SANDIA REPORT

SAND2015-8872
Unlimited Release
Printed October 2015

Interim Status Report for Risk Management for SFRs

Zachary K. Jankovsky, Matthew R. Denman, Katrina M. Groth, Timothy A. Wheeler

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Sandia National Laboratories

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone:	(865) 576-8401
Facsimile:	(865) 576-5728
E-Mail:	reports@adonis.osti.gov
Online ordering:	http://www.osti.gov/bridge

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd
Springfield, VA 22161

Telephone:	(800) 553-6847
Facsimile:	(703) 605-6900
E-Mail:	orders@ntis.fedworld.gov
Online ordering:	http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online

SAND2015-8872
Unlimited Release
Printed October 2015

Interim Status Report for Risk Management for SFRs

Zachary K. Jankovsky, Matthew R. Denman, Katrina M. Groth, Timothy A. Wheeler
Risk & Reliability Analysis
Sandia National Laboratories
P.O. Box 5800
Albuquerque, NM 87185-MS0748

Abstract

Accident management is an important component to maintaining risk at acceptable levels for all complex systems, such as nuclear power plants. With the introduction of passive, or inherently safe, reactor designs the focus has shifted from management by operators to allowing the system's design to take advantage of natural phenomena to manage the accident. Inherently and passively safe designs are laudable, but nonetheless extreme boundary conditions can interfere with the design attributes which facilitate inherent safety, thus resulting in unanticipated and undesirable end states. This report examines an inherently safe and small sodium fast reactor experiencing a variety of beyond design basis events with the intent of exploring the utility of a Dynamic Bayesian Network to infer the state of the reactor to inform the operator's corrective actions. These inferences also serve to identify the instruments most critical to informing an operator's actions as candidates for hardening against radiation and other extreme environmental conditions that may exist in an accident. This reduction in uncertainty serves to inform ongoing discussions of how small sodium reactors would be licensed and may serve to reduce regulatory risk and cost for such reactors.

Acknowledgments

The authors would like thank Jeff Cardoni (Sandia), who conducted the SAS4A simulations referenced in this report, ANL for their assistance in providing an initial SAS4A model for the analysis, and DOE NE Advanced Reactor Technologies program for providing the funding for this project.

The authors would also like to thank George Luger, Michael Darling, and Thomas Jones, all of the University of New Mexico, for their contributions to ALADDIN, which was vital for loading SAS4A data into GeNIe.

Contents

Figures

Tables

Executive Summary

Severe accidents are extremely rare in the nuclear power industry. However, as demonstrated by the Fukushima accident, rare events are not impossible events, and responding to these accidents can be extremely difficult. Severe Accident Management Guidelines (SAMGs) serve as a critical resource that would help operating crews respond to severe accidents. Sandia is investigating whether dynamic PRA could improve SAMGs and thus human reliability.

Dynamic, simulation-based Probabilistic Risk Assessment (PRA) methods can provide a scientific basis for supporting the diagnosis and response planning for current and future reactor designs. Recent advances in computing enable simulation-based PRA approaches to explore thousands of accident scenarios. Coupling these scenarios with plant simulations allows prediction of plant parameters and consequences associated with each accident scenario. In effect, running thousands of advanced PRA simulations allows experts to explicitly map out the relationship between known accident scenarios and observable reactor parameters. Dynamic PRA offers a comprehensive understanding of the accident scenarios and the associated plant states.

The methodology proposed in [1, 2] would allow the results of dynamic PRA to be harnessed to provide comprehensive, science-based support to operators facing severe accidents that fall beyond the scope of existing procedures, training, and experience. By formally encoding advanced PRA knowledge in SMART (Safely Managing Accidental Reactor Transients) SAMGs, we could reduce the socio-technical challenges associated with responding to severe accidents, and provide an additional line of defense against events which have traditionally been related to beyond design basis accidents or residual risk.

This report describes a severe accident management tool and demonstration for a small sodium reactor subjected to transient overpowers resulting from 0.3g and 0.5g earthquakes, transient overpower due to control rod removal, loss of primary coolant flow, and loss of operating and shutdown heat removal. SNL's PRA research under DOE's Advanced Reactor Technology program had three goals. A primary goal was to analyze a spectrum of accidents and explore the impact both of initial uncertainties and of human intervention, in order to provide operational insights to reduce potential for reactor damage. That goal was advanced and documented in ”Advance Liquid Metal Reactor Discrete Dynamic Event Tree/Bayesian Network Analysis and Incident Management Guidelines (Risk Management for Sodium Fast Reactors)” [3]. Another goal was to use the insights from the accident analyses to create a Dynamic Bayesian Network (DBN) which can learn from the instrumented variables with the objective of inferring key states of the reactor. This objective was also advanced in Reference [3]. A third goal was to use the DBN to provide example inferences throughout the accident progression tree. This goal was accomplished and documented in this report.

Key Results

The key results documented in this report include:

- Thousands of SAS4A simulations can be structured in a series of tractable dynamic event trees that explore uncertainties in reactivity coefficients, electromagnetic pump performance, reactor protection system performance, and operator modifications of the direct reactor auxiliary heat removal system and primary and secondary loop electromagnetic pump operation.
- Temporal information, including indicators that may be available to the operators such as temperatures, can be extracted from SAS4A, transformed into a reduce order data set that is compatible with a DBN, and combined with component end states from the simulation to train a DBN.
- Key observations, number less than 10 separate observations, can be applied to the trained DBN to diagnose some accidents that do not involve failure of the reactor protection system.
- Diagnosis of accidents that involve failure of the reactor protections system will require an automated process for feeding in accident parameters to the trained DBN due to the large amount of evidence needed to overcome the low probability of reactor protections system failure.

Once the DBN can be utilized across the accident space, the initial SAS4A simulations can be integrated to inform optimal accident management procedures.

Nomenclature

ADAPT Analysis of Dynamic Accident Progression Trees

ALADDIN Automatic Loader of Accident Data for a Dynamic Inferencing Network

BOP Balance of Plant

DBN Dynamic Bayesian Network

COL Combined License

CR Control Rod

DC Design Certification

DDET Discrete Dynamic Event Tree

DHRS Decay Heat Removal System

DRACS Direct Reactor Auxiliary Cooling System

EMP Electro-Magnetic Pump

EPRI Electric Power Research Institute

IHTS Intermediate Heat Transport System

KL Kullback-Leibler

LOF Loss of Flow

LOHR Loss of Operating Heat Removal

LWR Light Water Reactor

NRC United States Nuclear Regulatory Commission

PRA Probabilistic Risk Assessment

PSID Preliminary Safety Information Document

RCS Reactor Coolant System

RPS Reactor Protection System

SAMG Severe Accident Management Guideline

SFR Sodium Fast Reactor

SMART Safely Managing Accidental Reactor Transients

SMR Small Modular Reactor

TOP Transient Overpower

ULOF Unprotected Loss of Flow

UTOP Unprotected Transient Overpower

1 Introduction

1.1 Challenge/Problem Description

In this manuscript, we develop a proof-of-concept Dynamic Bayesian Network accident management model for a sodium fast reactor (SFR) which could be used to assist nuclear reactor operators to infer the state of the reactor with a subset of information that would be available during the accident as illustrated in [1]. We then use the model to investigate whether such a model is capable of providing insight into which reactor parameters provide the most valuable information for diagnosis. In the near term, the results could be used to determine which reactor parameters should be instrumented in the control room. In the longer term, the results would be a first step toward a full SMART procedures system.

1.1.1 Relevance to Industry

Generally, SFRs have several technological advantages that can affect the operation and safety of the plant. Two examples are passive safety features and low reactor pressure. Passive safety features utilize gravity-driven or natural convection systems rather than active, pump-driven systems to supply heat removal during upset conditions. An example of this is the sodium pool. The combination of large thermal mass and high conductivity helps to cool the reactor if pumps are not operational. The relatively low reactor vessel pressure (when compared to traditional LWRs) can reduce the cost of the vessel itself and simplify the overall system. On the other hand, SFRs also have unique technological challenges, which include positive coolant void and coolant temperature coefficients. While the overall net reactivity does decrease with increased temperature, the individual positive coefficients may pose a regulatory challenge. The fast-spectrum fuel is also not in its most reactive configuration when operating, which can lead to recriticality concerns during a severe accident.

SFR designs face regulatory challenges since regulations have been built up around the popular LWR technology. Nuclear Regulatory Commission (NRC) regulations specify the operator, security, and emergency response requirements for licensed nuclear reactors. Many of the existing regulations may not be applicable to SFR designs, especially considering the different coolant material and fast neutron spectrum.

Though no power SFRs have been licensed by the NRC and built, there is significant sodium reactor history in the USA. EBR-II [4] was operated for approximately 30 years, and was involved in numerous successful safety and fuel breeding tests. Its operational history has led to it being heavily influential in subsequent sodium reactor design. The Clinch River Breeder Reactor [5] was proposed as a joint effort between the Atomic Energy Commission and industry, but funding was canceled before construction began. The Fast Flux Test Facility (FFTF) [6] was run successfully as a fuel and material testing reactor for approximately 13 years. Finally, significant ground work toward licensing of a modern sodium power reactor was laid by the PRISM reactor project [7]. The severe accident management tool documented in this report would provide a critical component to the demonstration of the technical safety basis for licensing of power SFRs and the potential for reduced control room staffing for SFRs as compared to large LWRs.

1.2 Relationship with Collaborators

Figure 1 sketches the contributions of various actors in the overall Advanced Reactor Technologies (ART) project. Argonne National Laboratory has contributed through SAS4A support and expert elicitation for defining accidents to be considered. In return, Sandia has provided source code modifications to SAS4A to enable creation of discrete dynamic event trees (DDETs). A DDET differs from a traditional event tree in that events are given discrete timing. This allows DDETs to better capture complex behavior, as different timing of events may change the response of the system and thus branching of the event tree. Oak Ridge National Laboratory and Idaho National Laboratory have contributed probabilistic multi-physics insights that further inform accident definition. In addition, Idaho has launched an online portal that, when completed, will host databases to store the results of simulations and PRA documents to be used by all laboratories. Finally, the University of New Mexico has provided computer science support to translate SAS4A output data into populated Dynamic Bayesian Networks (DBNs).

1.3 Methodology

The theoretical framework for developing SMART procedures involves coupling dynamic PRA, system simulations codes, and Dynamic Bayesian Networks (DBNs) to provide fast-running diagnostic support [2, 1]. The methodology, as shown in Figure 2, takes outputs from an advanced PRA and aggregates them into a DBN to provide decision support. This coupled approach provides a process for extensive and comprehensive modeling of both the accident space and the plant response, in a fast-running framework. The research team develops and executes a full spectrum of runs using Discrete Dynamic Event Trees (DDETs) coupled to a simulation code (e.g., MELCOR, SAS4A); these runs are designed to cover the possible state-space of the accident. DBNs are used to synthesize and reduce this information into a framework that can be used for faster-than-real-time decision support. This information is used in combination with PRA information, e.g. system failure probabilities, to provide a detailed, probabilistic model of the accident sequence space. The resulting DBN model is an extensive knowledge base covering a wide spectrum of possible accidents, encoding the best-available knowledge from PRA to be used when needed.

The SMART procedures framework is implemented using a combination of tools. The DBN models are generated in GeNIe [8], which is a development environment for graphical decision-theoretic models developed by the University of Pittsburgh Decision Systems Laboratory. The structure of the model is built by the analyst. The model is built as a plate-based model containing nodes for accident states and reactor systems/components (outside of the temporal plate) and for

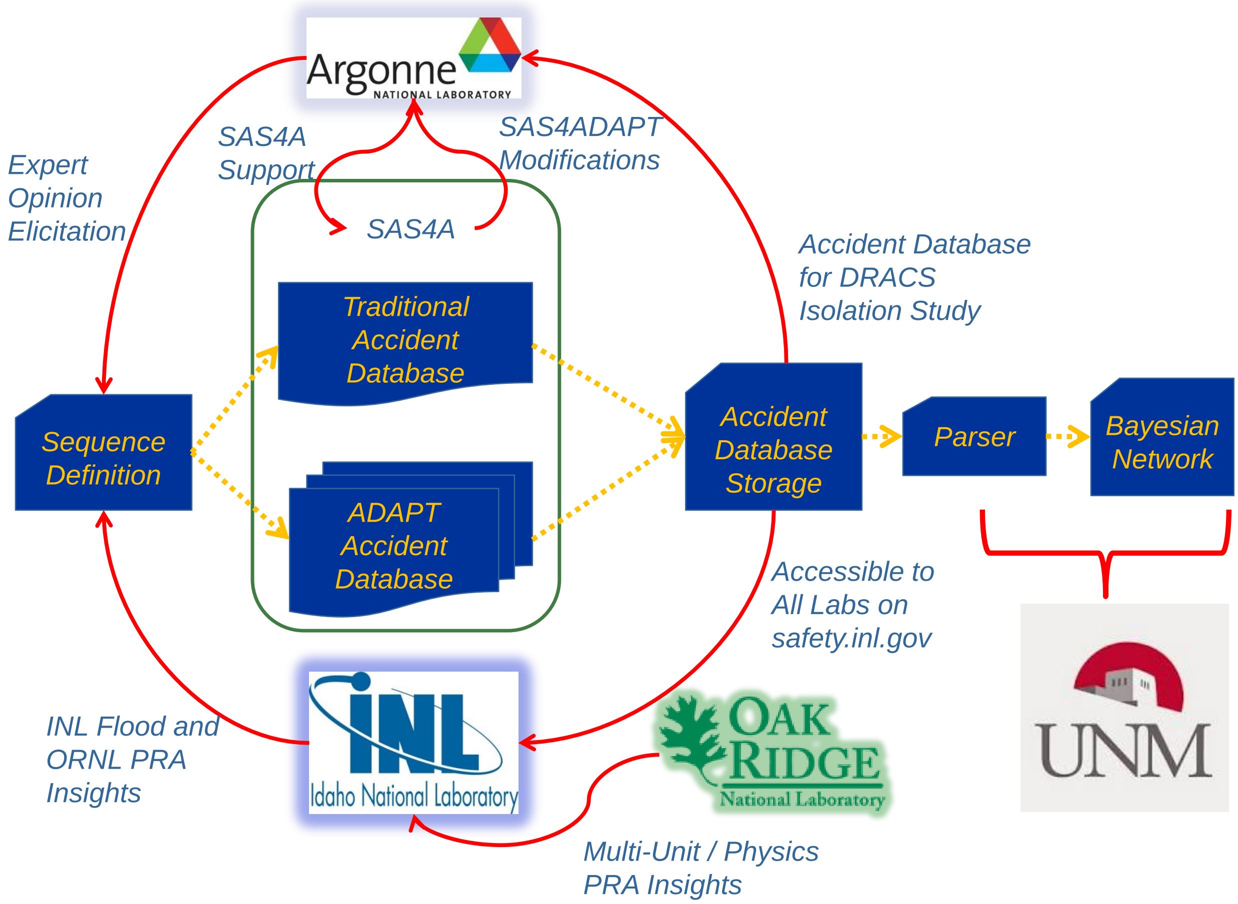

Figure 1: Relationship with Collaborators.

plant parameters (inside the temporal plate). Arcs are developed based on known causal relationship between the accident sequences, the reactor system components, and the plant parameters. The accident nodes are modeled as target nodes in GeNIe. The number and size of the time steps in the DBN are selected by the analyst.

