

International Journal on Artificial Intelligence Tools
© World Scientific Publishing Company

Intelligent Modeling for Nuclear Power Plant Accident Management

Michael C. Darling

Data-driven & Neural Computing
Sandia National Laboratories
and
Electrical & Computer Engineering
University of New Mexico
Albuquerque, New Mexico, USA
michael.darling@sandia.gov

George F. Luger

Department of Computer Science
University of New Mexico
Albuquerque, New Mexico, USA
luger@cs.unm.edu

Thomas B. Jones

Department of Computer Science
University of New Mexico
Albuquerque, New Mexico, USA
thomasbJones2@gmail.com

Matthew R. Denman

Risk & Reliability Analysis
Sandia National Laboratories
Albuquerque, New Mexico, USA
mrdenma@sandia.gov

Katrina M. Groth

Center for Risk & Reliability
Department of Mechanical Engineering
University of Maryland
College Park, Maryland, USA
kgroth@umd.edu

Received (16 May 2017)
Revised (28 August 2017)
Accepted (25 October 2017)

2 *M.C. Darling, G.F. Luger, T.B. Jones, M.R. Denman, K.M. Groth*

This paper explores the viability of using counterfactual reasoning for impact analyses when understanding and responding to “beyond-design-basis” nuclear power plant accidents. Currently, when a severe nuclear power plant accident occurs, plant operators rely on Severe Accident Management Guidelines. However, the current guidelines are limited in scope and depth: for certain types of accidents, plant operators would have to work to mitigate the damage with limited experience and guidance for the particular situation. We aim to fill the need for comprehensive accident support by using a dynamic Bayesian network to aid in the diagnosis of a nuclear reactor’s state and to analyze the impact of possible response measures.

The dynamic Bayesian network, DBN, offers an expressive representation of the components and relationships that make up a complex causal system. For this reason, and for its tractable reasoning, the DBN supports a functional model for the intricate operations of nuclear power plants. In this domain, it is also pertinent that a Bayesian network can be composed of both probabilistic and knowledge-based components. Though probabilities can be calculated from simulated models, the structure of the network, as well as the value of some parameters, must be assigned by human experts. Since dynamic Bayesian network-based systems are capable of running better-than-real-time situation analyses, they can support both current event and alternate scenario impact analyses.

Keywords: Counterfactual Reasoning; Decision Support; Dynamic Bayesian Networks

1. Introduction

The goal of this research project is to make a real-time analysis and prognostic system for the production of electric power through the use of a nuclear reactor. There are three critical aspects to the project: 1) to provide modern tools to augment the paper-based diagnostic methods currently used by reactor operators, 2) to offer real-time diagnosis and analysis, and 3) the ability to generate and validate remediation strategies in the case of problems arising. All three of these goals require the use of Artificial Intelligence technology. This paper chronicles our work in this domain.

1.1. *Modernizing Response Tools*

Though nuclear power plant accidents are extremely rare, the effects can be harmful for people, the environment, and the economy¹. The steps taken in the immediate aftermath of an accident are critical for limiting the extent of the damage to the plant and its surrounding area. Nuclear power plant operators currently follow paper manuals with step-by-step *Emergency Operating Procedures* for each type of anticipated accident.

For unanticipated “beyond-design-basis” accidents there are Severe Accident Management Guidelines (SAMGs) which are developed from expert judgment and best-estimate analyses. If the plant’s monitoring instruments fail, which was the case during 2011 Fukushima accident¹, operators must act with severely limited information about the current state of the reactor.

In order to understand the possible range of responses of nuclear power plants during severe accidents, researchers simulate accidents with various conditions using nuclear reactor accident simulation software, such as MELCOR² or SAS4A³.

To condense the large amount of information these simulators output into a fast and human-understandable knowledge base, we have used the simulated data to build dynamic Bayesian networks to simulate each process scenario. We call this framework “SMART procedures” (where “SMART” was originally an acronym for “Safely Managing Accidental Reactor Transients”).

SMART procedures is an ongoing project aimed at providing diagnostic tools in the event of a beyond-design-basis nuclear power plant accident. Currently, the SMART procedures framework provides inference on two types of accidents given partial information about key plant parameters.

1.2. Real-time analysis and diagnosis

Building a nuclear reactor for producing electric power is a complex engineering process. An explanation driven analysis tool is critical for understanding the interactions of its components as well as generating diagnostic recommendations. Dynamic Bayesian Belief Network (DBN) technology is an important generalization of the hidden Markov model, and offers a factored form of the full Bayesian representation^{4,5}. Besides offering a more expressive representation for the components of a complex causal system, it also supports tractable inference⁶.

The transparency of the DBN technology is an important aspect of its use in the nuclear domain. Our DBN is composed of both probabilistic and knowledge-based processes. The conditioned aspect of our system comes from both testing individual components, such as sensors, as well as from evaluating the ascribed relationships between components. The probabilistic aspect of these known relationships comes from thorough testing where expected outcomes can be quantified.

The knowledge-based side of the DBN design comes from the initial construction methodology where domain experts describe causal processes and assign specific probabilities to these component relationships as they are understood. With a flexible time discretization process available for the DBN model, the power generation system may be monitored intensely, by every second, for example, when situations are critical, as well as more slowly, by minutes or hours, when the health state of the reactor supports this.

