ORNL/TM-2015/373

Development of a First-of-a-Kind Deterministic Decision-Making Tool for Supervisory Control System

Sacit M. Cetiner
Roger A. Kisner
Michael D. Muhlheim
David L. Fugate

July 2015

OAK RIDGE NATIONAL LABORATORY
MANAGED BY UT-BATTELLE FOR THE US DEPARTMENT OF ENERGY

DOCUMENT AVAILABILITY

Reports produced after January 1, 1996, are generally available free via US Department of Energy (DOE) SciTech Connect.

Website http://www.osti.gov/scitech/

Reports produced before January 1, 1996, may be purchased by members of the public from the following source:

National Technical Information Service
5285 Port Royal Road
Springfield, VA 22161
Telephone 703-605-6000 (1-800-553-6847)
TDD 703-487-4639
Fax 703-605-6900
E-mail info@ntis.gov
Website http://www.ntis.gov/help/ordermethods.aspx

Reports are available to DOE employees, DOE contractors, Energy Technology Data Exchange representatives, and International Nuclear Information System representatives from the following source:

Office of Scientific and Technical Information
PO Box 62
Oak Ridge, TN 37831
Telephone 865-576-8401
Fax 865-576-5728
E-mail reports@osti.gov
Website http://www.osti.gov/contact.html

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

ORNL/TM-2015/373

Reactor and Nuclear Systems Division

DEVELOPMENT OF A FIRST-OF-A-KIND DETERMINISTIC DECISION-MAKING TOOL FOR SUPERVISORY CONTROL SYSTEM

Sacit M. Cetiner
Roger A. Kisner
Michael D. Muhlheim
David L. Fugate

Date Published: July 31, 2015

Prepared by
OAK RIDGE NATIONAL LABORATORY
Oak Ridge, Tennessee 37831-6283
managed by
UT-BATTELLE, LLC
for the
US DEPARTMENT OF ENERGY
under contract DE-AC05-00OR22725

CONTENTS

LIST OF FIGURES

Figure **Page**

LIST OF TABLES

Table **Page**

ACRONYMS

ALMR	advanced liquid-metal reactor
ART	Advanced Reactor Technologies
DOE	US Department of Energy
EM	electromagnetic
ESFAS	engineered safeguards features actuation system
FW	feedwater
ICHMI	Instrumentation, Controls, and Human-Machine Interface
IHTS	intermediate heat transport system
KT	Kepner-Tregoe
MAUT	multi-attribute utility theory
ORNL	Oak Ridge National Laboratory
PHTS	primary heat transport system
PNNL	Pacific Northwest National Laboratory
PRA	probabilistic risk assessment
PRISM	Power Reactor Innovative Small Module
RPS	reactor protection system
SCS	supervisory control system
SG	steam generator

ACKNOWLEDGMENTS

This project is funded by the U.S. Department of Energy Office of Nuclear Energy under the Instrumentation, Controls, and Human-Machine Interface (ICHMI) technical area of the Advanced Reactor Technologies (ART) program.

ABSTRACT

Decision-making is the process of identifying and choosing alternatives where each alternative offers a different approach or path to move from a given state or condition to a desired state or condition. The generation of consistent decisions requires that a structured, coherent process be defined, immediately leading to a decision-making framework. The overall objective of the generalized framework is for it to be adopted into an autonomous decision-making framework and tailored to specific requirements for various applications. In this context, automation is the use of computing resources to make decisions and implement a structured decision-making process with limited or no human intervention. The overriding goal of automation is to replace or supplement human decision makers with reconfigurable decision-making modules that can perform a given set of tasks reliably.

Risk-informed decision-making requires a *probabilistic assessment* of the likelihood of success given the status of the plant/systems and component health, and a *deterministic assessment* between plant operating parameters and reactor protection parameters to prevent unnecessary trips and challenges to plant safety systems.

The implementation of the probabilistic portion of the decision-making engine of the proposed supervisory control system was detailed in previous milestone reports. Once the control options are identified and ranked based on the likelihood of success, the supervisory control system transmits the options to the deterministic portion of the platform.

The deterministic multi-attribute decision-making framework uses variable sensor data (e.g., outlet temperature) and calculates where it is within the challenge state, its trajectory, and margin within the controllable domain using utility functions to evaluate current and projected plant state space for different control decisions. Metrics to be evaluated include stability, cost, time to complete (action), power level, etc.

The integration of deterministic calculations using multi-physics analyses (i.e., neutronics, thermal, and thermal-hydraulics) and probabilistic safety calculations allows for the examination and quantification of margin recovery strategies. This also provides validation of the control options identified from the probabilistic assessment. Thus, the thermal-hydraulics analyses are used to validate the control options identified from the probabilistic assessment.

Future work includes evaluating other possible metrics and computational efficiencies.

EXECUTIVE SUMMARY

This technical report was generated as a product of the Supervisory Control of Multi-Modular Advanced Reactor Plants project within the Instrumentation, Controls and Human-Machine Interface (ICHMI) technical area under the Advanced Reactor Technologies (ART) Research and Development program of the US Department of Energy.

The supervisory control system (SCS) includes a *probabilistic* portion and a *deterministic* portion to identify control options, select an option, and initiate a command. The project members at the Oak Ridge National Laboratory selected probabilistic risk assessment (PRA) as the analytical method and tool for accomplishing the probabilistic aspect of the decision-making process. The technical details of the probabilistic portion of the decision-making engine are documented in the previous milestone report.

This report documents the technical basis of the *deterministic* decision-making function for the SCS. *Deterministic behavior* means that the systems operate in a predictable and repeatable manner. That is, deterministic systems will produce the same outputs for the same set of input signals. *Decision-making* is the process of identifying and choosing alternatives based on an agreed-upon set of metrics and preferences established by the decision maker. Indirectly implied in decision-making is that there are alternative options to be considered. Each option offers a different approach or trajectory to move from a given state or condition to a desired state or condition.

This report presents a framework by which deterministic considerations can be taken into account for decision-making. More specifically, this report adapts classic *utility theory* as its mathematical method to realize decision-making for accomplishing the deterministic portion of supervisory control. Utility theory offers a unified conceptual framework for uniformly quantifying the consequences of decision alternatives, defining their weights, and dealing with associated uncertainties. In the context of supervisory control, *utility* is simply defined as a measure of preferences over a given set of actions. The consequence of a decision is the result that follows from that decision.

The notion that mathematical analysis should guide rational choice under uncertainty was formulated as early as the seventeenth century. Since its introduction and early use in economics, utility theory has been adapted for multiple uses in portfolio management, business, government, and the military. Practical improvements include inclusion of the capability to compare multiple attributes as would be encountered in complex organizational decisions. Only recently has utility theory been considered for use in decision-making for real-time systems such as automated automobile driving. Others are integrating utility functions with optimal control so that the cost function is more reflective of risk than it would otherwise have been. A possible application of a quadratic cost function is electric utility load-generation control.

This report describes several new developments in the application of utility functions to real-time decision-making. Several accomplishments of this work contribute distinctive properties to supervisory decision-making:

1. *Special utility functions:* Special utility functions are defined that express utility values for set-point type process variables. Specifically, a Gaussian distribution is adapted as the mapping mechanism of the utility function. A bundle of utility functions may incorporate both Gaussian (optimizable) and monotonic functions to provide a uniform means of comparing preferences.

2. *Wide range utility functions:* Utility functions are defined in this work that range from negative infinity to positive unity in contrast with the zero to unity range of previously defined utility functions. With this arrangement, all attributes are uniform for positive contributions, but the bundle sum can be penalized for a particular variable in a seriously degraded region.

3. *Diversely integrated:* Integration of diverse attributes into a uniformly calculated (single) utility value permits robust decision-making across a selection of alternative recovery solutions. The diversity of attributes includes process variables that capture system dynamics, economics, maintenance, and reliability/risk parameters from PRA output.

4. *Integration of rapid simulation*: Both static and dynamic variable simulations are used to estimate minimum and maximum expected process variables from alternative solution plans. These min/max values are used to generate the utility values for conducting a weighted comparison.

The developments discussed in this report combine PRA-based decision-making with deterministic decision-making in a formal manner to achieve a risk-informed comparative decision-making process that has not been presented in the literature. Based on the results described in this report, the next steps are to execute a series of examples that involve simple piping and fluid transport components. The examples demonstrate the combined PRA and deterministic decision-making capabilities over a range of system and component failures.

1. INTRODUCTION

This document is an interim report providing an update on the progress being made on the autonomous decision-making component of the proposed supervisory control system (SCS). This system is the key capability for autonomous control. This report documents the deterministic portion of the autonomous risk-informed decision-making process within the SCS, which combines probabilistic assessments with deterministic rules and insights.

Automation refers to the use of computing resources to perform repeatable tasks based on a predetermined set of rules and actions.

Autonomy, on the other hand, refers to the use of computing resources to make decisions and implement a structured decision-making process with limited or no human intervention. The overriding goal of autonomy is to replace or supplement human decision makers with reconfigurable decision-making modules that can perform a given set of tasks reliably.

Decision-making is the process of identifying and choosing alternatives based on an agreed-upon set of metrics and preferences established by the decision maker. Indirectly implied in decision-making is that there are alternative options to be considered. Each option offers a different approach or trajectory to move from a given state or condition to a desired state or condition.

The generation of consistent decisions requires that a structured, coherent process be defined, which immediately leads to a decision-making framework. The generalized framework for autonomous decision-making can be adopted and tailored to specific requirements for various applications.

This section provides an outline of the generalized decision-making framework that was presented and described in greater detail in previous reports [1-1, 1-2]. It also summarizes the key concepts of a risk-informed decision-making process for an SCS.

1.1 GENERALIZED FRAMEWORK FOR AUTONOMOUS DECISION-MAKING

Ultimately, the objective of a decision-making process is to consider uncertainties and evaluate options for the current component and system status. Hence it is quite possible that evaluation and assessment steps will require consideration of multiple attributes of a system, components, or elements of a system, or their future states. This is especially true for large-scale, complex systems such as a nuclear power plant.

While there are minor differences in the literature about the necessary and sufficient steps for decision-making, the decision-making process for the SCS is based on three fundamental elements:

1. identification: define decision alternatives,

2. evaluation: assess alternative decisions,

3. resolution: generate a single solution or a single trajectory, and collect the steps needed to finalize an action, and

4. action: execute control actions.

These elements, as illustrated in Fig. 1, define the generalized autonomous decision-making framework.

Fig. 1. Elements considered within the generalized framework for autonomous decision-making.

The steps shown in Fig. 1 form a generalized framework within which various decision-making methods can be implemented. This framework can be applied to a variety of engineering problems.

1.2 AUTONOMOUS DECISION-MAKING FOR SUPERVISORY CONTROL

The generalized framework provides a conceptual structure that only includes abstract rules, elements, and the relationships between them. Adoption of this framework for application to an SCS requires that a specific implementation be created to define how the individual objectives will be accomplished. This section provides a functional definition and some generic specifications for the proposed autonomous decision-making framework for an SCS. Details of this implementation and the functionality of the architecture were previously reported [1-1].

1.2.1 High-Level Description of the SCS

The SCS shall comply with the following high-level requirements:

1. The SCS shall be implemented as a non-safety-related system.

2. The SCS shall follow all the applicable rules and regulations regarding the separation and isolation of safety- and non-safety-related systems.

3. The SCS shall not perform any safety-related function.

4. The SCS shall not interfere with the functionality and operation of any safety system.

5. The SCS shall not override operator directives.

These requirements are enforced to define the domain of operation of the SCS. Implementing the SCS as a non-safety–related system avoids placing an undue regulatory burden on the vendor and the owner—especially considering the complexity of the system.

The fundamental assumption that goes into the design of the SCS is that, should the SCS fail to act during a transient, then the safety system will independently initiate and bring the plant to a nominal or acceptable shutdown state.

1.2.2 Definition of Terms

The following terms are used throughout the report. A brief glossary is provided below to avoid misinterpretation and to maintain consistency throughout this report.

Risk
In safety analysis, risk is defined as the product of frequency and consequence. However, in the context of the proposed supervisory control architecture and the autonomous decision-making framework, risk is defined as the probability of challenging a safety system, or the probability of safety actuation. The goal of the SCS is to avoid challenging a safety system.

Controllable Domain
An SCS is required to support human decision-making under normal operating conditions and to make autonomous decisions. All of the possible states that the plant can assume constitute the controllable domain. The boundary of the controllable domain is primarily defined by the trip setpoints of the reactor protection system (RPS) or the engineered safeguards features actuation system (ESFAS). This domain is illustrated in light blue and orange in Fig. 2.