The SAS4A [9] safety analysis code is used to simulate SFR accident characteristics. SAS4A is a system-level code that is capable of simulating SFR thermal-hydraulics in the core and reactor coolant system (RCS), neutronics, and liquid metal reactor accident phenomena. The reactor is nodalized in SAS4A as shown in Figure 3 [10]. It is noted in Reference [3] that SAS4A was incompatible with the DDET code ADAPT. While the event trees being constructed are simple enough that they are functionally identical to DDETs with similar branching rules, the eventual use of DDETs will allow more flexibility in order and timing of events. The incompatibility was resolved late in FY15 and test DDETs have been created and run with satisfactory results. It is expected that the next set of studies will incorporate the use of ADAPT.

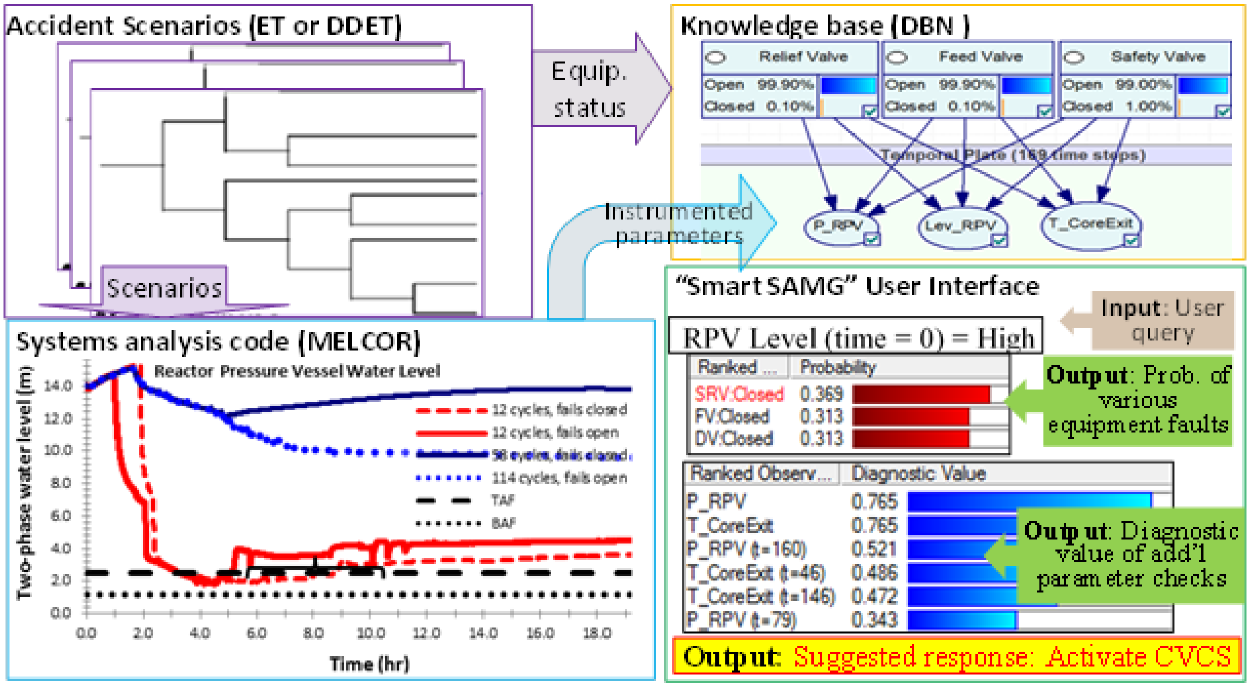

Figure 2: Conceptual process to develop risk-informed SMART SAMG procedures for nuclear power plant diagnostic support.

The data from the SAS4A simulations are processed through a data processing system, which is described in a draft SAND report titled "ALADDIN: The Automatic Loader of Accident Data for Dynamic Inferencing Networks" [11]. This process automates the quantification of the DBN model by filling the conditional probability tables in GeNIe with conditional probabilities based on external data. The system discretizes the SAS4A results and uses the discretized data to build a conditional probability tree that records the conditional probabilities of each observed variable given each combination of target states. The nodes are assigned a conditional probability at each time step. These probabilities are conditioned on the state of the reactor component/system and accident state variables.

1.4 Structure

This report is divided in the following way:

- Chapter 1: Outline of the motivation for the work and the DBN methodology
- Chapter 2: The particular reactor system and accidents to be evaluated
- Chapter 3: The accident simulation outputs in terms of the inputs that will be presented to the DBN for diagnoses
- Chapter 4: The structure of the particular DBN

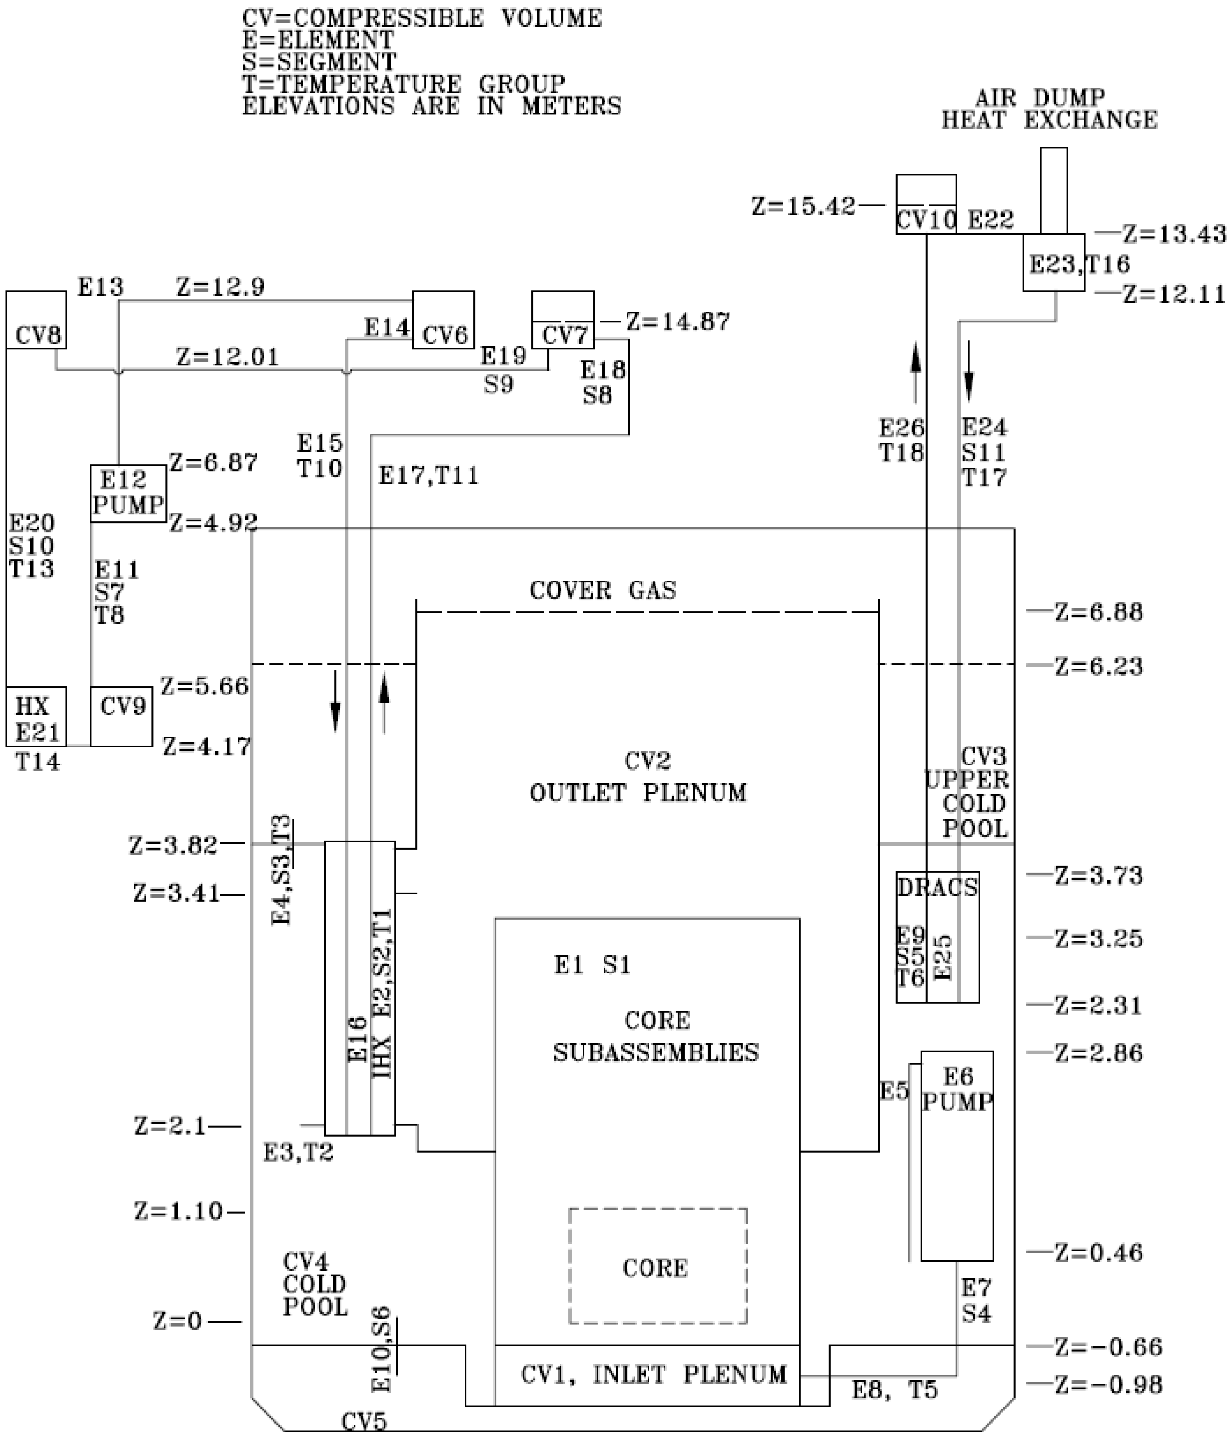

Figure 3: SFR SAS4A Nodalization

- Chapter 5: Insights drawn from the DBN as an example of the methodology
- Chapter 6: Planned and future work in the area

- Chapter 7: Summary of the types of insights that can be gained with the methodology and the overall applicability of the particular case

2 SFR Case Study Event Trees

This chapter describes the notional event trees of SAS4A simulations that will be used to train the DBN reasoning model. The SAS4A model used in this analysis is functionally equivalent to the model used in a previous report for this work [3]. The reactor model used in this study is a generic, small modular metallic fueled SFR with features adopted from the Advanced Liquid Metal Reactor (ALMR) design [12]. Some key design features which are relevant to modeling the selected accident sequences are:

- Four Electro-Magnetic Pumps (EMPs) $\rightarrow$ Provide forced circulation in the primary system to cool the reactor core. These pumps may experience thermal damage above 500 $^{\circ}$C operating temperature.
- Direct Reactor Auxiliary Cooling System (DRACS) $\rightarrow$ Passive decay heat removal system (DHRS) which uses natural circulation to transfer heat to air.
- Inherent reactivity shutdown $\rightarrow$ the reactor system exhibits strong negative reactivity feedback to increases in overall system temperature; thus the reactor can move from fission to decay heat levels without control rod insertion.

2.1 Purpose of Notional Event Trees

The notional event trees described herein are intended only to provide the probabilistic backbone for the DBN reasoning model. The generic small sodium reactor is a notional, pre-conceptual design and as such many assumptions were made in the assessment of failure probabilities. When required, probabilities were taken from the PRISM PSID [13], especially for initiating event annual occurrence probabilities. The potential for some events, such as transient overpowers (TOP), are highly design specific and thus the PRISM TOP magnitudes and initiating frequencies may not intuitively translate to a smaller reactor core with a 30 year refueling lifetime.

Due to these issues, the overall risks calculated with the simplified event trees provided in this report are not intended to accurately estimate the risk of a small sodium reactor. Instead, the following notional event trees are only intended to provide a rough estimate of the shape of the probability surface and associated reactor response in order to train and validate the DBN reasoning model.

The prototype model is intended to focus on four fundamental accidents: seismic-induced transient overpower, non-seismic transient overpower (TOP), loss of flow (LOF), and loss of operating heat removal (LOHR). Each accident sequence may be of the type ”protected” or ”unprotected”, depending on whether SCRAM occurs. Each accident sequence has the potential for long-term reduction in heat removal, such as degraded auxiliary cooling functionality or primary pump trip or damage.

For ease of visual representation and explanation, these major branches of the event tree have been broken down into sub-trees. Between the four sub-trees, 7,188 total branches were created, run in SAS4A, and the resulting data loaded into the DBN to be used for inferencing. See Chapter 3 for a discussion of the SAS4A results and how they were processed by ALADDIN [11] for inclusion into the DBN.

2.2 General Description of Branching Parameters

The notional event trees were created to provide the DBN with a wide array of accident sequences over which to reason and can be divided into two major categories. The first set of event trees depends on internal events. These event trees covers combinations of TOP, Loss of Flow (LOF), and Loss of Operating Heat Removal (LOHR) accident sequences. The second set is referred to as the earthquake event trees and is an expansion of the 0.5g earthquake event tree analyzed in [3].

The internal events event tree has the following branch conditions:

- TOP magnitudes
- LOF magnitudes
- BOP availability (LOHR)
- Scram state
- DRACS state
- Secondary (intermediate loop) pump power
- Inherent Reactivity Response
- Operator induced trip of the primary coolant pumps at 525K
- High temperature operating efficiency of the primary pumps

Reactor Protection System The RPS system is designed to bring the reactor to a safe configuration upon achieving the following set points:

- Flux (power) rises above 119% of its nominal value
- Core inlet temperature drops below 400 °C
- Core outlet temperature rises above 570 °C
- Normalized power to flow ratio drops below 1.13

When the RPS reaches any of the above conditions, it sends a signal to the reactor control rods to drop into the core. If the gravity release of the control rods fails, the RPS will actuate the control rod drive motors to push the control rods into the core. Once the normalize flux drops to 0.27 of its initial value, the RPS will signal the primary EMPs to trip. The EMP trip is treated as concurrent to SCRAM for this analysis. The RPS will also isolate operating power heat removal.

2.3 Establishment of Branching Probabilities

The reliability of the RPS was informed by the event tree in Figure 4, from the PRISM Preliminary Safety Information Document [13].

