and maintenance records and further to establish reference data as a basis for initiating event frequencies used.

CONCLUSIONS

The ambitions of the ongoing phase 3 program are high. The transfer of working methods for PSA application will continue and a larger part of the work will be performed by the plant staff who will also step by step take over the responsibility for the work. Support from the Russian chief designer, especially in the field of deterministic analyses, will be necessary in the long term even if the plant take over the responsibility for safety. In this context the PSA play an important role the most important use of the PSA will be to prioritized among safety issues and to use the PSA as one element in the long term safety management and operation of the plant.

RESULTS AND INSIGHTS FROM THE LOW POWER AND SHUTDOWN LEVEL 1 PRA OF THE GRAND GULF NUCLEAR POWER STATION*

Donnie W. Whitehead[1], John Darby[2], Jeff Yakle[3], Bevan D. Staple[1], Steve Miller[3], John Forester[3], and Sharon L. Daniel[1]

[1]Reactor Systems Safety Analysis Division
Sandia National Laboratories
Albuquerque, New Mexico 87185

[2]Science & Engineering Associates, Inc.
Albuquerque, New Mexico 87110

[3]Science Applications International, Corp.
Albuquerque, New Mexico 87185

INTRODUCTION

As a result of the Chernobyl accident and other precursor events, such as the 1990 loss of AC power event at Vogtle, the U.S. Nuclear Regulatory Commission's Office of Nuclear Regulatory Research has undertaken a two-phase project to analyze the frequencies, consequences, and risk of accidents occurring during modes of operation other than full power (i.e., low power and shutdown (LP&S) conditions). The Phase 1 work (Whitehead et al., 1991) represented a "coarse screening" of all identified modes of operation. The results were used to identify which set of plant operating conditions would be subjected to a detailed analysis during Phase 2.

This paper presents an overview of the results from the current Phase 2 effort to identify and quantify the various accident sequences that can lead to severe core damage. This overview encompasses the detailed analysis of potential accidents that could occur at the Grand Gulf Nuclear Power Plant while in the Cold Shutdown Phase [Plant Operational State (POS) 5] of a refueling outage. From the Phase 1 analysis, it was

*This work is supported by the U.S. Nuclear Regulatory Commission and is performed at Sandia National Laboratories, which is operated for the U.S. Department of Energy under Contract Number DE-AC04-76DP00789.

determined that on a relative basis, POS 5 contributed approximately 59.5 percent of the total core damage frequency. The only other POS to contribute a significant fraction to the total core damage frequency was POS 6, where POS 6 represents the plant refueling configuration when the vessel head is off and the water level is raised to the Main Steam Lines. In addition to contributing the largest fraction to the total core damage frequency, POS 5 also contributed the most sequences to the most important "simplified" plant damage state analysis (i.e., open containment and early core damage). Based upon these results, POS 5 was selected for detailed examination during the second phase of the project.

After selecting POS 5, an extension of the models and data developed during the screening study was conducted. Based on a more detailed examination of resource materials, the list of initiating events was expanded to include new events that could possibly occur during POS 5 that had not been originally identified in Phase 1. An example of such initiators include those that could occur while the plant undergoes a system hydro test. The frequencies used for the initiating events were updated using new information sources. Based upon updated plant-specific information and insights gained from the Phase 1 work, changes were made to the existing event trees and fault trees, and new event trees were constructed to account for the new initiating events. In addition, a multiple initiating event analysis was performed to identify the valid double and triple combinations of initiating events. The resulting combinations were then examined to determine if the event trees adequately accounted for the multiple initiating event combinations that survived the screening criterion.

The ability of systems to mitigate an accident initiating event depends on the initial status of the systems in POS 5 (i.e., Cold Shutdown). The event trees used in the detailed analysis consider the major variations in systems availabilities, such as forced or natural circulation and suppression pool inventory. These changes, plus the incorporation of additional human actions (many of which represent recognition/decision points) into the trees, form the basis of the accident sequence analysis task performed in this study.

The following sections describe the current status of the Phase 2 work for the LP&S analysis of Grand Gulf. The first section will provide a list of the initiating events and their associated frequencies. The second section will briefly describe the event trees constructed during the project. The third section will outline the approach used to quantify the diverse human actions incorporated into the event trees. Finally, the fourth section will present results from the accident sequence quantification task along with a brief discussion of the dominant accident sequence categories.

INITIATING EVENTS

This section of the paper summarizes the identification and quantification of initiating events. The reanalysis of the initiating events was undertaken for two main reasons. First, as POS 5 was modeled in detail, new initiating events were identified. For example, modeling of the Hydro condition in Cold Shutdown led to the consideration of Loss of Coolant Accidents (LOCAs) at full system pressure; and the modeling of the possibility that forced recirculation is initially available in POS 5 led to the consideration of loss of forced recirculation as an initiating event. Second, a trend identified in the screening study indicated that certain initiating events, such as loss of instrument air, should be quantified more accurately.

Accident initiating events were classified into three groups: Transients, LOCAs, and Decay Heat Removal Challenges. A Transient was defined as any non-LOCA initiating

Table 1. Initiating Events for POS 5

Initiating Event Nomenclature	Description	Mean Frequency per Year for POS 5
T_1	Loss of Offsite Power (LOSP) Transient	0.13
A_3	Large LOCA at Low Pressure	3.62E-05
A_{HY}	Large LOCA during Hydro Test (High Pressure)	1.25E-04
S_1	Intermediate LOCA at Low Pressure	3.62E-05
S_{1H}	Intermediate LOCA during Hydro Test (High Pressure)	1.25E-04
S_2	Small LOCA at Low Pressure	3.62E-05
S_{2H}	Small LOCA during Hydro Test (High Pressure)	1.25E-04
S_3	Small-small LOCA at Low Pressure	3.62E-05
S_{3H}	Small-small LOCA during Hydro Test (High Pressure)	1.25E-04
H_1	Diversion to Suppression Pool via RHR	6.1E-02
J_2	LOCA in connected system (RHR)	1.56E-02
E_{1B}	Isolation of SDC loop B only	5.7E-02
E_{1C}	Isolation of RWCU as DHR	1.57E-03
E_{1D}	Isolation of ADHRS only	5.7E-02
E_{1T}	Isolation of SDC common suction line	0.356
E_{1V}	Isolation of common suction line for ADHRS	0.356
E_{2B}	Loss of operating RHR shutdown system	6.5E-02
E_{2C}	Loss of RWCU as DHR	1.57E-03
E_{2D}	Loss of ADHRS only	6.5E-02
E_{2T}	Loss of SDC common suction line	3.8E-02
E_{2V}	Loss of common suction line for ADHRS	3.8E-02
T_{3A}	Loss of all Standby Service Water (SSW)	2.4E-02
T_{3B}	Loss of all Turbine Building Cooling Water	2.4E-02
T_{3C}	Loss of all Plant Service Water (includes Radial Well)	2.4E-02
T_{3D}	Loss of all Component Cooling Water	2.4E-02
T_{AB}	Loss of 1E 4160 V AC Bus B	1.66E-03
T_{DB}	Loss of 1E 125 V DC Bus B	6E-03
T_{IA}	Loss of Instrument Air	0.18
T_{ORV}	Inadvertent Open Relief Valve at Shutdown	7.2E-02
T_{IOP}	Inadvertent Overpressurization (makeup greater than letdown)	1.57E-03
T_{IHP}	Inadvertent Pressurization via spurious HPCS actuation	1.4E-02
T_{IOF}	Inadvertent Overfill via LPCS or LPCI	2.2E-02
T_{RPT}	Loss of Recirculation Pump	7.2E-02
T_{LM}	Loss of Makeup	8E-03

event, excluding Decay Heat Removal Challenges. Decay Heat Removal Challenges are initiating events in which the operating shutdown cooling system is directly lost. As a result of the reanalysis, thirty-four initiating events were identified and quantified. Table 1 provides a list of the initiating events and their frequencies as used in Phase 2.

EVENT TREES

The event trees used in this study of POS 5 have been extensively updated from those used in the original screening analysis (Whitehead et al., 1992). Two major reasons for the extensive modifications to the trees are (1) a variety of initial conditions were considered in the detailed analysis beyond those considered in the screening study, and (2) updated information obtained from the staff at Grand Gulf indicated that auto isolation of the low pressure shutdown cooling suction piping on high pressure (135 psig) was available, where in the original screening analysis it was considered unavailable.

To simplify the development of the necessary event trees, three types of event trees were developed. These included (1) Generic Functional Event Trees, (2) Generic System-Level Event Trees, and (3) Specific System-Level Event Trees. Generic functional event trees are event trees at a functional level that apply to any transient. Generic system-level event trees are trees at the mitigating systems level that form the basis for the event tree models for each specific transient initiating event (see Forester, et al., for a more detailed discussion of generic event trees). Specific system-level event trees are the event trees that model the mitigating systems response to each of the 34 specific initiating events. For each transient initiating event, the specific system-level event trees were produced from the generic system-level event trees. For LOCA initiating events, the specific system-level event trees were created without using any generic trees.

The event tree modification effort produced more than 200 unique event trees. These event trees incorporated 81 different initial conditions, numerous operator action and/or decision points, and many potential alternatives for restoring a lost function. While this effort resulted in event trees that are more realistic from the point of view of containing many of the possible operator responses to any given initiating event, it did significantly increase the complexity and the quantity of the accident sequences that had to be quantified.

HUMAN RELIABILITY ANALYSIS

This section outlines the Human Reliability Analysis (HRA) performed for the detailed study of POS 5. The general methodology used for identifying and quantifying human actions was the Accident Sequence Evaluation Program Human Reliability Analysis Procedure (ASEP HRAP) (Swain, 1987). This methodology was selected because (1) the human error probabilities (HEPs) obtained using the procedure are considered to be conservative relative to those that would be obtained from other methodologies, (2) it was used in the full power PRA performed at Grand Gulf (Drouin, et.al, 1989) as part of NUREG-1150, (3) it is straightforward to use and can produce internally consistent HEPs that appear to reflect at least the relative potential of human failures in the nuclear power plant environment, and (4) it allows for straightforward adjustments in HEPs as a function of the results from interviews with plant personnel.

In general, the HRA data collection and analysis process outlined in the ASEP HRAP was followed except for pre-accident human actions. The pre-accident human actions included in this analysis used the same HEP values that were used in the full power PRA of Grand Gulf. For the post-accident human actions, interviews with operators and other

plant personnel were conducted over a two-day period. While all the scenarios analyzed in the PRA could not be covered in the interviews, most of the critical operator actions analyzed were discussed with the operators. Follow-up information and additional data regarding operator actions and time requirements for specific operator actions (both inside and outside the control room), were obtained through telephone conversations with relevant plant personnel.

In deriving the HEPs, the ASEP HRAP was closely adhered to. Basic human error probabilities for each human action event were determined with the methodology and adjusted according the rules for applying the Performance Shaping Factors (PSFs). Deviations from the prescribed methodology were taken only when it was felt that the LP&S environment created a situation that was not well fitted to the ASEP HRAP, and each deviation was documented.

RESULTS AND DISCUSSION

As a result of having 34 initiating events and highly complex event trees, literally tens of thousands of accident sequences required quantification. Using a truncation value of 1E-8 for each cut set of a sequence, 4,309 sequences survived the initial screening pass. The cut sets of these sequences were then examined for validity and appropriate recovery actions were applied. After recovery, 192 sequences containing 296 cut sets survived. The total point estimate for POS 5 during a refueling outage at Grand Gulf for the 192 sequences is 3.476E-05.

Two classes of initiating events dominate the preliminary results. As can be seen in Table 2, loss of support systems (SUPPORT) (e.g., DC Power Bus B, Component Cooling Water, and Instrument Air) and loss of offsite power (LOSP) constitute approximately 74% of the total core damage frequency.

Table 2. Contribution to Point Estimate by Class of Initiating Event

IE Class	Point Estimate Core Damage Frequency	% Contribution To Core Damage Frequency
SUPPORT	1.8E-5	51.9%
LOSP	7.7E-6	22.0%
DHR	3.2E-6	9.3%
OTHER	3.2E-6	9.1%
LOCAs	2.6E-6	7.6%

Figure 1 provides a graphical representation of the percent contribution of the dominant initiating events to the point estimate core damage frequency. The figure indicates that four initiating events (i.e., T1: Loss of Offsite Power, TDB: Loss of DC Power Bus B, T5D: Loss of Component Cooling Water, and TIA: Loss of Instrument Air) contribute approximately 67% of the total core damage frequency. For the important initiating events identified in Figure 1, three general types of accident sequences are among those that are dominant. These are (1) Blackout - Loss of all AC power, (2) Steam at

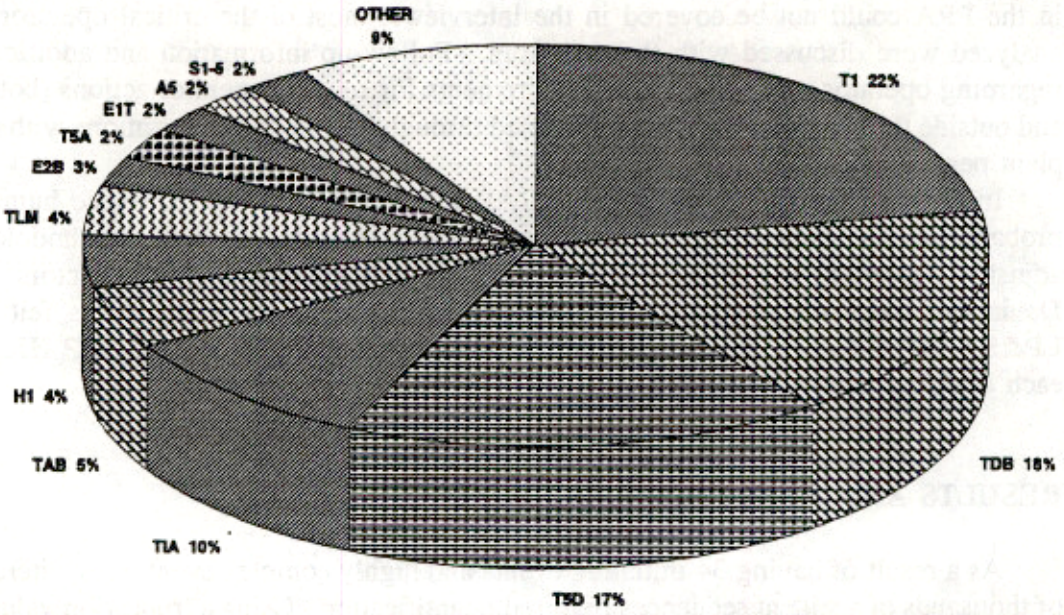


Figure 1. Percentage Contribution to Point Estimate Core Damage Frequency by Initiating Event

High Pressure - Steaming at high pressure with no makeup available, and (3) Flooding Containment - Injection of water into vessel, out the SRVs to the suppression pool, and finally out the open lower containment personnel lock.

The preliminary results from this study appear to indicate that the core damage frequency associated with operating in POS 5 is comparable to that obtained for a full power analysis. However, the reader is cautioned about placing too much emphasis on the "bottom line numbers" resulting from this study. As with any new work in research, the most important insights come from what is learned by doing the work. For example, the HRA process, which even for a full power analysis is complex, is even more so for a study such as this one. Thus, improved methods in the HRA field might provide additional insights from future studies of low power and shutdown conditions. One other example of an insight gained from this work is the increased importance of the time-dependent nature of the progression of an accident sequence and the operators' response to the unfolding events. While it is possible to incorporate some aspects of time-dependence into the event tree structure, a more robust and practical method should be developed. Such a method would also probably require improvements in the HRA methods currently used.

REFERENCES

D.W. Whitehead et al. "BWR Low Power and Shutdown Accident Sequence Frequencies Project, Phase 1 - Course Screening Analysis," Vol. 1, Draft Letter Report, Sandia National Laboratories, November 23, 1991 update, copies available at the NRC Public Document Room.

D.W. Whitehead et al. "Level 1 Risk Analysis of Low Power and Shutdown Operations at a BWR: Phase 2 Results," SAND92-1839C, presented at the 20th Water Reactor Safety Information Meeting, October 1992.

J. A. Forester et al., "Generic Event Trees and the Treatment of Dependencies and Non-Proceduralized Actions in a Low Power and Shutdown Probabilistic Risk Assessment, this volume.

A.D. Swain. "Accident Sequence Evaluation Program Human Reliability Analysis Procedure," NUREG/CR-4772, February 1987.

M.T. Drouin et al. "Analysis of Core Damage Frequency: Grand Gulf, Unit 1 Internal Events," NUREG/CR-4550, SAND86-2084, Vol. 6, Rev. 1, Part 1, September 1989.

COMPARATIVE EVALUATION OF THE PERFORMANCE OF ADVANCED SEMI-PASSIVE PRESSURIZED WATER REACTOR IN THE LOSS OF OFFSITE POWER EVENT

Meng Ouyang and Michael Golay

Department of Nuclear Engineering
Massachusetts Institute of Technology
Cambridge, MA 02139

INTRODUCTION

In the work reported here, a conventional PWR power plant is compared with a newly proposed advanced passive PWR design on their performances in response to LOOP event within the internal initiating event group. The purpose of the work is to evaluate the approaches and clarify the results achieved through simplification of advanced passive PWR design with respect to the two ways of reducing human errors. It builds upon the prior work of Golay(1988) and Seong(1988).