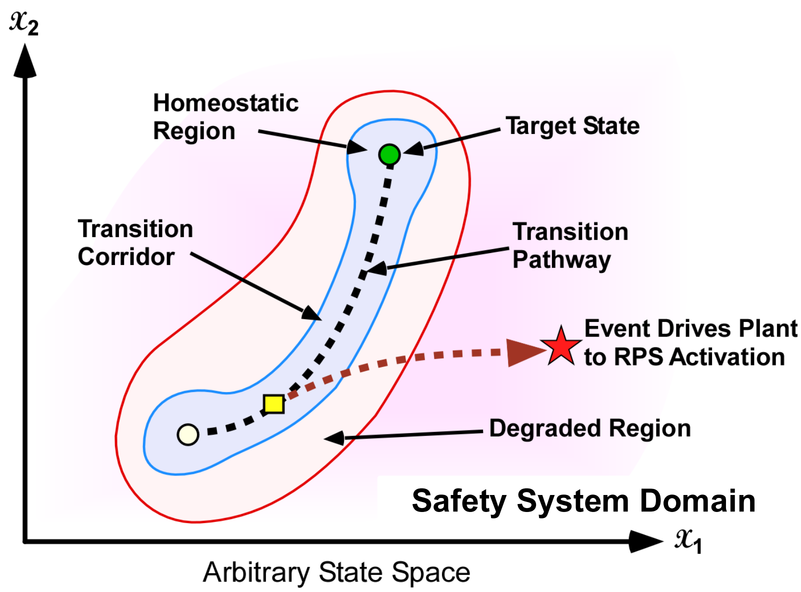

Fig. 2. A conceptual state space formed by arbitrary state variables x_1 and x_2.

Challenge Surface
The surface of the controllable domain is called the challenge surface, beyond which a safety system actuation is warranted by the design of the plant. The challenge surface is illustrated with the red line in Fig. 2.

Safety System Domain
This is the domain outside the challenge surface of the plant state space. The safety system domain is illustrated in fading purple in Fig. 2. Because this region represents the safety functions (e.g., protection system functions), it is outside the scope and capabilities of the control system.

Probability of Departure from Controllable Domain

This metric is an indication of proximity of the plant state to the challenge surface. While there might be numerous ways to define this probability metric, it can be simply defined as a distance function between the current plant state and the closest point on the challenge surface. The closer the plant gets to the surface, the higher the probability of the protection system actuation. Higher order moments of the states can also be considered, such as the rate of approach.

1.3 SUPERVISORY CONTROL SYSTEM ARCHITECTURE

The proposed architecture for autonomous decision-making implements the general framework using two methods. In the first method, the probabilistic method is implemented using probabilistic risk assessment (PRA) techniques to identify decision options. In the second method, the deterministic portion is implemented using utility theory to evaluate the alternatives identified by the probabilistic portion and to generate a single solution, or the resolution of the autonomous decision-making process. These methods are shown in Fig. 3. The cost function for finding the optimal or desired decision is determined by the evaluation metric. Additional constraints, such as regulatory rules and operating guidelines, can be enforced in the deterministic evaluation phase.

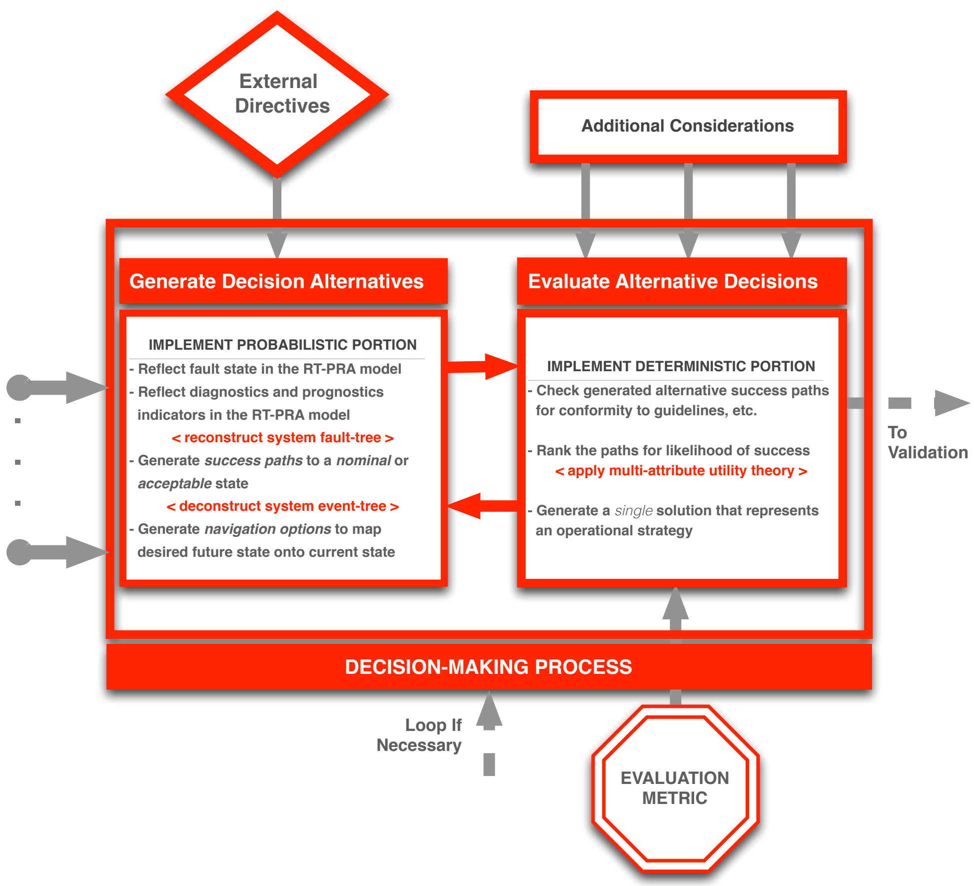

Fig. 3. The proposed framework for autonomous decision-making adopted for the supervisory control system.

Fig. 4 shows the functional architecture of the SCS and illustrates how the decision-making block in Fig. 3 relates to the overall architecture.

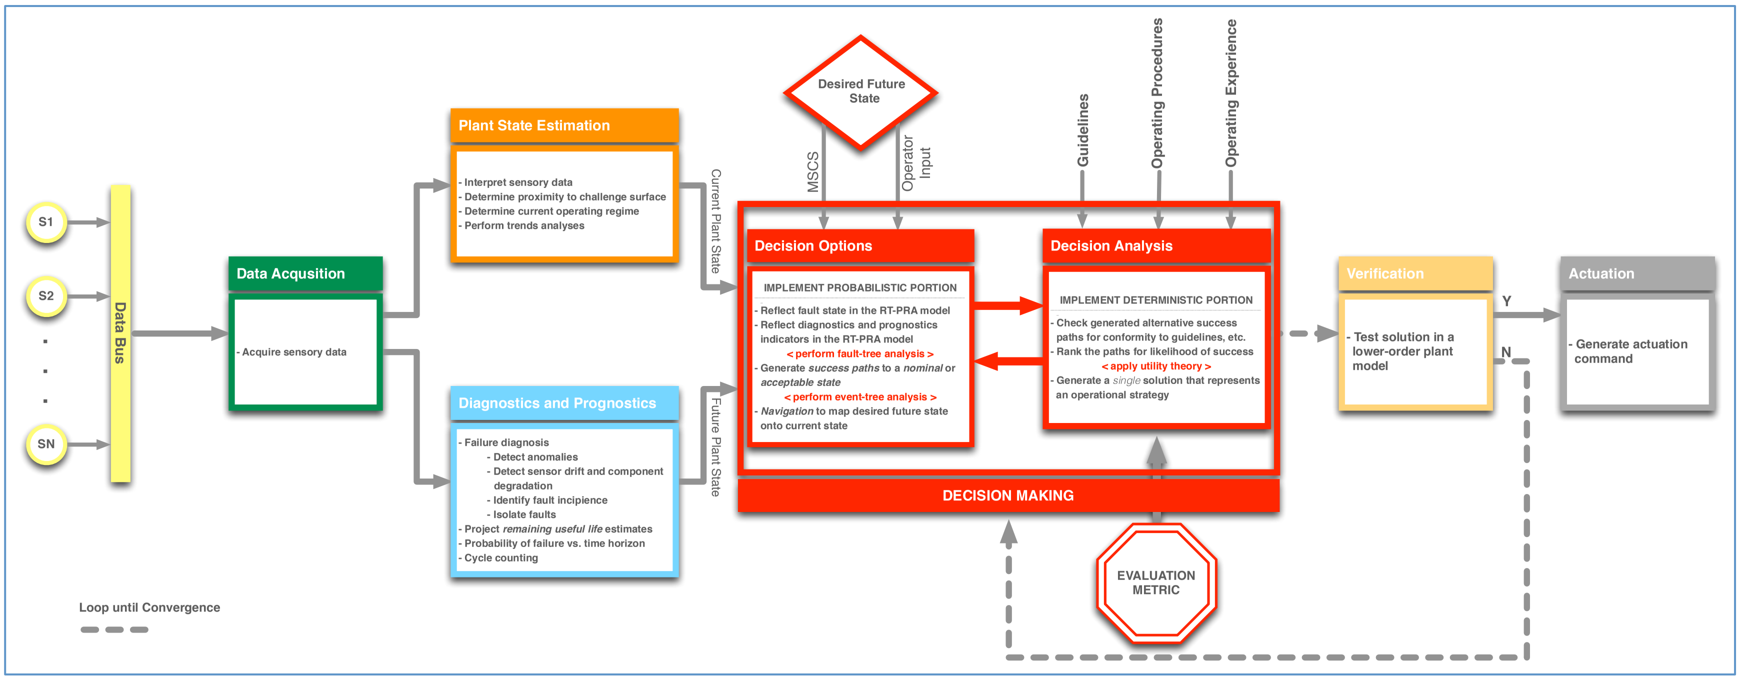

Fig. 4. Functional architecture of the SCS with a specific implementation of the generalized decision-making framework.

1.4 RISK-INFORMED DECISION-MAKING FRAMEWORK

Control systems should be "capable of maintaining system variables within prescribed operating ranges" [1-3]. More specifically, plant control systems in general and reactor control systems in particular in nuclear power plants are designed to maintain system variables (such as reactor power, coolant flow rate, power-to-flow ratio, reactor outlet temperature, coolant level, and turbine status) within prescribed operating ranges. Deviating outside these operating ranges (e.g., exceeding the setpoint of a control system variable) results in a plant transient and a challenge to plant mitigating systems, including a potential challenge to plant safety systems.

The use of a risk-informed decision-making framework for the control system (i.e., one that couples a probabilistic assessment with a deterministic assessment) (1) allows consideration of a broader set of potential challenges, (2) provides a logical means for prioritizing these challenges, (3) allows consideration of a broader set of resources to prevent or mitigate these challenges, (4) explicitly identifies and quantifies sources of uncertainty in the analysis, and (5) leads to better decision-making by providing a means to test the sensitivity of the results against key assumptions [1-4].

1.4.1 Probabilistic Interface to Risk-Informed Decision-Making

The purpose of this document is to define a framework for conducting the deterministic multi-attribute decision-making portion of the risk-informed framework. The probabilistic assessment portion of the risk-informed decision-making framework provides the control options for successfully avoiding trip setpoints given the status of the plant/systems and the likelihood of success for each of those options. The probabilistic portion of the decision-making engine was detailed in a previous milestone report [1-2].

Based on plant operating status, component health, and equipment failures, the decision-making capabilities for the supervisory control system uses the probabilistic analyses to identify a set of control options. These options, if taken, should prevent the actuation of the protection system. The possibility for one or more outcomes distinguishes probabilistically informed decision-making implemented in real time from more traditional decision-making.

The probabilistic portion of the decision-making algorithm ranks the likelihood of success of each decision path based on the current system/plant status and component health. Based on the likelihood of success metric under these conditions, the decision-making algorithm automatically chooses the top candidate control options as decision alternatives for the execution of the corresponding set of corrective actions. Selecting any of the control options would allow operations to continue by maintaining system status within the acceptable region.

A probabilistically based decision-making algorithm would simply select the option with the greatest likelihood of success; however, this may not be the best choice based on other criteria. For example, the most likely option for avoiding a trip setpoint probabilistically could be to manually shut down the reactor, but deterministic factors such as reduced generation of heat (i.e., power reduction) may re-rank this option to the least favorable of the choices.

Once the control options are identified and ranked, the supervisory control system transmits the options to the deterministic portion of the platform.

1.4.2 Deterministic Interface to Risk-Informed Decision-Making

The probabilistic portion of the risk-informed decision-making framework identifies those actions that would maintain system status within the acceptable region if taken. However, knowing the likelihood of avoiding a trip setpoint does not inform the system of where within the challenge surface the system is operating, how close it is operating to the degraded region, and whether the operating trajectory is toward the safety system domain. The deterministic analyses of those control options selected probabilistically provides the necessary information to optimize the selection of the control action to be taken. The metric for the deterministic analysis is based on maintaining and/or controlling the heat balance from the reactor core to the ultimate heat sink and must also reflect the success of maintaining this heat balance. Thus, the deterministic assessment between plant operating parameters and reactor protection parameters is used to identify the margin between operating conditions and trip setpoints.

Operation anywhere within the homeostatic region is considered normal. The plant control systems employ appropriate feedback control strategies, provided that the system parameters are maintained within the homeostatic region. Should operation be driven into the degraded region, the control objectives become (1) to maintain continuous and uninterrupted delivery of principal products of the system, if possible; (2) to prevent or minimize equipment damage; and (3) to preclude initiation of the plant safety and protection systems. Transitioning out of the degraded region back into the homeostatic region may require faster response control options to maintain system variables within the challenge surface. If a system variable transitions into the safety space region, it enters the domain of the protection system, which is independent of and isolated from the control system.