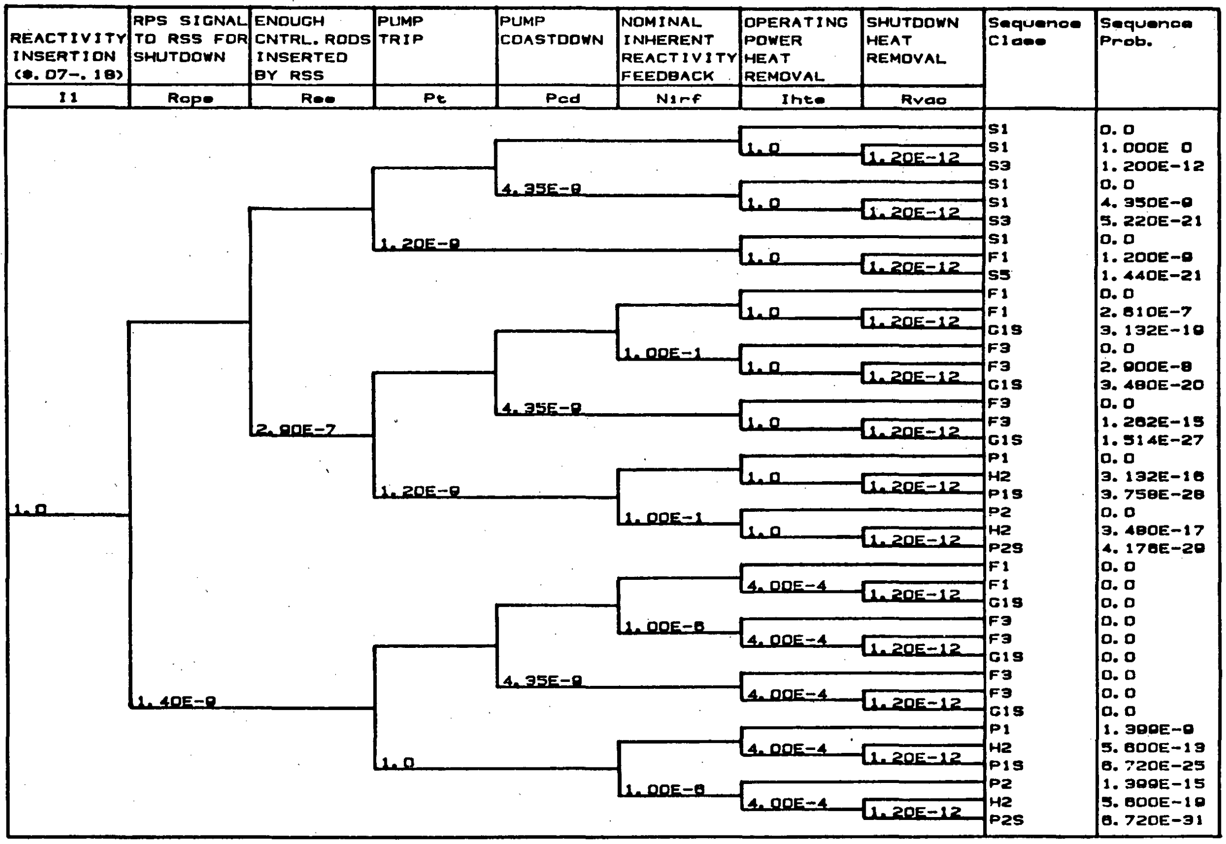

Figure 4: PRISM RPS Event Tree

These probabilities are summarized in Table 1, where SHR refers to Shutdown Heat Removal and OHR refers to Operating Power Heat Removal. An event with a bar over its name indicates that the event has not occurred. For example, the likelihood that SCRAM does not occur when RPS has actuated is 2.9×10^{-7}. The event tree analysis conducted in this study does not differentiate between failures that occur because RPS failed or due to other failure modes because the difference is not seen in the dynamic simulations. Thus, Table 2 shows the probabilities of the four different outcomes when the RPS system is commanded to actuate. It should be noted that, unless the initiating event is loss of heat removal, it is assumed that if scram occurs then operating heat removal will be unavailable but shutdown heat removal will be available. Furthermore, it is assumed that if scram does not occur then operating heat removal is available.

Table 1: Conditional Probability for RPS

Event	Probability
$\overline{RPS}$	1.4×10^{-9}
$\overline{SCRAM}$|RPS	2.9×10^{-7}
$\overline{EMP\_Trip}$|RPS	1.2×10^{-9}
$\overline{OHR\_Trip}$|RPS	4×10^{-4}
$\overline{SHR\_Trip}$	1.2×10^{-12}
$\overline{EMP\_Trip}$|$\overline{RPS}$	1.0

Frequencies for various initiating events were obtained from the PRISM PSID, and are presented in Table 3. These are used to build the prior probabilities in the DBN.

Table 2: Overall Probability for RPS

RPS Signal Reached	Scram	No Scram
Pump Trip	0.9999997	1.4×10^{-9}
No Pump Trip	1.4×10^{-9}	2.9×10^{-7}

Table 3: Initiating Event Probability for PRISM

Initiating Event	Probability per Year
Reactivity Insertion \$0.07 - \$0.18	1×10^{-4}
Reactivity Insertion \$0.18 - \$0.36	1×10^{-4}
Reactivity Insertion > \$0.36	1×10^{-6}
Earthquake 0.3g - 0.375g	1×10^{-4}
Earthquake 0.375g - 0.825g	2×10^{-5}
Earthquake > 0.825g	7×10^{-7}
Vessel Fracture	1×10^{-13}
Local Core Blockage	2×10^{-6}
Reactor Vessel Leakage	1×10^{-6}
Loss of One Primary Pump	1.6×10^{-1}
Loss of Substantial Primary Coolant Flow	5×10^{-2}
Loss of Operating Power Heat Removal	8×10^{-2}
Loss of Shutdown Heat Removal via BOP	8×10^{-3}
Loss of Shutdown Heat Removal via IHTS	1×10^{-2}
IHTS Pump Failure	5×10^{-2}
Station Blackout	3×10^{-5}

2.3.1 TOP Initiating Event Frequencies

From the PRISM PSID [13], four reactivity insertion values (0, 6, 30, 50 cents) were selected for analysis. For a point of reference, at Middle of Cycle (MOC) conditions the highest reactivity rod has 30 cents of reactivity worth and all of the control rods combined have $2.68 of reactivity worth. The PRISM design incorporated rod stops into the control rod drive design to limit the distance the control rods could be pulled from the core at a given time. PRISM allowed for 50 cents of reactivity to be available after each rod adjustment. Since the small SFR modeled in this study initially breeds reactivity for the first 15 years and then burns reactivity for the last 15 years, rod stops, if included, would be applied in a more complicated procedure than was proposed for PRISM. For example, reactivity changes more at the beginning and end of the 30 year reactor cycle, and so requires more operator actions to reposition rod stops. This increased activity leads to an increased probability of inadvertent reactivity insertion. This activity is at a minimum at MOC.

It should be noted that power adjustments, and thus opportunities for TOPs, would be much less frequent at MOC than at other portions of the refueling interval. The purpose of this project is to provide a probabilistic accident space to support DBN reasoning, not to design a reactivity control scheme for a notional small SFR, thus the initiating frequencies for the TOPs were simply taken from the ALMR PRA [12] with this note of the caution. The PRISM PRA TOP probabilities were nearly identical to the ALMR PRA TOP probabilities, with the lowest TOP bin ($0.11 - $0.35) assigned an annual probability 10^{-4} and the second lowest TOP bin ($0.35 - $1.75) also assigned an annual probability of 10^{-4}. In all, the ALMR frequencies seemed to provide a greater risk-tradeoff potential for this nominal study. The final TOP probabilities are given in Table 4.

Table 4: TOP Initiating Event Frequencies

TOP Magnitude (cents)	Probability per Year
0	$1 - (1x10^{-3}+1x10^{-4}+1x10^{-5})$
6	$1x10^{-3}$
30	$1x10^{-4}$
50	$1x10^{-5}$

2.3.2 LOF Initiating Event Frequencies

The modeled small SFR notionally has four Electromagnetic Pumps (EMPs) to provide forced circulation through the primary circuit. Two stages of partial LOF, in addition to the no LOF and complete LOF cases, were analyzed in the DBN. It was assumed for this analysis that the DBNs ability to infer a LOF accident could be achieved with only 0.75 (3 of 4 working EMPs), 0.5 (2 of 4 EMPs functional), and 0 (0 of 4 EMPs functional) flow branch conditions. The annual occurrence probabilities were derived from the PRISM PSID, with an annual failure rate for two pumps of 0.05. This probability is dominated by loss of an electrical bus which stops forced flow to two EMPs. Loss of all four EMPs is calculated as the loss of both electrical busses with a common cause failure factor β of 0.1. When the RPS functions as designed, a trip signal is sent to the pumps. Thus, the probability of pump trip is near 1 (i.e., 1 - $1.2\text{x}10^{-9}$) when the reactor protection system functions. The remaining probability $1.2\text{x}10^{-9}$ is divided by 2 for the 100% and 50% flow branches. In reality, these branches likely represent operator actions, not stochastic failures. If the reactor protections system fails, the probability of independent pump trip concurrent to the initiating event is simply Pr(IE)*Pr(X% Flow) under the rare event approximation. The final LOF probabilities are given in Table 5.

Table 5: LOF Initiating Event Frequencies

Flow %	Probability per Year | β
100	1
75	0.16
50	0.05
0	0.0051

2.3.3 LOHR Initiating Event Frequencies

The loss of heat sink initiating event frequency is divided into situations where operating power heat removal is needed (i.e., events where scram did not occurs) and events where shutdown heat removal is needed (i.e., events were scram did occur). The PRISM PSID assigns an annual probability of 0.08 to the failure of the operating power heat removal and 0.182 to loss of shutdown heat removal through a combination of balance of plant (BOP) and intermediate heat transport system (IHTS) failure events. It should be noted that DRACS is available in varying degrees in all accident conditions.

2.4 Seismic-Induced Accident Branches

For earthquake cases, the first branching point was the magnitude of the earthquake. The options for magnitude were 0.3g and 0.5g. Next, one of twenty-five sets of reactivity coefficients (RCs) was applied. It was reasoned that the earthquake will induce some level of control rod chatter proportional to peak ground acceleration, and that the magnitude of reactor power changes will depend on reactivity coefficients. Figure 5 shows the specific example of 0.3g ground acceleration and the first set of reactivity coefficients. The next branching point allowed the RPS to scram the reactor, trip the pumps, both, or neither. There was also a possibility for half of the primary pumps to fail, with or without scram, and without a trip of the other pumps. Next, the operator may override a pump trip or lack of trip by restarting or tripping pumps, respectively. If the pumps are running when they reach the thermal damage point, there is a branch for multiplying torque by (1.0, 0.5, or 0.0). For example, if half of the pumps failed initially and thermal damage fraction was 0.5, the reactor is left with 0.25 of nominal pumping torque. Finally, the secondary pump speed is branched out to 1.0, 0.5, or 1.5 times its nominal speed. The branches are identical for

0.5g acceleration and the other sets of reactivity coefficients. The numbers on the right side of the figure are the number of branches for each sub-tree. For example, if SCRAM occurs and the pump trips, there are 6 ways the scenario can branch: the operator may or may not restart the pumps, and there are three secondary pump speeds.

The total number of branches is in the upper right corner of Figure 5: 63 in this case. Each earthquake acceleration sub-tree produced 1,575 branches, as there were 63 branches per set of reactivity coefficents and 25 sets of reactivity coefficients. A total of 1,800 branches were run, however, because branches that were considered to have extremely low probabilities were still run. For example, there were no branches where a SCRAM occurred and the cold pool reached a high enough temperature for thermal pump damage. This accident scenario was not represented in the event tree, but was still modeled in SAS4A and used in the DBN. This decision was made for all accident types and magnitudes, and the number of branches given in the text represents the total including extremely low probability branches.

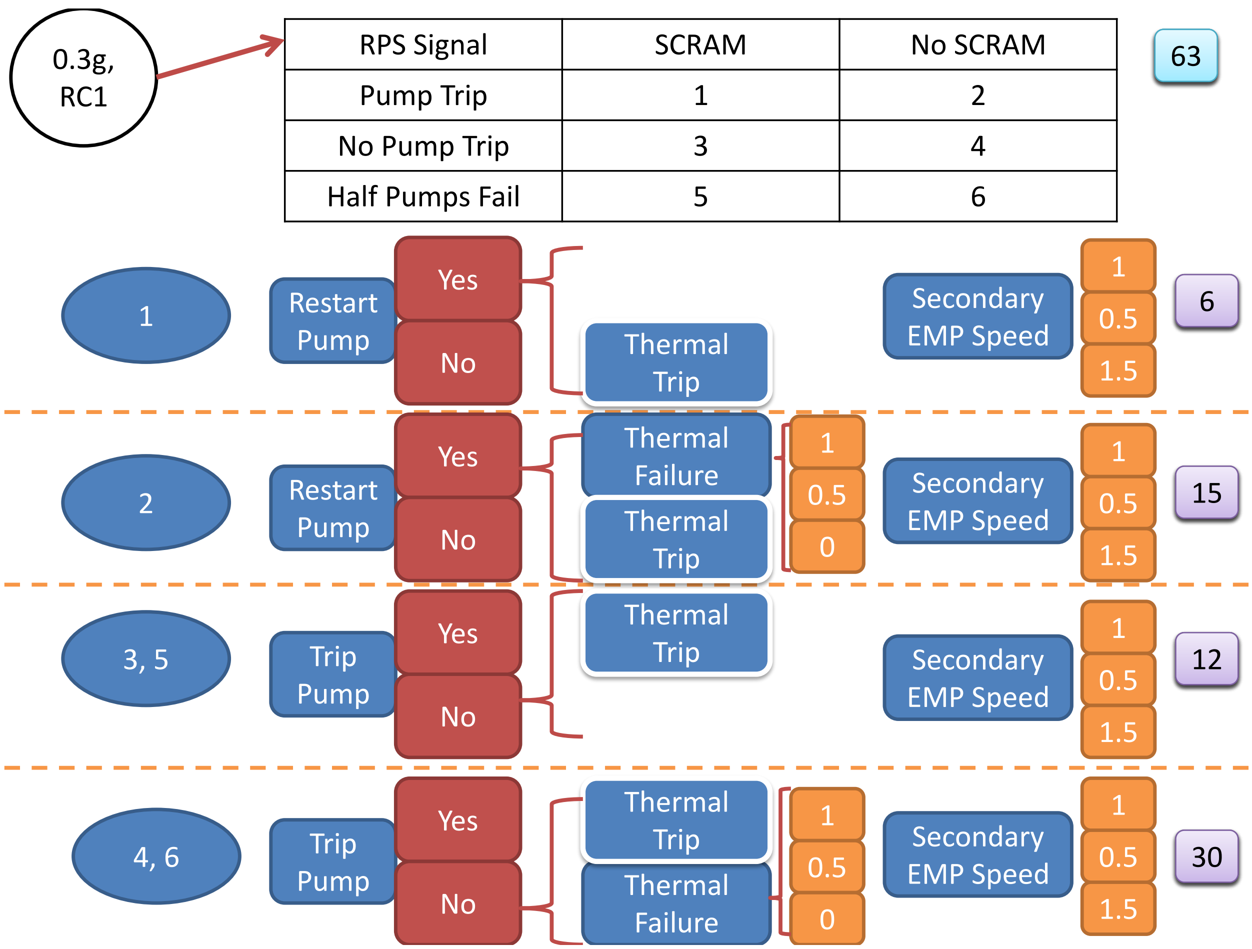

RPS Signal	SCRAM	No SCRAM
Pump Trip	1	2
No Pump Trip	3	4
Half Pumps Fail	5	6

Figure 5: Earthquake (0.3g, RC Set 1): Accident Branching

2.5 TOP Accident Branches

The first branching point of the TOP accident tree is the presence or absence of heat transfer from the core to the balance of plant (BOP). The next major branching is the degree of reactivity insertion. Figure 6 shows the branches for TOP with BOP and 6 cents insertion. The next branching point is the scram and pump trip combination. If the pumps do trip, there is a branching for the operator to restart them. If the pumps do not trip, there is a branching for the operator to trip them. Some cases do not branch further. For example, if the RPS scrams the reactor and does not trip the pumps, and the operator does not trip the pumps, the reactor is assumed to be in a safe configuration. Similarly, the pump thermal failure fraction is not branched if the pumps are tripped. In this case, however, reactor damage may still occur. The branching is identical for 30 and 50 cent insertions. Altogether, TOP with BOP produces 1,512 cases.