The DBN system is able to run situation analyses in better-than-real-time, supporting both current event and alternate scenario impact analyses. Judea Pearl⁵ offered suggestions on how this prognostic process might be automated when he proposed causal “counterfactual algorithms”. For Pearl, “counterfactual” means to test situations that are not “currently true”. We use this logic to explore the impact of alternative scenarios, where, given a particular critical state, possible solution alternatives can be explored. These prognostic scenarios can be tested using a Bayesian model in better-than real time.

1.3. Generation and validation of remediation strategies

A critical component of nuclear power generation monitoring is the need to know how to respond to certain critical states that might occur. The DBN is able to automate the complex processes for analysis of critical accident scenarios. This automatization process is very important as, once it is verified as optimum, it will allow the human reactor monitors to see in real time what the alternatives to particular situations are and what steps are recommended (rather than having to look these situations up in complex paper manuals as is the current protocol). It also allows operators, given the previous states of the system, to infer the likely states of reactor parameters in the event of instrument or control-room failure.

There has been use of AI technology⁷ to analyze the components of the nuclear power generation process. However, there has been no work to date, other than our own, analyzing the use of the dynamic Bayesian network technology to monitor the full complex processes involved in nuclear power generation.

Our previous work demonstrated proof-of-concept using small models increasing in size and complexity, starting with a model for light water reactors with three observable variables⁸, and several conceptual models for sodium reactors with up to fourteen observable variables^{9,10}. In this paper, we investigate the feasibility of impact analysis to provide the ability to base accident mitigation decisions on data-driven probabilistic investigations of possible outcomes. We present a proof-of-principle tool and demonstrate its use of counterfactual reasoning for prognostic accident remediation.

Section 2 discusses the motivation of SMART procedures. Section 3 provides a detailed analysis of the engineering behind the sodium cooled reactor and the expressiveness and transparency afforded by modeling this technology with a DBN. Further, we discuss the development of our prototype model: the generation of data through simulating accident sequences, the construction of the DBN, and then demonstrate its interface.

Section 4 explains the validation of the Bayesian model through analysis of the Kullback-Leibler divergence of its parameters, and through its accuracy and F-scores. In section 5, we explore the viability of analyzing the impact of critical prognostic decisions in several accident response scenarios. We conclude with a summary of the project and suggest further research.

2. The Motivation of SMART Procedures

Existing nuclear power plant accident management guidelines rely on expert judgment and best-estimate analyses in order to capture the physical responses of the plant. However, procedure developers cannot anticipate every possible accident scenario. This limitation can be addressed with the use of *Dynamic Probabilistic Risk*




Fig. 1. Risk-informed “Smart SAMG” development process for nuclear power plant diagnostic support. Accident scenarios are generated using event trees, the scenarios are simulated with a system analysis program such as MELCOR¹¹ or SAS4A, the data is used to generate a Bayesian network.

6 M.C. Darling, G.F. Luger, T.B. Jones, M.R. Denman, K.M. Groth

Assessment simulations^a coupled with *Discrete Dynamic Event Trees*^b to provide comprehensive coverage of the potential accident scenario space.

Dynamic Probabilistic Risk Assessment simulations can explore thousands of scenarios that form the basis for comprehensively learning the values of observable reactor parameters during accidents. These simulators also model the possible actions of plant operators using the plant's procedure manuals. The result of this approach is a comprehensive training data set. However, the amount of data generated provides too much information to process in real-time during the occurrence of a severe accident. Thus, these simulators provide detailed insight into specific accidents after they happen, yet cannot be used to support real-time accident diagnosis. For example, in the aftermath of the Fukushima power plant accident, researchers ran simulations in order to assess their modeling capabilities in comparison to the actual events¹³. The results of these simulations provided useful information. The aim of this paper is to take the insights provided by such simulations and construct a knowledge base capable of supporting decisions during the actual progression of an accident.

The SMART procedures framework was developed in order to condense the nuclear accident simulation data into a real-time analysis tool⁸. This is accomplished with a dynamic Bayesian network. After the simulation of thousands of permutations of possible scenarios, we populate the Bayesian network with conditional probabilities calculated from the data. Therefore, SMART Procedures, as described in Figure 1, can be used to support diagnosis of the likely state of a nuclear reactor given the values of the observed plant parameters. This will enhance operators' decision making abilities, especially during beyond-design-basis accidents.

3. Building a Sodium-Cooled Reactor Model

Sections 1 and 2 discuss the primary goals of the research: the creation of a better-than-real-time modeling system for diagnosis and prognosis. In the present section we go into the methodology for creating the full original DBN model. In Section 4, we describe how we have validated the model.

3.1. Accident Data Generation and Processing

To build a large probabilistic model it is necessary to have both the knowledge of how the components of the reactor relate to each other as well as probabilities describing component and component relationship failures.

The simulated accident data was generated using the SAS4A¹⁵ liquid metal reactor simulator which performs deterministic analysis on nuclear accident scenarios.

^a*Probabilistic Risk Assessments* are used by the United States Nuclear Regulatory Commission to quantify the causes, likelihood, and consequences of nuclear accidents. *Dynamic Probabilistic Risk Assessment* studies a system's dynamics often by employing reactor state simulators.