LOOP EVENT SEQUENCES COMPARISONS

The reason for examining LOOP event sequences for comparison of two design is that they are risk-dominant in conventional PWR plants. PLG(1983) shows that in a conventional PWR plant, LOOP sequences constitute about 39% of the expected CDF of internal initiating events.

Responses to LOOP Event in Two Stations

Conventional PWR Station

In conventional PWR station, the LOOP event modeling includes top events involving specific component (train) failures of four important accident mitigating systems, as well as failure of one group of important operator recovery actions. The systems are (1) Diesel Generators, (2) Service Water System, (3) Primary Component Cooling Water System, and (4) Emergency Feedwater System; and the group of operator Electric Power Recovery (ER) actions.

The risk dominant LOOP sequences show that most important failure phenomenon is the reactor coolant pump seal LOCA. Station blackout and failure to recover the power supply within a limited time period will definitely lead to SWS and PCCWS failures and then to the failure of reactor coolant pump thermal barrier. The operator recovery actions are assessed to be most important in mitigating LOOP event sequences.

Advanced Semi-Passive PWR Station

The major plant design features which are different from those of the conventional PWR plant are:

1) A simplified Reactor Coolant System (RCS) design, with the features of
   - Hermetically sealed canned motor pumps mounted in the inverted position
   - Close-coupled steam generator and pump configuration

   These features help eliminate the RCP seal LOCA events which can occur

easily in the conventional PWR plant.

2) Simplified, highly redundant non-safety-related systems for normal operations and for first response to initiating accidents, except for LOOP event; While in LOOP event, passive safety systems are mainly relied upon to mitigate the event consequences and thus reduce the importance of operator recovery actions and the importance of supporting systems.
3) Passive Safety Systems which include:
   - Passive Core Cooling System (PXS) which itself consists of (1) Passive Residual Heat Removal Heat Exchangers (PRHR HXs), (2) Core Makeup Tanks (CMTs), (3) four stages of valves of Automatic Depressurization System (ADS), (4) In-containment Refueling Water Storage Tanks (IRWSTs), and (5) Coolant Accumulators
   - Passive Containment Cooling System (PCS)

In the LOOP event, this plant may response differently according to how the LOOP event occurs and proceeds. Table 1 summarize these cases. This plant has a full load rejection capability that permits, upon loss of the offsite power supply, running the reactor and main turbine generator at a reduced power level sufficient to supply power to the plant auxiliary loads. Restoration of the grid within one-half hour makes the LOOP event similar to a loss of feedwater flow coincident with a loss of reactor coolant system flow event. However if the grid is not recovered and batteries are available, the plant must rely upon passive systems, or systems powered by diesel generators. The loss of offsite power (LOOP) coincident with the common cause failure of the batteries, leads to a station blackout event in response to which only the passive residual heat removal system is actuated. In the event of failure to restore the grid within 24 hours, it is conservatively considered that the diesel generators have also failed. The limited battery capacity (24 hours for safety systems and 2 hours for non-safety systems) does not influence the system operation since all safety systems are actuated within the 24 hours period, and do not need a power supply after that time.

There are two major offsite power recovery human action, namely within 30 minutes OTH-R05 and within 24 hours OTH-R24. In the semi-passive plant, recovery of offsite power may be treated as having equivalent difficulties as in the conventional PWR plant.

Risk Importance Ranking of Systems and Human Actions

The LOOP event sequences are studied in detailed in order to identify three importance measures of LOOP event sequence core damage frequencies:

The Fussell-Veselly (FV), Risk-Reduction (RR) and Risk-Increase (RI) importance measures for LOOP events area presented in Figures 1 to Figure 6.

Comparing the F-V importance ranking results between two designs (Figures 1 and 2), we see that for advanced passive PWR design, LOOP event mitigation relies more upon several passive safety systems while conventional PWR LOOP event mitigation relies most upon power recovery actions and the diesel generator system. In other words, the advanced semi-passive design scatters the mitigation function among several passive systems rather than relying upon recovery action and diesel generator system success solely as in conventional design.

Comparing the R-R importance ranking results between two designs (Figures 3 and 4), we see that for advanced passive PWR design, improvements in recovery actions and in the availability of four systems (SFWP, PRHR, CMT and ADS) are critical in reducing risk; while for the conventional design, only improvements in the diesel generator system availability and in two kind of power recovery actions (ER1, ER5) are critical in reducing risk. This means that if an effort is taken to improve the performance of the plants, the advanced passive PWR design will more effectively reduce the magnitude of the LOOP CDF than will the conventional PWR design.

Comparing the R-I importance ranking results between two designs (Figures 5 and 6), the trend observed is just opposite to the trend in R-R importance ranking. When the degradation of performance of plants occurs, advanced semi-passive design has fewer systems than with conventional design which are critical to the increases of LOOP CDFs. The Figures show that CMT system is the most critical for LOOP CDF increase in advanced semi-passive design, while primary component cooling water system is the most critical to LOOP CDF increase in the conventional design.

CONCLUSIONS AND REMARKS

Results obtained from this work are the followings:

- The total internal LOOP event sequences Core Damage Frequencies of the two plants differ by two orders of magnitude, with that of the advanced passive PWR being smaller and negligibly significant in the total plant Core Damage Frequency
- The Reactor Coolant Pump Seal LOCA event is completely eliminated from the advanced design (That event is typical in conventional PWR designs)

- The numbers of typical mechanical components are few in the advanced design, but their individual failure rates are similar with those in conventional PWR plant
- The number of required human actions and failure rates modeled with the advanced design are compared with those of conventional plant, with former being simpler, less stressful and more successful.
- The system risk importance rankings are graphically presented. For the three importance measures, the advanced passive PWR design is always better than conventional PWR design.

Table 1. Advanced Semi-Passive PWR LOOP Cases and Corresponding Systems

LOOP CASES	CORRESPONDING SYSTEMS
1. LOOP + Battery Failure = Station Blackout	Only the Passive Residual Heat Removal System is available and actuated.
2. LOOP Occurs, but the grid is restored within 0.5 hour	There is no need for diesel generator actuation. All safety and and non-safety systems are available; This has the same effect as the Lossf of Feedwater flow combined with Loss of Reator Coolant System flow;
3. LOOP occurs, grid is not recovered within 0.5 hours, but is recovered within 24 hours	Onsite power has to be provided (by diesel generators and /or batteries)to support non-safety system and safety -related systems
4. LOOP occurs and grid is not recovered within 24 hours.	One Conservatively takes the Diesel Generators as having failed. Only batteries support non-safety related system for 2 hours and/or safety-related systems for 24 hours.

Figure 1. Advanced Passive PWR design Fussell-Vesely System Importance Ranking In terms of Contribution to LOOP Event CDF (in units of %)

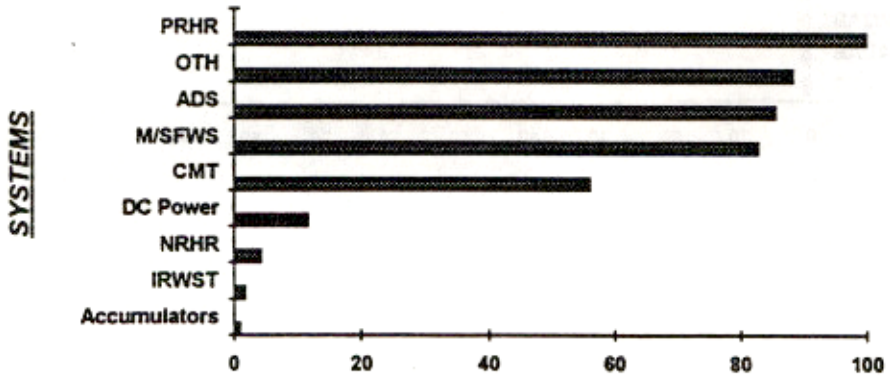


Figure 2. Conventional PWR Plant Fussell-Vesely System Importance Ranking
In terms of contribution to LOOP Event CDF (in unit of %)

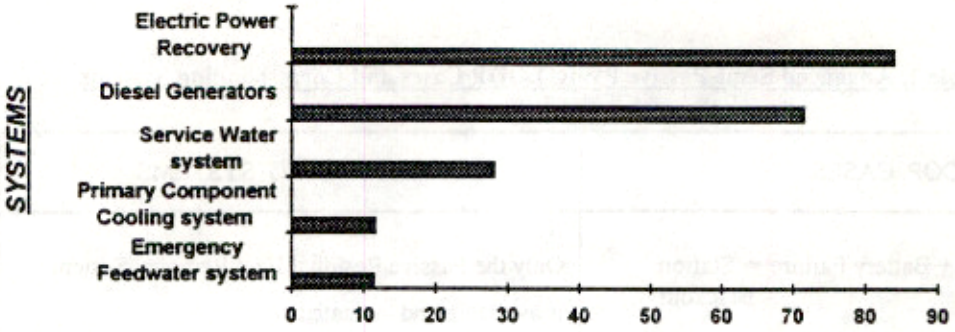


Figure 3 Advanced Passive PWR Design Risk Reduction Importance Ranking of Events in Units of % of LOOP Event CDF

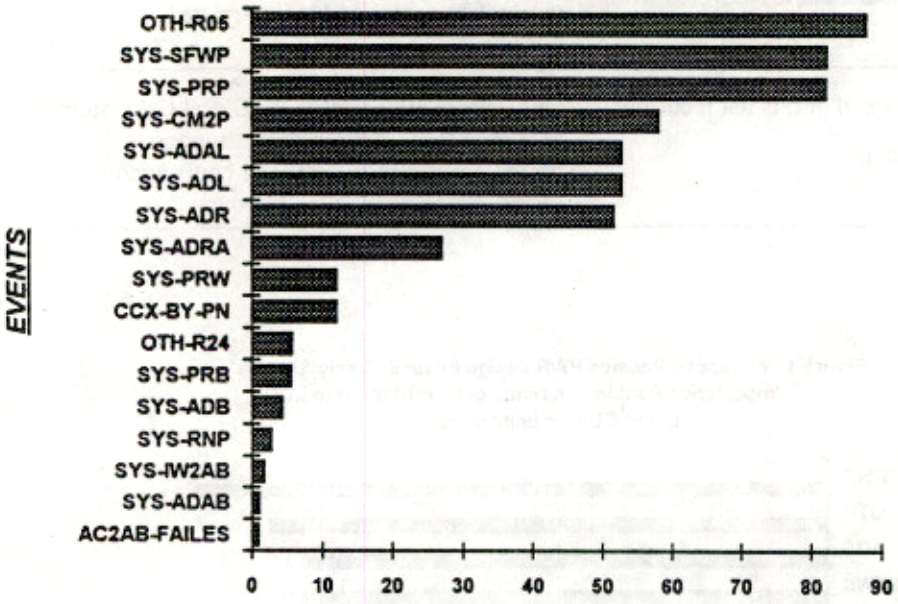


Figure 4 Conventional PWR Plant Risk Reduction Importance Ranking of Events in Units of % of LOOP Event CDF

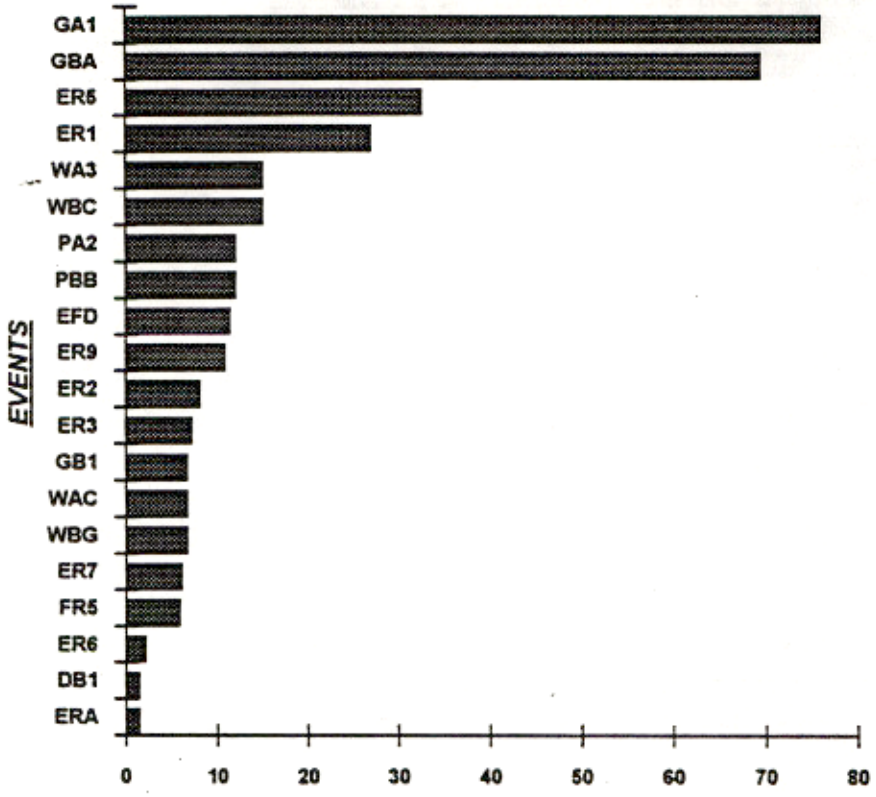


Figure 5. Advanced Passive PWR Design Risk Increase Importance Ranking of Events in Units of Risk Increase multiplier to LOOP Event CDF

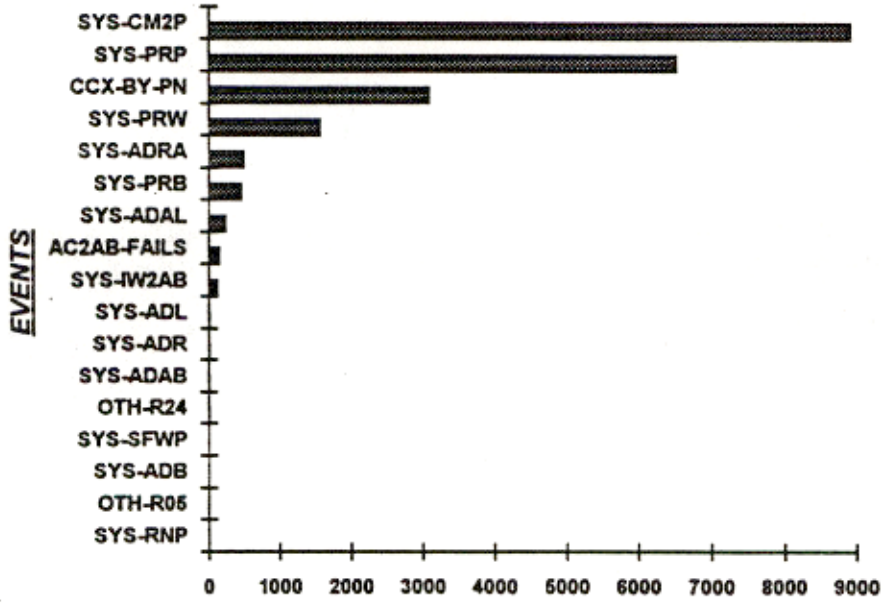
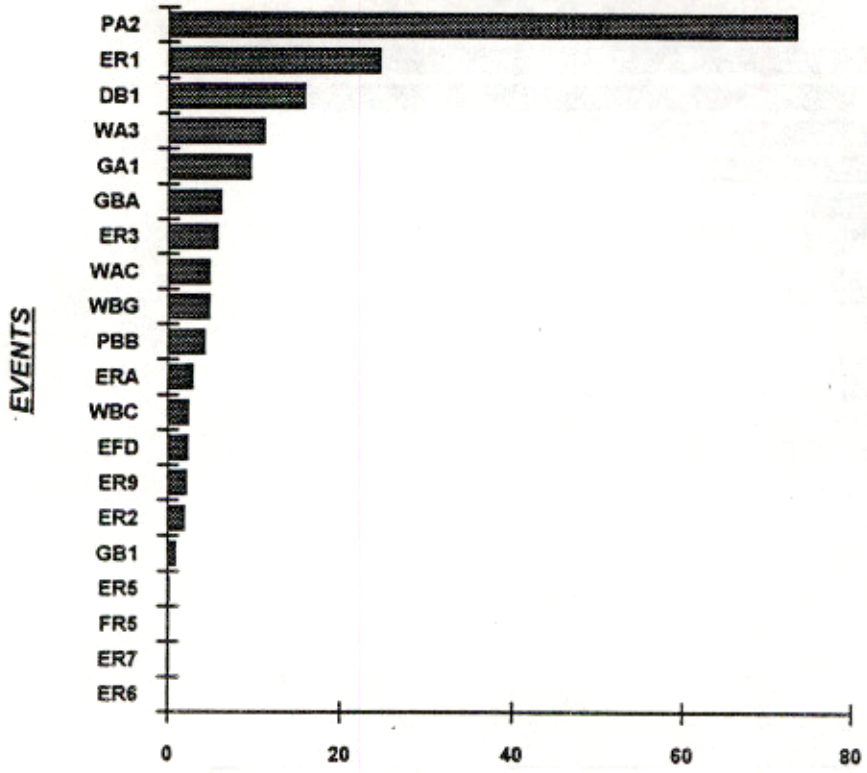


Figure 6. Conventional PWR Plant Risk Increase Importance Ranking of Events in Unit of % of LOOP Event CDF



ACKNOWLEDGMENT

This study was supported by the Advanced Light Water Reactor program of the Electric Power Research Institute.