The next major branch is TOP without BOP. This leads to substantially fewer branches, as the secondary pump speed becomes irrelevant. All other branching points are identical to TOP with BOP. Figure 7 shows the branching for TOP without BOP in the specific case of a 6 cent insertion. There are 396 cases for TOP without BOP.

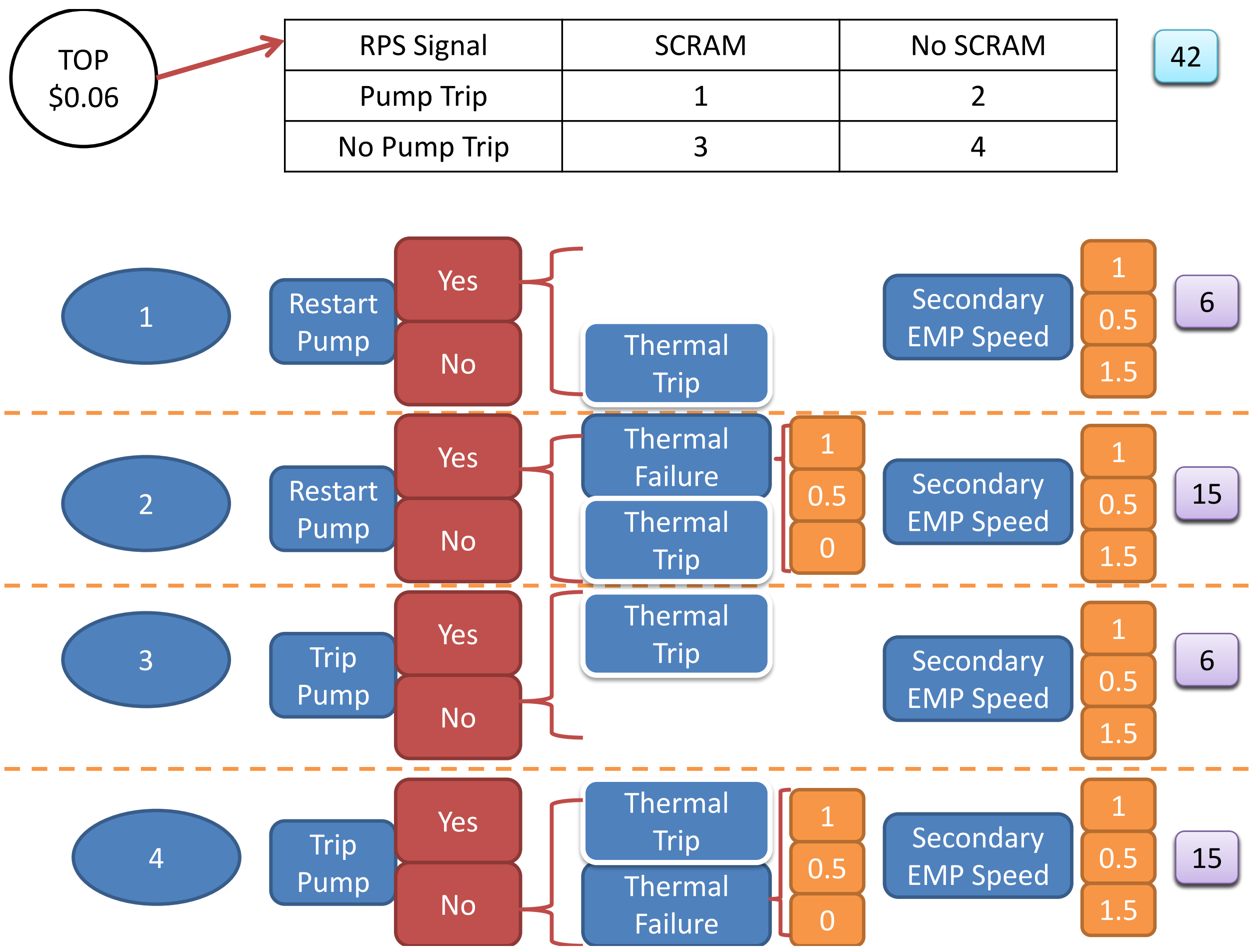

Figure 6: TOP with BOP Present: Accident Branching

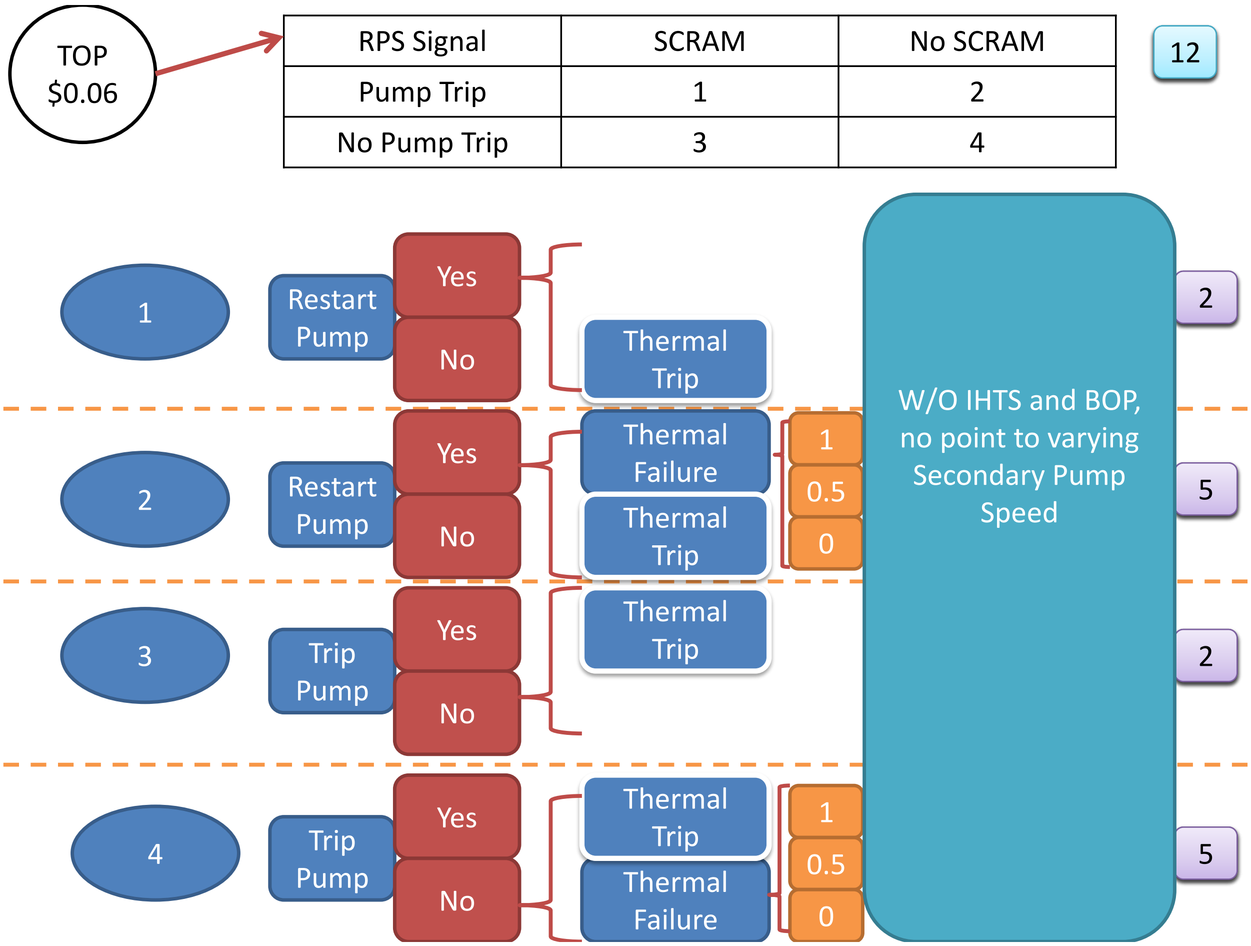

Figure 7: TOP with BOP Absent: Accident Branching

2.6 LOF Accident Branches

The first branching point of the LOF accident tree is the presence or absence of heat transfer from the core to the BOP. The next major division is the number of pumps failed. Options are one, two, or four pumps failed out of four total. In cases with one or two pumps failed, branching continues much as in the TOP cases. See Figure 8 for an example. There are again options for the RPS to scram, trip the remaining pumps, both, or neither. There is also the option for operator intervention to either trip or restart the remaining pumps. In the case of thermal failure in a LOF accident, the failure fraction (1, 0.5, or 0) is applied to the total pumping capability of the remaining pumps. This is to simulate all pumps experiencing some thermal damage, not necessarily any single pump failing completely. The branching is identical for 2 failed pumps. If all pumps have failed (Figure 9, the branching is simplified because there is no ability to trip or restart primary pumps. In this case, the only further branching point is secondary pump speed. There are 837 cases for LOF with BOP.

If the BOP is not available in a LOF accident, branching is again considerably simplified. Figure 10 shows the branching for LOF without BOP in the specific case of one pump failure. Scram, trip, operator action, and thermal failure are all still in play. Secondary pump speed is now irrelevant. The branching is identical for two failed pumps. Figure 11 shows the branching when all pumps have failed and BOP is unavailable. The only options are to scram or not scram. There are 279 cases for LOF without BOP.

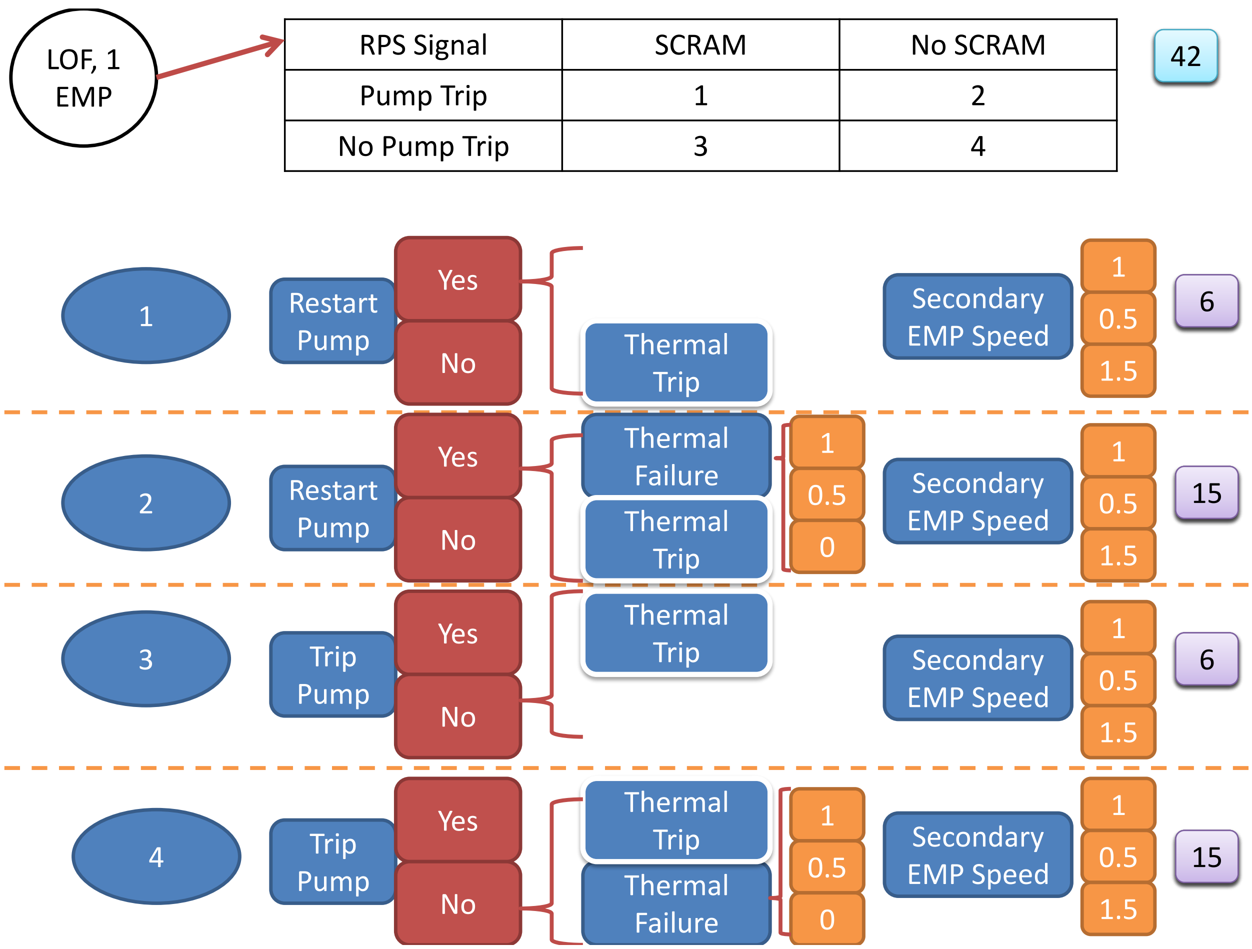

Figure 8: LOF (1 Pump Failed) with BOP Present: Accident Branching

2.7 LOHR Accident Branches

The first major branching condition for the loss of operating heat removal (LOHR) accident is the presence or absence of heat transfer from the core to the balance of plant (BOP). Figure 12 shows the branches for LOHR with BOP. First, the RPS may scram, trip pumps, both, or neither. Then, as in other accident scenarios, the pumps may be tripped or restarted by the operator, and may experience thermal damage. There are 396 cases for LOHR with BOP.

If the BOP is lost, secondary pump speed is irrelevant. Scram, trip, operator action, and thermal failure are still in play. Figure 13 shows the possible branches for LOHR with no BOP. This reduces the number of cases to 168.