Historically, safety margins have been set conservatively (for example, based on design and operational limits) to compensate for uncertainties in measured values and system responses. Because safety is critical to the successful operation of the nuclear power plant fleet, motivations are strong to better characterize and manage safety and its associated margins. Among these motivations is having improved knowledge of both the qualitative and quantitative aspects of safety margins to provide for enhancements and improvement in plant performance. Magnitude and speed can be important if the parameter of interest is close to, or moving rapidly toward, a reactor trip setpoint. The integration of deterministic calculations using multi-physics analyses (i.e., neutronics, thermal, and thermal-hydraulics) and probabilistic safety calculations allows for the examination and quantification of margin recovery strategies. This also provides confirmation of the control options identified from the probabilistic assessment. Thus, the thermal-hydraulics analyses are used to assess the control options identified from the probabilistic assessment by providing the following information:

1. How far is (are) the variable(s) of interest from the preferred transition corridor (magnitude of correction)?
2. Is the variable within the allowable variance along the Transition Pathway (tolerance)?
3. What is the path within the homeostatic or degraded region (direction)?
4. How fast a correction must be made (speed of correction)?

Most decisions require more information than risk alone. The use of a risk-informed approach in a control system allows probabilistic insights to be coupled with deterministic factors of concern, such as magnitude of deviation from a nominal setpoint and speed of parameter adjustment needed. Each control option has a different probability of success and can be linked to magnitude of deviation, speed of adjustment, and other metrics of interest. Chapter 2 introduces the theory, and provides the technical bases for calculating the metrics for the deterministic for deterministic assessments. Chapter 3 describes how the methodology is applied within the SCS, and defines the interfaces and plant variables used to address the questions above.

1.5 CHAPTER 1 REFERENCES

1-1. S. M. Cetiner, M. D. Muhlheim, G. F. Flanagan, D. L. Fugate, and R. A. Kisner, "Development of an Automated Decision-Making Tool for Supervisory Control System," ORNL/TM-2014/363 (SMR/ICHMI/ORNL/TR-2014/05), Oak Ridge National Laboratory, Oak Ridge, Tenn., Sept. 2014.

1-2. S. M. Cetiner and M. D. Muhlheim, *Implementation of the Probabilistic Decision-Making Engine for Supervisory Control*, ORNL/SPR-2015/140, Oak Ridge National Laboratory, Oak Ridge, Tenn., March 2015.

1-3. NUREG-0800, Rev. 5, *Standard Review Plan*, Chapter 7.7, "Control Systems," March 2007.

1-4. Risk-Informed Security Regulations Workshop, Sponsored by the U.S. Nuclear Regulatory Commission's Office of Nuclear Regulatory Research and Office of Nuclear Security and Incident Response, Albuquerque, N.M., Sept. 14–15, 2010.

2. ELEMENTS OF DETERMINISTIC DECISION-MAKING

An overview of the decision framework is provided to allow for a clear presentation of the mathematical and statistical concepts, notation and structure involved in decision modeling. The necessary data interfaces to be able to process the output generated by the PRA calculation for a decision node is provided in Chapter 3.

2.1 UTILITY THEORY

The notion that mathematical analysis should guide rational choice under uncertainty was formulated as early as the seventeenth century. The concept of maximizing expected monetary return dates to the eighteenth century, when the principle of utility was conceived by Jeremy Bentham (1748–1832), Thomas Bayes (1702–1761), and Daniel Bernoulli (1700–1782) as a means of maximizing the greatest happiness or pleasure when making decisions from among alternative choices by applying rational analysis [2-1]. The bases of their early work can be summed up by the standard model of human motivation: given a person has a desire Y, and if they believe that by performing action X, they can achieve Y, then they will choose X [2-2]. The early work, which merged logic and moral principles, was expanded over the centuries by numerous philosophers and intellectuals to encompass economics, statistics, and psychology.

Bernoulli offered the approach that any gain brings a utility inversely proportional to the total wealth of an individual [2-3]. Mathematically, the relationship can be shown as Eq. (2-1a), in which a utility u with an increase in wealth Δz, relative to the current wealth z, is related to a positive constant c. For a small Δ, the expression can be written as the differential Eq. (2-1b). Integrating the differential equation yields Eq. (2-1c), where z_0 is the constant of integration—interpreted as the wealth necessary to obtain zero utility.

$$u(z + \Delta z) - u(z) = c\frac{\Delta z}{z} \tag{2-1a}$$

$$du(z) = c\frac{dz}{z} \tag{2-1b}$$

$$u(z) = c\log(z) - \log(z_0) \tag{2-1c}$$

From this derivation, Daniel Bernoulli determined that the utility of wealth is not linear but logarithmic. Thus, utility theory became analytical. It is interesting to note that Bernoulli's formulation has the characteristic of being unbounded. Treatment of bounded utility does not appear until the twentieth century with the work of Savage [2-4].

Further evolution of utility theory involves the concept of *marginal utility*, which emerged not long after Bentham's work. Marginal utility refers to the gain resulting from an increase in the consumption of a good or a service or, conversely, the loss resulting from a decrease in consumption. The marginal utility of a good is derived from its most important use to a person. The work of Gregory King (1648–1712), refined by the work of Charles Davenant (1656–1714), led to the King-Davenant law of demand that describes the inverse relationship between price and quantity [2-5]. Jules Dupuit, chief engineer of Paris, used water consumption in a city as an example of marginal utility. Dupuit, the first economist to present a cogent discussion of marginal utility, argued that if it were difficult to obtain water and consumers had to pay large use fees and they purchased it, it had to provide the household with at least that much utility.

However, he further argued that as more water was introduced to the city, the point would come when households would not require more. The relationship is simply illustrated in Fig. 5.

Utility theory and numerous associated corollaries were refined in the nineteenth and twentieth centuries, resulting in progressively more organized methods to rank alternative decisions in their *relative* order of consumer preference. (Note that it is impossible to reliably measure the *absolute* economic utility derived from a good or service.) Since a consumer's choice is constrained by price as well as disposable income, the rational consumer, it is assumed, will not spend money for a unit of goods or services unless the marginal utility is perceived as at least equal to or greater than that of a unit of another good or service. Therefore, the price of a good or service may be associated with its marginal utility; the consumer will rank personal preferences accordingly. This relationship is essentially a restatement of the standard model of human motivation previously described above.

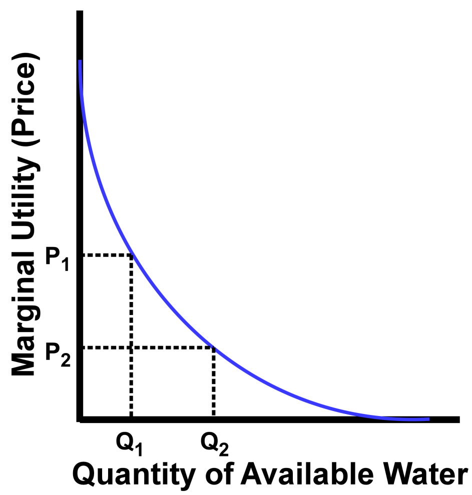

Fig. 5. Dupuit's argument for the law of diminishing marginal utility.

2.1.1 Modern Utility Theory

Modern utility theory for decision-making originated in 1944 with von Neumann and Morgenstern, who postulated a set of axioms using objective probabilities [2-6]. They demonstrated that a utility value could be assigned to each possible decision outcome in a manner that allows the decision-maker to always select the outcome with the highest expected utility. This result, denoted the expected utility hypothesis, lends true credibility to decision outcomes derived using utility theory. Savage provided a major contribution in 1954 when he presented the first axioms concerning subjective expected utility, which incorporates subjective probability estimates made by the decision maker. Numerous other researchers have investigated a variety of decision-making characteristics beginning mid-twentieth century [2-7–2-15]. These investigations have provided a basis for developing additive and multiplicative multi-linear utility functions. One important finding is that the form of the utility function applicable in a given decision-making circumstance depends on the independence conditions that exist among the attributes. Another discovery is that consumer attitude towards risk is reflected in the shape of the utility. With a Bernoulli utility function representation of risk preferences, an individual is risk averse if and only if the inequality of Eq. (2-2) is true.

$$\int_{-\infty}^{+\infty} u(x)\, dF(x) \leq u\left(\int_{-\infty}^{+\infty} x\, dF(x)\right) \text{ for all } F(\cdot) \tag{2-2}$$

where $u(x)$ is the utility function, and $F(x)$ is the distribution function for outcome variable x. The effect, called Jensen's inequality function, is shown in Fig. 6 [2-16].

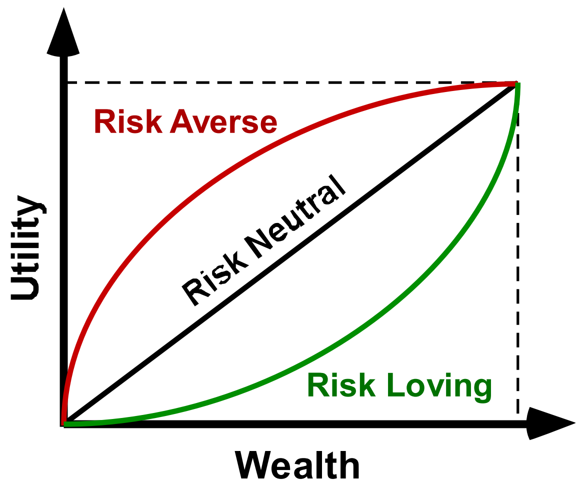

Fig. 6. Three different expected utility maximizers reflect different risk preferences.

Utility theory as applied in the eighteenth and nineteenth centuries dealt with one or two variables and hence one or two utility attributes. Humans possess a limited capacity to perform complex information processing in terms of quantifiable characteristics they can process. The assignment of utility to the attributes also tended to be strongly subject to user interpretation. The desire to expand utility theory and decision-making to complex decisions for large-scale organization systems involving purchases, for example, led to multi-attribute utility theory (MAUT). MAUT is commonly applied to those decision-making problems in which the possible decision consequences or alternatives are characterized by many attributes.

2.1.2 Kepner-Tregoe

The Kepner-Tregoe (KT) model was developed by Charles Kepner and Benjamin Tregoe in the mid-twentieth century for the purpose of systematically evaluating a quantity of qualitative issues that have significant trade-offs between them. The KT model was developed as a result of observing methods used successfully by military officers to determine the best course of action among multiple alternatives. The KT method is very similar to utility theory and MAUT in particular and can be adapted to evaluation of quantitative issues. The steps in KT methodology are as follows:

1. Clearly define the problem at a high level.
2. Determine strategic requirements (musts), operational objectives (wants), and constraints (limits).
3. Rank objectives and assign relative weights.
4. Generate alternatives.
5. Assign a relative score for each alternative on an object-by-objective basis.
6. Calculate the weighted score for each alternative and identify the top two or three alternatives.

7. List adverse consequences for each top alternative, and evaluate the probability and severity.
8. Make a final, single choice between the top alternatives.

Although KT literature makes no reference to utility theory, it is clear that this methodology is a practical but simplified implementation. The KT method has been taught for several decades as a process to help business organizations make critical (usually high-cost) decisions.

Risk is handled only as an input to the KT decision process when identifying alternatives; however, the weighting process assumes linear utility functions, which have no predilection for risk (the linear case of Fig. 2). In fact, the generation of a utility function is not part of the KT process; hence the implied utility function simply becomes a linear multiplier.

2.1.3 Multi-Attribute Utility Theory

Multi-attribute utility theory is a quantitative comparison method used to evaluate dissimilar measures of costs, risks, and benefits. The underpinning of MAUT is the direct use of utility functions that transform diverse criteria to one common, dimensionless scale (0 to 1) known as a multi-attribute utility. An alternative's quantitative data (values) can be converted to a utility score through utility functions that ranges between zero (unacceptable) and unity (excellent). Each criterion is weighted according to importance. The procedure to identify the preferred alternative requires multiplying each alternative's utility score results across all assessed criteria, then summing these products. The preferred alternative will have the highest summation score.

Utility functions (and MAUT) are used when quantitative information is known about each alternative—a distinction over KT. Better estimates and comparison of the alternative performance are possible with MAUT than with KT because of the utility functions. Utility functions (for every decision criterion) are created on the basis of the data for each criterion. The utility functions transform an alternative's raw score [which may be in engineering dimensions, discrete state information (e.g., Celsius, kilograms, liters per minute, dollars), or a dimensionless utility score ranging between 0 and 1]. Criteria are weighted according to importance. Each normalized alternative's utility score results are multiplied for all of an alternative's criteria to identify the preferred alternative. The preferred alternative will have the highest total score.

The MAUT evaluation method works well for complex decisions with multiple criteria and many alternatives. Additional alternatives can be added to a MAUT analysis with minimal effort. Any number of alternatives can be scored against them once utility functions have been developed.

2.2 UTILITY THEORY AS A DECISION-MAKING TOOL

Utility theory has not yet been integrated into decision-making for real-time automated industrial processes. Utility theory has not found significant support among control scientists and engineers, mainly because of the following limitations [2-31]:

1. *maximizing behavior* cannot be consistently rational in any non-zero-sum multi-agent situation,
2. the *axiom of independence* cannot be justified in real life situations, and
3. value/utility functions are static (i.e., they do not change in time).

Although the first two limitations have restricted the widespread use of utility theory in control applications, the third limitation can be remedied by creating a system to upgrade the utility functions periodically to remain cognizant with changing plant conditions—something not considered in classical utility theory. An application that has recently applied utility theory in a control system is automated driving systems (see Section 2.3). The automated driving systems combine rule-based systems with utility theory to overcome the limitations cited above.