REFERENCES

Michael W. Golay, V.P. Manno, and C.Vlahoplus, Jr., Nonprescriptive nuclear safety regulation: the example of loss of offsite power, *Nuclear Safety, Vol.* 29, *No.* 1. January-March 1988

Poong H. Seong, Vincent P. Manno and Michael W. Golay, Application of a power plant simplification methodology: the example of the condensate feedwater system, *Nuclear Engineering and Design*, 110 (1988)33-46

PLG "Seabrook Station Probabilistic Safety Assessment" Main report Volume 1-6, December, 1983

Westinghouse Electric Corporation "AP 600 Probabilistic Risk Assessment", DE-AC03-90SF18495, June 26, 1992 Volume 1-3

071 Accident Precurssors

Chair: V.M. Bier, U. Wisc.-Madison

Accident Sequence Precursor Methodology for Maritime Safety
L.H.J. Goossens (Delft U. Technol.); C.C. Glansdorp (Marine Analytics)

Developing Organizational Safety Information System for Monitoring Potential Dangers
M. Tamuz (U. Texas)

The Performance of Precursor-Based Estimator
V.M. Bier, W. Yi (U. Wisc.)

ACCIDENT SEQUENCE PRECURSOR METHODOLOGY FOR MARITIME SAFETY

Louis H.J. Goossens[1] and Cees C. Glansdorp[2]

[1] Safety Science Group, Delft University of Technology, Delft, the Netherlands
[2] Marine Analytics, Maritime Consultancy Bureau, Rotterdam, the Netherlands

INTRODUCTION

Maritime safety has always been given high priority in the Netherlands. This led to the use of various types of marine safety systems like markings, piloted waterways, patrol boats and so on. It also led to the development of traffic guidance systems like VTS (vessel traffic services). Policy evaluations indicate that advanced, newly developed safety systems are added on to the existing system configurations without sufficient in-depth knowledge of the added-on effectiveness of the new systems in relation to already existing systems.

From an inventory of marine safety highlights an in-depth study of incidents in relation to safety systems appeared to be highly relevant in order to contribute to the cost-effectiveness of (combined) marine safety systems, given an accepted level of safety. A basic element in such a study was the analysis of ship incident data, for which the Accident Sequence Precursor Methodology (ASP Method) has in principle a suitable framework. This paper presents provisional results from the ongoing research program EURET 1.3 (Commission of the European Communities) preceeded by a pilot study, sponsored by the Dutch Ministry of Traffic and Waterways (Goossens et al, 1992).

ACCIDENT SEQUENCE PRECURSOR METHODOLOGY - GENERAL ASPECTS

The ASP Methodology originated at the Oak Ridge National Laboratory shortly after the accident at Three Mile Island as a part of the Accident Sequence Precursor program for light water nuclear reactors in the United States (Minarick, 1990). As a general conclusion, the precursor program has made a very significant contribution to the field of risk assessment.

The Dutch Ministry of Housing, Physical Planning and Environment has sponsored, taking its responsibility for developing risk management methodologies, a research program into the feasibility of the ASP Method for the chemical process industries in the Netherlands (Cooke et al, 1987 and 1990).

Whereas the traditional probabilistic risk analysis methodology was developed to predict risks in systems for which no or scarce operational experience was available, the ASP methodology is specifically designed for situations for which operating experience is available, though insufficient to allow statistical predictions of serious accident probabilities. This is typically the case in the chemical process (and nuclear) sector.

The ASP modelling is backed by generic event trees describing typical accident sequences at an aggregated level such that sufficient data might be available for the probability calculations. A precursor is formally defined as:

- Either an initiating event, perhaps accompanied by one or more system failures; and
- An unavailability in one or more safety systems, without an occurrence of an initiating event.

For marine systems, one might apply the ASP method for predicting accident probabilities of severe release accidents of hazardous materials from ships based on precursor data of small scale incidents (traffic and ship's integrity incident data), no major hazards accident having occurred yet. This attempt will not be further elaborated in the forthcoming sections. The focus in here lies on the evaluation of safety systems' performances.

Apart from using generic event trees the ASP Method calculations of the overall accident probability is done by summing conditional probabilities for accidents, given accident sequence precursors. Otherwise stated, the overall accident probability (known from incident registrations) reveals the accident sequence precursor data contributing to the occurred accidents. This attempt is explored by defining generic accident sequence paths within the generic event tree. Each path consists of a subset of relevant precursors, indicative for the safety systems' effectiveness.

ROLE OF THE NAVIGATOR

It appears that these sequence paths are heavily centered around the actions of the navigator (with or without pilotage). This intrinsic modelling dependency leads to complications in the modelling of the sequences of the individual sub-events in the event tree. In order to provide a sensible set of generic sub-events two modelling assumptions are made:

1. The task cycle of the navigator consists always of three distinct phases: observing, processing information, and decision making on course and speed. Although in practice a complete task cycle may be executed within a very short period, for sake of modelling it is felt to be necessary to distinguish between these three subtasks. One additional reason is that the three subtasks are differently influenced by stimuli which the navigator receives either from the ship (signals or alerts from crew members), or from the surroundings (visual cues, alerting signals or VTS-information).

2. Distinction is made between actions which determine course and speed of the ship and actions from the surroundings, which may influence the course and speed of the ship eventually. This is of importance, because in principle the navigator sails the ship based on his skills and knowledge about the subsystem (ship), sustained by information from outside the subsystem. In this respect active and passive aids-to-navigation are designed.

APPROACH TO MARITIME SAFETY

The following constraints are present in maritime situations:

- complex maritime routes (like estuaries of big rivers and harbour entrances, or dense sea routes like the English Channel and parts of the North Sea) are considered
- course and speed of the ships are important factors in determining the cost-effectiveness of ship manoeuvres and the risk-benefit-ratio
- the expected behaviour of the navigator with respect to course and speed and his response to deviations determine the failures and recovery actions made by the navigator
- the expected behaviour of the ship (mechanical failures) and its surrounding elements (other ships and fixed objects) determine the physical constraints
- passive safety measures have a function in increasing the attention level of the navigator and to provide information on the conditions of the maritime route
- active safety measures also have a function in increasing the attention level of the navigator and to provide information on the conditions of the maritime route, plus they have the ability of providing concrete time-constrained advices and communications.

As indicated above, the incidents are identified as specific paths in the generic event tree, shown in figure 1. The basic elements (branches of the generic event tree: see Goossens et al, 1992) are

- the identification of the initiating event
- the behaviour of the navigator in terms of establishing the right or wrong set of observables with respect to the critical situation prior to the incident
- the corrective effect of the safety systems with respect to the set of observables of the navigator (he should adjust the set)
- whether or not the point-of-no-return has been reached (all incidents are a result of having no ability of prevention after the point-of-no-return has been reached, otherwise there would have been no incident at all).

In figure 1 only the identified deviations are taken into account. If an identified deviation occurs, an initiating event (E) is present. The next step is whether the navigator is able to determine the right set of observables: if wrong then N, if right then N'. The B between brackets means that the environment has an ability to respond to the navigator by means of active safety systems; this part is relevant for the analysis indicated by B.

The next step is whether or not an active safety system is able to correct or divert the navigator's informational and processing procedure and avoids a critical situation: if wrong (or present) then C, if right then C'. Whether or not an incident can be prevented eventually, depends on the point-of-no-return and the abilities of the navigator: if not prevented then I, if prevented then I'.

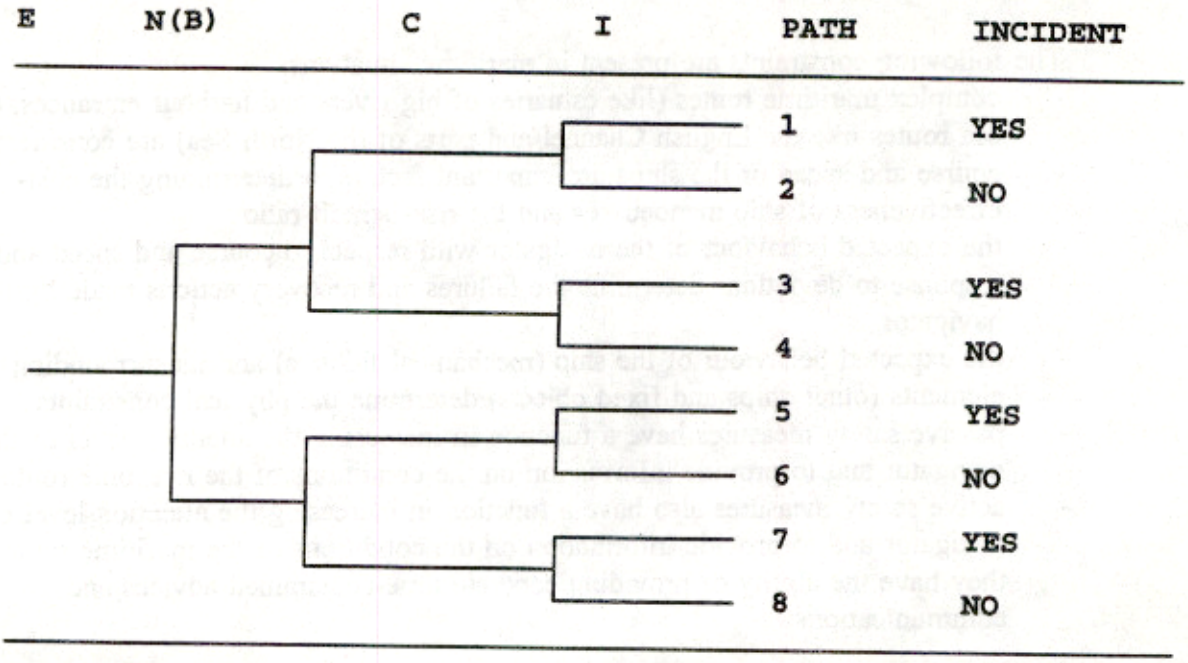


Figure 1. Generic event tree representing navigation through safeguarded waterways for identified deviations only

So any incident path and no-incident path can be represented by E, N, C and I. For instance, incident path 1 (see figure 1) is represented by:

$$p(\text{inc path } 1) = p(ENCI)$$

which can be written as,

$$p1 = p(I|ENC).p(C|EN).p(N|E).p(E)$$

in which p1 = p(inc path 1). It should be noted that the other three incident paths (path 3, path 5 and path 7) have a similar representation as p1, by which N and C are alternately used as N' and C'. By marginalizing the conditional probabilities on the right hand side, one can rewrite p1 into

$$p1 = \Phi_E \cdot \Phi_B \cdot \Phi_C \cdot \Phi_I$$

PRELIMINARY RESULTS

Let us now consider the most important terms for this analysis: Φ_B and Φ_C.

The term p(N|E) represents the navigators' ability to assess a set of observables, as defined in terms of situation images. If the navigator is able to do so, we talk about non-complex situations indicated by p(N'|E), if the navigator is unable to assess the right set, we talk about complex situations: p(N|E). For practical use, p(N|E) may be represented by the prior safety system parameter Φ_B. However, this parameter only represents the communication between safety systems and the navigator. This communication may be anything and Φ_B may depend on the type of initiating event.

For the analysis the Western Scheldt data analyzed in Goossens et al (1992) over the period 1985-1988 have been used. In 1986 the VTS was built in parts of the Western Scheldt. Φ_B represents complex situations shown via incident paths 1 and 3, Φ_B' non-complex situations shown via incident paths 5 and 7. Their ratio indicates the influence of the new VTS. It is all based on the assumption that the non-incident paths (2 + 4) and (6 + 8) behave comparably.

Table 1. Comparison of complex and non-complex situations

year(s)	Φ_B / Φ_B' collisions	strandings
1985-1988	6	5.2
1985-1986	4.8	3.9
1987-1988	12	8.3

Overall Table 1 is shown that complex situations give six times more incidents than non-complex situations. Given the fact that non-complex situations occur much more often (e.g., 100 times more), complex situations have a probability of occurrence 600 times (say, between 100 and 1000 times) more likely than non-complex situations.

The term p(C¦EN) represents the avoidance capabilities once a critical situation has occurred. The effect is solely determined by active safety systems (communication from outside the ship). If there is no communication possible (VHF, VTS or whatever) the p(C¦EN) = 1 and p(C'¦EN) = 0. The same is true for p(C¦EN') = 1 and p(C'¦EN') = 0. For practical use, p(C¦EN) may be represented by the avoidance parameter Φ_C, which indicates the direct effect of active safety systems (VTS) and the like.

From the pilot study data (Goossens et al, 1992) the following calculations can be made for the ratio of incidents for complex situations (VTS-influence thru path 1) / (no-VTS-influence thru path 3); see Table 2.

Table 2. Comparison of possible VTS-influence on maritime incidents

ratio in year (path 1)/(path 3)	collisions	strandings
1985-1986 (no VTS)	0.6	8
1987-1988 (with VTS)	0.4	7.3

These results indicate some influence of VTS on collisions and almost no influence at all on strandings.

CONCLUSIONS

The Accident Sequence Precursor Methodology has been shown to provide insight into the mechanisms under which ship incidents in densely navigated waterways may occur. The defined incident sequence paths lead to typical representations of the navigator's behaviour and the response of active safety systems like VTS to that. Preliminary data show the limited effects of such active safety systems.

ACKNOWLEDGEMENTS

This paper reports preliminary results under the CEC research program EURET 1.3 and from a pilot study sponsored by the Ministry of Traffic and Waterways in the Netherlands [Marine Analytics report MA-1250].

REFERENCES

Cooke, R. and L. Goossens, "The accident sequence precursor methodology for the European Post-Seveso era", Reliability Engineering & System Safety, 1990, Vol.27, No.1, p.117-130

Cooke, R.M., L.H.J. Goossens, A.R. Hale and J. van der Horst, "Accident Sequence Precursor methodology - A feasibility study for the chemical process industries", Report for the Dutch Ministry of Environment, TUDelft/TNO Apeldoorn, March 1987

Goossens, L.H.J., C.C. Glansdorp and C. Deutsch, "Pilot Accident Sequence Precursor Method - Western Scheldt", Project MA-125, Report MA-1250, Marine Analytics b.v., Rotterdam and Delft University of Technology, Delft, March 1992

Minarick, J.W., "The US NRC Accident Sequence Precursor Program: Present method and findings", Reliability Engineering & System Safety, 1990, Vol.27, No.1, p.23-51

DEVELOPING ORGANIZATIONAL SAFETY INFORMATION SYSTEMS FOR MONITORING POTENTIAL DANGERS

Michal Tamuz

University of Texas School of Public Health
1200 Herman Pressler Street, W230
Houston, Texas 77030

INTRODUCTION

Organizations have limited opportunities to learn from high risk environments, where disastrous accidents rarely occur and the threat of disastrous consequences hinders trial-and-error learning. This paper examines how such organizations supplement their sparse experience with accidents by collecting information about and learning from "near accidents," events which, under slightly different circumstances, could have resulted in an accident (NRC, 1980). The main argument is that an organization's interpretation of near accidents influences how it collects information about near accidents and thus, its capacity to learn from them. Specifically, how an organization interprets near accidents affects its choice of information-gathering methods which, in turn, affects the quantity and contents of the information it collects.

This paper is based on organizational research on aviation safety information systems. It also draws on the preliminary findings of an organizational study that compares how potential dangers are monitored in the nuclear power, space shuttle and air transportation industries. The research builds on interviews with key participants in safety-monitoring systems, analyzes reports generated within these systems, and examines the organizational procedures for learning from these incident reports. To illustrate its arguments, the paper uses examples from studies of five aviation safety information systems (SIS) that collect information about potential dangers.

Organizational Learning

Guided by their own experience and the experience of others, organizations learn by constructing and changing their rules of behavior. Organizational learning is defined as "...learning by encoding inferences from history into routines that guide behavior" (Levitt & March, 1988:320). This perspective focuses on learning as a process rather than as an outcome. Thus, organizational learning does not necessarily result in beneficial outcomes.

Organizational learning is hampered when organizations have little past experience to guide their actions. For example, an airline rarely has fatal accidents. Yet it wants to learn from experience how to lessen the chances that such a catastrophe will occur. When experience is limited, the usual methods of learning from experience, imitating others, or making rational cost-benefit calculations prove ineffective (Tamuz, 1987). To enrich their actual experience with catastrophic accidents, organizations gather information about near accidents to simulate "near histories". In constructing a "near history," an organization not only generates information about the specific events which occurred but also about the possibility that other events may have occurred under somewhat different circumstances (March, Sproull & Tamuz, 1991).