^b*Discrete Dynamic Event Tree* based software dynamically branches the accident simulations whenever there is more than one possible outcome.




Fig. 2. High level diagram of a Sodium Fast Reactor (SFR). The primary system pump flows relatively cold sodium over the reactor core. The fuel in the core (at a rate managed by control rods) heats the sodium which then transfers energy through the heat exchanger to secondary sodium. This secondary sodium heats water in a steam generator which drives a steam turbine. The secondary pumps return the cooled sodium from the steam generator to the heat exchanger for reuse¹⁴.

We executed simulations varying the states of four target variables: the functional capacity of the nuclear power plant's main coolant system (differential pressure), Direct Reactor Auxiliary Cooling System (DRACS), the Balance of Plant (BOP) systems, and control rod insertion (SCRAM). These systems, seen in Figure 3 and Table 1, were chosen as target variables since compounded failures of these systems have the potential to cause extensive core damage.

The differential pressure parameter has three possible states: 100%, 50%, and 0%. This variable describes the capacity of the pumps removing heat from the core. The DRACS parameter has three possible states: *available*, *degraded*, and *enhanced*. Balance of Plant (BOP) has three states: *operational*, *decay*, and *shutdown*. The *SCRAM* parameter has three states: control rods *nominal*, *fully in*, or *withdrawn*. The *SCRAM* parameter describes the level of the operators' attempt at an emergency core shutdown^c. The states of these reactor systems are initial conditions for the simulations.

^cThe control rods in a nuclear power plant are used to control the fission rate of the core.

The nuclear power industry uses *Discrete Dynamic Event Trees* to capture the changing of complex situations across time. Discrete Dynamic Event Trees capture the situation when a fixed initial condition evolves into an array of possible end states, when decisions that create the end states are context specific. Nuclear power plant accidents almost always start from steady power operation and, as events and decisions change over time, they are captured as branches in the tree.

The simulations provided for the prototype focus on two types of accidents: both seismic and non-seismic-induced *Transient Overpower* and *Loss of Flow* accidents. The *Discrete Dynamic Event Tree* was designed to branch on multiple conditions including the magnitude of an earthquake, balance of plant availability, the SCRAM state, the DRACS state, secondary pump power, reactivity response, and coolant pump status. Some branching conditions were determined dynamically by SAS4A. Further description of the simulations, branching conditions, and branching probabilities, is provided in¹⁶. The event tree also included nominal scenarios with no earthquake and with various combinations of states of the variables. Such scenarios are investigated in order to provide baseline conditional probabilities in the model. The event tree is comprised of 7189 distinct SAS4A simulations. Each simulation contains 2588 time steps corresponding to the first 48 hours of its scenario.

3.2. Data Processing

Each simulation offers a permutation of the initial conditions. These initial conditions are drawn from the *Discrete Dynamic Event Tree* and are tied to a particular state of the target variables. The resulting states of the observation variables — variables that could be inferred by an operator’s instrumental observations — are returned by the SAS4A simulator. Each variable provided 2558 times steps of data for each of the 7189 simulations. With 12 observed variables, we had generated 2.3 gigabytes of data.

The next challenge was to process the data in order to construct the Bayesian network. We previously hand-quantified simple Bayesian networks using the GeNIe^{d12} probabilistic modeling software. We augmented these simple hand-made models with the SAS4A data. This required us to implement a data processing program which we call ALADDIN^{e17}. ALADDIN discretizes the data, calculates the conditional probabilities of each of the observation variable states given each combination of possible target variable states at each time step, and builds the dynamic Bayesian network using SMILE^{d12}.

The proper discretization of data values can improve the performance of supervised learning algorithms¹⁸. Therefore we discretized the conditional probability states of each variable into equal width distributions. This binning method calculates the maximum and minimum values for each variable’s probability distribution

^dGeNIe is a graphical interface to the University of Pittsburgh’s Structural Modeling, Inference, and Learning Engine or “SMILE”

^eAutomatic Loader of Accident Data for Dynamic Inferencing Networks

and distributes the values into k bins of equal width. We divided the probability values for each variable into 3 bins.

After parsing and discretization, ALADDIN calculates the conditional probabilities for each of the nodes in the DBN. If we let P be the number of plant parameter nodes, T be the number of time steps, N be the number of bins for each plant parameter node, and S be the number of reactor system state combinations, then the number of conditional probabilities is $P \cdot N \cdot S \cdot T$. For this example model where $P = 12$, $T = 96$, $N = 3$, and $S = 108$, we have 373,248 conditional probabilities.

3.3. Construction of The Network

ALADDIN reads a provided dynamic Bayesian network outline which the user constructs with GeNIe. This pre-built network contains a node for each of the plant's parameters, including the target and observation variables, and defines the relationships between them (see Figure 3). The observation variables, along with their states are given in Table 1. The system then populates the model with the conditional probabilities of each observation node at each time step.

Besides providing an interface for constructing the Bayesian network's structure, GeNIe also provides a platform to analyze the network by propagating evidence and diagnosing the states of the plant's parameters. This functionality is used as a decision support system. Users can input a set of known conditions, which propagate evidence to the unobserved target variables. The posterior probability is then used to predict the evolution of important reactor systems.