In classical utility theory, the utility function can be curved to reflect the risk attitude of the consumer as shown in Fig. 6 (i.e., from risk-averse to risk-loving). As the concepts of utility functions are applied in real-time automation, the risk preferences take on a new meaning in automation space that is related to the significance of anticipating the approach to a threshold. For a control system, the current state can transition into an infinite number of states—within the homeostatic, degraded, or safety space regions. This would imply an infinite number of end states. Fortunately, the probabilistic assessment limits the number of alternatives. The deterministic assessment evaluates the top probabilistic alternatives by mapping the state space onto a utility function. How this is done is provided below.

The current and future states of a plant system are represented by the two-dimensional state-space diagram, as illustrated in Fig. 7. Represented in this example are two arbitrary component temperatures. It is desired to maneuver the entire system's operation from the current state to a new target state. The maneuver is accomplished by the actions of local controllers that cause the temperatures to move from the current value to a new (target) value along a trajectory that lies within a transition corridor shown in the figure. The planned trajectory and associated transition corridor avoid passing through unfavorable regions (e.g., any protection system challenges). However, in the example, an event occurs partway through the maneuver that changes the system's trajectory. Without intervention, temperatures would be forced out of the transition corridor into the degraded region crossing T_{UB} (degraded state) eventually causing actuation of the RPS at T_{max} (trip state).

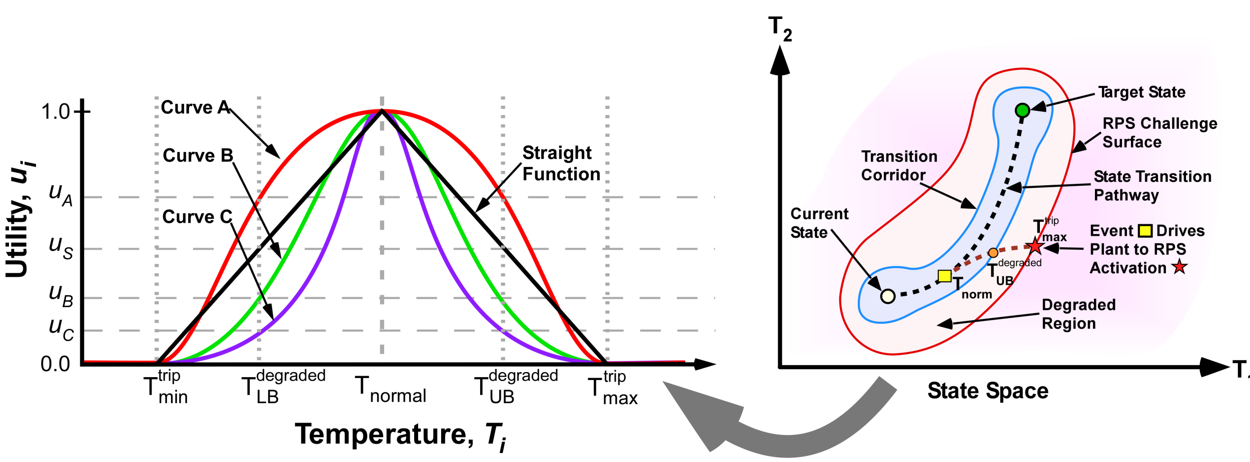

Fig. 7. How temperature movement in state space maps to a utility function.

A collection of utility functions as shown in the left portion of Fig. 7 can be formulated for the temperatures. The utility function ordinate value was chosen to observe the customary range of 0 to 1 from economics. The abscissa represents the temperature space with minimum- and maximum-temperature bounds for the challenge surface, T_{min} and T_{max} for the RPS, and boundary between the corridor and degraded region, T_{LB} and T_{UB}. A given temperature generates a specific utility value depending on which curve is employed. Curves A, B, and C are utility functions $u_i(T)$ based on Gaussian distributions around temperature T having three different variance parameters, σ_i, around the expected temperature value T_μ, which is T_{nom} for the example.

$$u_i(T) = e^{-\frac{(T-T_\mu)^2}{2\,\sigma_i^2}} \tag{2-3}$$

The three Gaussian curves (curves A, B and C) plus the straight-line curve (S)[1] illustrate varying tolerance in anticipating the approach to a threshold. Table 1 shows that for normal operating temperatures, as defined by the planned trajectory, all curves generate a perfect value (unity). Likewise, at the threshold crossing into the RPS challenge surface, all curves generate zero values, meaning that there is no utility for that temperature and beyond. The significance of curve shape becomes apparent as the temperature departs from normal and heads toward the degraded region: shape determines how important the distance from normal becomes. Curve A with a higher variance is tolerant of deviation, whereas curve C with a lower variance is not and reduces utility value rapidly away from normal. The straight line is considered tolerance neutral.

Table 1. Comparison of utility values having different shapes for the same temperature

Temperature	Curve	Utility Value (arbitrary)
T_{nom}	A, B, C and S	1.00
T_{UB}	A	0.66
T_{UB}	S	0.45
T_{UB}	B	0.25
T_{UB}	C	0.12
T_{max}	A, B, C and S	0.00

For this example, the utility values were restricted to the range zero to one as customarily restricted in economics applications. However, utility values do not necessarily need to be limited to that range. Later in this report examples will be given for utility functions with ranges from $-\infty$ to 1.Utility functions are traditionally summed as shown in Eq. (2-4) to obtain bundled utility values:

$$U = \sum_{i=1}^{N} \omega_i\, u_i(x_i) \tag{2-4}$$

where x_i is the independent utility variable, $u_i(x_i)$ is the utility function that describes the shape of the utility variable x_i, ω_i is the weight of the utility function u_i, and U is the total utility value of the bundle. Other functions can be applied, including multiplication (product sum) as well as unique functions that further define the weight of options.

Utility functions do not have to be monotonically increasing as typically adopted for game theory and economic applications. Furthermore, dynamic utility functions can be constructed such that they respond

[1] The straight-line (curve S) is represented by straight lines between T_{min} to T_{nom} and T_{nom} and T_{max}.

to varying plant conditions. However, the stability, reliability, and safety of this application have not yet been demonstrated.

2.3 ENGINEERING APPLICATIONS OF UTILITY THEORY

The preponderance of utility theory applications involves decision-making at an organizational level (e.g., governmental economics, business, and military). It has only been recently that the underlying concepts of utility theory have been considered for automated system application.

2.3.1 Advanced Driver Assistance Systems and Highly Automated Driving

The literature related to automated automobile driving systems suggests two leading decision-making processes. Many mathematical driving system models incorporate rule-based decision-making processes [2-18–2-20]; in comparison with the more sophisticated models that favor utility functions [2-21–2-24]. The principal advantage of utility functions over rule-based approaches is that multiple criteria can be weighed and compared quantitatively. Moreover, a utility function can be modified and extended with less effort than rule-based structures [2-25].

Autonomous driving has become a practical research topic over the last decade [2-26–2-28]. In intelligent transportation systems, most of the research work has focused on lane change assistant systems, which are inherently more difficult than finding the shortest or fastest path and have significant, immediate safety implications [2-29]. Recent work by BMW stands out in the application of utility functions [2-30]. This work specifically considers uncertainty by establishing a utility for each lane—each lane is evaluated based on driver settings (e.g., desired velocity); host vehicle properties; and environment information (e.g., road, lanes, and objects). The lane utility is described by means of a probability distributed stochastic variable, U, instead of a regular variable, u, as shown in Eq. (2-5).

$$U \sim N(\mu_U, \sigma_U^2) \tag{2-5}$$

An overall utility, U_k, for time step k is determined by using a weighting factor $\omega_{k,i}$. The formulation of the utility function is shown in Eq. (2-6).

$$U_k = \sum_{i=1}^{N} \omega_{k,i}\, u_{k,i} \tag{2-6}$$

where $u_{k,i}$ is the utility function for the ith attribute at time step k.

The utility is calculated for state vectors that represent the traffic situation development from the past to the future. These vectors contain the position and the velocity of surrounding vehicles. The future state has to be predicted; current and past state vectors can be stored and reused for past time steps. The method that BMW uses to determine the utility of a lane change is based on combining utility function and a rule-based system, as illustrated in Fig. 8. In this example, the utility of each lane (U_k) on a highway at any discrete time step (k) is calculated using a range of utility attributes as a time series within the *Discretionary Lane Changes* block. The utility of each lane varies as a function of road and traffic conditions obtained by sensor readings. The results of the calculation are then reassessed by the *Mandatory Lane Changes* block, which is a rule-based system, to assure that the decisions from the discretionary decision-making block would not violate the traffic rules or create safety issues. The approach by the BMW team was successfully tested August 2011on an actual public freeway.

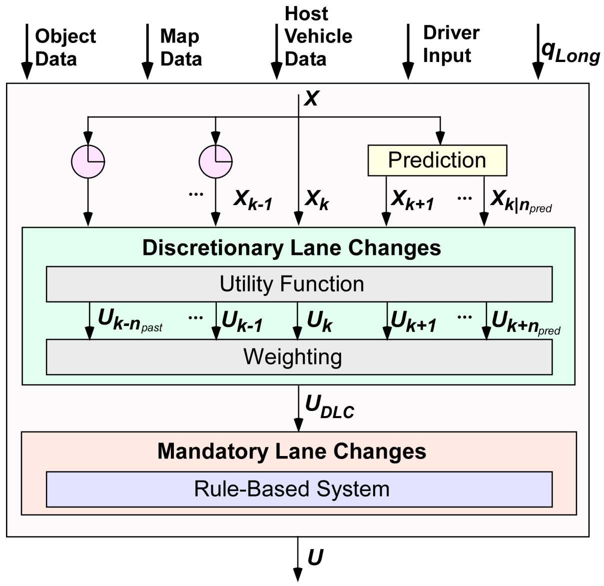

Fig. 8. BMW utility determination method.

2.4 ADOPTION OF UTILITY THEORY FOR SUPERVISORY CONTROL DETERMINISTIC DECISION-MAKING

The concept of utility is applied to decision-making for supervisory control applications with the intention of establishing a *uniform* scale for measuring the overall value of a choice. Utility becomes a true measure of value to the decision maker. Utility theory as described previously provides a consistent method to compare and measure values. Therefore, the choice with the highest utility value will be preferred. There will always be an interpretation or characterization of the utility functions made by designers. This interpretation will always have an aspect of subjectivity that is no different from human operators interpreting operating procedures at a power plant (which also have been devised with some level of subjectivity derived from engineering experience).

The decision-making sequence almost always comes down to the simplified progression depicted in Fig. 9. Each box can be expanded into a subsystem that together comprises a supervisory decision-making system. Previous work at Oak Ridge National Laboratory developed architectures and decision tools for building an automated decision-making supervisory control system [2-32]. In a related project, a decision-making engine based on PRA was devised [2-33]. The developments described in these reports are combined in the architecture shown in Fig. 10. This architectural interconnection becomes the functional blocks of a generalized decision-making framework employing probabilistic and deterministic methods to arrive at decisions.

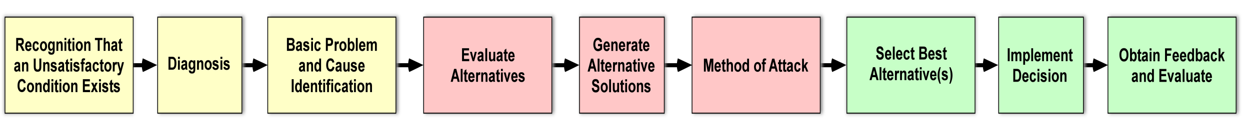

Fig. 9. Most basic sequence for arriving at a decision and implementing it.

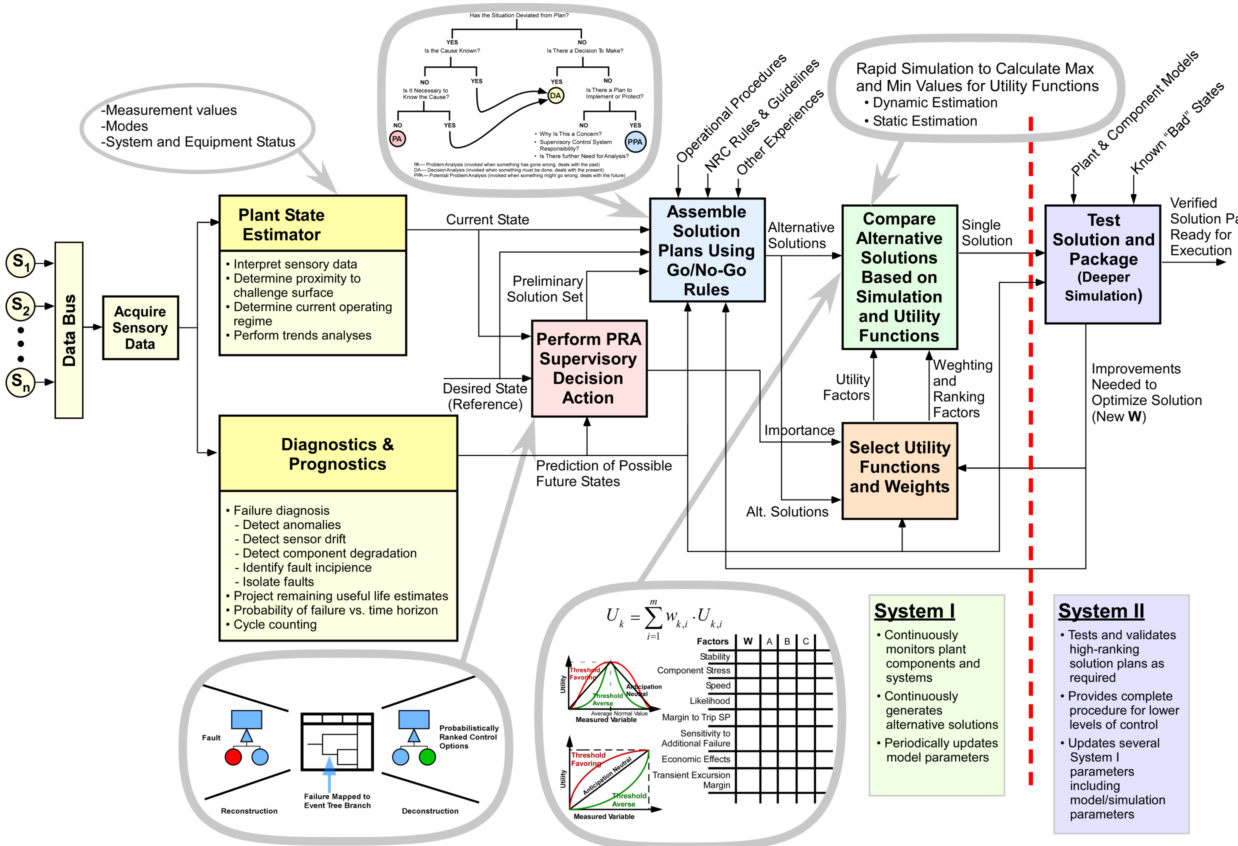

Fig. 10. Architectural interconnection of functional blocks of the generalized decision-making framework showing probabilistic and deterministic modules.