Interpreting Near Accidents

Organizations vary in how they define and interpret potential dangers. Some organizations interpret near accidents as an opportunity for learning. They seek to prevent accidents by discovering their precursors. In these organizations, people assume that under slightly different circumstances, the conditions leading to a near accident could also result in an accident. They consider a near accident almost as if it had been an accident in that the organization investigates why it occurred.

Organizations concentrate on learning because of the importance of discovering accident precursors and the difficulty in obtaining information about them. First, accidents rarely occur, their causes are complex, and it is difficult to untangle causal factors obscured by accident investigations and lawsuits (Tasca, 1989). Second, it is hard to weigh the value of near accident information against the usefulness of maintaining accountability. Its value may only become apparent in hindsight, for example, as an accident spurs investigators to seek near accidents that occurred under similar conditions.

Other organizations interpret near accidents as a means of controlling employee behavior. By deterring individuals from engaging in behavior associated with near accidents, they attempt to avert real accidents. They treat a near accident almost as if it had been an accident in that the organization penalizes the individuals held responsible. While these organizations focus on curbing potentially hazardous behavior, they may be charged with accomplishing a dual mission: not only to prevent errors but also to learn from them. All too often these organizations use the same source of information to accomplish their dual mission.

Organizations face inherent problems in using the same information as a means of policing behavior and as a source of learning (Tamuz, 1988). The incentives that organizations use to control individual behavior can hinder their efforts to learn. By offering rewards for desired behaviors and penalties for undesired ones, organizations also give individuals incentives to withhold information that may reflect poorly on their performance (Lawler & Rhode, 1974). When individuals withhold negative information from the organization, they can still learn from their own mishaps, but they reduce the organization's capacity for learning.

Effects of Near Accident Interpretations

It is critical to understand whether an organization collects information solely for learning or primarily for policing behavior because this choice influences the methods by which an SIS gathers information which, in turn, affect the quantity and quality of information it collects. The safety information systems and other organizations in the industry depend on this near accident information to learn how to avert accidents. Although learning-centered and dual-mission safety information systems both analyze near accident data to learn about accident precursors, they differ in the source of information they use. The following discussion illustrates how this process unfolds in safety information systems as they allocate incentives, define potential dangers, set reporting requirements, and determine confidentiality standards.

INCENTIVES AND THE DUAL USE OF INFORMATION

Gathering Information Focused On Learning

When learning is the focus, Van der Schaaf advises organizations to use near accident information solely for study and analysis (Van der Schaaf, Lucas & Hale, 1991). To gather scarce information about near accidents, organizations, in particular, safety information systems do not penalize individuals for being involved in or reporting near accidents. For example, pilots report near accidents without fear of punishment to the Air Line Pilots Association. Similarly, in the nuclear power industry, no penalties are associated with reporting non-consequential events or potential problems to the Human Performance Enhancement System. Safety information systems may also provide immunity from punishment. The Aviation Safety Reporting System (ASRS), for example, offers pilots eligibility to obtain limited immunity from prosecution by the Federal Aviation Administration (FAA). For immunity-based incentives to be effective, individuals must be convinced that the organization will not use the information they provide to judge their performance. Consider the example cited by Lucas (Van der Schaaf, Lucas & Hale, 1991). When a space

shuttle contractor punished several employees, despite its policy to forgo punishment in exchange for hazard-related information, technicians lost confidence in the "forgiveness" policy and refrained from reporting hazardous incidents.

Organizations that use positive incentives (or withhold negative ones) encourage individuals 1) to report near accidents and 2) to provide unbiased information about them. By increasing the number of near accident reports filed, the SIS expands its database, and thus, augments its capacity for learning. By removing the incentives for individuals to cover up their mistakes, it lessens the self-interested bias that distorts information. For example, in excerpts of ASRS reports published anonymously in the newsletter "Callback," pilots give candid accounts of how their actions contributed to potentially dangerous situations.

Gathering Information Designed to Control Behavior

When deterrence is key, organizations often define involvement in a near accident as being in violation of company rules, punishable by penalties, public censure, and the obstruction of career advancement. Because a rule violation may be involved, those who provide information about such events are subject to punishment. Consequently, when the organization gives individuals incentives to refrain from engaging in safety violations, it simultaneously provides them with disincentives to disclose information about such potentially hazardous activities (Tamuz & Sitkin, 1993). Thus, if organizations use the same source of information for control and learning, the incentives that enable organizations to discipline individuals for their mistakes can impede organizational efforts to learn how to correct the system conditions contributing to these "operator errors". The following examples illustrate how disincentives for information disclosure led to the under reporting of near accidents in one SIS and how the punitive actions of one organization had unanticipated repercussions for the information gathering activities of another.

Under-Reporting Near Accidents. The Federal Aviation Administration (FAA) collects reports of near midair collisions (NMACs) to monitor and discipline pilot behavior. It also attempts to learn from the information it collects for disciplinary purposes. The FAA holds pilots who report NMACs subject to prosecution, giving them disincentives to provide information. The results of an FAA experiment suggest that these disincentives contribute to the under-reporting of NMACs. From 1968-1971, the FAA provided immunity from prosecution for pilots who reported NMACs. As expected, pilots increased the number of NMACs they reported. Indeed, the FAA's NMAC safety information system recorded a remarkable upsurge in the number of reports filed, increasing from 559 in 1965 up to 2,230 in 1968, 1,444 in 1969, 1,456 in 1970, and 1,350 in 1971. But when the FAA retracted its offer of immunity in 1972, the number of NMACs reported dropped to 231 (FAA, 1987). This experiment illustrates how the use of negative sanctions designed to alter pilot behavior also gave pilots incentives to refrain from information-reporting, resulting in the under-reporting of near midair collisions and a reduced pool of information from which to learn.

Repercussions for other SIS. When the FAA intensified its detection efforts in air traffic control centers, it had unanticipated repercussions on the Aviation Safety Reporting System. The ASRS, which offers eligibility for immunity from FAA prosecution, recorded a surge in pilot reports of possible air traffic violations. The dramatic increase occurred after the FAA began computerized surveillance of air traffic control centers, increasing the perceived likelihood that pilots would be prosecuted. It suggests that if individuals perceive that a regulatory agency is clamping down on offenders, by improved detection or tougher penalties, they have greater incentives to seek immunity and thus, to report near accidents to an SIS offering immunity. Surprisingly, the pilots increased the total number of incidents they reported to the ASRS, providing more information not only about incidents subject to prosecution but also about other safety-related incidents (Tamuz, 1987). This example illustrates how one organization's efforts to tighten its control of employee behavior by improving its information-gathering methods unexpectedly influenced another organization's capacity to gather information as a basis for learning.

Consider yet another repercussion. As the FAA intensified its efforts to detect and prosecute air traffic violations, it inadvertently created information overload for the ASRS. While pilot incident reports streamed into the Aviation Safety Reporting System, its level of funding remained unchanged. The information overload sorely taxed the resources of the ASRS, straining its capacity to process and analyze incident reports. As a result, the ASRS

had to modify its information processing routines, reducing the depth and detail of analysis for selected incidents.

Biased Information. Disciplinary investigations, designed to determine accountability, tend to result in biased information, ill-suited for learning. The adversarial nature of disciplinary investigations tends to generate biased information because each side presents evidence that supports its allegations. In particular, organizations give individuals strong incentives to withhold potentially damaging information. Consider the incentives the FAA uses to control computer-detected operational errors. An air traffic controller is subject to mandatory retraining and a blemished record; a pilot is subject to FAA prosecution for violating federal aviation regulations. Under these conditions, neither controllers nor pilots are likely to reveal self-incriminating information during disciplinary investigations. When organizations seek to learn from the biased information generated in disciplinary hearings, they may draw the wrong conclusions and choose a misguided course of corrective action. For example, as Perrow (1984) would suggest, they may attribute hazardous conditions to operator error rather than to system-based problems and thus, mistakenly attempt to improve operator training.

Undermining Trust in the System. While the FAA's use of negative sanctions limits its ability to collect valid information, its use of legalistic procedures to allocate sanctions undermines bonds of trust, further obstructing its information collection efforts. Paradoxically, the use of legalistic procedures to resolve differences between conflicting explanations of problems may have been adopted to foster trust by protecting the due process rights of the participants who are subject to FAA sanctions. However, as Sitkin & Roth (1993) would predict, the use of such legalistic procedures may indirectly undermine the trust relations between employees and the organization and between subgroups of employees. First, it shifts the arena of inquiry from learning what went wrong to establishing blame and allocating punishment. Second, the use of adversarial procedures positions the FAA investigators opposite the controllers and pilots. Third, by creating a zero-sum game between the controllers and pilots in which someone has to be held accountable, the use of legalistic procedures may erode trust in the FAA for pitting the groups against one another.

DEFINING POTENTIAL DANGERS

Definitions of potential dangers vary in specificity. In the air traffic industry, definitions of near accidents range from general statements about aviation safety to precisely-measured deviations from performance standards (Tamuz, 1988). SIS that gather near accident information solely for learning are not constrained in how they define a near accident; they are free to use general definitions of potential dangers. But if an organization considers a near accident as a violation of organizational rules, it must precisely define what behavior constitutes a violation. The following discussion explores how ambiguous definitions expand the capacity for organizational learning while precise definitions constrict it.

Definitions that Foster Learning

The enactment of an official organizational definitions legitimates the reporting of potentially dangerous incidents. It recognizes that "close calls" can yield substantial information, relevant to learning, even though they did not result in significant consequences. Consider an example from the nuclear power industry. The Human Performance Enhancement System (HPES), founded by the Institute for Nuclear Power Operations (INPO), has developed official definitions for "potential problems" and "non-consequential events" (INPO, 1990).

By defining potential dangers broadly, a safety information system fosters near accident reporting by appealing to professional norms. The use of broad, indeterminate definitions indicates respect for the professional judgment of information-reporters because it enables them to exercise discretion in reporting incidents. For example, both the Aviation Safety Reporting System and the SIS operated by the Air Line Pilots Association rely on the judgment of individual pilots who identify and report safety-related incidents. Consistent with professional norms, pilots report a variety of potential dangers to these systems even when they do not benefit directly from filing such reports (Tamuz, 1987).

The use of a broad near accident definition enables a SIS to gather information about relatively unknown potential hazards. General definitions encompass categories of well-known and readily-measured hazards (including deviations from performance standards), broad classifications of common mishaps (i.e., misinterpretation of oral instructions), and hitherto unknown potential dangers. By using an ambiguous definition, the SIS enhances its capacity to discover potentially dangerous circumstances that go beyond existing definitions.

Definitions Designed to Guide Behavior

Safety analysts apply measures of severity to the conditions which precipitate near accidents, rather than to their consequences. They differentiate precursor conditions by their potential for damage. To illustrate, consider how the FAA calibrates its definitions of near midair collisions for distance and ranks them by the need to take evasive action. The FAA defines a near midair collision as a "potential" hazard, if two aircraft come within 500 feet of one another and if one of the pilots takes evasive action to avoid a collision. If the two aircraft come within 100 feet of one another and a collision was avoided due to chance, the FAA defines the near midair collision as "critical" (FAA, 1987).

The FAA also differentiates its definitions by accountability. When two aircraft fail to maintain the prescribed distance from one another, the FAA defines the incident as a "loss of separation". It adds the definition "operational error", if it holds the controller responsible for the incident, for example, if the controller failed to instruct the pilot. But if the pilot strayed from her course and failed to heed (or hear) the controller's instructions, the FAA categorizes the same incident as a "pilot deviation" (FAA, 1987).

While an organization may rely on precisely-defined, accountability-based definitions to police unsafe actions, these definitions restrict organizational information gathering. They focus attention on clearly-defined problems that are also under the jurisdiction of the organization. The use of such definitions tends to define away other potentially dangerous situations as irrelevant. To illustrate, consider the FAA Quality Assurance Program for the automatic surveillance of air traffic control centers in which a computer detects every time two aircraft lose their allotted separation. If the aircraft are flying under the guidance of air traffic control (IFR), then the air traffic control center identifies the event as a possible controller operational error, thoroughly investigates the loss of separation incident, and keeps statistics on it. But even though the computer detects an event, if one or more of the aircraft are not flying under the jurisdiction of the air traffic control center (VFR), the incident does not "count" as a possible rule violation. Consequently, control center administrators do not study it or record the incident in their statistics (Tamuz & Sitkin, 1993).

REPORTING REQUIREMENTS AND CONFIDENTIALITY

A learning-centered safety information system may choose to enact voluntary reporting and this choice influences the information it collects. For example, the Aviation Safety Reporting System and the Air Line Pilots Association SIS enact voluntary reporting. Because individuals provide information voluntarily and the definition of a potentially dangerous situation depends on the perception of the information-reporter, the number of incidents recorded may reflect fluctuations in reporting behavior rather than changes in underlying safety conditions. The number of reports filed with the ASRS, for example, varied with publicity, politics and perceive reporting incentives (Tamuz, 1987).

Organizations tend to require the mandatory reporting of near accidents if these events constitute a rule violation. The choice of mandatory reporting influences information gathering, although not always as expected. Since many near accidents go unnoticed except by those directly involved, organizations find it difficult to enforce mandatory reporting. Consequently, even if an organization requires reporting, individuals can refrain from disclosing potential dangers and near accidents remain under-reported. (See NMAC reports in FAA, 1987). By using computerized surveillance, the FAA has eliminated the need to enforce mandatory reporting in air traffic control centers. But not all potentially hazardous situations can be readily-measured or detected by a computer.

A learning-centered SIS that protects confidentiality trades accountability for the capacity to gather information. By protecting the identity of individuals and organizations, it can gain the support of the industry and promote near accident reporting. To illustrate, the ASRS assures confidentiality by eliminating any information that could identify a particular flight,

airline, or aircraft manufacturer. Although these provisions preclude holding people and corporations accountable for unsafe conditions, they allow the ASRS to gather valuable information about near accidents, especially human factor data, that is difficult to obtain from other sources.

By interpreting near accidents as rule violations, an organization may limit its ability to protect the identity of information reporters, and thus, discourage individuals from reporting the violations it seeks to monitor. If an individual reports information that incriminates others, the organization is caught in the dilemma of protecting the identity of the reporter or upholding the due process rights of the accused (Tamuz, 1988). Further, guarantees of confidentiality may prove ineffective if organizational conditions enable supervisors or co-workers to deduce who reported a potentially hazardous situation. Consider the FAA Safety Hotline confidentiality guarantee. Because the hotline prosecutes violations of Federal Aviation Regulations, it must maintain a degree of specificity in its incident investigations which may enable employees to guess the identity of information providers.

SUMMARY

This paper explores how the dual use of information, for learning and controlling behavior, influences the capacity for organizational learning. The interpretation of near accidents as a learning opportunity or primarily as means of deterring unsafe behavior shapes how organizations collect information about near accidents. The organization's choice of information-gathering methods, in turn, affects the quantity and contents of the information it collects. Specifically, if an organization collects information for disciplinary purposes, it lessens the amount of near accident reporting and increases the self-interested bias that distorts accounts of potential dangers. Further research is needed to study these learning processes in depth and to determine whether they are specific to the air traffic industry or can be generalized to other industries.

ACKNOWLEDGMENTS

I am indebted to the National Science Foundation, Decision, Risk and Management Sciences Division, for funding the research on which this paper is based. My thanks to Sim B Sitkin for his helpful comments on the paper.

REFERENCES

Federal Aviation Administration, Office of Aviation Safety (1987). Unpublished statistics.

Institute of Nuclear Power Operations (1990). Human Performance Enhancement System. Atlanta: INPO 90-005.

Lawler, E. E., & Rhode, J. G. (1976). Information and Control in Organizations. Pacific Palisades, CA: Goodyear.

Levitt, B.& March, J.G. (1988). Organizational learning. *Annual Review of Sociology*. 14:319-40.

March, J. G., Sproull, L. S., & Tamuz, M. (1991). Learning from samples of one or fewer. *Organization Science*. 2,1:1-14.

National Research Council, Committee on FAA Airworthiness Certification Procedures. (1980). Improving Aircraft Safety: FAA Certification of Commercial Passenger Aircraft. Washington, D.C.: NAS.

Perrow, Charles (1984). Normal Accidents: Living with High-Risk Technologies. New York: Basic Books.

Sitkin, S. B. & Roth, N. L. (1993). Explaining the limited effectiveness of legalistic remedies for trust/distrust. *Organization Science*. 4,3:367-92.

Tamuz, M.(1988). Monitoring Dangers in the Air: Studies in Information and Ambiguity. Ph.D. Dissertation. Stanford University.

Tamuz, M. (1987). The impact of computer surveillance on air safety reporting. *Columbia Journal of World Business. 22,1:69-77.*

Tamuz, M. & Sitkin, S.B (1993). The invisible muzzle: Organizational and legal constraints on the public disclosure of information about health and safety hazards. Unpublished manuscript.

Tasca, Leo (1989). The Social Construction of Human Error. Ph.D. Dissertation. SUNY-Stony Brook.

Van der Schaaf, T.W., Lucas, D. A., & Hale, A.R., Eds. (1991). Near Miss Reporting as a Safety Tool. Oxford: Butterworth-Heinemann.