Target	States	Prior Probabilities
Balance Of Plant	Operational	1.19×10^{-12}
	Shutdown	0.9999
	Decay	3.97×10^{-13}
Differential Pressure	0% flow	3e-13
	50% flow	1.04×10^{-4}
	100% flow	.9999
SCRAM	Fully In	0.0150
	Nominal	0.985
	Withdrawn	3.04×10^{-6}
DRACS	Enhanced	0.9850
	Operational	0.0150
	Degraded	7.95×10^{-12}

4. Model Performance

In this section we study the effectiveness and performance of the DBN and its components. We perform model and variable validation using two traditional ap-

10 *M.C. Darling, G.F. Luger, T.B. Jones, M.R. Denman, K.M. Groth*




Fig. 3. User-defined DBN Model Structure: the user creates the DBN model structure to define the relationships between observation nodes, target nodes, and intermediate nodes. The oval-shaped nodes, which are placed on a temporal plate, represent observed variables. The rectangular nodes represent reactor systems. The rectangular nodes on the far right are targets representing the accident states. The system reads this hand-made model and populates its conditional probability tables based on the data.




Fig. 4. The accuracies and F-scores of all target variable states with varied data sets. The baseline value is when the most frequent target combination is always chosen. Out of all the observed variables, $t_{coolant}$ and $reactivity$ have the highest and lowest KL divergence values respectively. Evaluating the BN while dropping each of these variables from the data set illustrates the change in performance of the BN's predictive power. For all of the target variables, dropping $reactivity$ has no effect. Whereas removing the coolant temperature reduces the BN's performance.

proaches, a traditional AI technique for evaluating model adequacy in stochastic systems^{19,20}. First we test for variable dependencies using Kullback-Liebler divergence. We then conduct model cross-validation with analysis of F-scores and pre-

diction accuracies.

4.1. *KL Divergence*

In order to measure the pertinence of each of the plant parameters, we implemented a variable evaluator based on Kullback-Leibler (KL) divergence²¹. KL divergence measures the distance between two probability distributions (in this case, between two BN models). In information theory, the Kullback-Leibler divergence equation is used to measure the amount of information lost when the distribution function Q is used to approximate the actual distribution function P :

$$D(P||Q) = \sum_i P(i) \log \left(\frac{P(i)}{Q(i)} \right) \quad (1)$$

where $P(i)$ represents the true probability distribution and $Q(i)$ represents a theoretical distribution. The equation defines divergence D between P and Q (i.e., the information lost when $Q(i)$ is substituted for $P(i)$).

In this application, KL divergence is used to compare the BN model that includes all of the plant parameters with BN models that have removed one of the parameters. The divergence between two models shows how much information is lost by the elimination of one of the variables. If there is a large amount of information lost when a node is removed, then the node is highly pertinent. If the information loss is minimal, then the node may be unnecessary and thus a candidate for removal.

In calculating the KL divergence of an arc, $P(i)$ is the model that includes a particular node to be measured while $Q(i)$ is the model without that node. The values summed over i are combinations of possible observed and target states. KL divergence is calculated for each arc between the observation and target nodes in a method similar to²². Joint KL divergence calculations are conducted over all the target nodes for each observation node. In calculating the joint KL divergence, we treat each combination of possible target states as a single state in a joint target node that collects all targets into a single node.

After calculating the KL divergence value of each of the observation variables we found that the coolant temperature ($t_coolant$) had the highest KL divergence value at 6.793×10^{-12} bits while *reactivity* had the lowest KL divergence value at 3.234×10^{-18} bits. Thus out of all observed variables, the reactor's coolant temperature and reactivity provide the highest and lowest amount of information gain respectively. A plot of the KL divergence values for all of the variables can be seen in Figure 5.

4.2. *Cross-Validation*

When investigating predictive models, cross-validation is a standard technique to measure the predictive power of a model while accounting for the possibility of data "overfitting" that model²³. To perform the cross-validation analysis, we construct




Fig. 5. KL divergence values for the plant parameters. (Time steps have been converted into accident time). KL divergence is measured as the number of bits of information lost when the variable is removed from the model.

the model on a subset of the data and test with the remaining data. The goal is to measure how well the Bayesian network models the SAS4A simulation data.

To measure model performance, we use the F-score and accuracy of each of the target variable states predicted by the Bayesian network when fed the discretized observation variable states produced by the SAS4A simulator. The F-score is defined as:

$$F_{\beta} = (1 + \beta^2) \cdot \frac{\text{precision} \cdot \text{recall}}{(\beta^2 \cdot \text{precision}) + \text{recall}}$$

where β is a weighting parameter between precision and recall (it is usually set to 1 for even weighting), and where precision and recall are defined as:

$$\text{Precision} = \frac{\text{TruePositives}}{\text{TruePositives} + \text{FalsePositives}}$$

$$\text{Recall} = \frac{\text{TruePositives}}{\text{TruePositives} + \text{FalseNegatives}}$$

The F-score is useful for measuring the behavior of a machine learner on an asymmetric data set, while the accuracy is generally used for symmetric data sets.