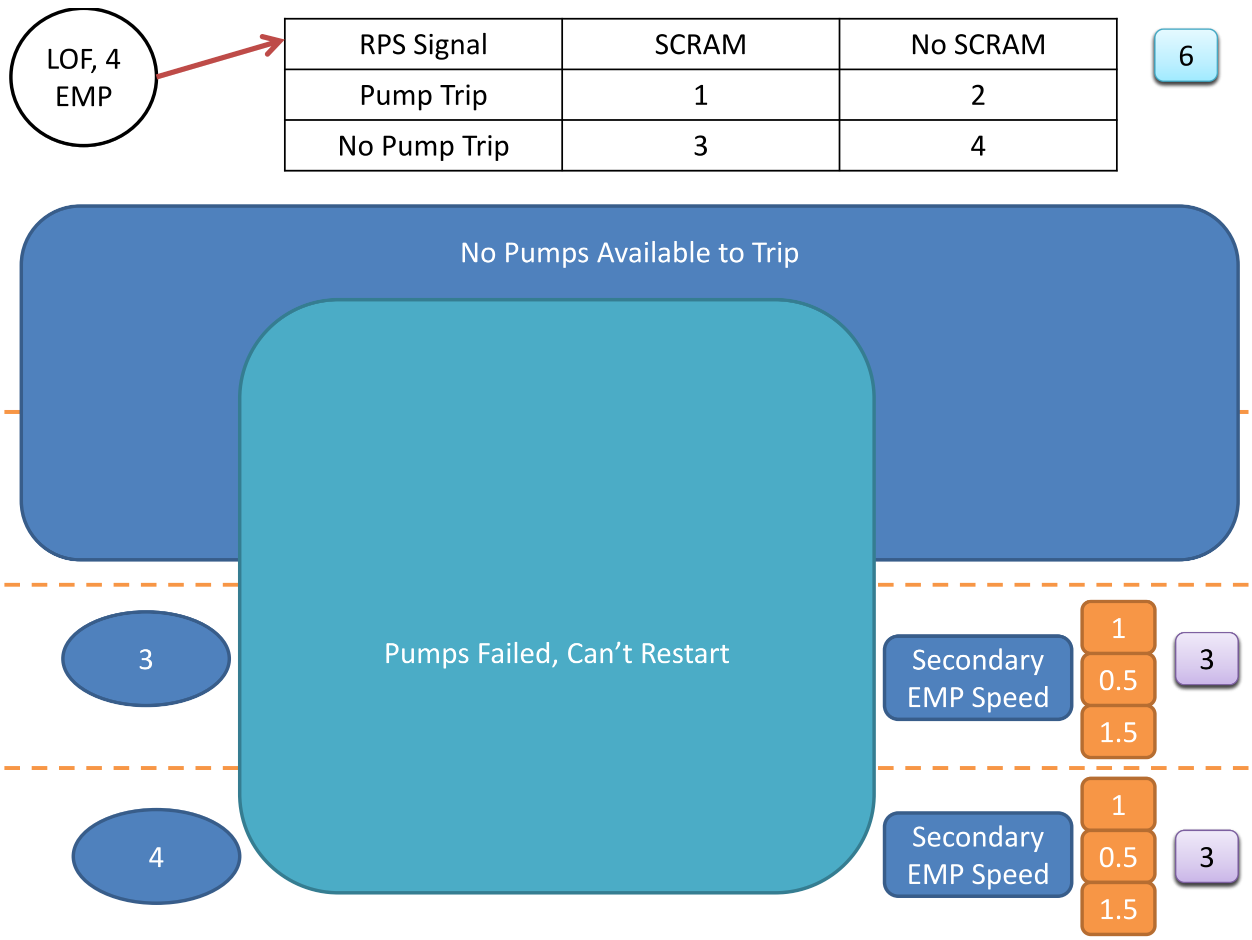

Figure 9: LOF (4 Pumps Failed) with BOP Present: Accident Branching

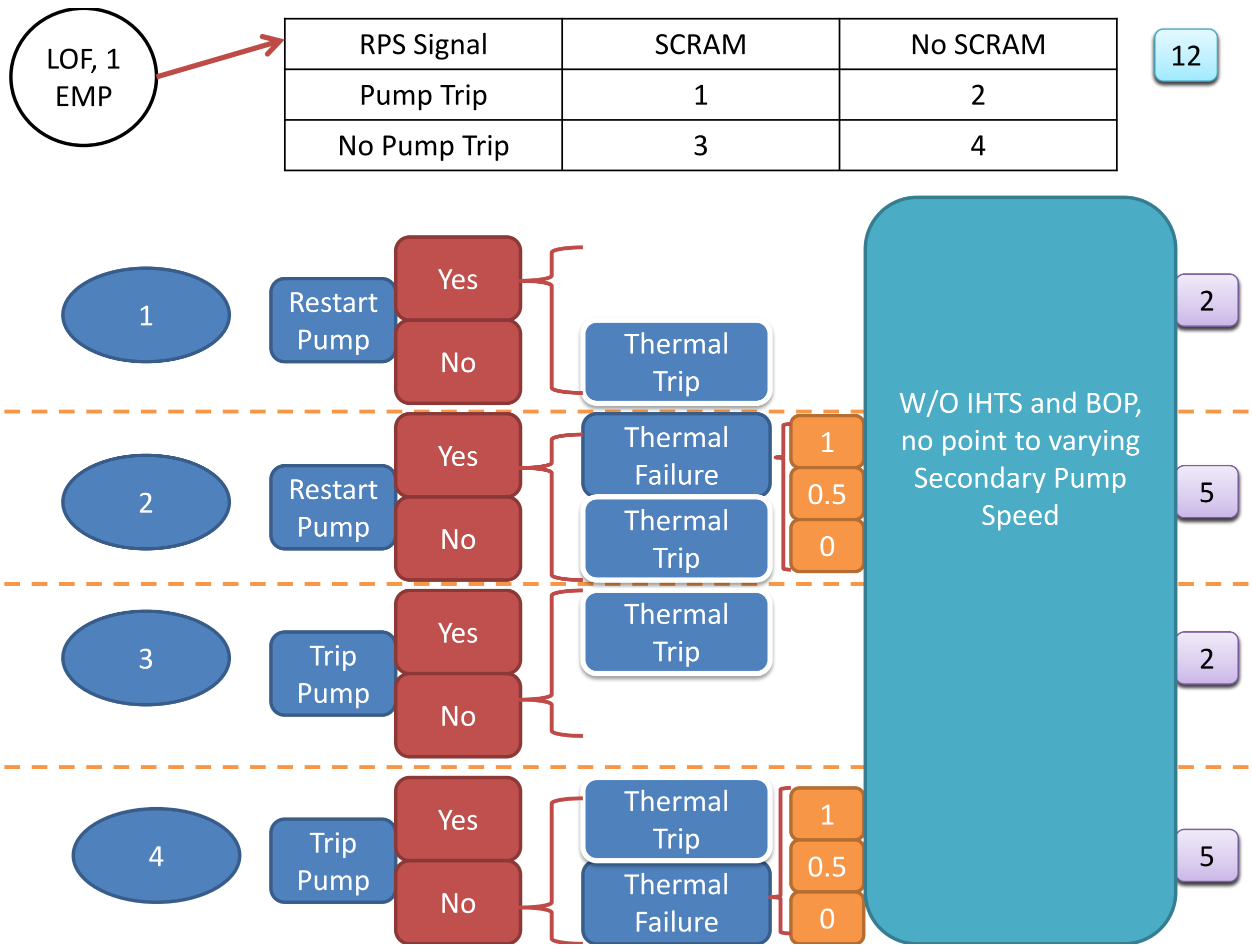

Figure 10: LOF (1 Pump Failed) with BOP Absent: Accident Branching

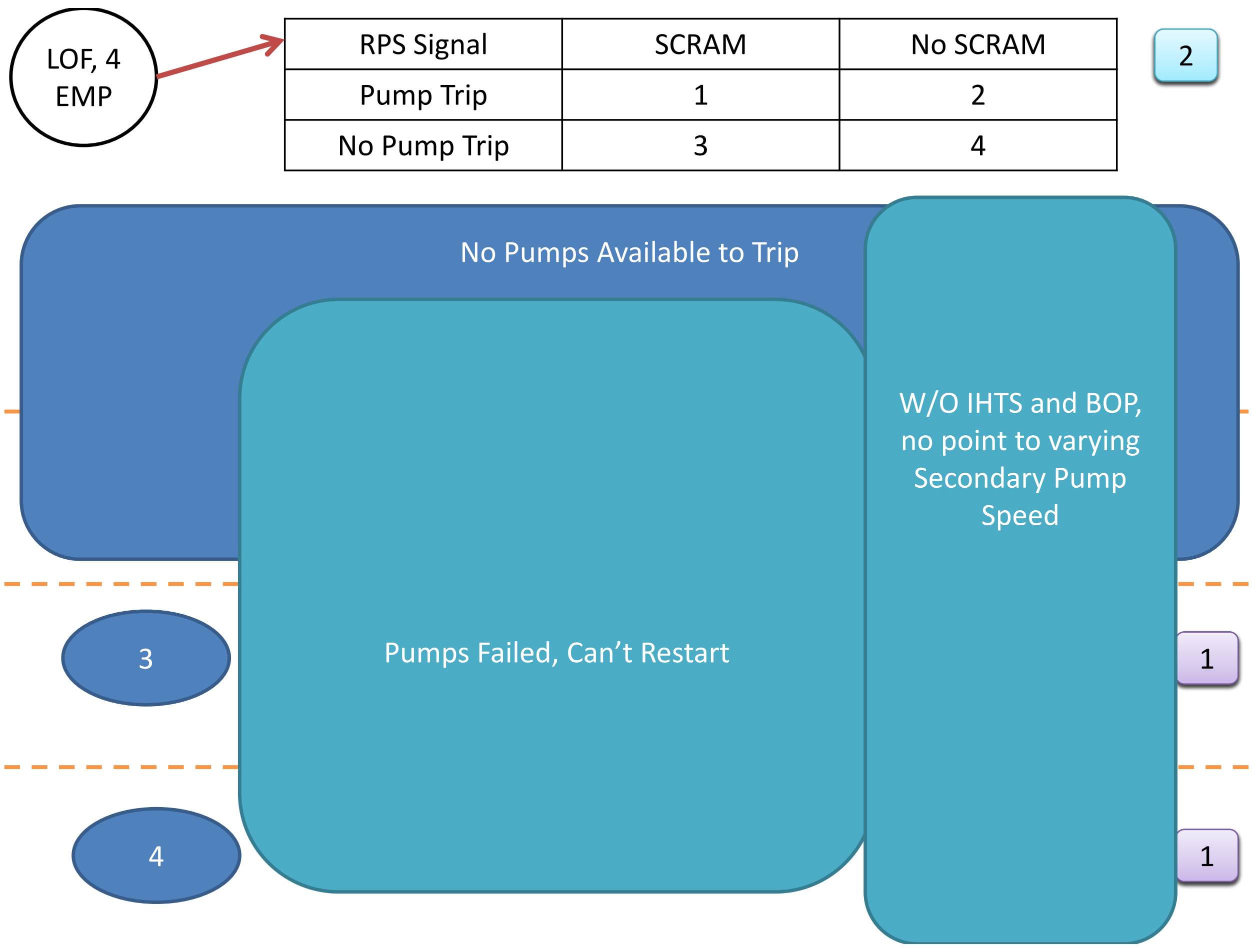

RPS Signal	SCRAM	No SCRAM
Pump Trip	1	2
No Pump Trip	3	4

Figure 11: LOF (4 Pumps Failed) with BOP Absent: Accident Branching

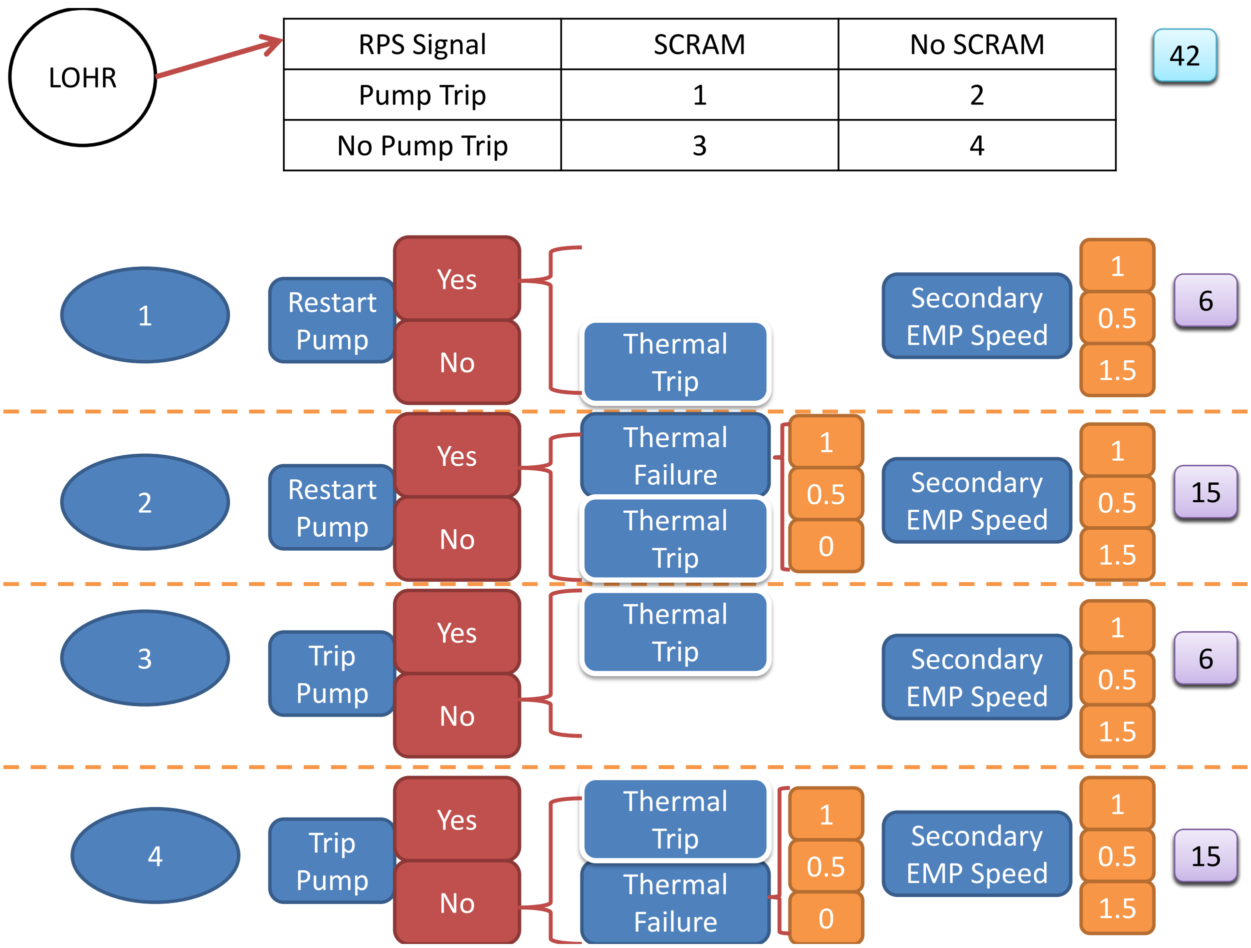

Figure 12: LOHR with BOP Present: Accident Branching

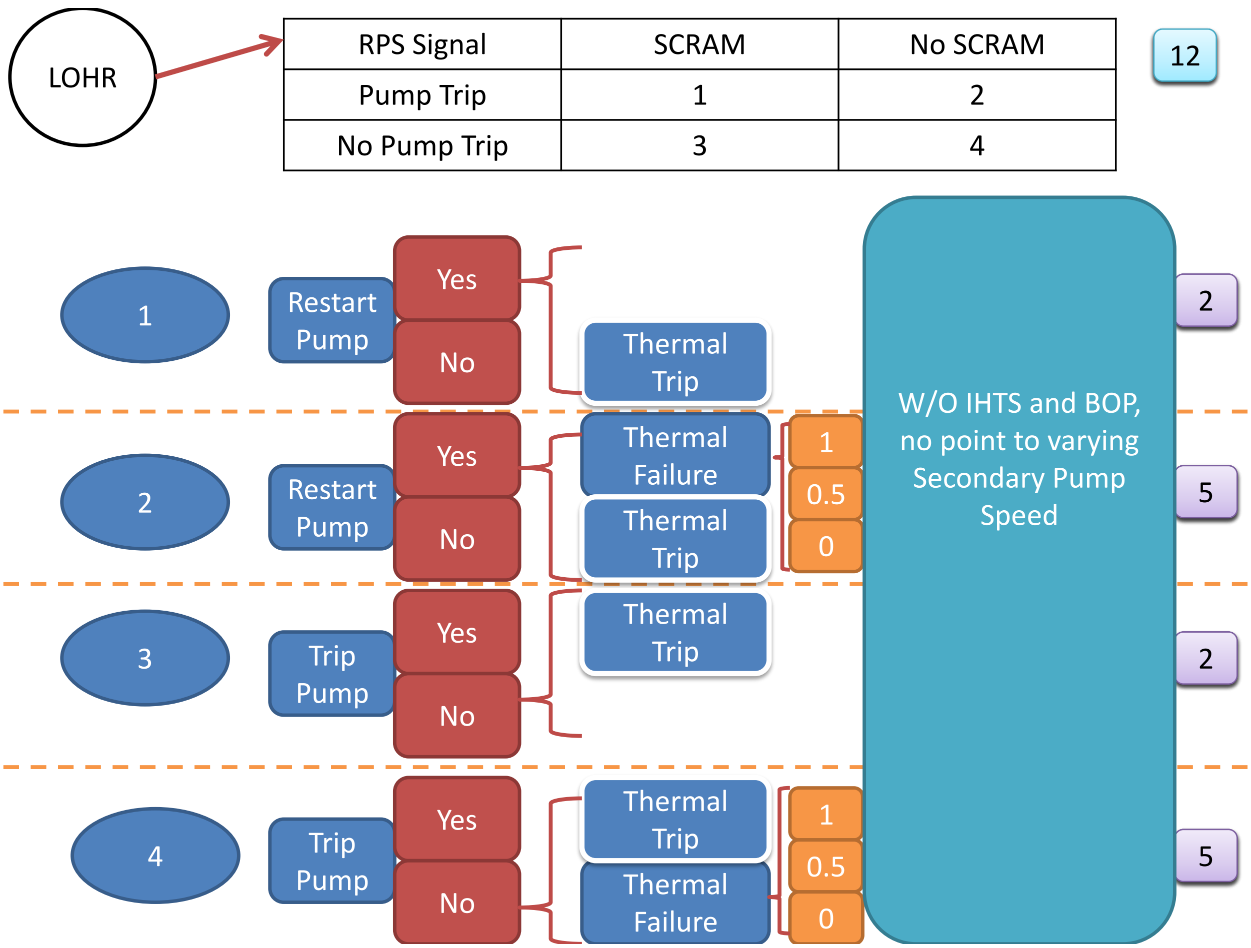

Figure 13: LOHR with BOP Absent: Accident Branching

3 Processing of SAS4A Results

This section describes the SAS4A data as it is presented to the DBN. The processing involved and results of important choices are briefly discussed. The data output from the SAS4A runs create a three dimensional matrix of 7,188 by 2,588 by 14. These dimensions are defined by the 7,188 SAS4A scenarios described in Chapter 2, the 2,588 time steps returned by SAS4A as it ran from start to finish through the branches, and the 14 SAS4A variables which were selected a priori to represent either those monitored plant parameters observable in the main control room or those unmonitored plant parameters which could be inferred by operators. This three dimensional matrix of SAS4A output was processed by the data translator ALADDIN [11]. The data was sampled across the time step dimension such that the results of 96 of the 2,588 time steps were retained for informing the DBN. The data retained (96 values for each of the 14 plant parameters for each of 7,188 branches) are further reduced through by creating bins of time-dependent data for informing the DBN model. This binning process is discussed below.

3.1 SAS4A Data Binning Process

For this study, an equal-width binning scheme was used with three bins. That is, the range of each variable is calculated and each bin covers 1/3 of that range. Bin 0 covers the lower 1/3 of the range, Bin 1 the middle 1/3, and Bin 2 the top 1/3 of the range of any particular variable. This method is susceptible to the effects of a single outlier, as evidenced by reactor power (Figure 24). In that case, a single time step of a single scenario reaches 10^{54} times nominal. This causes each bin to be very wide, and nearly all data to be placed in Bin 0. It is unlikely that the particular time step would be used by the DBN, and so the most likely condition is that every DBN node for Power will be in Bin 0. This renders the variable useless for making inferences, as it never changes. This also affected the binning of peak fuel temperature (Figure 18), cold pool level (Figure 23), and power to flow ratio (Figure 26). For future analyses, different binning strategies will be explored. For example, each bin may be set to a desired percentage of the range.