THE PERFORMANCE OF A PRECURSOR-BASED ESTIMATOR

Vicki M. Bier[1] and Woojune Yi[2]

[1]Department of Industrial Engineering
[2]Department of Nuclear Engineering and Engineering Physics
University of Wisconsin
Madison, WI 53706

INTRODUCTION

In estimating the frequency ϕ of a rare event (e.g., a core melt accident at a nuclear power plant), the number of observed events is usually too small to support the development of accurate estimates by means of the usual statistical estimator, $\phi^*=N/T$, where N is the number of observed events and T is the number of years of experience. This estimator is unbiased (i.e., its expectation is equal to the true accident frequency ϕ), and also consistent (i.e., it converges to ϕ as T gets large), but has extremely large variance for small T.

While we certainly want an estimator that is consistent, we also want it to be reasonably accurate in the short term, not only as T becomes large. The use of data on accident "precursors" or "near misses" has been suggested as an approach to achieve this goal. A number of precursor-based estimators have been suggested, and a comparison of their properties (particularly consistency and bias) is provided by Bier (1993). Unfortunately, each of these estimators has its weaknesses. Some are not even consistent, and all are biased. This paper discusses the performance of selected estimators in practice, based on the results of a simulation study.

BACKGROUND

Cooke and Goossens (1990) proposed a precursor-based estimator:

$$\phi_{CG} = \frac{1}{T}\max_j [N_j \hat{P}_j] \qquad (1)$$

where N_j is the number of precursors of type j occurring at any plant, $\hat{P}_j$ is the estimated probability that a precursor of type j will progress to an accident, and T is the number of plant years of experience. Cooke and Goossens suggested that ϕ_{CG} will tend to use data

from "precursors of maximal length in the data base"; i.e., the most severe observed precursors. This might be a desirable property, since severe precursors provide information on a large fraction of the overall accident process, thereby reducing the room for error in estimating conditional accident probabilities.

In fact, Bier (1993) proposed a modification to ϕ_{CG} that is less biased, and *always* uses data from the most severe observed precursor. This estimator simply takes the maximum over the P_j themselves (i.e., the conditional accident probabilities for all observed precursor types), instead of over the product N_jP_j:

$$\phi'_{CG} = \frac{N_i\hat{P}_i}{T}, \; i = \underset{j}{\operatorname{argmax}}(\hat{P}_j \mid N_j > 0) \qquad (2)$$

Essentially, the estimator ϕ'_{CG} looks only at the precursor of maximal severity (or "length"), while ϕ_{CG} selects the precursor that yields the largest estimate. Thus, $\phi_{CG} \geq \phi'_{CG}$. In addition, ϕ'_{CG} always reverts to ϕ^* as soon as an accident has been observed. Therefore, it is trivially consistent. However, ϕ'_{CG} is still biased high over finite time periods, since ϕ'_{CG} will equal ϕ^* when $\phi^* > 0$, but will frequently exceed ϕ^* when $\phi^* = 0$.

Despite its consistency and relatively small bias compared to ϕ_{CG}, ϕ'_{CG} can still yield large errors in some circumstances. In particular, the bias in ϕ'_{CG} will increase as the observation time decreases. To understand why, note that ϕ'_{CG} selects one precursor-based estimator from a pool of candidate estimators. However, a given precursor-based estimator will be included in this pool only if at least one precursor of the corresponding type has actually occurred. This means that the estimators from which ϕ'_{CG} is selected are all conditioned on the occurrence of at least one precursor. Therefore, the estimated frequency of precursors of type i will be biased high by an amount equal to $1/(1-e^{-\lambda_i T})$, where λ_i is the actual frequency with which precursors of type i occur.

For any given precursor type, the extent of bias will decrease rapidly with T, at least when T is small. Similarly, the bias also decreases with λ_i; in other words, the bias will be greater for rare precursors than for frequent precursors (especially when T is small). However, this is a potential problem. Since ϕ'_{CG} chooses the precursor with the highest conditional accident probability, it will tend to favor extremely rare precursors, for which the bias in the estimated precursor frequency will be greatest. Therefore, when T is small (e.g., substantially less than $1/\lambda_i$), ϕ'_{CG} can overestimate the actual accident frequency by a large margin.

Despite the potential for extremely large bias when T is small, ϕ'_{CG} may still be a desirable estimator. One reason is that we may actually wish to incur some conservative bias, especially given the catastrophic nature of the accidents being analyzed. Another factor supporting Cooke and Goossens' suggestion to "look only at precursors of maximal length in the data base" is the possibility that estimates of conditional accident probabilities may be biased in the non-conservative direction. In particular, as pointed out by Cooke and Goossens, system failures are likely to be associated in the sense of Barlow and Proschan (1975). However, this association may be difficult to model. If we instead incorrectly model subsequent system failures as if they are independent, then conditional accident probabilities will be systematically underestimated. Using data on the most severe precursors may reduce the extent of this bias by increasing the fraction of the accident process that is addressed statistically (rather than by PRA modeling), and hence reducing the room for error in estimating conditional accident probabilities.

However, as seen above, reducing non-conservative bias in estimated conditional accident probabilities by using data on more severe precursors will also increase the conservative bias in our estimates of precursor frequencies, possibly by a large margin. In fact, while errors in conditional accident probabilities will likely be minimized for the rarest (and hence most severe) precursors, overestimation of precursor frequencies is minimized

for the most frequent precursors. Thus, using data from severe precursors will tend to reduce the underestimation of conditional accident probabilities, but while increase the overestimation of estimated precursor frequencies. Moreover, these are precisely the precursors favored by the proposed estimator.

It is therefore not immediately clear that the approach reflected in ϕ'_{CG} (i.e., selecting the most severe observed precursors) is necessarily the best for use in practice. In fact, Bier (1993) even speculated that it might be helpful to modify ϕ'_{CG} to avoid use of extremely rare precursors, for which the bias in estimated precursor frequency will be greatest. In this paper, the performance of ϕ'_{CG} is assessed in a simulation study, to determine whether such a modification is warranted.

A MEASURE OF ERROR

In assessing the performance of precursor-based estimators, the first problem is to decide on an appropriate measure of error. Mean squared error would appear to be an obvious choice, and has the advantage that for many estimators it is analytically tractable. However, on further consideration, mean squared error does not turn out to be a very helpful measure for assessing the performance of precursor-based estimators. Given the sparsity of precursor data, and also the substantial uncertainties entailed in estimating conditional accident probabilities, order of magnitude errors in precursor-based estimates of rare event frequencies are not unlikely, and mean squared error is not good at capturing the impact of such large uncertainties.

For example, an order of magnitude overestimate of the true accident frequency ϕ would yield an absolute error of 9ϕ, or a contribution of $81\phi^2$ to the computation of mean squared error. By contrast, an order of magnitude underestimate would correspond to an absolute error of 0.9ϕ, or a contribution of $0.81\phi^2$ to mean squared error (a factor of 100 smaller). In fact, based on the criterion of mean squared error, an estimate of 0 (corresponding to an absolute error exactly equal to ϕ) would be preferable to a factor of 3 overestimate, even though the overestimate would at least indicate the correct order of magnitude of the accident frequency, while an estimate of 0 would be totally uninformative.

In practice, we are likely to be at least equally concerned about underestimates and overestimates of accident frequencies. (In fact, given the catastrophic nature of the accidents being analyzed, we may if anything be more concerned about underestimates; underestimates of accident frequencies can lead to a false sense of security, and hence to catastrophic accidents, while overestimates will likely yield only larger expenditures on safety improvements.) For this reason, mean squared error was abandoned as a basis for assessing the performance of precursor-based estimators.

Instead, a measure of error based on the ratio of the estimated accident frequency to its true value was adopted as more appropriate; namely, root mean squared logarithmic error (RMSLE),

$$RMSLE(\hat{\phi}) = \sqrt{E[\log^2(\hat{\phi}/\phi)]} \qquad (3)$$

where $\hat{\phi}$ is an estimator of the true accident frequency ϕ. At least this measure of error will weight order of magnitude over- and underestimates of risk comparably. In addition, it weights infrequent but large errors (e.g., an order of magnitude or more) more heavily than more common minor misestimations of accident frequency. Thus, an order of magnitude error in estimating accident frequency (in either direction) will be weighted 100 times more heavily than an error of only 0.1 orders of magnitude (which might well be considered negligible in light of the large uncertainties typically present in risk analysis). Also, using RMSLE, large errors in opposite directions do not cancel each other out, which

might give a false impression of accuracy. Finally, this measure is readily interpretable. Thus, for example, an RMSLE of 0.5 is equivalent to half an order of magnitude.

RMSLE measure does have a competing weakness. In particular, it is undefined (or infinite) when an estimator takes on a value of 0, which can in principle occur if no precursors of the appropriate type have been observed. However, this is not a significant problem for ϕ'_{CG}, as long as at least one precursor of any type has been observed).

PERFORMANCE OF ϕ'_{CG}

Since RMSLE cannot be computed in closed form, errors were instead estimated using simulation. A single simulation trial was defined to represent the equivalent of 100 years of actual observation. (The U.S. commercial nuclear power industry includes roughly 100 operating reactors, so a single calendar year of data provides the equivalent of 100 years worth of observation.) Thousands of simulation trials were run in each case, as necessary to ensure accurate simulation results.

In any given 100-year trial of this simulation, some number N_0 of precursors with frequency 1 per year was generated, according to a Poisson process. Subsequently, each of these N_0 precursors was assumed to progress to the next more severe precursor (with frequency 10^{-1} per year) according to a Bernoulli process with probability 0.1; let the number of these more severe precursors be N_1. Precursors of even greater severity and smaller frequency were simulated in a similar manner. Thus, N_{i+1}, the number of precursors with frequency $10^{-(i+1)}$, was simulated by random selection from the N_i precursors of next lower severity level (unless the hypothesized accident frequency ϕ is greater than $10^{-(i+1)}$, in which case N_{i+1} was set equal to 0, since precursors cannot occur less frequently than actual accidents). Therefore, N_{i+1} can never exceed N_i. Moreover, if $N_{i+1}=0$ but $N_i>0$, then the most severe observed precursors will be of frequency 10^{-i}, and the portion of ϕ'_{CG} that estimates the precursor frequency (rather than the conditional accident probability) will equal $N_i/100$.

With regard to conditional accident probabilities, we have noted, as initially pointed out by Cooke and Goossens (1990), that there is likely to be non-conservative bias in the estimation of these probabilities. Since conditional accident probabilities are generally estimated using PRA, the likely magnitude of this source of bias cannot be rigorously predicted based on a simple Poisson or Bernoulli model (as was used for precursor frequencies); in the absence of empirical data on the accuracy of PRA, therefore, some other assumption was needed. In this study, we used a simple parametric model to explore the effects of differing degrees of error in conditional accident probabilities via sensitivity analysis.

In particular, we assume that no underestimation of conditional accident probabilities occurs when the true conditional accident probability is equal to one, since in this case the observed precursor is simply an accident, which should be readily identifiable. When the true conditional accident probability is equal to P (for $0<P<1$), the conditional accident probability is assumed to be underestimated by -K log(P) orders of magnitude; in other words,

$$\hat{P} = P^{1+K} \qquad (4)$$

Thus, $\hat{P}$ decreases by K orders of magnitude for every order of magnitude by which P is less than one. A value of K=0.5 obviously indicates the potential for very large errors in $\hat{P}$, since $\hat{P}$ will lose half an order of magnitude of accuracy for every order of magnitude by which P is less than 1. When K is equal to 0.0375, by contrast, $\hat{P}$ will underestimate P by less than 10% per order of magnitude. While this may still sound like a substantial

error, it is actually quite small considering the large uncertainties typically encountered in risk analysis. Note, by the way, that the error in P is not assumed to be random, but instead is simply a deterministic function of P itself. This is obviously a simplistic assumption, but nonetheless provides a preliminary basis for studying the effects of differing magnitudes of error.

In practice, of course, we will not know the parameter K characterizing the errors in our estimates of P. (In fact, if we knew K, we would simply correct our estimates of P accordingly!) However, sensitivity analysis that systematically varies K over reasonable ranges can still be helpful. In particular, the interesting question is whether overestimation of precursor frequencies will tend to dominate underestimation of conditional accident probabilities, or vice versa. Simple parametric analysis can help to clarify this question. The lines labeled "Base case" in Table 1 present the results of such analyses for both $\phi=10^{-4}$ and $\phi=10^{-6}$. (All RMSLE's given in Table 1 have standard errors less than 0.5% of the reported value, showing that sufficient samples were run to yield accurate results.)

As can be seen from the base case results, RMSLE(ϕ'_{CG}) is substantial--never much less than half an order of magnitude, and as much as two orders of magnitude in some cases. Note also that the true errors are likely to be if anything greater than those shown in Table 1, since the error in P has been assumed to be deterministic, thus eliminating one source of noise in ϕ'_{CG}. However, the largest errors occur when K (the underestimation of conditional accident probabilities) is large, suggesting that at least in these cases overestimation of precursor frequencies is not the dominant cause of error.

Because RMSLE eliminates the signs of all errors, however, it is impossible to tell from the base case results whether the errors when K is small reflect primarily overestimation or underestimation of accident frequencies. We answered this question by empirically studying the effects of modifications to ϕ'_{CG}. In particular, we explored the suggestion initially made in Bier (1993), that errors in ϕ'_{CG} might be reduced by avoiding the use of extremely rare precursors. In other words, rather than basing our estimator on the most severe observed precursor, we ignored those precursors with frequency less than or equal to 10^{-3} (i.e., those events for which the errors in estimated precursor frequency would tend to be greatest). As shown in Table 1, this modification yields only negligible changes in RMSLE (i.e., no more than about 0.06 orders of magnitude) when K is large, but more substantial reductions in RMSLE (nearly 0.2 orders of magnitude) for smaller values of K.

Thus, it appears that the cases when ϕ'_{CG} is the least accurate on average (i.e., when K is large) cannot be substantially improved by ignoring the most severe precursors. In fact, reducing the error in such cases would if anything require precursors of greater rather than lesser severity, in order to reduce reliance on inaccurate estimates of the conditional accident probability P. But of course precursors of greater severity can only be obtained by increasing the observation period, since ϕ'_{CG} already uses data from the most severe observed precursor. Thus, lengthening the observation period to ensure the occurrence of more severe precursors (with frequencies of 10^{-2} or less, for example) might yield substantial improvements in RMSLE. By contrast, while ignoring rare (and hence severe) precursors can achieve reductions in RMSLE in some cases, those reductions occur only when RMSLE(ϕ'_{CG}) is already relatively small (less than half an order of magnitude). Moreover, since the true frequencies of observed precursors will generally not be known, ignoring allegedly low-frequency precursors is likely to be controversial, and would violate the principle of "rational consensus" (Cooke, 1991).

Therefore, based on the results of the simulation study reported here, ϕ'_{CG} appears to be a reasonable estimator for use at present. While it can yield large overestimates of risk when extremely rare precursors are observed, this happens only infrequently. Modest improvements in accuracy could in principle be obtained by ignoring such precursors, but this procedure would likely be controversial, violating the principle of "rational consensus"

Table 1. RMSLE(ϕ'_{CG}) as a function of K.

	ϕ	K = 1/2	K = 1/4	K = 1/8	K = 1/16
Base case:	10^{-4}	1.18	0.66	0.47	0.41
Ignoring precursors with frequencies of 10^{-3} or less:	10^{-4}	1.18 (+0.00)	0.60 (-0.06)	0.34 (-0.13)	0.24 (-0.17)
Base case:	10^{-6}	2.09	1.05	0.59	0.44
Ignoring precursors with frequencies of 10^{-3} or less:	10^{-6}	2.15 (+0.06)	1.06 (+0.01)	0.54 (-0.05)	0.31 (-0.13)

(Cooke, 1991). In fact, improving the accuracy of ϕ'_{CG} when it is least accurate would require precursors of greater severity (to avoid large errors in estimated conditional accident probabilities), rather than lesser severity. Therefore, reducing error in ϕ'_{CG} could perhaps best be accomplished by increasing the length of the observation period.

Further work may still be advantageous, however. First, our analysis of ϕ'_{CG} has been less than fully realistic in some respects. For example, the true errors are likely to be greater than those reported here, since the error in $\hat{P}$ has unrealistically been assumed to be deterministic, thus eliminating one source of noise in ϕ'_{CG}. Correcting such simplifying assumptions might provide greater confidence in the results presented here. In addition, the error in ϕ'_{CG} may still be greater than we desire, with RMSLE's ranging from 0.4 to 2.1 orders of magnitude for 100 plant-years of observation. Alternative approaches may permit the development of improved estimators.

ACKNOWLEDGEMENTS

This paper was prepared with the support of the U.S. Nuclear Regulatory Commission (NRC) under Award No. NRC-04-92-089. The opinions, findings, conclusions and recommendations expressed herein are those of the authors and do not necessarily reflect the views of the NRC.

REFERENCES

Barlow, R. E. and F. Proschan, 1975, Statistical Theory of Reliability and Life Testing (Holt, Rinehart and Winston, New York).

Bier, V. M., 1993, Statistical methods for the use of accident precursor data in estimating the frequency of rare events, Rel. Eng. Sys. Safety 41.