The results of the F-score and accuracy measurements can be found in Figure 4. The figure shows that the F-score of the BN with full simulation data tends to outperform the baseline. In a few instances, however, this is not the case – specifically, the F-score for DRACS *operational* and the *SCRAM nominal* state. Both of these states are the chosen state for the baseline Bayesian network. This indicates that the Bayesian network is occasionally incorrectly predicting that these two states are not the correct state when they, in fact, are. This is to be expected when moving from a static baseline that always chooses the most likely target combination while ignoring the observation variables.

The accuracy of the DRACS operational and enhanced states are the only states where the baseline guess is better than the BN with full observational data or observation data with dropped variables. There were also a few instances where dropping the variable with the highest KL divergence value actually improved results in the BN. These included two of the *SCRAM* states — *Nominal* and *Withdrawn* in accuracy and the *SCRAM nominal* and BOP *Decay* on the F-score. This points to the importance of choosing the correct variables to predict the target accuracies on each of these variables. The highest KL variable, coolant temperature ($t_{coolant}$), was useful for predicting most variable states, but it caused a decrease in accuracy and F-score for some *SCRAM* and BOP states.

5. Impact Analysis Using Counterfactuals

Impact analysis using counterfactual reasoning involves exploration of the events that led up to the current state of affairs when they do not coincide with expectations, and the possible likely outcomes of pending decisions. This exercise allows operators to determine whether an unexpected outcome is due to a fault in planning, or in an unexpected external event. It also supports reasoning over possible and probable consequences of future actions. Applications of counterfactual analysis include fault diagnosis planning and policy analysis. Balke and Pearl define the query in the abstract as: "If A were true, would C have been true?", where A specifies an event that is contrary to one's real-world observations, and C is known as the counterfactual consequent²⁴.

The ultimate goal of our data-driven decision-support system is to be able to perform impact analysis both on the decisions that have led to the current state of affairs, as well as on real-time decisions that would influence the future state of the reactor. This analysis tool allows operators to determine the most likely outcomes of actions under consideration. In the remainder of this section we present three examples of counterfactual reasoning, demonstrating how, given a current "danger" state of the sodium-cooled nuclear power generator, human operators can determine which decisions are best for the "health" of the system. Two of the situations are described in general detail, outlining states of the model and possible decisions. The third situation we describe in detail, demonstrating the justification data that support the remediating decision.

5.1. *Description of Problems and Decisions*

Even though sodium fast reactors are designed to be inherently safe across a wide array of accident conditions without intervention, operators may still feel compelled to act to reduce potential stresses to reactor components. Most accident management plans are focused on ensuring that more heat is removed from the reactor core than is generated in the reactor core. If the reactivity control system fails and the power does not decrease upon initiation of the accident, the most direct method for accident management is to fix the fault which prevented the reactivity control system from activating. This fault may be a function of the instrumentation and control system, a relatively easy fix, or it may be a mechanical failure of the control rod drive mechanisms, a more difficult fix due to access restrictions under accident conditions. If the reactivity control system cannot be activated, the operators might decide to reduce flow through the primary pumps to reduce the heat generation loads from those pumps in the primary system. This action would reduce forced circulation over the fuel which might cause undesirable overheating. A third action may be to attempt to increase the energy removed through the emergency heat removal system. This is the action which is explored in the following subsection.

5.2. *Impact Analysis of Enhancing DRACS*

For a detailed impact analysis, we analyze the decision of whether or not the operators should choose to enhance the Direct Reactor Auxiliary Cooling System (DRACS). This decision is conditioned by the operators' failure to fully engage the SCRAM system and the possibility of complete DRACS failure. When the core begins to overheat, the operators attempt to control its reactivity by further inserting the control rods (this is the *SCRAM* system); if the control rods are fully inserted, the core's reactivity will rapidly decrease. However, if the control rods fail to further insert, due to being knocked out of alignment or some other mechanical failure, then alternative measures for cooling the core will be explored. Figure 6 shows the coolant temperatures in all simulations where the operators take no action on the DRACS system and the control rods are either somewhat or completely withdrawn. In these scenarios, there is a chance of the pool of sodium coolant boiling (liquid sodium boils at 1173.15 Kelvin), which is an emergency the operators must prevent.

"Enhancing" the DRACS system is one hypothetical emergency option for cooling core and preventing the coolant from boiling. This enhancement implies modifying the DRACS system in a manner beyond its design (i.e., adding water to a chamber meant only for air) and can result in complete failure of the system. This would leave the operators with no further recourse for cooling the core if the primary coolant system is also damaged. Figure 7 shows all simulations where the attempt to enhance DRACS results in its degradation.

It is imperative to maintain the coolant temperature below its boiling point. The operators of the power plant would have to decide whether it is worth the risks to enhance their DRACS before the coolant temperature approaches 1173.15




Fig. 6. The liquid sodium temperatures in all simulations where the operators take no action on the DRACS system. The graph on the left shows the cases where the control rods are in their nominal position which approximately results in an 11.34 percent chance of the sodium boiling; the figure on the right shows the simulations where control rods are completely withdrawn: there is approximately a 68.71 percent chance of the sodium reaching its boiling point.