In Fig. 10, plant sensors are sampled (far left) and feed software analyzes those signals to interpret sensory data to determine the following:

- Proximity to RPS challenge surface
- Current operating regime/state
- Exceptional trends over specific time blocks

Additionally, sensor data (as measured and historical) are examined for diagnostic and prognostic purposes:

- Failure diagnosis
  - Detect anomalies
  - Detect sensor drift
  - Detect component degradation
  - Identify fault incipience
  - Isolate faults

- Project remaining useful life estimates
- Probability of failure vs. time horizon
- Cycle counting

A probabilistic decision-making engine (the block: *Perform PRA Supervisory Decision Action*) acts on failed component information as well as sensory and state information to identify and rank control restoration actions. A list of possible actions is ranked based on the potential for success based on real-time plant equipment and state information. These actions are turned into a set of alternative solution plans by the block *Assemble Solutions Plans Using Go/No-Go Rules*, which accesses plant operation procedures, NRC rules and guidelines, and other databases such as collective experience.

The block *Compare Alternative Solutions Based on Simulation and Utility Functions* is in effect the primary block for accomplishing the deterministic decision-making function. This block applies the utility functions and other weighting schemes to derive a bundled utility value for all the possible solutions from the *Assemble Solutions Plans* block. Basic and simplified simulations for crucial points in each of the alternative plans will be accomplished in the *Compare Alternative Solutions* block. This block receives weighting and utility function anticipation values from the block *Select Utility Functions and Weights*. The relative importance of the various utilities can be modified by information provided by diagnostics and prognostics, PRA, and the feedback results of the more rigorous simulation and analysis performed in the *Test Solution and Package* block.

The architecture is divided into two distinct systems characterized by a difference in their depth of simulation. System I (left of the red dashed line) continuously monitors plant variables, components, and systems. It may regularly generate possible solutions for minor deviations that may not ultimately be considered sufficiently important to pass on for actual implementation. There may be some continuous learning and calibration so that the ability to generate solutions and achieve near optimal comparative ranking is always robust. In contrast, System II, which receives the highest ranked solution from System I, tests and validates solution plans by a more detailed simulation exercise. A successful simulation proceeds on to become packaged as a complete procedure for implementation by lower levels of control in the hierarchy. As mentioned, this block updates several System I parameters including its simpler model/simulation parameters.

Several enhancements have been made to the classic utility theory that make the functionality of the *Compare Alternative Solutions Based on Simulation and Utility Functions* block distinctive and effective for the task of supervisory decision-making:

1. *Special utility functions*: Special utility functions are defined that express utility values for setpoint type process variables. Specifically, a Gaussian distribution is adapted as the mapping mechanism of the utility function. A bundle of utility functions may incorporate both Gaussian (optimizable) and monotonic functions to provide a uniform means of comparing preferences.

2. *Wide range utility functions*: Utility functions are defined in this work that range from negative infinity to positive unity in contrast with the zero to unity range of previously defined utility functions. With this arrangement, all attributes are uniform for positive contributions, but the bundle sum can be penalized for a particular variable in a seriously degraded region.

3. *Diversely integrated*: Integration of diverse attributes into a uniformly calculated (single) utility value permits robust decision-making across a selection of alternative recovery solutions. The diversity of attributes includes process variables that capture system dynamics, economics, maintenance, and reliability/risk parameters from PRA output.

4. *Integration of rapid simulation*: Both static and dynamic variable simulations are used to estimate minimum and maximum expected process variables from alternative solution plans. These min/max values are used to generate the utility values for conducting a weighted comparison.

2.5 CHAPTER 2 REFERENCES

2-1. J. Bentham, *An Introduction to the Principles of Morals and Legislation*, Oxford: Clarendon Press (1789).

2-2. D. Read, *Utility Theory from Jeremy Bentham to Daniel Kahneman*, The London School of Economics and Political Science, LSEOR 04-64, ISBN. 07530-1689-9 (2004).

2-3. D. Bernoulli, "Exposition of a New Theory on the Measurement of Risk," *Econometrica*, **22**(1), pp. 23–36 (Jan. 1954; an English translation of Bernoulli 1738).

2-4. L. J. Savage, *The Foundations of Statistics*, New York: Wiley; and New York: Dover Publications (1954 and 1972).

2-5. D. Charles, "An Essay upon the Probable Methods of Making a People Gainers in the Ballance of Trade," Printed for James Knapton, at the Crown in St. Paul's Church-yard, 1699, accessed 7/2015 (http://quod.lib.umich.edu/e/eebo/A69897.0001.001?view=toc).

2-6. J. Von Neumann and O. Morgenstern, "Theory of Games and Economic Behavior," Princeton University Press, 2007 reprint.

2-7. J. W. Pratt, "Risk aversion in the small and in the large," *Econometrica* **32**, pp. 122–36 (1964).

2-8. R. O. Schlaifer, *Analysis of Decisions under Uncertainty*, McGraw-Hill, New York (1969).

2-9. R. R. Novick and D. V. Lindley, "Fixed-state assessment of utility functions," *Journal of the American Statistical Association* **7**(4), pp. 306–11 (1979).

2-10. A. Tversky, and D. Kahneman, *Judgment under Uncertainty: Heuristics and Biases*, The Hebrew University, Jerusalem, Israel (1973).

2-11. P. C. Fishburn, *Decision and Value Theory*, Wiley, New York (1964).

2-12. P. C. Fishburn, "Independence in utility theory with whole product sets," *Operations Research* **13**, pp. 28–45 (1965).

2-13. P. C. Fishburn, *Bernoullian Utilities for Multiple Factor Situations: Multiple Criteria Decision Making*, J. L. Cochrane and M. Zeleny, eds., University of South Carolina Press, Columbia, South Carolina (1973).

2-14. R. L. Keeney and H. Raiffa, *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*, Wiley, New York (1973).

2-15. P. C. Fishburn, "Additivity in utility theory with denumerable product sets," *Econometrica* **34**, pp. 500–503 (1966).

2-16. J. L. W. V. Jensen, "Sur les fonctions convexes et les inégalités entre les valeurs moyennes," *Acta Mathematica* **30** (1), pp. 175–93 (1906).

2-17. C. R. White, "Performance Improvement in Aircraft Maintenance," *Production Planning and Control*, **5**(5), pp. 475–84 (Sept. 1994).

2-18. U. Sparmann, “Spurwechselvorgänge auf zweispurigen BAB- Richtungsfahrbahnen,” in *Forschung Straßenbau und Straßenverkehrstechnik*. Straßenbau, Germany : Bundesministerium für Verkehr, (1978).

2-19. P. G. Gipps, “A model for the structure of lane-changing decisions,” *Transp. Res. B, Methodol*. **20**(5), pp. 403–14 (Oct. 1986).

2-20. A. Niehaus and R. Stengel, “Probability-based decision making for automated highway driving,” *IEEE Trans. Veh. Technol*. **43**(3), pp. 626–34 (Aug. 1994).

2-21. K. I. Ahmed, “Modeling drivers’ acceleration and lane changing behavior,” Ph.D. dissertation, Mass. Inst. Technol., Cambridge, MA (1999).

2-22. T. Toledo, “Integrated driving behavior modeling,” Ph.D. dissertation, Mass. Inst. Technol., Cambridge, MA (2003).

2-23. D. Ehmanns, “Modellierung des taktischen Fahrerverhaltens bei Spurwechselvorgängen,” Ph.D. dissertation, RWTH Aachen, Forschungsgesellschaft Kraftfahrwesen, Aachen, Germany (2003).

2-24. A. Kesting, M. Treiber, and D. Helbing, “General lane-changing model MOBIL for car- following models,” *J. Transp. Res. Board*, vol. 1999, no. 1999, pp. 86–94 (2007).

2-25. R. Sukthankar, J. Hancock, S. Baluja, D. Pomerleau, and C. Thorpe, “Adaptive intelligent vehicle modules for tactical driving,” in *Proc. AAAI Workshop Intell. Adaptive Agents*, pp. 13–22 (1996).

2-26. A. Furda and L. Vlacic, “Enabling Safe Autonomous Driving in Real-World City Traffic using Multiple Criteria Decision Making,” *IEEE Intelligent Transportation Systems Magazine* **3**(1), pp. 4–17 (2011).

2-27. Robin Schubert, “Evaluating the Utility of Driving: Toward Automated Decision Making under Uncertainty,” *IEEE Trans. on Intelligent Transportation Systems* **13**(1) (Mar 2012).

2-28. A. Kesting, M. Treiber, and D. Helbing, “General Lane-Changing Model MOBIL for Car-Following Models,” Transportation Research Record: Journal of the Transportation Research Board, No. 1999, Transportation Research Board of the National Academies, Washington, D.C., pp. 86–94 (2007).

2-29. S. O’Hara, “Vehicle Path Optimization of Emergency Lane Change Maneuvers for Vehicle Simulation,” Master of Science Thesis, University of Maryland, College Park (2005).

2-30. M. Ardelt, C. Coester, and N. Kaempchen, “Highly Automated Driving on Freeways in Real Traffic Using a Probabilistic Framework,” *IEEE Trans. on Intelligent Transportation Systems,* Vol. 13, No. 4 (Dec. 2012).

2-31. Ljubo Vlacic, Daniel Thomas, Joshué Pérez Rastelli, “Real-Time Co-Operative Decision Making & Control Systems,” *XXIII International Symposium on Information, Communication and Automation Technologies*, Sarajevo, Bosnia and Herzegovina (Oct. 2011).

2-32. S. Cetiner, M. Muhlheim, G. Flanagan, D, Fugate, and R. Kisner, *Development of an Automated Decision-Making Tool for Supervisory Control System*, ORNL/TM-2014/363, SMR/ICHMI/ORNL/TR-2014/05, Oak Ridge National Laboratory, Oak Ridge, Tenn., September 2014.

2-33. S. Cetiner and M. Muhlheim, *Implementation of the Probabilistic Decision-Making Engine for Supervisory Control*, ORNL/LTR-2015/140, Oak Ridge National Laboratory, Oak Ridge, Tenn., March 2015.

3. APPLICATION OF UTILITY THEORY IN DETERMINISTIC DECISION-MAKING FOR SUPERVISORY CONTROL SYSTEM

The objective of deterministic decision-making is to capture the physical behavior of a system (i.e., the time evolution of physical variables for a known disturbance). A system is said to be *deterministic* if its future states does not involve random behavior. Hence, a *deterministic model* is a representation of a system behavior that will produce the same set of outputs for a given set of inputs and an initial state. Typically, the deterministic behavior of a system is represented by a set of differential, difference, or algebraic equations.

The deterministic decision-making framework is intended to provide the necessary interfaces for the probabilistic portion and to generate a resolution (i.e., a single solution) of the decision-making process.

3.1 INTERFACES TO PROBABILISTIC DECISION-MAKING

As documented in the previous milestone report (Ref. 3-1), the outcome of the probabilistic module is to generate a set of decision alternatives, each of which may have a varying number of control actions, as illustrated in Fig. 11. The probabilistic module provides a probabilistic ranking of these alternatives with a metric called *likelihood of success*, which is retrieved from the *Frequency* column of a PRA calculation, as shown in Fig. 12.