Cooke, R. M., 1991, Experts in Uncertainty: Opinion and Subjective Probability in Science (Oxford University Press, New York).

Cooke, R. and L. Goossens, 1990, The Accident Sequence Precursor methodology for the European post-Seveso era, Rel. Eng. Sys. Safety 27:117.

072 Aging Analysis: Methodology & Applications

Chair: D.H. Johnson, PLG

Sensitivity and Uncertainty Analyses in Aging Risk-Based Prioritizations
M. Hassan, S. Uryas'ev (BNL); W.E. Vesely (SAIC)

Bayesian Aging Analysis of Hanford Waste Storage Tank Leak
W.G. He, S.B. Rao (PLG); D.R. MacFarlane (LANL)

Lifetime Prediction and Physical Asset Management
R. Menich (Failure Analysis Assoc.)

SENSITIVITY AND UNCERTAINTY ANALYSES IN AGING RISK-BASED PRIORITIZATIONS[1]

Mahbubul Hassan,[2] Stanislav Uryas'ev,[2] and William E. Vesely[3]

[2]Brookhaven National Laboratory
Building 130
Upton, NY 11973

[3]Science Applications International Corporation
655 Metro Place South
Suite 745
Dublin, OH 73016

INTRODUCTION

Aging risk evaluations of nuclear power plants using Probabilistic Risk Analyses (PRAs) involve assessments of the impact of aging structures, systems, and components (SSCs) on plant core damage frequency (CDF). These assessments can be used to prioritize the contributors to aging risk reflecting the relative risk potential of the SSCs. Aging prioritizations are important for identifying the SSCs contributing most to plant risk and can provide a systematic basis on which aging risk control and management strategies for a plant can be developed. However, these prioritizations are subject to variabilities arising from uncertainties in data, and/or from various modeling assumptions. The objective of this paper is to present an evaluation of the sensitivity of aging prioritizations of active components to uncertainties in aging risk quantifications. Approaches for robust prioritization of SSCs also are presented which are less susceptible to the uncertainties.

The methodology used to perform aging risk-based prioritization of active components is that described in NUREG/CR-5510.[1] Using this methodology, the PRA models are separated from component aging models, allowing any PRA to be used and any aging models to be applied for the components. Minimal cutsets obtained from standard PRAs are used to calculate single and multiple component aging risk-importance coefficients. The aging effects are incorporated into component aging models through changes in component failure rates with age for a given maintenance, and the consequent changes in component unavailabilities. The age-dependent component unavailabilities then are combined with component aging risk-importance coefficients to obtain the increase in CDF due to aging. In this study,

[1]Work performed under the auspices of the U.S. Nuclear Regulatory Commission.

sensitivity and uncertainty analyses approaches were used to evaluate aging prioritization sensitivity to a) uncertainties in component aging risk-importance coefficients, b) uncertainties in aging failure rate data, c) uncertainties in test and maintenance data, d) component aging failure models, d) plant age, and e) selection criteria for risk-significant components. Sensitivity and uncertainty analyses results are also used to define aging prioritizations in terms of component groups which are relatively stable, even when there are large uncertainties in input data. To demonstrate these approaches, two NUREG-1150 plants, a PWR and a BWR, were used in the evaluations.

APPROACH

The methodology to prioritize risk from aging components is based on some average or midpoint values for the input parameters (base case). The prioritized list of risk-significant components is selected based on some cutoff value for the increase in CDF due to aging. Uncertainties in component aging risk quantifications can impact prioritization results in the following ways, a) by altering the individual rankings of the components, and b) by modifying the list of risk-significant components, since determinations of both the selection and the ranking of components are based on criteria related to component aging risk effects.

To address these effects, the approach taken in this study is that for each component j, the aging rate a_j, sensitivity coefficient S_j, replacement interval L_j, and test interval T_j are considered as random variables with corresponding midpoint values $\hat{a}_j$, $\hat{S}_j$, $\hat{L}_j$, and $\hat{T}_j$ over possible ranges. For component j, the contribution to CDF due to aging, $\Delta\hat{C}_j$, is estimated. If this contribution is less than some cutoff value ϵ, the component j is not included in the list of important contributors. While in an average sense, contribution $\Delta\hat{C}_j$ may be smaller than cutoff value ϵ, the real contribution ΔC_j can be larger than ϵ due to variabilities in input. With 90% probability, random variables a_j, S_j, L_j, and T_j are assumed to belong to the intervals

$$\underline{a}_j \le a_j \le \bar{a}_j,\ \underline{S}_j \le S_j \le \bar{S}_j,\ \underline{L}_j \le L_j \le \bar{L}_j,\ \underline{T}_j \le T_j \le \bar{T}_j$$

where

$$\underline{a}_j = \hat{a}_j/10,\ \bar{a}_j = 10\hat{a}_j,\ \underline{S}_j = \hat{S}_j/5,\ \bar{S}_j = 5\hat{S}_j,$$

$$\bar{L}_j = 2\hat{L}_j,\ \underline{L}_j = \hat{L}_j/2,\ \bar{T}_j = 2\hat{T}_j,\ \underline{T}_j = \hat{T}_j/2$$

The bounds on the parameters were obtained from Reference 2. The error ranges were assigned either to account for plant specific variations in aging and maintenance data, or to include the PRA uncertainties in aging risk estimates.

The increase in core damage frequency due to aging of component j, ΔC_j, may be calculated, for example, with the formula

$$\Delta C_j = S_j \Delta q_j = S_j a_j \left[\frac{1}{4}(L_j - T_j)T_j + \frac{1}{6}T_j^2\right]$$

considering only single component contributions and assuming a linear aging behavior.[1] Δq_j is the unavailability of component j due to aging.

Component j does not belong to the list of important contributors if

$$\Delta\hat{C}_j = \hat{S}_j \Delta\hat{q}_j = \hat{S}_j \hat{a}_j \left[\frac{1}{4}(\hat{L}_j - \hat{T}_j)\hat{T}_j + \frac{1}{6}\hat{T}_j^2\right] \le \epsilon$$

Nevertheless, because of uncertainties in the parameters, contribution ΔC_j can be larger then cutoff value ϵ. The probability that component j belongs to the list of important contributors

$$P\{\Delta C_j \ge \epsilon\} = P\left\{S_j a_j \left[\frac{1}{4}(L_j - T_j)T_j + \frac{1}{6}T_j^2\right] \ge \epsilon\right\}$$

can be calculated as a four-dimensional integral with respect to a_j, S_j, T_j, and L_j.

By definition

$$P\{\Delta C_j \ge \epsilon\} = \int_{S_j a_j\left[\frac{1}{4}(L_j - T_j)T_j + \frac{1}{6}T_j^2\right] \ge \epsilon} \mu_j(a_j, S_j, L_j, T_j)\, da_j\, dS_j\, dL_j\, dT_j$$

where $\mu_j(a_j, S_j, L_j, T_j)$ is a joint density of random values a_j, S_j, L_j, T_j. We consider the case with independent random variables a_j, S_j, L_j, and T_j. Therefore density

$$\mu_j(a_j, S_j, L_j, T_j) = \phi_j^a(a_j)\,\phi_j^S(S_j)\,\phi_j^L(L_j)\,\phi_j^T(T_j)$$

is a product of the densities ϕ_j^a (a_j), ϕ_j^S (S_j), ϕ_j^L (L_j), and ϕ_j^T (T_j) for the variables a_j, S_j, L_j, and T_j. For the evaluation of the integral, appropriate densities (lognormal, loguniform) for the specified random variables may be used. The integral can be evaluated numerically using standard computer codes.

Data sensitivity evaluations examined the sensitivity of component rankings by minimizing relevant parameters for the dominant contributors from the base case (average values for the parameters) and maximizing these parameters for the rest of the components. The effects of individual parameters were evaluated as well as their combined effects. Minimal maintenance practice was assumed for the components. Under minimal maintenance, tests do not alter the aging rates or the age of the component and renewal is done only at failure. In evaluating the sensitivity of aging prioritization to component aging models, linear and nonlinear time-dependence were considered in component aging failure rate behavior. Aging prioritizations based on these models were compared by defining equivalent mean time to failure (MTTF) for the components.

SELECTED RESULTS

Table 1 shows component prioritization results summary for a PWR (Plant A) and a BWR (Plant B). For each plant, the base case utilizes midpoint values for the variables over possible ranges. The sensitivities of number of risk-significant single aging contributors to uncertainties in different variables are shown. The combined effect refers to the case where uncertainties in the variables are correlated to produce the maximum impact on the number of contributors. Components are considered risk-significant if CDF contributions due to aging are at least 1% of C (ϵ = .01C), where

C is the baseline CDF estimated by the PRA. For a single variable, aging rate has the maximum impact on the number of risk-significant single contributors for both plant A and Plant B. Uncertainties in test and replacement intervals had no impact on the number of significant contributors to aging risk for Plant A.

Table 1. Number of Risk-Significant Single Aging Contributors

	Base Case	Additional Single Contributors Due to Uncertainties in				
		Aging Rate	Importance Coefficient	Replacement Interval	Test Interval	Combined Effects
Plant A	27	4	2	0	0	16
Plant B	12	10	9	3	6	13

While only a few additional components are added to the list of risk-significant aging contributors through the sensitivity analyses, the detailed results (reported in Reference 3) show that individual component rankings can vary very widely. However, while the individual component orderings are subject to variations, it is possible to define component aging prioritizations in terms of broad groups which are relatively stable. For example, based on the sensitivity analyses for Plant A and Plant B, the aging risk contributors may be grouped into three categories, 1) components from the base case which consistently appear in the list of significant contributors in all cases when single parameter uncertainties are considered, 2) rest of the contributors from the base case, along with additional contributors from single-parameter sensitivity analyses, and 3) components which appear as risk-significant only when all parameter uncertainties are correlated. Table 2 shows the number of risk-significant single contributors in each category.

Table 2. Component Aging Prioritization Grouping Based on Sensitivity Analysis

Component Category	Plant A	Plant B
Group I	11	11
Group II	21	11
Group III	11	3

Figures 1 and 2 show the likelihood values for the components to be risk-significant for Plant A and Plant B, respectively, when uncertainties in $\{a_j\}$, $\{a_j$ and $S_j\}$, $\{a_j$, S_j, and $L_j\}$ and $\{a_j$, S_j, L_j, and $T_j\}$ were included. The probabilities in Figures 1 and 2 refer to the CDF estimates due to aging for the components to be larger than ϵ (where $\epsilon = .01\hat{C}$ and $\hat{C}$ is the baseline CDF calculated by PRA without the aging effects). For example, in Figure 1, for a component with a base case aging CDF estimate of 1e-07, there is approximately a 25% chance of exceeding ϵ. Components with probabilities of 0.5 or higher are included in the list of risk-significant contributors from base case evaluation. There are small changes in these probabilities when curves for a_j only and $\{a_j$ and $S_j\}$ are compared. However, introduction of uncertainties in L_j and T_j did not appreciably change the probabilities. This observation applies to both plants.

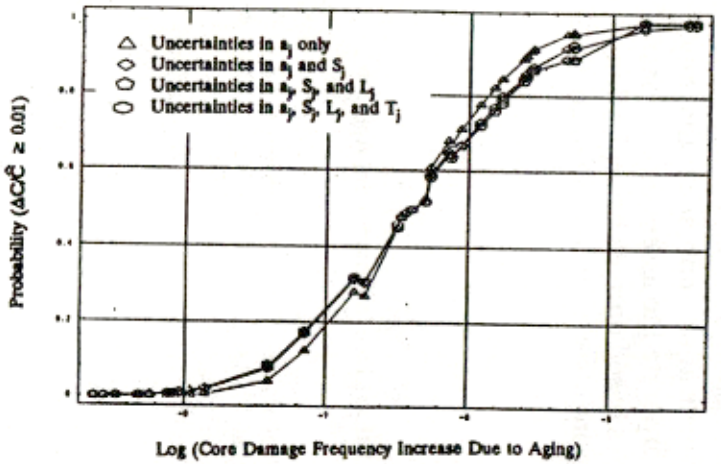


Figure 1. Plant A: Prioritization sensitivity to parameter uncertainties
Probability that aging core damage frequency contribution is larger than cutoff (ϵ=1% $\hat{C}$)

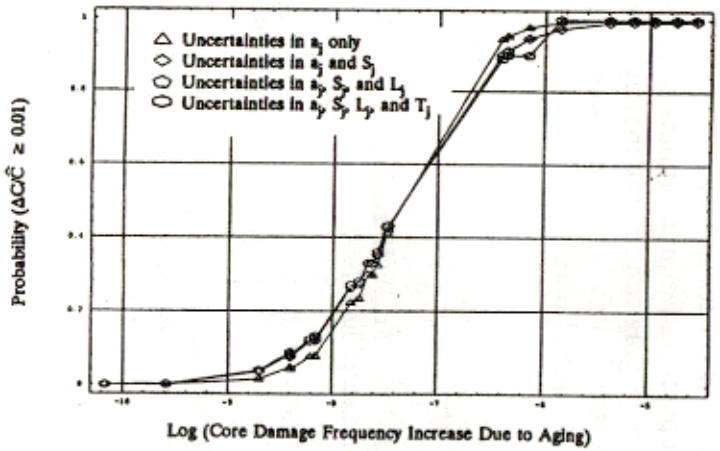


Figure 2. Plant B: Prioritization sensitivity to parameter uncertainties
Probability that aging core damage frequency contribution is larger than cutoff(ϵ=1% $\hat{C}$)

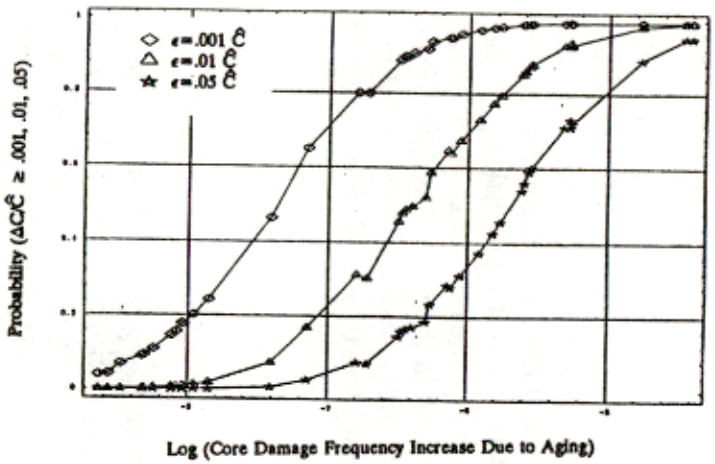


Figure 3. Plant A: Prioritization sensitivity to cutoff criteria
Probability that aging core damage frequency contribution is larger than cutoff

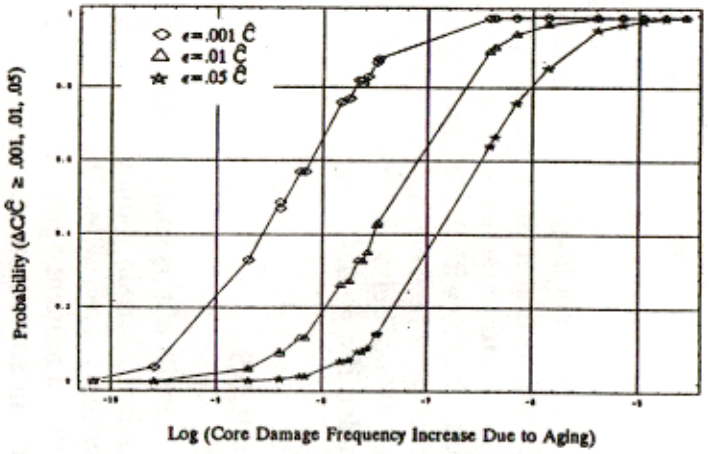


Figure 4. Plant B: Prioritization sensitivity to cutoff criteria
Probability that aging core damage frequency contribution is larger than cutoff

Figures 3 and 4 present the probability distributions for the CDF estimates due to aging for the components to be larger than ϵ, for ϵ = .05C, ϵ = .01C, and ϵ = .001C for Plant A and Plant B, respectively. Contributors with probabilities $\geq$0.5 when ϵ = .01C correspond to the list of prioritized components in the base case. Figure 3 for Plant A shows that components which were not included in the base case for ϵ = .01C have less than 20% chance to be included if the requirements were ϵ = .05C. This means that if the ϵ = .01C criteria is used to select aging risk-significant components, it can be assured, with a high degree of confidence, that most of the possible contributors with $\Delta C \geq .05C$ have been included in the prioritized list. If the requirements were relaxed to ϵ = .001C, the components excluded have probabilities of less than 10% of being included in the base case, and less than a 4% probability of being included if ϵ = .05C.

CONCLUSIONS

The results of this study indicate that PRA-based aging risk evaluations, along with sensitivity-uncertainty analyses, can be effective in prioritizing active components in safety systems, even when there are large variabilities in basic inputs. Another finding is that dominant aging contributors are generally quite stable and only a few components, in addition to the ones identified from the base case, are deemed to be potentially risk-significant even when all major uncertainties are accounted for. Individual rankings of the components may change significantly with plant age. These rankings also are sensitive to the assumed aging behavior for the components. However, the list of risk-significant contributors is not easily affected. The larger uncertainties in aging failure rates of the components drove the uncertainties in aging risk estimates, and consequently, the aging prioritizations for the components.