Fig. 7. The liquid sodium temperatures in all simulations where the operators attempt to enhance the DRACS system but it ultimately fails (degrades). The graph on the left shows the cases where the control rods are in their nominal position which results in an 27.97 percent chance of the sodium boiling; the figure on the right shows the simulations where control rods are completely withdrawn: there is approximately a 69.39 percent chance of the sodium reaching its boiling point.

K. Therefore, we measure the probability of coolant boiling as a function of its temperature and the time at which the temperature occurs to support the decision of whether or not to enhance the system. Figure 8 shows the coolant temperatures in all simulations where the operators choose to enhance the DRACS system.

The probability of the sodium boiling increases as its temperature rises. However, since the temperature is dependent on the state of the reactor and the mechanical functioning of the coolant systems, the surge in temperature indicates the likelihood that there is a malfunction. Therefore, the decision to enhance the DRACS can be based on the probability of the sodium boiling given its temperature and the time

16 *M.C. Darling, G.F. Luger, T.B. Jones, M.R. Denman, K.M. Groth*




Fig. 8. The liquid sodium temperatures in all simulations where the operators attempt to enhance the DRACS system. The graph on the left shows the cases where the control rods are in their nominal position which results in an approximately 18.06 percent chance of the sodium boiling (note that some of these simulations result in a degraded DRACS); the figure on the right shows the simulations where control rods are completely withdrawn: there is an approximately 69.05 percent chance of the sodium reaching its boiling point.




Fig. 9. Simple example: The probability of the liquid sodium coolant boiling based on time and its temperature. The left panel displays the coolant temperatures for 3 simulations as they progress in time; the panel on the right displays the probability of the coolant boiling when its temperature is in the range of 892 – 898 K during the same time period. Only one of the three simulations will eventually reach the boiling point. The three simulations have a temperature of exactly 898 K at 4 points in the simulation. These points are annotated by numerical markers. At markers 1 and 2 the three simulations exhibit identical behavior. Therefore, at these two points the probability of the coolant boiling is exactly 1/3. However, at marker 3, the temperature of one of the simulations rises more rapidly than the other two, and is the only one to cross 898 K at this point in time. Therefore, at the time of marker 3 the right panel shows that the probability of boiling is 1. Finally, at marker 4, the other two simulations cross 898 K but since they will not eventually boil, the right panel shows a boiling probability of 0.

at which it occurs.

For the purpose of illustration, Figure 9 shows an example for three simulations. In this scenario, the control rods are in their nominal position but the operators




Fig. 10. The panel on the left shows the probability of sodium boiling when the operators choose not to enhance DRACS; the right panel shows the probability of boiling when they do enhance DRACS. These graphs indicate that, under these conditions, if the coolant temperature reaches 880 K soon after the accident begins, there is a greater-than 90% chance of the coolant boiling, which is slightly lower if they choose to enhance—therefore, they might choose to do so. However, if the coolant temperature reaches 880 later into the accident sequence, there is a greater probability of sodium boiling if they choose to enhance. Therefore, at this point in time, these sequences generate so much heat that the DRACS is ineffectual and thus operators should focus on other ways of managing the accident.

choose to enhance DRACS and are successful in preventing the boiling temperature two out of three times. This example is a snapshot of the times at which the temperature is in the range of 892 – 898 K. All three simulated coolant pools reach 898 two times simultaneously. This is why the panel on the right shows the probability of coolant boil is 1/3 at the two points of their simulation where temperature is 898 K. However, at marker 3, the temperature of one of the simulations rises more rapidly than the other two and is the only one to cross 898 K at this point in time. Therefore at the time of marker 3 the right panels shows that the probability of boiling is 1. Finally, at marker 4, the other two simulations cross 898 K but since they will not eventually boil, the right panel shows a boiling probability of 0.

Figure 10 shows the results of the same scenario at 880 K calculated using the full data set. The panel on the left shows the probability of sodium boiling when the operators choose not to enhance DRACS; the right panel shows the probability of boiling when they do enhance DRACS. These graphs indicate that, under these conditions, if the coolant temperature reaches 880 K soon after the accident begins, there is a greater-than 90% chance of the coolant boiling, which is slightly lower than if they choose to enhance—therefore, they might choose to do so.

However, based on these calculations, if the coolant temperature reaches 880 later into the accident sequence, there is a greater probability of sodium boiling if they choose to enhance. Therefore, at this point in time, these sequences generate so much heat that the DRACS is ineffectual and thus operators should focus on other ways of managing the accident. Figure 11 shows the same scenario at multiple temperature ranges. The information captured by these graphs provides the opera-

18 *M.C. Darling, G.F. Luger, T.B. Jones, M.R. Denman, K.M. Groth*




Fig. 11. The probability of the coolant boiling given the time steps at various temperatures. The panels on the left show the probability of sodium boiling when the operators choose not to enhance DRACS, the right panels show the probability of boiling when they do enhance DRACS. By comparing the probabilities of enhancing and not enhancing, the operators would be able to decide whether the decision is advantageous from a probabilistic perspective.

tors some measure with which to base their decision, whereas currently they would be considering this drastic choice of action using instinct and raw data (assuming properly functioning instruments).

5.3. Further Examples of Counterfactual Inference

While this paper focuses on the operator action relating to enhancing the DRACS performance, two other operator actions are also discussed: restoring the reactivity control system and reducing power to the main coolant pumps. These operator actions are not examined but are briefly described here.