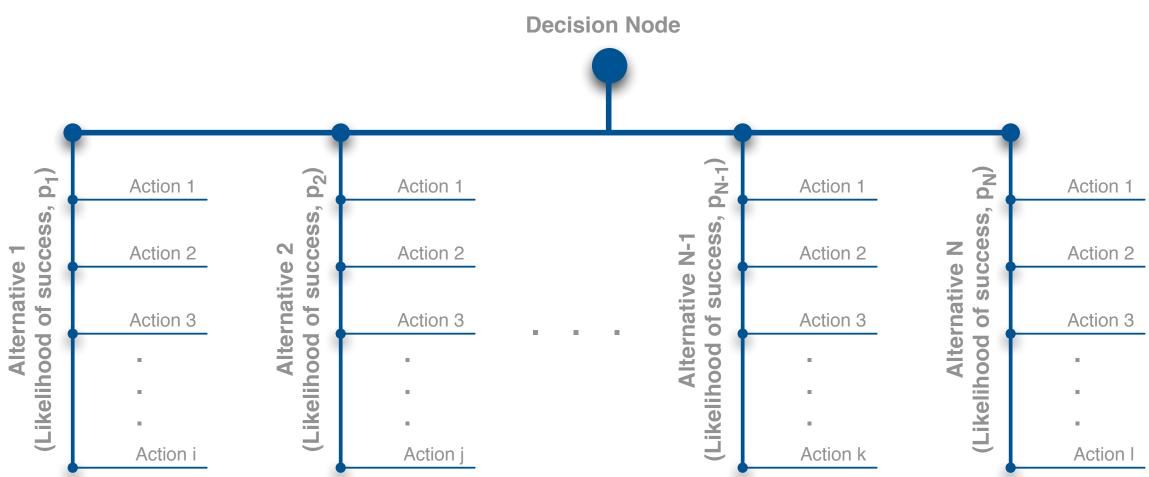

Fig. 11. A decision tree generated as a result of a probabilistic risk assessment.

The purpose of the decision-making method is to assess each alternative, considering its likelihood of success as a whole as well as by taking into account the implications of individual control actions. The outcome of this assessment must be a *singleton* measure that quantifies the favorability of an alternative.

Based on an extensive literature survey, control theory (classical and contemporary) does not offer analytical tools that can be employed to tackle the decision-making problem for the SCS. While *Bayesian networks* and *decision graphs* offer a wide spectrum of mathematical tools for decision-making [3-2], the scale of a complex system—such as a nuclear power plant—and the interdependence of its dynamic processes, makes these tools inefficient. However, these tools may prove useful in the future for certain subtasks of complex decision-making processes.

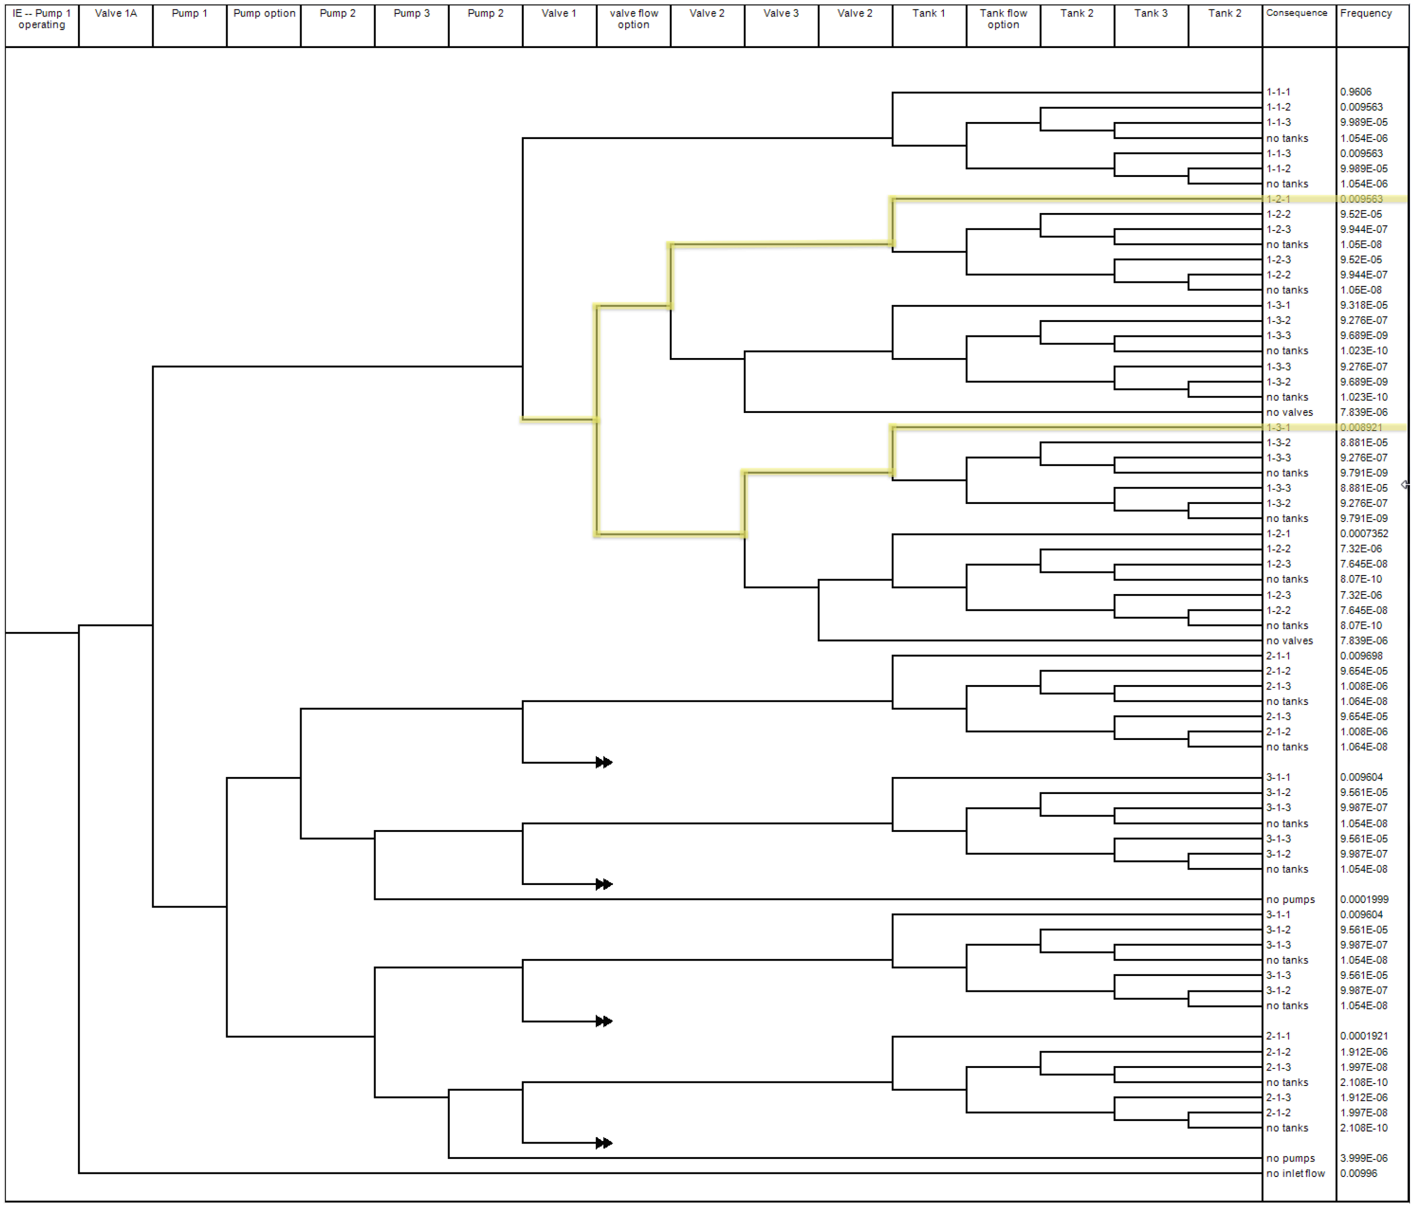

Fig. 12. Trajectories for each decision alternative and associated likelihood of success.

Conversely, utility theory provides a means by which a series of decision alternatives—called *utility attributes*—can be collectively evaluated, as briefly introduced in Chapter 2 for large-scale complex systems. This chapter describes the data interfaces that are needed to be able to process the output generated by the PRA calculation for a decision node.

3.2 SELECTION OF UTILITY VARIABLES

The objective of employing the utility theory is to create a framework by which the physical behavior of the system can be assessed along with the probabilistically ranked decision alternatives. The plant dynamics are captured by a set of state variables that are determined to be key actors in control.

3.2.1 Utility Attributes for Plant State

The objective of the deterministic decision-making module is to incorporate the physical behavior (current and projected) of the system. In order to achieve that capability, the utility variables must be selected such that the projected physical behavior of the system can be factored into the decision-making with the probabilistically ranked options from the PRA calculation. This is best accomplished by linking the desired utility attributes to key process variables (i.e., the ones that provide insight about the status of the system). A partial list of system design variables for Advanced Liquid-Metal Reactor (ALMR) Power Reactor Innovative Small Module (PRISM) and their nominal steady-state values are shown in Table 2.

Table 2. ALMR PRISM heat transport system design values

Variable	Description	Nominal Value	Unit
$\dot{Q}_{R_X}$	Reactor thermal power	425	MWt
$T_{R_{Xo}}$	Reactor outlet temperature	468.3	ºC
$T_{R_{Xi}}$	Reactor inlet temperature	321.1	ºC
ΔT_{R_X}	Reactor temperature difference	147.2	ºC
ω_p	Primary coolant mass flow rate (total)	2016	kg/s
$\omega_{p,\,disc}$	Primary pump discharge volumetric flow rate*	0.66	m^3/s
h_p	Primary pump head	96.3	m
T_{hl}	Intermediate hot leg temperature	426.67	ºC
ω_i	Intermediate coolant mass flow rate (total)	2268	kg/s
$\omega_{i,\,disc}$	Intermediate pump discharge volumetric flow rate	2.6	m^3/s
h_i	Intermediate pump head	95.7	m
$\dot{Q}_{SG}$	Steam generator thermal power**	432	MWt
$T_{SG,o}$	Steam generator outlet temperature	285	ºC
$p_{SG,\,o}$	Steam generator outlet pressure	6.895	MPa
$T_{SG,\,fw}$	Steam generator feedwater temperature	216	ºC
ω_{SG}	Steam flow rate	233.5	kg/s

* Volumetric flow rate per pump; total of four pumps.
** Including pump heating from primary loop, intermediate loop, and steam generator pumps (~ 6.82 MWt).

The selection criteria for utility variables must address the safety envelope of the controls domain. As illustrated in Fig. 2, the fundamental objective of the SCS is to maintain the plant state within the controllable domain, which is delineated by the red line in Fig. 2—also referred to as the *challenge surface*. In its simplest form, the challenge surface is formed by the *trip variables* that, if exceeded, initiate an RPS and/or ESFAS actuation. Reactor safety functions and associated trip variables for ALMR PRISM are listed in Table 3.

Table 3. Reactor trip variables and associated safety functions for ALMR PRISM

	Safety Function	Monitored Variable	Type
Flux	Monitor for insertion of reactivity (threshold function of operating power level)	Reactor core neutron flux	TRIP
Flow	Monitor for loss of flow*	Primary loop sodium level	TRIP
		Primary loop EM pump discharge inlet pressure	
Temperature	Monitor for loss of heat sink	Reactor core outlet temperature Cold pool temperature	TRIP
Level	Monitor for loss of sodium	Primary loop sodium level	TRIP
Pressure	Monitor for electromagnetic (EM) pump outlet duct failure	Primary loop EM pump discharge inlet pressure	TRIP

* The loss-of-flow measurement is indirect, using the EM pump discharge pressure as an indicator of the primary loop flow rate.

During normal operation, the SCS tries to confine the plant state within an even tighter domain, which is delineated by the blue line in Fig. 2—also called the homeostatic region. Similarly, to incorporate a broader snapshot of the plant state, additional utility attributes must be linked with key process variables.

ALMR PRISM RPS actuates on the following trip variables [3-3]:

1. measured reactor core neutron flux (φ),
2. reactor core outlet temperature ($T_{R_{Xo}}$),
3. cold pool temperature ($T_{pool,cold}$),
4. pump discharge inlet pressure (p_{disc}), and
5. primary heat transport system (PHTS) sodium level (y_{PHTS}).

In addition to the RPS trip variables identified in the ALMR *PRISM Preliminary Safety Information Document* [3-3], the following variables were identified as important decision variables:

1. reactor core coolant temperature difference (ΔT_{R_X}),
2. intermediate heat transport system (IHTS) sodium level (y_{IHTS}),
3. steam generator (SG) drum level (y_{SG}), and
4. steam generator feedwater (FW) inlet flow rate (ω_{fw}).

To maintain consistency among the attributes, the utility variables are derived from the process variables through a simple linear transformation:

$$x_i = \frac{p_i - (p_i)_{min}}{(p_i)_{max} - (p_i)_{min}} \tag{3-1}$$

where x_i is the utility variable for the *ith* attribute, and p_i is the process variable linked to x_i; subscripts min and max are the minimum and maximum values each process variable is allowed to take (red line in Fig. 2). For safety-related variables (i.e., trip variables) these values are obtained by the setpoints of their processes from plant technical specifications. Fig. 13 shows an example plot of a how a utility variable and a process variable are linked through the linear transformation.