The approaches and results of this study are intended to provide guidance on the performance and use of aging prioritizations for applications. The implications of this study are that stable groups of prioritized components can be developed with high confidence which are not very sensitive to modeling assumptions and data. The advantages of sensitivity-uncertainty analyses based prioritization approach is that any changes in uncertainties in the inputs can be easily incorporated to assess their impact on the results. Aging control and management strategies based on such prioritizations will be able to identify and focus attention on components whose age need to be controlled to control plant risk.

REFERENCES

1. NUREG/CR-5510,"Evaluations of Core Melt Frequency Effects Due to Component Aging and Maintenance", June 1990.

2. NUREG/CR-5587, "Approaches for Age-Dependent Probabilistic Safety Assessments with Emphasis on Prioritization and Sensitivity Studies," August 1992.

3. NUREG/CR-6045 (Draft), "Sensitivity and Uncertainty Analyses in Component Prioritization Methods Including Aging and Maintenance Effects," April 1993.

BAYESIAN AGING ANALYSIS OF HANFORD WASTE STORAGE TANK LEAK

Wei G. He[1], Shobha B. Rao[1], and Donald R. MacFarlane[2]

[1]PLG, Inc.
4590 MacArthur Boulevard, Suite 400
Newport Beach, CA 92660-2027

[2]Los Alamos National Laboratory
Engineering and Safety Analysis Group, N6
P.O. Box 990, MS K557
Los Alamos, NM 87545

ABSTRACT

In this paper, Hanford Waste Storage single-shell tank (SST) failure due to leaks is modeled as an aging process. The methodology used for this aging analysis was initially developed for the Idaho National Engineering Laboratory as embodied in the computer code DORIAN. This methodology is an integrated package for performing Bayesian aging analysis of reliability data; e.g., for identifying trends in component failure rates and/or outage duration as a function of time. DORIAN allows several possible aging models along with their prior probability indicating our belief that each trend is the correct one to be hypothesized and input into DORIAN. DORIAN then combines the prior aging models with the specific tank leak experiences to develop the posterior tank leak rates. The actual tank leak data are also used to "update" our confidence of aging models and to develop a "posterior" probability for each aging model, representing the probability that each model is the correct one. DORIAN output is three dimensional, with time and failure rate as variables. Using the SST leak data during 1944 to 1992, future SST leak frequencies have been predicted using DORIAN. For each given tank age, a distribution for the leak rates is generated. To predict the number of future leakers in any given year, a Binomial model is used with the probability generated by DORIAN output. For predicting the number of leakers in any given period, a time-dependent Poisson model is used. The results show the increasing tank leak frequency and the decreasing confidence when the prediction extends further and further into the future.

HISTORICAL BACKGROUND

Since 1944, radioactive wastes from the processing of irradiated uranium fuels have been stored as alkaline slurries in Hanford underground tanks. A total of 149 SSTs of various designs and sizes have been built in groups at different times since 1944, the latest being built around 1964. In addition, between 1970 and 1986, 28 double-shell tanks (DST) were constructed for this use. The capacity of these tanks range from approximately 55,000 gallons to over 1 million gallons. The basic containment employed in the design of SST tanks was the provision for a reinforced concrete shell with a carbon steel liner. These tanks are grouped in 15 tank farms located in the Hanford 200 East and 200 West areas. The first waste was routed to tank storage in December 1944.[1]

Over the years, almost every group of these SSTs has developed leaks. Other tanks have behaved as though they leaked and were categorized as "questionable integrity." Various leak sizes, from a few hundred gallons to hundreds of thousands of gallons, have been estimated. When a leaking tank was discovered, all transfers to that tank were stopped, and liquid wastes were pumped out; although this process may have taken years to accomplish. Due to the environmental and potential health consequences of leaks, several attempts have been made to estimate the tank leak rate and to predict the potential future leaks in order to develop strategies to control the damage.

EARLY LEAK ASSESSMENT

The first SST leak rate analysis was performed in 1980*.[2] At the time of the analysis, the oldest tank was about 400 months. The length of time (up to the month) that each tank has been used (its service life) was determined. The analysts plotted the percentage of tanks leaked against service life on a normal distribution grid. It was found that the points fell essentially on a straight line. The last point on the line indicates about 40% of the tanks developed leaks that corresponded to their service life of 400 months. This indicates that the tank leaks are normally distributed. A slight extrapolation of the data gives a mean leak at 425 months and the standard deviation (the slope of the line) of 174 months. Therefore, if the prediction were right, approximately 84% of the tanks would fail when they were 599 months old. For the oldest tank group commissioned in 1944, the age (in 1993) was about 599 months, and the total failure within that group is about 42%.

In 1983, another analysis was performed.[3] This time, no attempt was made to find an analytical expression of tank leaks, and survival analysis was used to examine the data. Two major parameters used in the calculations were the Cumulative Proportion Surviving at Beginning of Interval (CPS) and the (time) Interval. Tanks were divided into three groups: survival tanks, dead tanks, and lost tanks. The inactive tanks, which are tanks with no waste transfers to and from, were defined as the lost tanks. The lost tanks were removed from the tank population at the end of each time interval. The CPS is simply the ratio of failed tanks versus the total tank population in the interested time interval. The analyst plotted the CPS against the time (month). As was shown in the plotting[3], this analysis predicted that 25% of the tanks would leak during the first 300 months, and 50% of the tanks would develop leaks during the first 386 months of service life. The failure rate would increase sharply beyond 300 months.

A more recent attempt (1992, Westinghouse Hanford Company analysis) to assess leak rates of SST did not distinguish between in-service and inactive tanks. It was assumed that

---

*In this analysis, as all later analyses including our study, the confirmed leakers and the questionable integrity tanks were lumped together as failed tanks.

both types could develop leaks and be detected if they leaked. A very simple engineering approach was used in the analysis, and the focus was the potential future leak, and, thus, no estimation of past leak was given. The predicted future potential leaks have a lower bound of 1.5 tank/year and an upper bound of 2.77 tank/year. The lower bound was obtained by dividing total leakers (66) by the average service life (44). This corresponds to (for the lower bound) a mean failure rate of 596 months. The upper bound was obtained by assuming that the rest of the 83 SSTs would all leak in the next 30 years.

BAYESIAN AGING ANALYSIS

Although various analyses result in different tank leak profiles, one common characteristic is that all analyses predict increased number of leaks in the future. Therefore, it seems that the age of the tank may play an important role in this failure. A preliminary examination of the existing tank leak experience also indicates that the failure rate has been increasing over the years (see Figure 1). Therefore, the failure was modeled as an aging process, with tank types and leak sizes being two distinct variables.

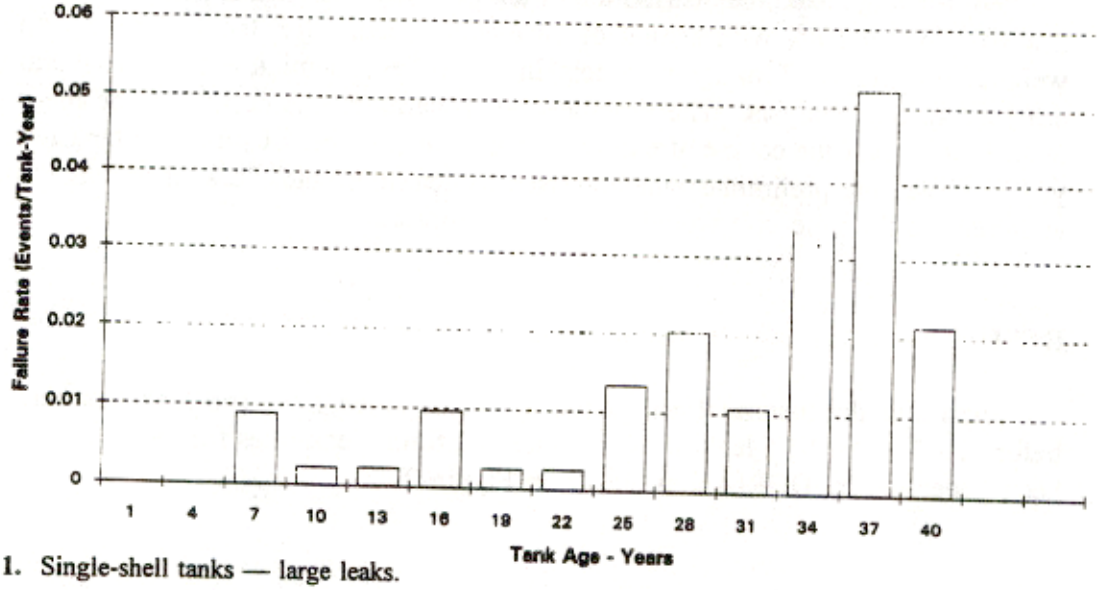


Figure 1. Single-shell tanks — large leaks.

There is little information available on root cause analysis of tank failure; thus, it is hard to identify any single cause or some fixed causes that could lead the tanks to develop leaks. The cause may be aging of the steel liner, accelerated by the contents of the tank, or in-leakage of water from the ground through the concrete. It could be the cyclic thermal stresses caused by the periodic addition of hot wastes, causing the steel liners to buckle and crack. It most likely is a combination of the above and other additional factors. Thus, we need an aging analysis tool that is able to incorporate all of the above thoughts and the uncertainties associated with them. The methodology suitable of this task is Bayesian aging analysis, as embodied in the computer code DORIAN. The code is an integrated package for performing aging analysis of reliability data; e.g., for identifying trends in component failure rates and/or outage duration as a function of time.[4]

To use the code, three types of input are required:

1. One or more possible aging models along with their prior probabilities indicating the analysts' belief that each trend is actually the correct one.
2. A postulated tank leak rate distribution during its first service year along with the analysts' judgment of the distribution uncertainties.
3. Actual tank leak experience.

The distribution of the first-year leak rate and aging models should be reasonably broad and flexible to include all possible failure rates and conceivable failure mechanisms. In this work, several test runs were executed to test and examine various assumptions of the first-year leak distribution and aging models before final results were generated.

1. Tank failure (leak) distribution as a function of time and leak rates.
2. A posterior probability for each aging model, representing the probability that each model is the correct one in light of the observed data.
3. A posterior probability for the first-year leak rates, representing the probability that each rate is the correct one in light of the observed data.

For this analysis, leak size of 5,000 gallons is used to separate leakers into small and large leak groups, and tank size of 55,000 gallons is used to separate tanks into small and large tank groups. Small and large leaks for the total SSTs are also performed. Similar to Reference 2, tanks are defined to be survival tanks, dead tanks, and lost tanks. However, the lost tanks are those incapable of leaks (no liquid to cause small leaks or few liquids to cause large leaks). To limit the scope of this discussion, the following discussion uses large leaks for the total SSTs as an example.[5]

For the large leak, the distribution of the first-year leak rate is postulated to range from 2×10^{-3} to 7×10^{-3}, with a total of 10 equally spaced (on the Log-space) and equally weighted leak rates. Two aging models, linear and exponential, each with several versions, are used in this analysis. The time-dependent parameter is defined to be the number of leaked tanks over the course of exposure tank-year (number of surviving tank in a calendar year). Given the preliminary method used to determine tank leak, it is judged that using tank-year as the time unit would be more appropriate.

RESULTS

Inputting the historical SST large leak data since 1951 (no leak has been detected before 1951), past SST leak rates and potential future leak rates for the next 10 years have been generated by DORIAN as shown in Figure 2.

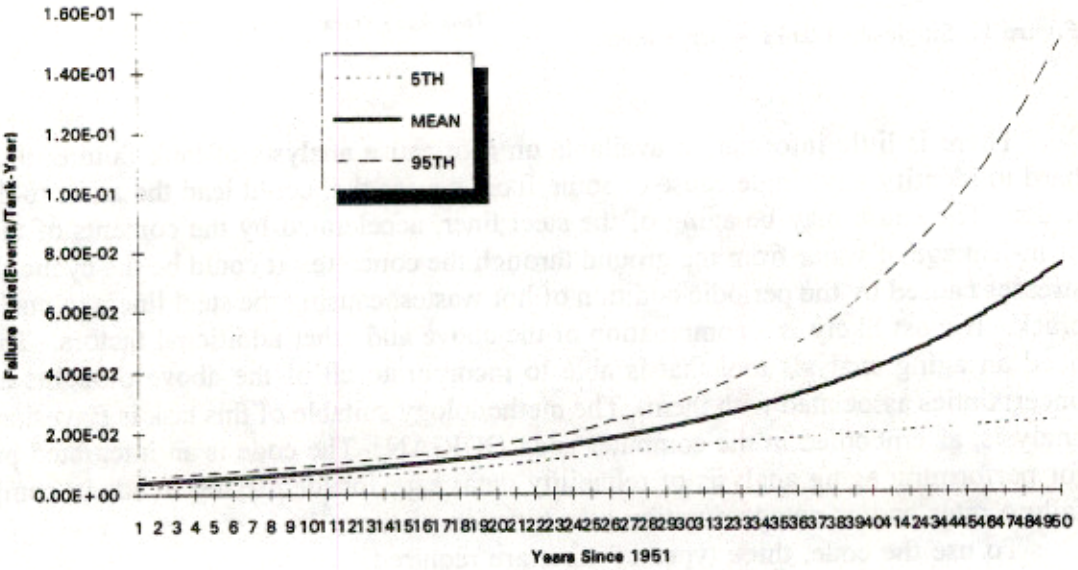


Figure 2. SST large leaks.

For each given tank age, this study gives a distribution for the leak rates, given that the tank has survived to that age. For instance, the results predict that the large leak rates

for an SST of 54 years old have a mean likelihood of 0.945 to survive another year. The results also indicate that we have 95% confidence that, for that tank, the likelihood of developing a large leak in its age of 55 would be less than about 0.101.

To estimate the likelihood of tank failure in each year, given the remaining population of 11 SSTs available for large leaks, the Binomial model is invoked.

$$p(K=k)=\begin{bmatrix} n \\ k \end{bmatrix} p^k(1-p)^{n-k}$$

where n is the total number of tanks, k is the number of tanks that developed a leak in the year, and p is the likelihood of leaking for a tank. For this problem, n equals 11, and p is supplied by the failure rate generated by DORIAN code*. The distribution of the failures in any year can be easily obtained by substituting the distribution of p from DORIAN output. In this study, a Monte Carlo simulation is used to calculate the distribution.

Using the above procedure, it can be easily shown, for tank group of age 55 (which is corresponding to about 1997 for tanks commissioned in 1944), that there will be about 34% chance that, on average, one tank will develop large leaks, and we have 95% confidence that the chance of one tank developing a large leak will be less than 39%. The same analysis determines that there will be about 54% chance that no tanks will develop large leaks. We have 95% confidence that the chances of no tanks developing large leaks for this tank group will be larger than 17%. (A similar approach can be used to predict the likelihoods of leak in any given year for small or large leaks.)

To estimate the likelihood of tank failures for a given period, such as the 5 years from 1993 to 1997, the Poisson model p(k,t; λ)is used. In principle, this is a time-dependent Poisson process since the failure rate is a function of time λ = λ(t). Since this failure rate is assumed to be a piece-wise constant, the integration can be replaced by summation.

$$p[K(t)=k] = \frac{\sum (\lambda_i t_i)^k \exp(-\lambda_i t_i)}{k!}$$

where i equals one to five for this example, and k equals zero to n, the number of remaining interested tanks**.

The Poisson model predicts that, on average, the likelihood of no tanks failure during this period is about 0.069. The results indicate that we have 95% confidence that the likelihood of no tank leak during this period will be no less than 0.042. The following table gives the Poisson model prediction of likelihood of number of tanks to failure, from 0 to 11, during this period.

Table 1. Estimated distributions of number of leakers during 1993 to 1997. (The decimal points reflect the cutoff from simulation.)

Leakers (k)	Mean	Variance	5th	50th	95th
0	0.0686	0.0003	0.0420	0.0663	0.1019
1	0.1814	0.0009	0.1330	0.1800	0.2327
2	0.2422	0.0003	0.2110	0.2441	0.2657
3	0.2177	0.0001	0.2023	0.2207	0.2240

*Note that the unit of output failure rate is 1 year.

**Since the λ is small, it is judged that the Poisson model is a good approximation.

Table 1 (continued). Estimated distributions of number of leakers during 1993 to 1997. (The decimal points reflect the cutoff from simulation.)

Leakers (k)	Mean	Variance	5th	50th	95th
4	0.1482	0.0004	0.1154	0.1498	0.1768
5	0.0815	0.0003	0.0527	0.0812	0.1122
6	0.0377	0.0001	0.0200	0.0368	0.0592
7	0.0151	0.0000	0.0065	0.0142	0.0268
8	0.0053	0.0000	0.0019	0.0048	0.0106
9	0.0017	0.0000	0.0005	0.0015	0.0037
10	0.0005	0.0000	0.0001	0.0004	0.0012
11	0.0001	0.0000	0.0000	0.0001	0.0003

The half-life of the tank is estimated to be 576 months by this study as compared to 386 months[3] or 425 months.[2] In general, smaller potential future leak rates are generated by this analysis compared with other studies. For instance, one of the analyses predicted that between 1.5 and 2.8 tanks would develop a leak in 1997, while the prediction of this study is that there will be about 50% chance each for one or more tanks to develop a large or small leak in 1997. For a tank of 200 months old, the probability of surviving another 36 months are respectively estimated to be

- Normal distribution grid: 0.056
- Survival table: 0.055
- DORIAN: 0.036
- Existing data: 0.031

The differences in leak rates between this study and other analyses are small for early tank leaks. However, the differences get larger when tanks become old.