3.2 Binned SAS4A Data

Figures 14 through 27 illustrate the results of the ALADDIN binning for each of the 14 plant parameters selected for the DBN. Figures 14, 15, 16, and 17 show binning of the reactivity contributions of various types of feedback, all of which had coefficients that were varied at the start of the scenarios. Note that these figures represent the bin populations of parameters with each scenario given equal weight, and do not reflect the overall likelihood of parameters being in specific bins on a ”typical day”. For example, the figures do not account for the reliability of the SCRAM system. It can be seen that the major changes in reactivity contribution came from axial and radial expan-

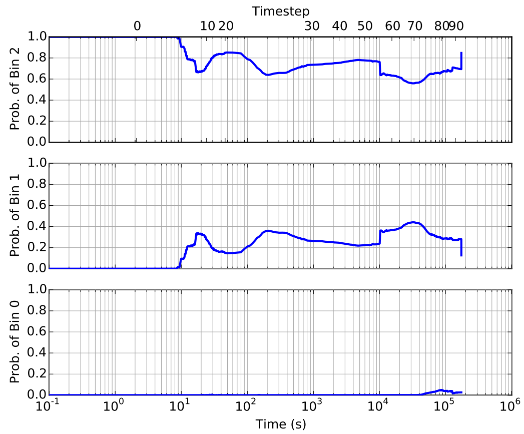

Figure 14: Axial Expansion Feedback Bins

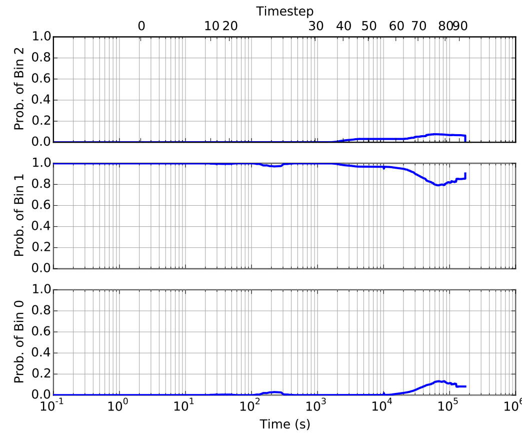

Figure 15: Radial Expansion Feedback Bins

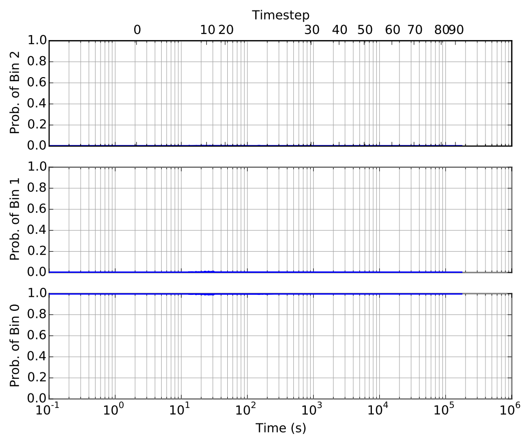

Figure 16: Coolant Feedback Bins

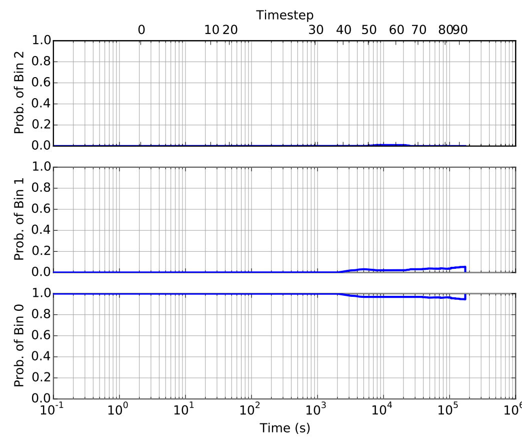

Figure 17: Doppler Feedback Bins

sion of the fuel. Figures 18 and 19 show binning of temperatures of fuel and coolant within the hottest core channel. Coolant temperature can be seen to rise in most scenarios with time. Figure 20 shows binning of cladding thickness, which can drop to Bin 0 suddenly as fuel pins rupture late in the accident.

It is seen in Figure 21 that coolant flow tends to decrease with time, which is to be expected as pumps are tripped or fail. Coolant flow is a physical state monitored by plant instrumentation that the operator can use as evidence to hypothesize the accident condition and make a decision to act. Cold pool temperature (Figure 22), cold pool level (Figure 23), power (Figure 24), and cover gas pressure (Figure 27) are also monitored, and the power-to-flow ratio (Figure 26) is calculated from power and flow. Total reactivity (Figure 25) is derived from changes in power, and so can be considered a known variable to the operator.

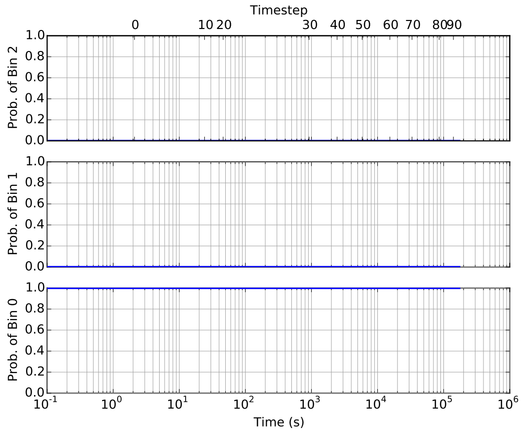

Figure 18: Peak Fuel Temperature Bins

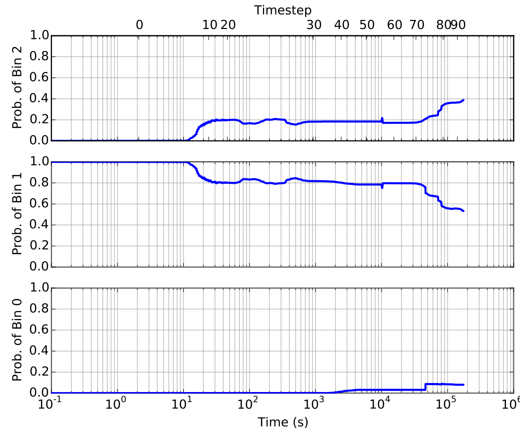

Figure 19: Channel Coolant Temperature Bins

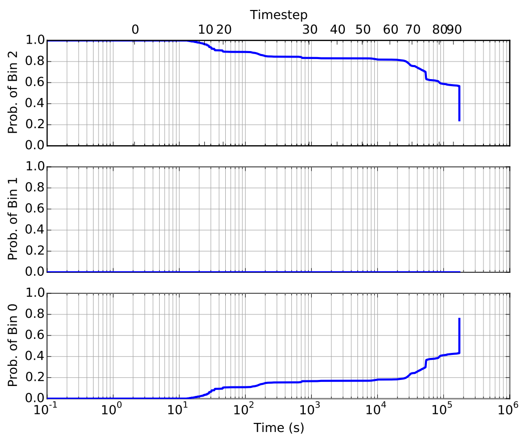

Figure 20: Cladding Thickness Bins

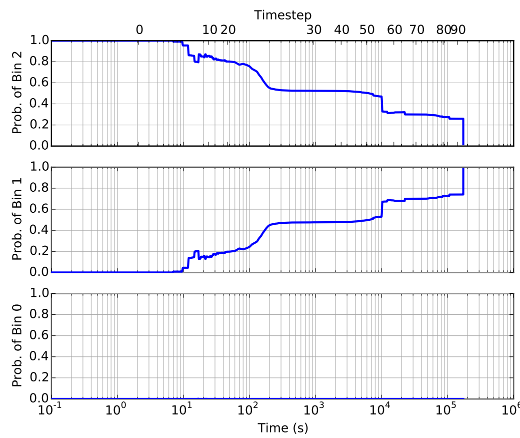

Figure 21: Channel Coolant Flow Bins (Instrumented)

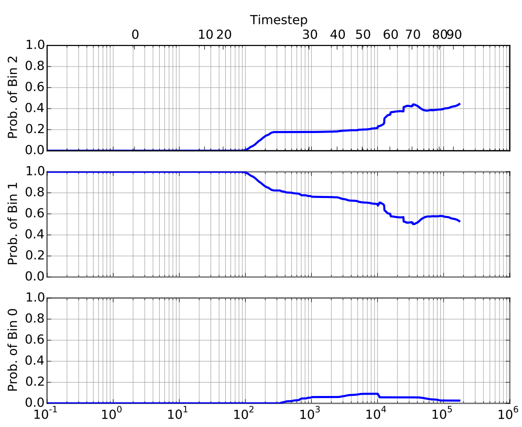

Figure 22: Cold Pool Temperature Bins (Instrumented)