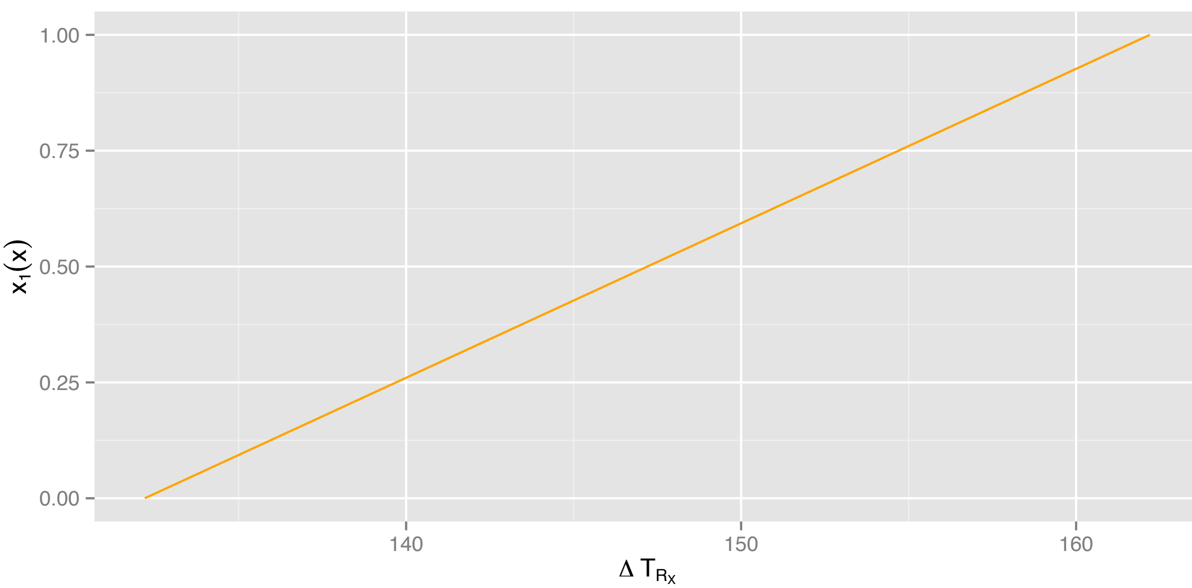

Fig. 13. The linear transformation maps the core differential temperature variable onto its utility variable.

A preliminary list of utility variables selected for the supervisory control system for the ALMR PRISM based on the nine variables identified above is shown in Table 4.

Table 4. Process utility variables for ALMR PRISM supervisory control system

Utility Variable	Safety Variable	Minimum	Lower Bound	Nominal	Upper Bound	Maximum	Linear Transformation
x_1	ΔT_{R_X} (ºC)	132.2	142.2	147.2	152.2	162.2	$x_1 = \frac{\Delta T_{R_X} - (\Delta T_{R_X})_{min}}{(\Delta T_{R_X})_{max} - (\Delta T_{R_X})_{min}}$
x_2	$T_{R_{Xo}}$ (ºC)	453.3	463.3	468.3	473.3	483.3	$x_2 = \frac{T_{R_{Xo}} - (T_{R_{Xo}})_{min}}{(T_{R_{Xo}})_{max} - (T_{R_{Xo}})_{min}}$
x_3	$T_{R_{Xi}}$ (ºC)	306.1	316.1	321.1	326.1	336.1	$x_3 = \frac{T_{R_{Xi}} - (T_{R_{Xi}})_{min}}{(T_{R_{Xi}})_{max} - (T_{R_{Xi}})_{min}}$
x_4	p_{dis} (kPa)	807	817	827	837	847	$x_4 = \frac{p_{dis} - (p_{dis})_{min}}{(p_{dis})_{max} - (p_{dis})_{min}}$
x_5	y_P (m)	9	10	12	14	15	$x_5 = \frac{y_P - (y_P)_{min}}{(y_P)_{max} - (y_P)_{min}}$
x_6	y_I (m)	2	3	5	7	8	$x_6 = \frac{y_I - (y_I)_{min}}{(y_I)_{max} - (y_I)_{min}}$
x_7	y_{SG} (m)	2	3	5	7	8	$x_7 = \frac{y_{SG} - (y_{SG})_{min}}{(y_{SG})_{max} - (y_{SG})_{min}}$
x_8	φ (n/cm^{2}s)	TBD	TBD	TBD	TBD	TBD	TBD
x_9	ω_{fw} (kg/s)	242	262	272	282	302	$x_9 = \frac{\omega_{fw} - (\omega_{fw})_{min}}{(\omega_{fw})_{max} - (\omega_{fw})_{min}}$

3.2.2 Utility Attributes for Decision Alternatives and Interface to Probabilistic Portion of Decision-Making

This utility attribute is the interface to the probabilistic portion of the decision-making; it integrates the probabilistically informed assessments with deterministic calculations.

As presented in Fig. 11 and Fig. 12, the probabilistic calculation module outputs a number of decision alternatives, with each alternative having a quantitative measure, $p_s \in [0, 1]$, indicating its likelihood of success. Since the attribute is defined within the same input domain, the utility variable is defined as:

$$x_{10} = p_s \tag{3-2}$$

3.3 SELECTION OF UTILITY FUNCTIONS

The functional characteristics of utility functions are a subject of research. As briefly discussed in Chapter 2 and illustrated in Fig. 6, the shape of a utility function has implications for its effect on the overall decision-making. While utility functions presented in the literature—mostly in the field of economics—are expected to satisfy certain criteria such as *monotonicity*, *non-decreasing,* or *strictly increasing* properties, these rules result from the field of the application.

Engineering applications of utility theory expand the classical definition of utility functions to address specifics needs and requirements. For instance, Ref. 2-30 employs Gaussian distributions for representing the relationship between a utility attribute and its functional form.

3.3.1.1 Utility Functions for Plant State

The proposed selection scheme of utility functions greatly expands its definition:

1. utility variables, x_i, are defined in $\mathbb{R} \in [0, 1]$, which maps an engineering variable operating range between its minimum and maximum value,
2. utility functions, $u_i(x_i)$, have a mean value of $\mu = 0.50$ (symmetry rule),
3. utility functions intersect the abscissa at a lower-bound and an upper-bound value of an engineering variable, and
4. utility functions are positive within the domain delineated by the lower- and upper-bound, and negative elsewhere.

This scheme allows for rewarding a particular utility for being contained within the operations domain while penalizing it for being outside. Depending on the other parameters used, the penalty for not being contained within the domain can be significant, as will be illustrated shortly.

The probability density of the Gaussian distribution is represented as

$$p(x|\mu, \sigma) = e^{-\frac{(x-\mu)^2}{2\sigma^2}} \tag{3-3}$$

where μ is the mean and σ is the standard deviation of the distribution. Some examples from the family of distributions are shown in Fig. 14.

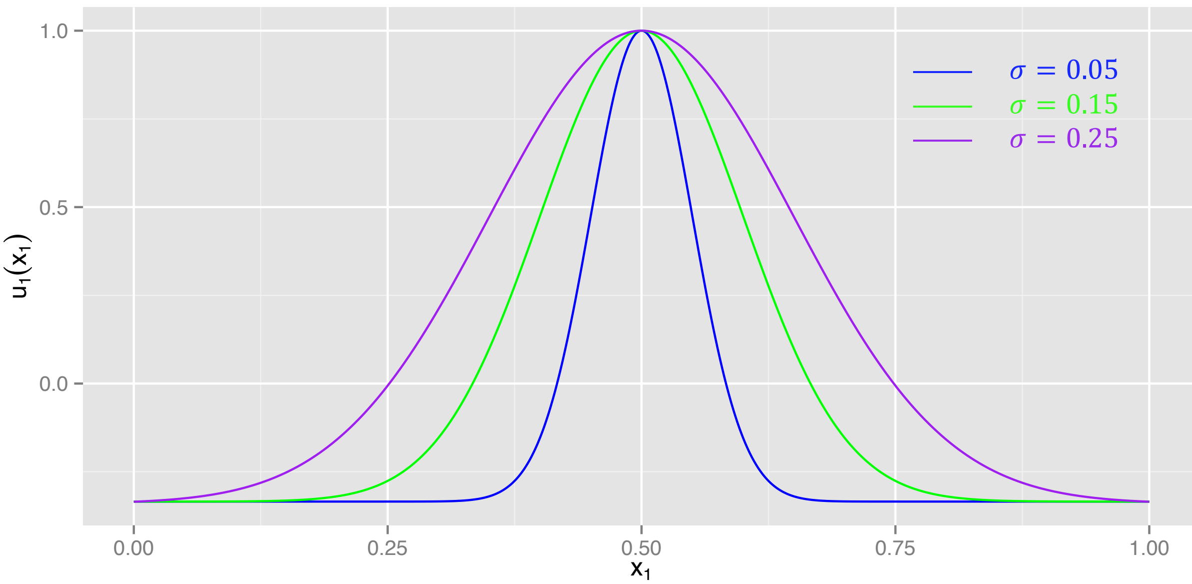

Fig. 14. Family of Gaussian probability density distributions for various standard deviations.

The utility functions are selected from the family of Gaussian distributions through a linear transformation (called *affine transformation*) represented by Eq. (3-4),

$$u(x|\mu, \sigma) = a\, e^{-\frac{(x-\mu)^2}{2\sigma^2}} + b \tag{3-4}$$

where a and b are the coefficients of the transformation. This transformation essentially determines the point where the curve intersects the abscissa. Fig. 15 shows the Gaussian family of curves after the linear transformation.

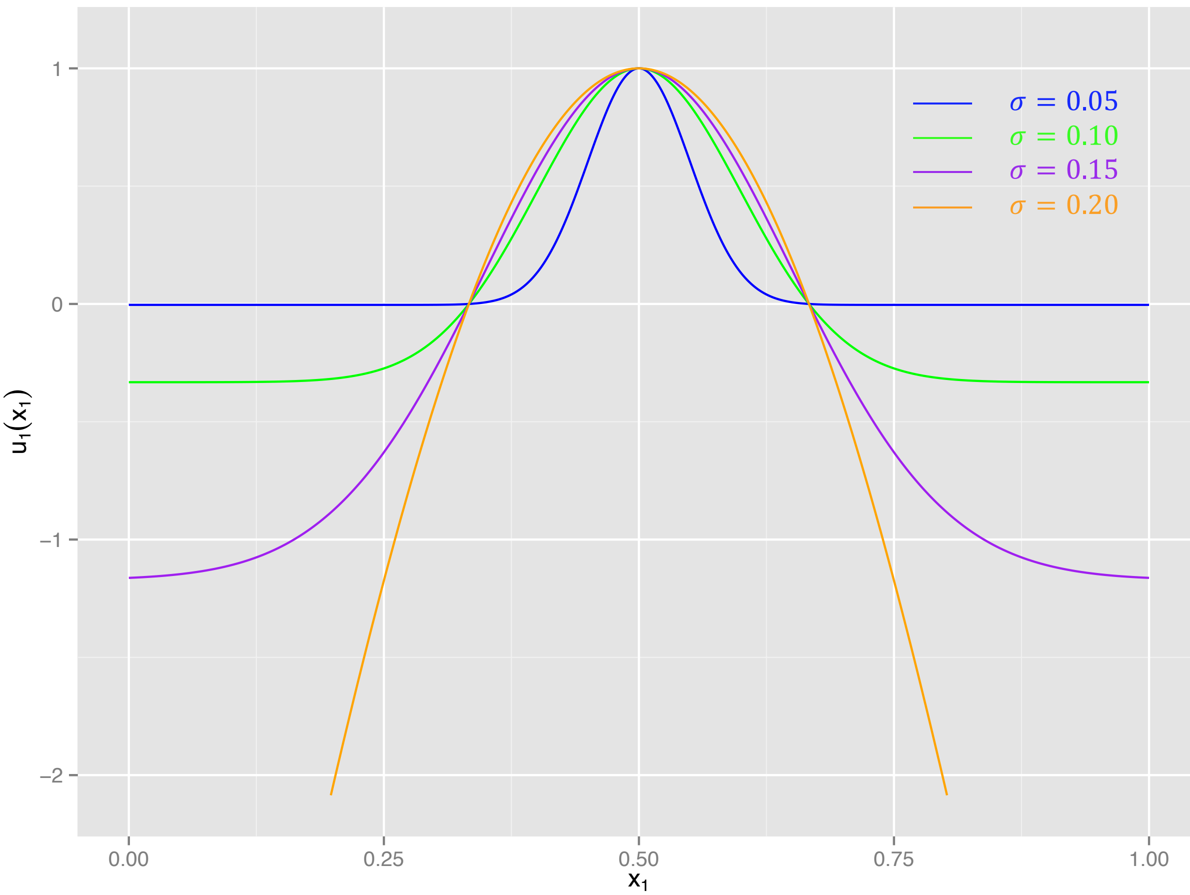

Fig. 15. Utility functions with varying standard deviations after linear transformation.

The intersection points are determined based on the lower- and upper-bound values of a safety variable. For instance, the lower- and upper-bound values for x_1 are determined as follows:

$$(x_1)_{LB} = \frac{\left(\Delta T_{R_X}\right)_{LB} - \left(\Delta T_{R_X}\right)_{min}}{\left(\Delta T_{R_X}\right)_{max} - \left(\Delta T_{R_X}\right)_{min}} \tag{3-5a}$$

$$(x_1)_{UB} = \frac{\left(\Delta T_{R_X}\right)_{UB} - \left(\Delta T_{R_X}\right)_{min}}{\left(\Delta T_{R_X}\right)_{max} - \left(\Delta T_{R_X}\right)_{min}} \tag{3-5b}$$

Based on the values given in Table 4, these are calculated as

$$(x_1)_{LB} = \frac{1}{3}$$

$$(x_1)_{UB} = \frac{2}{3}$$

$(x_i)_{LB}$ and $(x_i)_{UB}$ values are symmetrical about $x_i = 0.5$, as illustrated with $(x_1)_{LB}$ and $(x_1)_{UB}$.