Accident sequences that result in a loss of reactivity control are caused by either a failure of the digital control system or physical failure of the control rod insertion mechanisms. Failures in the digital control system can possibly be bypassed through redundant activation commands or may be physically bypassed through manual activation of circuit breakers that power the control rod insertion mechanisms. Physical failures of these control rod insertion mechanisms are located on the top of the reactor vessel, a space which may be both thermally hot and radioactive and thus difficult to access. Successfully completing fixing or bypassing the reactivity control system will always lead to improved end states. The difficulties associated within quantifying the likelihoods and timings associated with these actions prevented further study of this action in the dynamic event tree.

The primary system reactor coolant pumps force cold sodium over the hot reactor core, thus cooling the core. While the cooling function of these pumps is desired, the sodium reactors are generally designed to be cooled via natural circulation during accident conditions and with a negative temperature coefficient which reduces the power generation rate when the reactor heats up. Eventually the operation of the reactor coolant pumps artificially lowers the temperature of the core which allows for higher equilibrium power generation and worse long term impact. These pumps also dump heat into the primary system that will need to be removed by the DRACS²⁵. Unfortunately, loss of pumps during transient overpower conditions may cause the reactor to overheat and fail the fuel cladding. This decision balance between when the pumps are needed and when they are a hindrance is worthy of future study.

5.4. Constraints of Current Data

The probabilities of the coolant boiling provide the operators with a criterion with which to base their decision on whether to enhance the DRACS system. This is an improvement over the current practice where operators have no such criterion. However, the data used in this case study though dynamically generated for the plant parameters, is not dynamic with respect to the reactor systems (which includes DRACS).

This is significant since the decision to enhance the DRACS should not only be based on temperatures, it should also be based on the time it takes to initiate the enhancement and for the time it would take to actually affect the system. The decision would have to be made early enough in the accident sequence to have any hope of averting core damage. We leave it as future work to generate the needed dynamic data in order to perform this analysis on the DRACS enhancement

20 M.C. Darling, G.F. Luger, T.B. Jones, M.R. Denman, K.M. Groth

decision.

6. Conclusions and Future Work

SMART procedures is an ongoing project aimed at providing diagnostic tools in the event of a beyond-design-basis nuclear power plant accident, given the values of observed plant parameters. This framework is meant to condense the rich data output from nuclear reactor simulators into a model that can be used to probabilistically analyze an accident in faster-than-real time. This supports a greater understanding of the state of a reactor during accidents when only a subset of information might be available.

This paper presents a prototype implementation for the SMART procedures framework. This entails simulating accidents with a comprehensive set of conditions and using the resultant data to build a dynamic Bayesian network. Once built, the DBN can assist diagnosis of the states of observed power plant parameters that can lead to *Transient Over Power* and *Loss Of Flow Accidents*.

We evaluated the performance of the model in two ways: calculating its internal consistency and, using F-scores, its accuracy. First, we used Kullback-Leibler divergence measures to test the importance within the model of the individual nodes and arcs in determining the current state of the reactor. Second, we used the traditional F-score measure to evaluate the overall accuracy of the model in predicting the current state of the system. The results of these analyses are presented in Section 4.

The most important new result of this paper was to test counterfactual reasoning algorithms to determine possible next states of the system should certain decisions be taken, given an accident situation. Balke and Pearl²⁴ suggested counterfactual inference for exploring possible DBN system states that could occur in the future, given decisions made in the present. In Section 5, we used counterfactual reasoning to calculate the probability that the plant's coolant would boil based on the choice of whether or not operators enhanced its auxiliary coolant system beyond its designed capacity.

Although our current prototype shows promising results, there are still many components in need of improvement. Currently, we discretize the data using *Equal Width* binning; we plan to implement alternative, more sophisticated, techniques such as *Entropy* or dynamic *Minimum Descriptive Length* partitioning²⁶, as we suspect that these can improve the model's predictive accuracy. We also must continue to expand and calibrate the model parameters.

Another important route forward is measuring the effect of adding or subtracting variables so that only those that have the best predictive power for each target system are used. We also need to generate more data in order to model other types of accidents. In addition, we plan to implement a full impact analysis tool capable of analyzing multiple decisions under varying conditions.

Definitions

DRACS: Direct Reactor Auxiliary Cooling System

KL: Kullback-Leibler

LOF: Loss of Flow Accident

SCRAM: Control rod insertion

TOP: Transient Overpower Accident

Acknowledgements

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy’s National Nuclear Security Administration under contract DE-NA0003525. This project was funded in part by the DOE Office of Nuclear Energy Advanced Reactor Technologies Program under work package number AT-15SN200304.

The authors would like thank Jeff Cardoni and Zac Jankovsky, who conducted the SAS4A simulations referenced in this paper; Argonne National Laboratories for their assistance in providing an initial SAS4A model for the analysis.