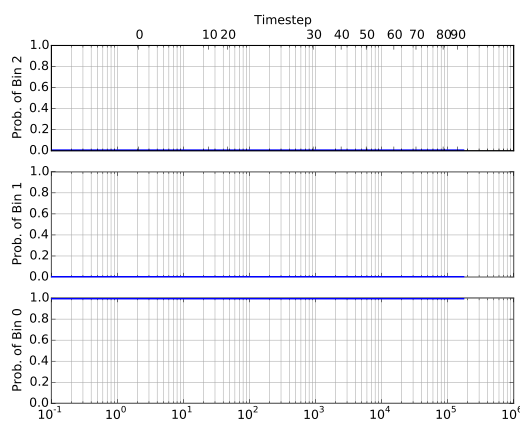

Figure 23: Cold Pool Level Bins (Instrumented)

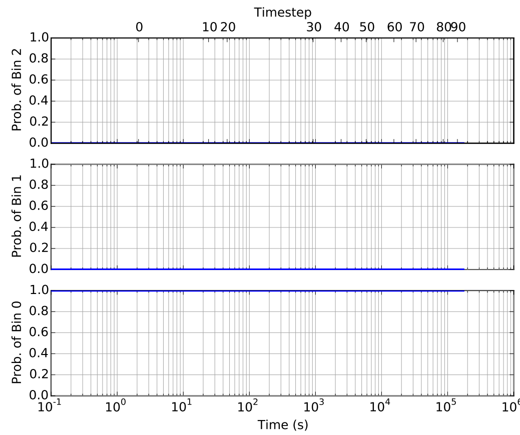

Figure 24: Reactor Power Bins (Instrumented)

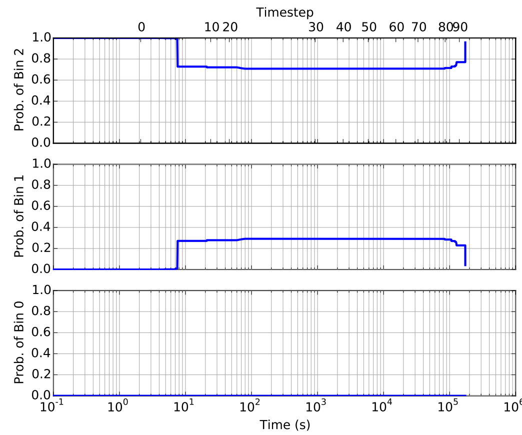

Figure 25: Total Reactivity Bins (Instrumented)

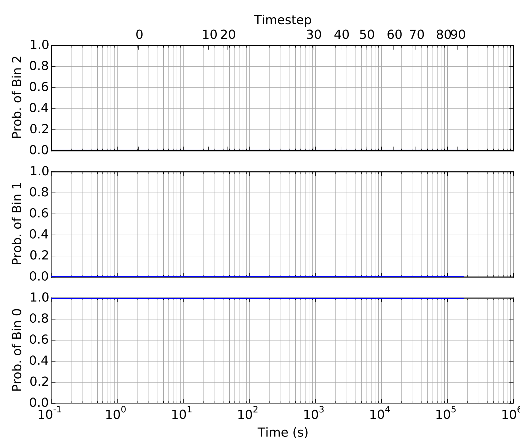

Figure 26: Power-to-Flow Ratio Bins (Instrumented)

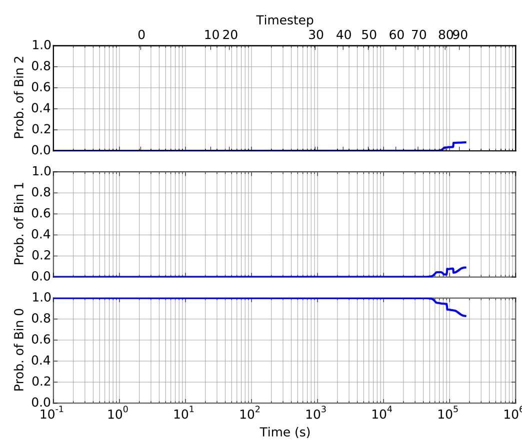

Figure 27: Cover Gas Pressure Bins (Instrumented)

3.3 DBN Target Node Values

The DBN target node states (differential pump pressure, balance of plant status, DRACS status and SCRAM status) were calculated from the SAS4A input and output data. Differential pressure states were determined by the pump pressure at the end of each simulation. The pressure would indicate whether 0, 1, 2, or all 4 pumps were still operating. Figure 28 shows the percentage of scenarios that end in each of the states under consideration. It can be seen that the great majority of cases end with all pumps lost or tripped. It should be noted that while the LOF accidents branched on 0, 1, 2, or all 4 pumps failed, the differential pressure target node represents 0, 1, 2, or all 4 pumps functioning at the end of the accident scenario. It was chosen to define the target node more finely at higher levels of pump outage (by failure or tripping), as fuel damage is generally more likely with more pumps failed. In this way, the relationship between higher levels of pump outage and fuel damage may be seen more clearly.

Next, BOP status is shown in Figure 29. The percentage associated with the Decay state indicates that there is zero or nearly zero heat transfer to the balance of plant in nearly 2/3 of cases. This is a parameter that is set in a priori in the event tree for each scenario. In a DDET the BOP status would be dynamic, dropping when a scram occurs or upon some other criteria for disconnection or damage of the IHTS.

The status of DRACS was also set a priori in this set of scenarios (Figure 30), but can be made dynamic. The operator may attempt to increase heat transfer by dumping water into the DRACS system. If the temperature is high enough, this may actually cause the heat transfer to be reduced as the system may be damaged by thermal shock.

Finally, the scram status (Figure 31) was determined from both the SAS4A input and output according to the following rules:

- If TOP and scram disabled → Control rods withdrawn
- If TOP and scram enabled → Control rods fully inserted
- If no TOP and scram disabled → Control rods nominal
- If no TOP and scram enabled → ...
  - If final reactivity `<-$10` → Control rods fully inserted
  - Otherwise → Control rods nominal

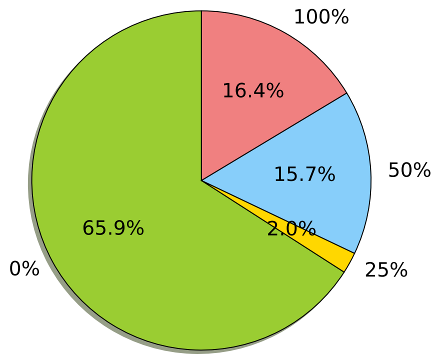

Figure 28: State Percentages for Differential Pressure

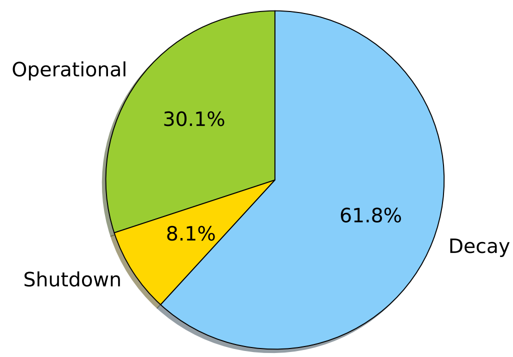

Figure 29: State Percentages for Balance of Plant Status

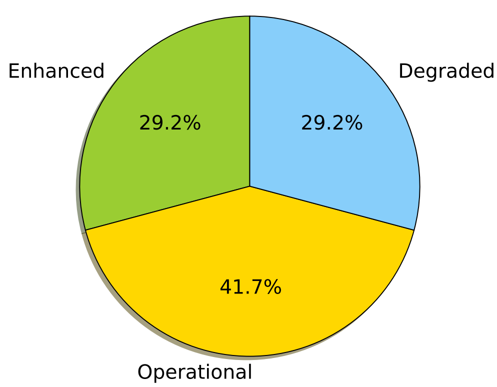

Figure 30: State Percentages for Direct Auxiliary Reactor Heating System Status

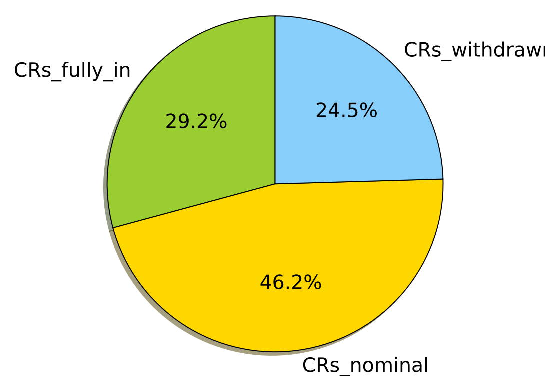

Figure 31: State Percentages for Scram Status

4 DBN Model

4.1 Model Structure

This chapter describes the structure of the DBN under study as well as the prior assumptions included within it, while Chapter 5 contains attempts to diagnose accident and plant parameters from the DBN with data loaded. Figure 32 contains the structure of the DBN generated for the case study. This figure contains a plate-based DBN modeling the relationship between reactor systems and components (denoted by gray nodes), one unmonitored physical state (denoted by blue nodes), plant parameters (denoted by green nodes), and accident types (denoted by yellow nodes). In this version only two accident types are recognized: TOP and LOF. LOHR will be added as an accident type in future versions. As such, in Chapter 5 only TOP and LOF accident are diagnosed. The same model populated with data is shown as Figure 33 with EM Pump 1 set as "operational". Given only this evidence (piece of data) the model reasons that there is a 99% chance that a TOP is not occurring. This largely reflects the prior probabilities of various plant states, and diverges as evidence is set that contradicts those probabilities.

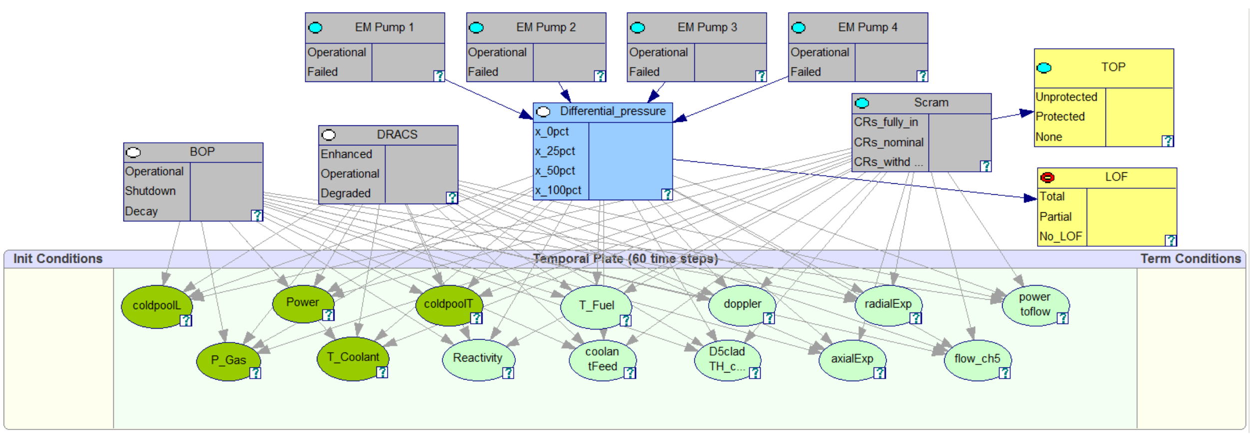

Figure 32: DBN Model Structure

The model in Figure 32 contains two accident states, TOP (transient overpower) and LOF (loss of flow). The model contains seven reactor systems and components (the DRACS, the BOP, four EMPs, and the scram system) and one unmonitored physical state (Differential Pressure, which is produced by operational EMPs). The model also contains fourteen plant parameters which may provide insight into the status of the reactor systems and the accident states. These plant parameters and their ranges from the SAS4A data are shown in Table 6. It should be noted that these values have not been edited. For example: pure sodium freezes at 370K, but SAS4A does not model the freezing phenomena of sodium, only of previously-molten fuel. SAS4A will therefore allow sodium to flow at temperatures significantly lower than 370K, which may not be valid.

The model structure shows that the four EMPs directly influence the amount of differential pressure; we assume each pump has the same influence on the differential pressure. The time-varying reactor parameters are duplicated once for each time step, which were distributed as follows:

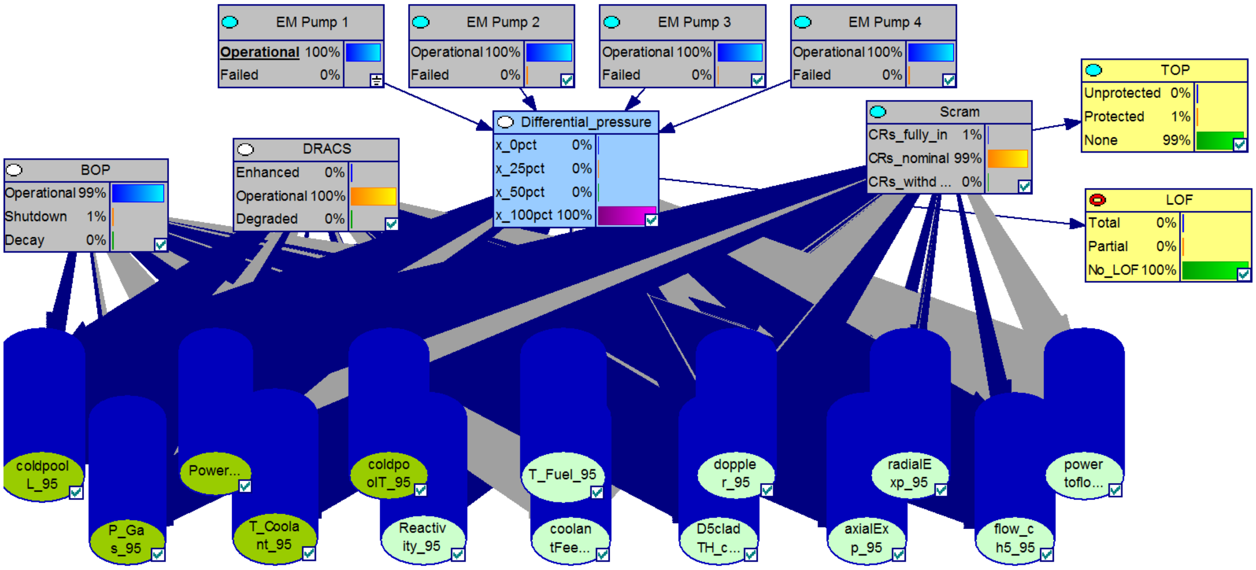

Figure 33: Unrolled DBN Structure

- 24 time steps for the first 0.1hr
- 24 time steps from 0.1hr to 1hr
- 24 time steps from 1hr to 10hr
- 24 time steps after 10hr

DRACS availability, scram status, BOP status, and pump differential pressure each influence the state of all fourteen (including five monitored) plant parameters at each time step in the model. In this example model, the status of the DRACS, BOP, scram system, and EMPs remain constant throughout the duration of the accident (i.e., they are modeled in the DBN to either have failed or remain operational a priori, they do not fail during the accident). The scram system influences the state of the TOP node; this represents the definitional relationship wherein an unprotected TOP is defined by failure of the scram system. Similarly the differential pressure influences LOF via a direct definitional relationship since a LOF accident is defined by loss of differential pressure.