The a and b values of the linear transformations as shown in Eq. (3-4) for a given utility function $u_i(x_i|\mu, \sigma)$ are calculated by solving the following set of equations:

$$\frac{b}{a} = \exp\left(-\frac{[(x_i)_{LB} - \mu]^2}{2\sigma^2}\right) \tag{3-6a}$$

$$a - b = 1 \tag{3-6b}$$

As an example, solving Eq.'s (3-6a) and (3-6b) for $\mu = 0.5$ and $\sigma = 0.15$ yields

$$a = 2.171117$$

$$b = 1.171117$$

The utility functions for the ALMR PRISM process variables based on the transformation and lower and upper limits are shown in Fig. 16.

3.4 COMPOUND UTILITY AND INTERFACE TO PROBABILISTIC PORTION OF DECISION-MAKING

Contributions from individual utilities $i \in [1, N]$ are combined into a compound utility metric, U_i, for the decision branch i on Fig. 11 by the following expression:

$$U_i = p_i \sum_{j=1}^{N} \omega_j \, u_j(x_j) \tag{3-7}$$

where p_i is the likelihood of success associated with the ith branch of the decision tree (Fig. 11), N is the total number of utility variables (e.g., nine utility attributes were identified for ALMR PRISM as shown in Table 4), ω_j is the weight of each utility function, and $u_j(x_j)$ is the utility function for the attribute x_j. The decision branch with the highest compound utility, U, is selected.

The compound utility is calculated based on the maximum or minimum values that process variables take, which is determined based on a detailed, end-to-end dynamic simulation for the plant. The compound utility may also be calculated as a time-varying variable. In that case, the lowest compound utility value must be used as the decision variable for each decision branch.

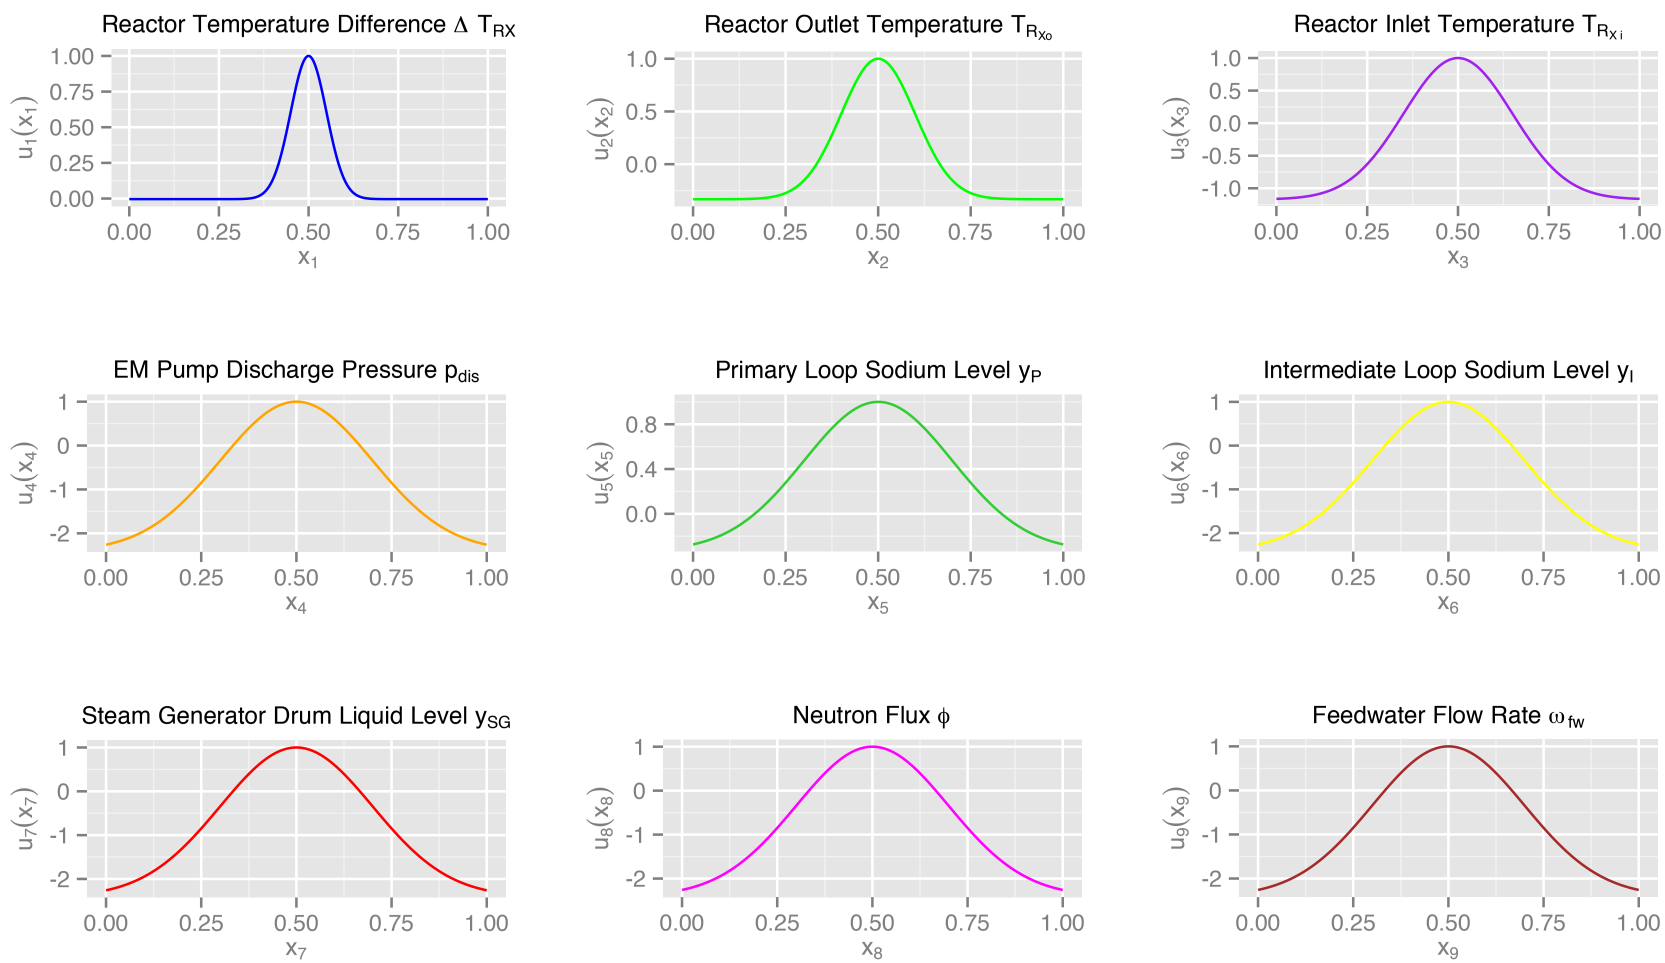

Fig. 16. Utility functions for ALMR PRISM process variables.

3.5 CHAPTER 3 REFERENCES

3-1. S. M. Cetiner and M. D. *Muhlheim, Implementation of the Probabilistic Decision-Making Engine for Supervisory Control*, ORNL/SPR-2015/140, Oak Ridge National Laboratory, Oak Ridge, Tenn., (March 2015).

3-2. F. V. Jensen and T. D. Nielsen, *Bayesian Networks and Decision Graphs*, Springer Information Science and Statistics Series, Second Ed. (February 2007).

3-3. *PRISM Preliminary Safety Information Document*, GEFR-00793, UC-87Ta, Prepared for US Department of Energy under Contract No. DE-AC03-85NE37937 (December 1987).

4. SUMMARY, CONCLUSIONS AND FUTURE WORK

This report documents the technical basis and a computational framework to accomplish the *deterministic* decision-making function for the SCS. The framework adopts *utility theory* as the mathematical method to perform the deterministic portion of the overall decision-making function. Utility theory offers a unifying measure that takes into account the value and potential consequences of individual control actions, which is reflected on the combined utility of a decision alternative. In the context of supervisory control and automated decision-making, *utility* is simply defined as a measure of preferences over a given set of actions.

The proposed *deterministic decision-making framework* accomplishes the following:

1. provides interfaces to the probabilistic decision-making function that allows incorporation of decision branches, which represent decision alternatives, as well as their likelihood of success;

2. provides a formal way to incorporate the physical behavior of the system with probabilistic assessments;

3. provides a theoretical foundation to combine the effects of individual physical dynamics, called *utility attributes* in utility theory, to ultimately achieve a metric that gives a singleton measure of its effectiveness—called its *utility*; and

4. provides a unifying method to incorporate other key considerations, such as plant life-cycle economics and predictive maintenance scheduling, into the final decision metric.

4.1 CONCLUSIONS

This report introduces several new developments in the application of utility theory to real-time decision-making:

1. *Special Utility Functions:* Special utility functions are defined that express utility values for process variables. Specifically, a Gaussian distribution is adopted as the mapping between the utility attributes and their utilities for process variables.

2. *Wide-Range Utility Functions:* Utility functions are defined in this work that range from negative infinity to positive unity for process variables in contrast with the zero-to-unity range of classical utility functions. This arrangement allows for rapid penalty of control actions that will take the plant to degraded operational domain.

3. *Diversely Integrated:* Integration of diverse attributes into a uniformly calculated (single) utility value permits robust decision-making across a selection of alternative recovery solutions. The diversity of attributes includes process variables that capture system dynamics, plant life-cycle economics, maintenance scheduling, and likelihood of success value from the probabilistic decision-making module.

4. *Integration of Rapid Simulation:* Faster-than-real time simulations are used to determine the effects of individual control actions for each decision branch. The violations of trip setpoints are noted, and those branches are automatically eliminated from the decision process. Utility calculations are performed based on the projected worst-case values of the tracked process variables following a disturbance until a steady state is reached.

The contributions in this report provide a formal way to combine probabilistic decision functionality with deterministic evaluations to achieve a risk-informed performance-based comparative decision-making process. This is a novel contribution that currently is not reported in the literature for an engineering application.

4.2 FUTURE WORK

The present work documents the progress made on incorporating a deterministic decision-making approach into the overall decision-making function. The ongoing work captures this approach in the simulation environment that was presented in previous milestone reports.

The simulation environment under development for demonstrating the supervisory control decision-making functionality makes use of a PRA simulation tool and a system-level dynamic simulation tool. These two diverse simulation approaches are coupled through a complex data-exchange scheme. Supervisory control capability leverages this capability and adds functionalities to perform the decision-making calculations.

4.2.1 Identification of Utility Function Weights

There is a large volume of publication in the scientific literature that addresses the problem of identifying the weights of individual utility functions. However, a great majority of these publications addresses the issue of maximizing the *marginal utility* or the *expected utility* of an economic or financial problem. It was observed that these solutions may not be directly applicable to the supervisory control decision-making problem. Currently, the weights are identical.

One potential solution being considered is to bring in an existing operating procedure for a given set of transients—similar to training an expert system. The designer would then set up the utility quadrature, run the simulations, and calculate utilities by performing a parametric scan over all attribute weights. Only those weights that lead to desired (or required) decisions would be selected; others that did not satisfy the condition would be eliminated from the parametric search.

This research activity is ongoing.

4.2.2 Incorporation of Power Runback Options

Power runback is a key functionality for SCS. While the research is still ongoing, it is considered that the most effective way to incorporate that functionality will be through multiple utility calculations performed at different reactor power levels such as 100%, 90%, 80%, etc. It is expected that the compound utility for the process variables will be more favorable as the power level is reduced—provided that the lower-level loop control functions are designed to maintain reactor operations at lower power levels than the nominal steady state power. The supervisory control would most likely accomplish this by changing the power-level setpoint of the control system that regulates the control-rod drive mechanisms, as would an operator. The highest power level that yields an acceptable level of utility would be selected.

This is an ongoing research activity.

4.2.3 Incorporation of System Diagnostic and Prognostic Information

An ongoing effort led by Pacific Northwest National Laboratory (PNNL) under the ART program within the ICHMI technical area is working on developing a methodology for enhanced risk monitors that integrate real-time information about equipment condition and probability of failure into risk monitors. The idea is to provide an instantaneous assessment of risk at a component level as plant equipment is used and ages. This methodology incorporates the sources of uncertainty into the enhanced risk monitors framework and addresses the question of how uncertainties propagate through the calculations. An uncertainty bound is estimated for predicted risk metrics [4-1].

While the content of the data that will be provided by the enhanced risk monitors module is not clear, based on the earlier interactions with the PNNL team, it is understood that the output of the prognostic calculation will include a multitude of projections, which would indicate likelihood of failures at different time horizons and associated confidence levels.

This area will be further investigated internally as well as with the research team at PNNL.

4.3 CHAPTER 4 REFERENCES

4-1. P. Ramuhalli, E. H. Hirt, G. A. Coles, C. A. Bonebrake, B. J. Ivans, Jr., D. W. Wootan, and M. R. Mitchell, *An Updated Methodology for Enhancing Risk Monitors with Integrated Equipment Condition Assessment*, PNNL-23478 Rev. 0 (SMR/ICHMI/PNNL/TR-2014/01), Pacific Northwest National Laboratory (July 2014).