References

1. FNI, Investigation committee on the accident at Fukushima nuclear power stations, Tech. rep., Final Report, July 23, 2012 (2012).
2. J. Cahalan, A. Tentner, E. Morris, Advanced LMR safety analysis capabilities in the SASSYS-1 and SAS4A computer codes, Tech. rep., Argonne National Lab., IL (United States). USDOE, Washington, DC (United States) (1994).
3. D. Chanin, J. Sprung, L. Ritchie, H.-N. Jow, MELCOR accident consequence code system (MACCS), Tech. rep., Nuclear Regulatory Commission, Washington, DC (USA). Div. of Systems Research; Sandia National Labs., Albuquerque, NM (USA) (1990).
4. J. Pearl, Probabilistic reasoning in intelligent systems: networks of plausible inference, Morgan Kaufmann, 1988.
5. J. Pearl, Causality, Cambridge university press, 2010.
6. K. P. Murphy, Dynamic Bayesian networks: representation, inference and learning, Ph.D. thesis, University of California, Berkeley (2002).
7. K. Kim, E. B. Bartlett, Nuclear power plant fault diagnosis using neural networks with error estimation by series association, Nuclear Science, IEEE Transactions on 43 (4) (1996) 2373–2388.
8. K. M. Groth, M. R. Denman, J. N. Cardoni, T. A. Wheeler, ‘Smart Procedures’: Using dynamic PRA to develop dynamic, context-specific severe accident management guidelines (SAMGs), in: Proceedings of the International Conference on Probabilistic Safety Assessment and Management (PSAM 12), Honolulu, HI, 2014.
9. K. M. Groth, M. R. Denman, T. J. Jones, M. C. Darling, G. F. Luger, Proof-of-concept accident diagnostic support for sodium fast reactors, in: European Safety and Reliability Conference ESREL 2015, Zurich Switzerland, 2015.
10. T. J. Jones, M. C. Darling, K. M. Groth, M. R. Denman, G. F. Luger, A dynamic bayesian network for diagnosing nuclear power plant accidents, in: The 29th International Florida Artificial Intelligence Research Society Conference, 2016.

22 *M.C. Darling, G.F. Luger, T.B. Jones, M.R. Denman, K.M. Groth*

11. R. Gauntt, R. Cole, S. Hodge, S. Rodriguez, R. Sanders, R. Smith, D. Stuart, R. Summers, M. Young, MELCOR computer code manuals, Sandia National Laboratories, NUREG/CR 6119.
12. R. Gauntt, D. Kalinich, J. Cardoni, J. Phillips, A. Goldmann, S. Pickering, M. Francis, K. Robb, L. Ott, D. Wang, et al., Fukushima daiichi accident study (status as of April 2012), Sandia National Laboratory Report, SAND2012-6173, Albuquerque, NM.
13. U. S. DOE, A technology roadmap for generation IV nuclear energy systems, in: Nuclear Energy Research Advisory Committee and the Generation IV International Forum, 2002.
14. Argonne National Laboratory, The SAS4A/SASSYS-1 Safety Analysis Code System, ANL/NE-12/4, Argonne, IL (2011).
15. M. R. Denman, K. M. Groth, J. Cardoni, T. Wheeler, Advance liquid metal reactor Discrete Dynamic Event Tree/Bayesian network analysis and incident management guidelines (risk management for sodium fast reactors), SAND2015-2484, Sandia National Laboratories, Albuquerque, NM (2015).
16. M. J. Druzdzel, SMILE: Structural modeling, inference, and learning engine and GeNIe: a development environment for graphical decision-theoretic models, in: Proceedings of American Association for Artificial Intelligence (AAAI-99), 1999, pp. 902–903.
17. M. C. Darling, T. B. Jones, G. F. Luger, K. M. Groth, ALADDIN: The automatic loader of accident data for dynamic inferencing networks, Tech. Rep. SAND2015-10401, Sandia National Laboratories, Albuquerque, NM (2015).
18. Y. Yang, G. I. Webb, A comparative study of discretization methods for naive-bayes classifiers, in: Proceedings of PKAW, Vol. 2002, Citeseer, 2002.
19. M. I. Jordan, Z. Ghahramani, T. S. Jaakkola, L. K. Saul, An introduction to variational methods for graphical models, Machine learning 37 (2) (1999) 183–233.
20. T. Minka, et al., Divergence measures and message passing, Tech. rep., Technical report, Microsoft Research (2005).
21. R. G. Cowell, Conditions under which conditional independence and scoring methods lead to identical selection of Bayesian network models, in: Proceedings of the Seventeenth conference on Uncertainty in artificial intelligence, Morgan Kaufmann Publishers Inc., 2001, pp. 91–97.
22. J. R. Koiter, Visualizing inference in Bayesian networks, Master’s thesis, Delft University of Technology (2006).
23. A. Y. Ng, Preventing “overfitting” of cross-validation data, in: ICML, Vol. 97, 1997, pp. 245–253.
24. A. Balke, J. Pearl, Counterfactuals and policy analysis in structural models, in: Proceedings of the Eleventh conference on Uncertainty in Artificial Intelligence, Morgan Kaufmann Publishers Inc., 1995, pp. 11–18.
25. V. Ho, G. Apostolakis, Dynamic operator actions analysis for inherently safe fast reactors and light water reactors, Reliability Engineering & System Safety 22 (1-4) (1988) 61–90.
26. E. J. Clarke, B. A. Barton, Entropy and MDL discretization of continuous variables for Bayesian belief networks, International Journal of Intelligent Systems 15 (1) (2000) 61–92.