CONCLUSION

A Bayesian probabilistic aging analysis, aided by the DORIAN code, is performed for a Hanford single-shell waste storage tank. The results of this analysis suggest that the overall tank leak rate is increasing with the time, although the rate of the increasing may not be as large as indicated by other researchers. In addition, the approach taken by this analysis provides two-dimensional probabilistic distributions for the leak rate. Results of this study can be continuously updated to include any new evidence as more SSTs develop leaks and more accurate data of leaks are determined. Finally, as was performed in this study, the approach used in this study can be extended to form the base to predict the potential double-shell tank leak.[5]

REFERENCE

1. J. D. Anderson, *A History of the 200 Area Tank Farms*, WHC-MR-0132 (1990).
2. R. J. Catlin, *Assessment of the Surveillance Program of the High-Level Waste Storage Tanks at Hanford*, U.S. Department of Energy (1980).
3. D. A. Reynolds, *Analysis of Single-Shell Tank Failure Data*, SD-WM-TI-032 (1983).
4. *Development and Demonstration of Methods for Nuclear Power Plant Aging Risk Analysis--DORIAN Users Manual*, PLG-0759 (1990).
5. *Risk Assessment for Hanford High-Level Waste Tank 241-SY-101*, PLG-0936 (1993).

LIFETIME PREDICTION AND PHYSICAL ASSET MANAGEMENT

Ronald P. Menich
Failure Analysis Associates
149 Commonwealth Drive
Menlo Park, CA 94025

INTRODUCTION AND SUMMARY

Economic vitality depends increasingly on sustained operation of highly mechanized and automated systems. These systems in turn depend on the integrity of physical assets such as power generation and electricity distribution, data and telecommunications networks, process industry facilities, aircraft, ships, trains, pipelines, and bridges. Cost-effective physical asset management is a critical factor for maintaining competitiveness, ensuring human safety, and avoiding consequences of environmental contamination.

An essential part of physical asset management involves estimation of remaining life for key system components. Deciding whether to inspect, maintain, repair, or replace based on what is "known" and what can be "predicted reliably" can be complicated. Inspections often provide incomplete information of the extent of actual degradation, and they typically involve shutting the system down, with attendant loss of revenue. Such considerations must be weighed against safety and the very substantial economic costs that can result if equipment fails while operating.

Predictions of remaining life are often based on probabilistic models for time to component failure, where the model is built up using statistical analyses of available data, including operating experience and/or typical industry data. This procedure has typified engineering practice in lifetime prediction for over thirty years. The probability model provides, or so it might appear, a tractable means for estimating rates of degradation in situations that can be extremely complicated from a process or operational standpoint.

In practice, however, use of probabilistic or statistical models for life prediction may cause critical issues to be overlooked. These include goodness-of-fit for parametric statistical models and the "structural uncertainties" inherent in the modeling approach adopted, whether that approach involves mathematics (e.g., probability and statistics) and/or engineering (Monte Carlo) simulation of degradation and failure.

In this paper we discuss the goodness-of-fit issues for an example case that involves degradation in a large piece of industrial equipment. The number of components subject to potential failure is quite large. The actual data is "interval-censored," that is, periodic inspections indicate the extent of degradation in component units. This inspection data is fit to a "galaxy" of probability distributions, using the method of maximum likelihood for parameter estimation, adapted to the interval-censored data. Goodness-of-fit is evaluated using the Pearson chi-square test, a theoretically sound and appropriate procedure for this data.

Parametric distributions investigated include the gamma, Weibull, and several others. Distributions based on exponential transformation of random variables investigated included the log normal and variations of the log Cauchy (i.e., distribution of $Y = e^X$, where X is a Cauchy random variable).

From a curve-fitting viewpoint, several of these distributions appeared to provide quite satisfactory models for the degradation data. For instance, a standard Weibull linear regression plot produced a coefficient of determination value of 0.9650. Yet the Weibull model was resoundingly rejected by the chi-square goodness-of-fit test, and none of the models tested were accepted by this test.

Most interesting, perhaps, is that the log Cauchy family provided a closer statistical fit to this data than any other distribution class investigated. The Cauchy distribution is not typically employed in lifetime prediction applications. Its density function is symmetric about the

median and its graph resembles the normal density, but approaches the horizontal axis so slowly that moments do not exist. It can be regarded as a model of "infinite variability."

There is a tendency in practice, enforced by what has heretofore been available in software, to fit failure data to well-known, well-behaved parametric families such as the Weibull, log normal, and gamma, and then select a model from this menu. Now, however, the ready availability of efficient numeric computation routines (e.g., *Excel Solver*) makes it feasible for practitioners to routinely investigate a broader class of more complicated probability distributions in their search for a better and more robust predictive model.

Our results suggest that caution should be exercised in using standard parametric statistical models as predictive mechanisms for complicated failure processes and as input for life-cycle decision making. An infusion of new methods and approaches for survival analysis is needed to support sound, cost-effective physical asset management.

A TYPICAL WEIBULL ANALYSIS

The source data in this example arises from a situation in which n = 10,164 identical mechanical units were placed into service at a particular point in time. At periodic intervals, the units were inspected to determine which had failed during the interval of time since the previous inspection. Unfailed units remain in operation indefinitely. Table 1 lists the inspection times and the number of units failed since the last inspection. For ease of exposition, we have adopted Lawless's (1982), notation for grouped (interval censored) data.

Table 1. Inspection times and number of failures since last inspection.

Inspection number (j)	Operational time in years at event of inspection (a_j)	Number failed since last inspection (d_j)
0	0	0
1	1.69	1
2	1.83	1
3	2.43	16
4	3.33	24
5	4.61	38
6	5.70	72
7	6.90	150
8 ($=k$)	7.97	760
9	∞	9102

Typically, engineers who analyze these units transform and plot the inspection values on Weibull paper and then use linear regression to estimate the distribution parameters. The form of the Weibull cumulative distribution function used is:

$$F(t)=\begin{cases}1-\exp\left[-\left(\frac{t}{\theta}\right)^{b}\right], & t\geq 0\\ 0, & \text{otherwise}\end{cases},$$

where the ϕ and b parameters are called the characteristic time and slope parameters, respectively. The input tuples for the regression are:

$$\left\{\left(\ln(a_j),\ln\left(-\ln\left(1-n^{-1}\sum_{i=0}^{j}d_i\right)\right)\right),\quad j=1,\ldots,k\right\}.$$

Under this transformation, the regression line slope serves as an estimator of b and the intercept as an estimator of $-b \ln(\phi)$. The procedure above yields the estimates $\hat{b} = 3.9408$, $\hat{\theta} = 15.1481$ for the data in Table 1, and produces a coefficient of determination value of 0.9560. With this high R-squared value in hand, the results of such an analysis are typically deemed acceptable, and it is time to move on to other problems.

PROBLEMS WITH STANDARD WEIBULL ANALYSIS

The method applied to the Table 1 data in the above example was ordinary least squares. In his discussion of least squares methods for survival data, Lawless (1982) states:

> This ordinary least squares estimate does not take into account the variance-covariance structure....

After introducing a generalized least squares estimator to approximate the covariance structure, Lawless continues:

> The [generalized least squares] estimates provided are suitable for some purposes, although where precise estimation is important, one should compute maximum likelihood or other more precise methods.

In addition to the conceptual caveats mentioned by Lawless, there are many other reasons why the application of standard Weibull (ordinary least squares) analysis to the grouped data of Table 1 is faulty. First, the application of standard Weibull analysis to grouped data results in the loss of information contained in the original untransformed data. The number of tuples upon which the ordinary least squares estimate is based is the number of inspections k: all information regarding n (the total number of units) has been lost.

There is also another loss of information in the sense that the standard Weibull analysis cannot represent the information provided by inspections prior to the discovery of the first failed unit. In the particular example of Table 1, there is only one such inspection, namely, at time zero. However, in other situations it has occurred that more than one inspection passes before the first failed unit is discovered. For such inspections, the cumulative number of failed units is zero, and hence the double logarithm transformation that prepares the input data for the linear regression is not possible.

BETTER ESTIMATION METHODS

There are statistical estimation methods available that do not have the drawbacks listed above for standard Weibull analysis. One is maximum likelihood estimation for grouped data: maximize over all positive b and θ the function

$$L(b,\theta) = \prod_{j=1}^{k+1} \left[S(a_{j-1}) - S(a_j) \right]^{d_j},$$

where $S(t) \equiv 1 - F(t)$ denotes the survivor function (Lawless, 1982). Once the maximum likelihood estimates have been computed, the Pearson X^2 goodness-of-fit test statistic (D'Agostino and Stephens, 1984),

$$X^2 = \sum_{j=1}^{k+1} \frac{\left[d_j - n\left(\hat{S}(a_{j-1}) - \hat{S}(a_j) \right) \right]^2}{n\left(\hat{S}(a_{j-1}) - \hat{S}(a_j) \right)},$$

may be computed and assessed against the percentage points of the $\chi^2_{(k+1-p-1)}$ distribution, where p is the number of parameters estimated ($p = 2$ in the Weibull example). Here the $\hat{S}(t)$ is the estimated survivor function value computed by substituting the estimated parameter values into the functional form for S(t).

OTHER DISTRIBUTIONS

The question naturally arises: "Are there any other distributional forms characterized by the same number of parameters that fit the data in Table 1 better than the Weibull?" Table 2 below lists the results of the maximum likelihood estimation method applied to various different distribution types, each having exactly two parameters. The table is sorted in decreasing order of the logarithm of the likelihood function at optimality: the higher in the table, the better the fit.

Table 2 shows that there are several distributions available that have lower values of the chi-square test statistic. In fact, there are at least three distributions, all derived from the standard Cauchy distribution, that have associated chi-square test statistics that are less than one-tenth that of the Weibull.

It is well known that a Cauchy random variable has no associated mean, variance, or other moments: the tails of the distribution are too heavy to allow convergence of the integrals that define these moments. It can be verified that, for all four of the distributions derived from the Cauchy in Table 2, none of the moments exist. Thus, these distributions that fit the data in Table 1 better than the Weibull have no means or variance. This important property of the data in Table 1 would have been missed had the analysis simply terminated after a standard Weibull analysis.

Table 2. Results of parameter estimation and goodness-of-fit tests.

Distribution	Functional form or reference	Max. log likelihood	Pearson χ^2 statistic
function #3 of Cauchy	Distribution of $W = \exp(\exp(\exp(aX + b))) - e$, where X is standard Cauchy distributed	-4473.07	39.78
function #2 of Cauchy	Distribution of $W = \exp(\exp(aX + b)) - 1$, where X is standard Cauchy distributed	-4475.26	42.24
function #4 of Cauchy	Distribution of $W = \exp(\exp(\exp(\exp(aX + b)))) - \exp(e)$, where X is standard Cauchy distributed	-4478.24	46.27
function #1 of Cauchy	Distribution of $W = \exp(aX + b)$, where X is standard Cauchy distributed	-4482.98	51.90
Gompertz	The distribution $1 - \exp[-a(\exp(ct)-1)]$	-4510.14	173.75
Dhillon	Problem 2.11 in Cox and Oakes	-4570.51	470.72
Weibull	WEIBULL function in Excel	-4577.41	516.29
log logistic	Page 21, Cox and Oakes	-4584.88	569.50
inverse power function	Distribution of X, where $P\{X > t\} = (1 + a\, t^b)^{-1}$	-4584.88	569.51
inverse log function	Distribution of Ξ, where $P\{X > t\} = (\ln(a\, t^b + e))^{-1}$	-4592.98	632.02
gamma	GAMMADIST function in Excel	-4638.92	838.58
lognormal	LOGNORMDIST function in Excel	-4684.13	1548.21
Birnbaum-Saunders	The distribution $\Phi\left\{\frac{1}{\alpha}\left[\left(\frac{t}{\beta}\right)^{0.5} - \left(\frac{t}{\beta}\right)^{-0.5}\right]\right\}$, where Φ denotes the standard normal distribution	-4728.59	6384.49
inverse Gaussian	Equation (2.14) in Chikkara and Folks	-4731.59	7629.58
function of extreme value	Distribution of $W = \exp[aX + b]$, where X has the standard extreme value distribution	-4800.39	7865.13

RECENT DEVELOPMENTS AND NEW DIRECTIONS

Mathematical modeling that takes explicit account of failure and wearout processes, survival in dynamic environments, and the implications for monitoring and "optimal" operation of complex systems are areas of recent and active investigation. Some of this work is described in the book of Sobcyzk and Spencer (1992) and the survey papers of Singpurwalla (1992, 1993). The template for a significant line of development in dynamic failure modeling was laid in Lemoine and Wenocur (1985) and carried forward in a series of subsequent papers.

A sample of work on monitoring, prognostic methods, and failure prevention includes Grove (1979); Shives and Mertaugh (1988); McDuff, Simpson, and Gunning (1989); Pusey and Pusey (1989); Feyock and Karamouzis (1991); and Ray, Wu, Carpino, and Lorenzo (1992). Widespread implementation in the field of prognostic methodologies has been limited heretofore by several factors, including on-board monitoring capability; portable, inexpensive computational power; pattern recognition, adaptive learning, and signal processing technologies; and data and record keeping systems. Dramatic advances in microprocessor power and versatility, combined with growing developments in pattern recognition, adaptive learning (neural networks), and signal processing, would promise substantial breakthroughs in the future for application of prognostic methodologies in real time or near real time situations. However, existing field equipment inventories must be utilized and managed employing the best applicable technologies and making maximum use of available experience and expertise.

Acknowledgements

I would like to acknowledge the significant contributions Dr. Austin Lemoine has made to this paper, and to thank him for his advice and encouragement.

BIBLIOGRAPHY

Chikkara, Raj S., and Folks, J. Leroy, 1989, *The Inverse Gaussian Distribution*, Marcel Dekker, Inc.

Cox, D.R., and Oakes, D., 1984, *Analysis of Survival Data*, Chapman and Hall.

D'Agostino, R.B., and Stephens, M.A., 1984, *Goodness-of-Fit Techniques*, Marcel Dekker.

Feyock, S. and Karamouzis, S., 1991, Design of an intelligent information system for in-flight emergency assistance, *Proceedings of Goddard Conference on Space Applications of Artificial Intelligence*, National Aeronautics and Space Administration, Washington, D.C., pp. 295-306.

Grove, R.C., 1979, Investigation of Advanced Prognostic Analysis Techniques, Technical Report Number NRTC-78-47-R (USARTL-TR-79-10), Army Research and Technology Laboratories, Fort Eustis, Virginia.

Hines, William W., and Montgomery, Douglas C., 1980, *Probability and Statistics in Engineering and Management Science*, 2nd. ed., Wiley.

Lawless, J.F., 1982, *Statistical Models and Methods for Lifetime Data*, Wiley, New York.

Lemoine, A.J. and Wenocur, M.L., 1985, On failure modeling, *Naval Research Logistics Quarterly*, Vol. 32, pp. 497-508.

McDuff, R. Simpson, P.K., and Gunning, D., 1989, An investigation of neural networks for F-16 fault diagnosis; I - System description, in *AUTOTESCON '89, IEEE International Automatic Testing Conference, IEEE Press, New York, pp. 351-357.*

Microsoft Excel for Windows, version 4.0a, copyright 1985-92, Microsoft Corporation.

Microsoft Excel Solver was developed by Frontline Systems, Inc., Incline Village, NV.

Pusey, H.C. and Pusey, S.C., Eds., 1990, Current practices and trends in mechanical failure prevention, *Proceedings of the 44th Meeting of the Mechanical Failures Prevention Group*, Vibration Institute, Willowbrook, Illinois.

Ray, A., Wu, M.K., Carpino, M., and Lorenzo, C.F., 1992, Fatigue damage control of aerospace systems for high performance and extended life, in: *Proceedings of Winter Annual Meeting of the American Society of Mechanical Engineers - Intelligent Control Systems*, ASME, New York, pp. 7-16.

Shives, T.R. and Murtaugh, L.J., Eds., 1988, *Detection , Diagnosis, and Prognosis of Rotating Machinery to Improve Reliability, Maintainability, and Readiness throught the*

Application of New and Innovative Techniques, Conference Proceedings, Cambridge University Press, New York.

Singpurwalla, N.D., 1992, Survival in Dynamic Environments, Technical Report, GWU/IRRA/Serial TR-92/1, George Washington University.

Singpurwalla, N.D., 1993, The warranty problem: Its statistical and game-theoretic aspects, SIAM *Review*, Vol. 35, pp. 17-42.

Sobczyk, K. and Spencer, B.F., 1992, *Random Fatigue: From Data to Theory*, Academic Press, New York.