4.2 DBN Node Probability Tables

4.2.1 Accident Type Nodes

The conditional probability table for the LOF node in the DBN is shown in Table 7. Since the LOF accident is defined by a loss of differential pressure, the conditional probability table for LOF is deterministic; meaning the state of LOF is completely determined by the state of differential pressure. If there is 0% of the required differential pressure, a Total LOF has occurred. If there is 25% or 50% of the required differential pressure, a Partial LOF has occurred. If there is approximately 100% of the required differential pressure, there is no LOF.

Table 6: Plant Parameter Nodes and Bins from ALADDIN, 96 Time Steps.

Node ID	Meaning	Minimum	Bin 0 Max	Bin 2 Min	Maximum
axialExp	Axial Expansion Reactivity Feedback ($)	-0.62	-0.34	-0.07	0.22
coldpoolT	Cold Pool Temperature (K)	48.99	438.89	828.79	1230.5
coldpoolL	Cold Pool Level (m)	5.68	3.3×10^4	6.6×10^4	1×10^5
coolantFeed	Coolant Reactivity Feedback ($)	-0.28	0.54	1.37	2.23
D5clad_TH	Cladding Thickness in Channel 5 (Fraction of nominal)	0.0	0.34	0.66	1.0
doppler	Doppler Feedback Reactivity ($)	-0.29	0.21	0.71	1.23
flow_ch5	Channel 5 Flow Rate (kg/s)	-96.47	-40.92	14.63	71.86
P_Gas	Cover Gas Pressure (Pa)	1.4×10^4	7.4×10^5	1.5×10^6	2.2×10^6
Power	Reactor Power Level (Fraction of nominal)	0.00	5.0×10^{53}	9.9×10^{53}	1.5×10^{54}
powertoflow	Power to Flow Ratio (Fraction of nominal)	-2.3×10^5	2.3×10^{55}	4.6×10^{55}	7.0×10^{55}
radialExp	Radial Expansion Reactivity Feedback ($)	-0.62	-0.22	0.18	0.60
Reactivity	Net Reactivity ($)	-24.60	-16.04	-7.47	1.35
T_Coolant	Peak Coolant Temperature in Channel 5 (K)	64.66	571.56	1078.45	1600.7
T_Fuel	Peak Fuel Temperature in Channel 5 (K)	65.75	5708.17	1.1×10^4	1.7×10^4

The conditional probability table for the TOP node is shown in Table 8. Since there is minimal available data on the reliability of SFR systems, the probability of transient overpower was assigned directly by the analysis team.

Table 7: Conditional Probabilities for LOF, given Differential Pressure

Diff. Pres.	x_0pct	x_25pct	x_50pct	x_100pct
Total	1	0	0	0
Partial	0	1	1	0
No_LOF	0	0	0	1

Table 8: Conditional Probabilities for TOP, given Scram State

Scram	CRs_fully_in	CRs_nominal	CRs_withdrawn
Unprotected	0	9.59×10^{-14}	2.9×10^{-7}
Protected	1	3.31×10^{-7}	0.9999997
None	0	0.9999997	0

4.2.2 Reactor Systems and Physical State Nodes

The marginal probability tables for the reactor systems (DRACS, EMPs, BOP, and the scram system) are shown in Table 9. Since there is no available data on the reliability of SFR systems, these values were directly assigned by the analysis team. The team will update these values if additional SFR reliability data becomes available.

The conditional probabilities for differential pressure are derived directly from the causal relationships between flow from the EMPs and differential pressure. The conditional probabilities for differential pressure are shown in Tables 10 and 11. The probabilities have been assigned based on expert judgment about the likely state of differential pressure given the status of the pumps. With all four pumps working, the differential pressure is expected to be 1.0. With one of four pumps in the failed state, the differential pressure is likely to be around 0.5. With three pumps failed, the differential pressure is likely to be at 0.25 of what is necessary. If all four pumps are failed, the DP is highly likely to be 0% of the necessary flow. These most likely states are thus assigned high probabilities. To accommodate the possibility that un-modeled factors could impact the relationship between EMPs and differential pressure, smaller probabilities have been assigned to other states that are possible.

Table 9: Marginal Probabilities for DRACS, the four EM Pumps, Scram, and BOP.

	State	Probability
DRACS	Enhanced	1.19×10^{-12}
	Available	0.999999999998
	Unavailable	3.97×10^{-13}
EM Pumps	Operational	0.9996
	Failed	4.38×10^{-4}
Scram	CRs_fully_in	0.0150
	CRs_nominal	0.985
	CRs_withdrawn	3.04×10^{-6}
BOP	Operational	0.985
	Shutdown	0.0150
	Decay	7.95×10^{-12}

Table 10: Conditional Probabilities for the Differential Pressure Node, given EM Pump State.

EMP 1	Operational							
EMP 2	Operational				Failed			
EMP 3	Operational		Failed		Operational		Failed	
EMP 4	Op.	Fail	Op.	Fail	Op.	Fail	Op.	Fail
x_0pct	0	0	0	0	0	0	0	0.05
x_25pct	0	0	0	0.05	0	0.05	0.05	0.9
x_50pct	0.0001	0.99	0.99	0.25	0.99	0.25	0.25	0.05
x_100pct	0.9999	0.01	0.01	0.7	0.01	0.7	0.7	0

Table 11: Conditional Probabilities for the Differential Pressure Node, given EM Pump State (Continued).

EMP 1	Failed							
EMP 2	Operational				Failed			
EMP 3	Operational		Failed		Operational		Failed	
EMP 4	Op.	Fail	Op.	Fail	Op.	Fail	Op.	Fail
x_0pct	0	0	0	0.05	0	0.05	0.05	0.9999
x_25pct	0	0.05	0.05	0.9	0.05	0.9	0.9	0.0001
x_50pct	0.99	0.25	0.25	0.05	0.25	0.05	0.05	0
x_100pct	0.01	0.7	0.7	0	0.7	0	0	0

5 Completed DBN Model Analysis

The methodology, as described in Chapters 2 through 4, results in the development of a DBN which has been informed from the results of 7,188 SAS4A accident simulations. This DBN (based on 96 time steps from the DDET) was exercised against simulated accident conditions (or evidence) to test its effectiveness for operators to infer which accident situation is likely taking place in the reactor (see Figure 34). The data points for each parameter at each time step are referred to in the DBN as pieces of evidence. This evidence may reinforce or contradict the prior probabilities. Initially, before any evidence is considered, the model reflects only the prior probabilities of various accidents and plant conditions. It is anticipated that in future research the DBN will be enhanced so that it will be dynamically tested by receiving real-time evidence from a SAS4A accident simulation as that simulation progresses through time. At this point in time, accident evidence is statically imposed upon the DBN by setting the values of specific target variables in the DBN (e.g., SCRAM status, BOP status, DRACS status) to simulate hypothetical outputs of a SAS4A simulation. Some simple inferences are made to show potential uses of the DBN.

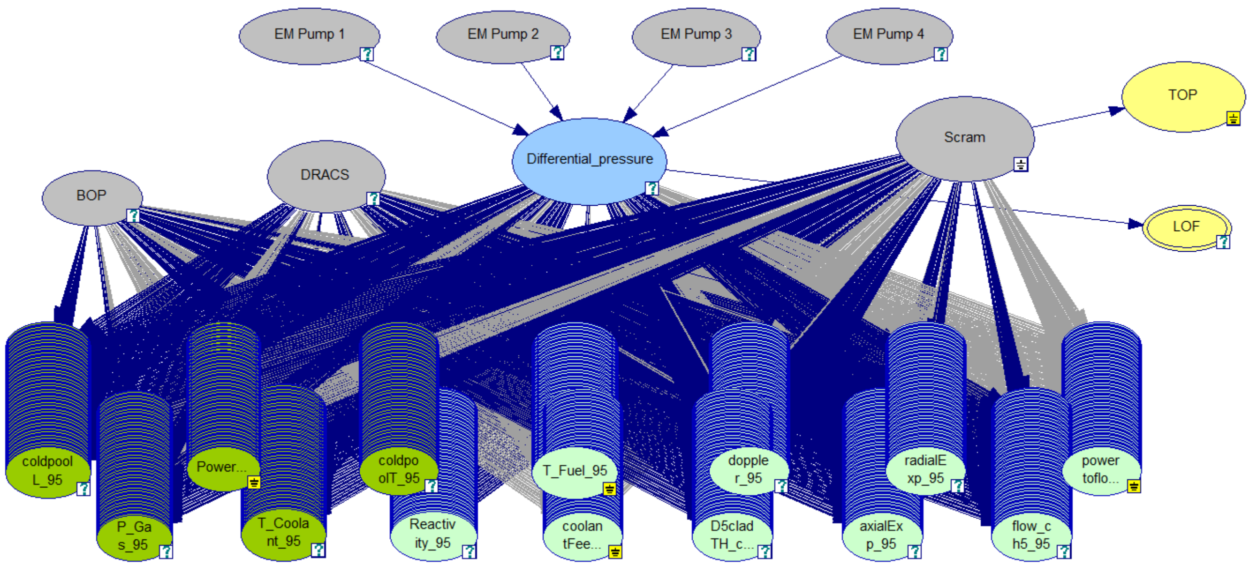

Figure 34: Populated DBN (96 Time Steps)

5.1 DBN Inference Directly from Target Node

First, a target node that is directly related to an accident will be exercised. The Scram target evidence is set to control rods fully in, which can happen in any accident scenario where power or temperature exceed certain set points and the RPS system inserts the control rods. The most common scram is for overpower.

When this evidence of control rod configuration is added, it has the effect (Figure 35) of setting all Power nodes to Bin 0, even before scram would have occurred. A consequence of n-ary binning with equal width bins is that a single outlier may greatly influence bin assignment. In this case, during one time step of one SAS4A scenario the power reached 10^{54} times nominal. This is considered to be an effect of a particular combination of reactivity coefficients and insertion which is highly unlikely and the numerical methods used in SAS4A. Another binning method is to allow the user to choose the fraction of the range for each bin, which can reduce the influence of an outlier if it is recognized. In this case every time step sampled for the DBN has power in Bin 0 (see Figure 24), and so setting evidence of Power in a higher bin would ”lock out” all accident scenarios at a probability of zero.

Peak fuel temperature (Figure 18) and power-to-flow ratio (Figure 26) are all set to Bin 0 as well, because of the outlier power spike. In all, 5 of 14 plant parameters are rendered useless for inferencing with this particular set of binned scenario data. Early reactivity nodes are set to Bin 2, which is expected for a case which would lead to a scram. The new probability of a protected TOP is 1. Other accident scenarios that end in scram have much smaller prior probabilities than protected TOP, and so it appears as the most likely accident.

5.2 DBN Inference Directly from Data

Next, an attempt will be made to deduce an accident scenario from a smaller set of evidence. In a LOF accident, it would be expected that coolant flow will be low and channel coolant temperatures will be high early in the accident. If channel coolant temperature evidence from approximately minutes 1 through 5 is set to Bin 2, and channel flow for the same time steps is set to Bin 1 (the lowest bin with a population), the model believes (Figure 36) that a protected TOP has happened with 1.000 probability. This is because protected TOP has a much higher prior probability and in most protected TOP scenarios the pumps also trip, reducing flow and increasing channel temperature.

In order to force the model to diagnose a loss of flow accident, evidence must be added to distinguish the plant state from that of a protected TOP. With a loss of flow, the cold pool temperature will likely stay low longer into the accident. This is because decay heat will not transfer as quickly from the core channels, and auxiliary heat removal systems will continue to cool the pool directly. With two later nodes of the cold pool temperatures set to Bin 0 the model is now certain (Figure 37) that a protected TOP and total LOF are both happening.

At this point the most useful piece of additional evidence would be channel coolant temperature later into the accident. With two later channel coolant temperature nodes set to Bin 2, there is no change in the diagnosis (Figure 38). Because the two accident conditions are not mutually exclusive and in fact often occur in tandem, this is seen as acceptable inferencing from the DBN.

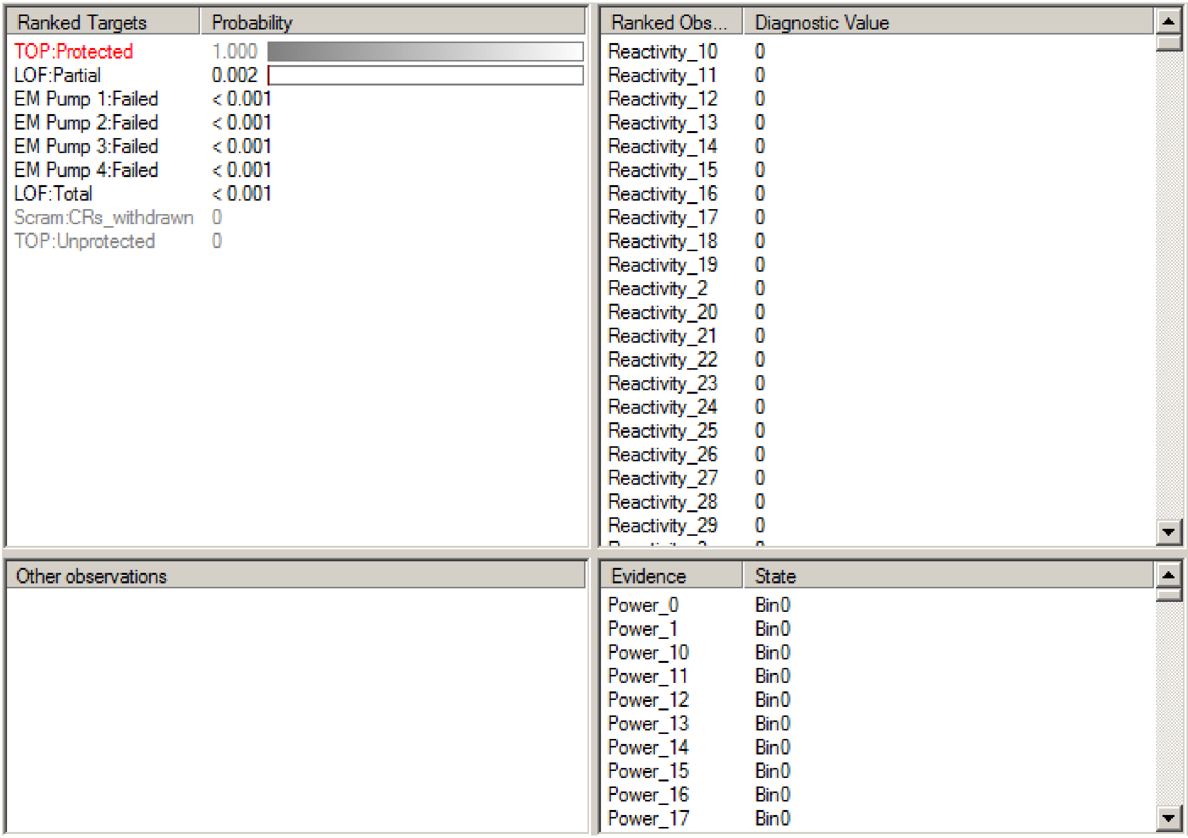

Figure 35: DBN Diagnosis, CRs Fully In

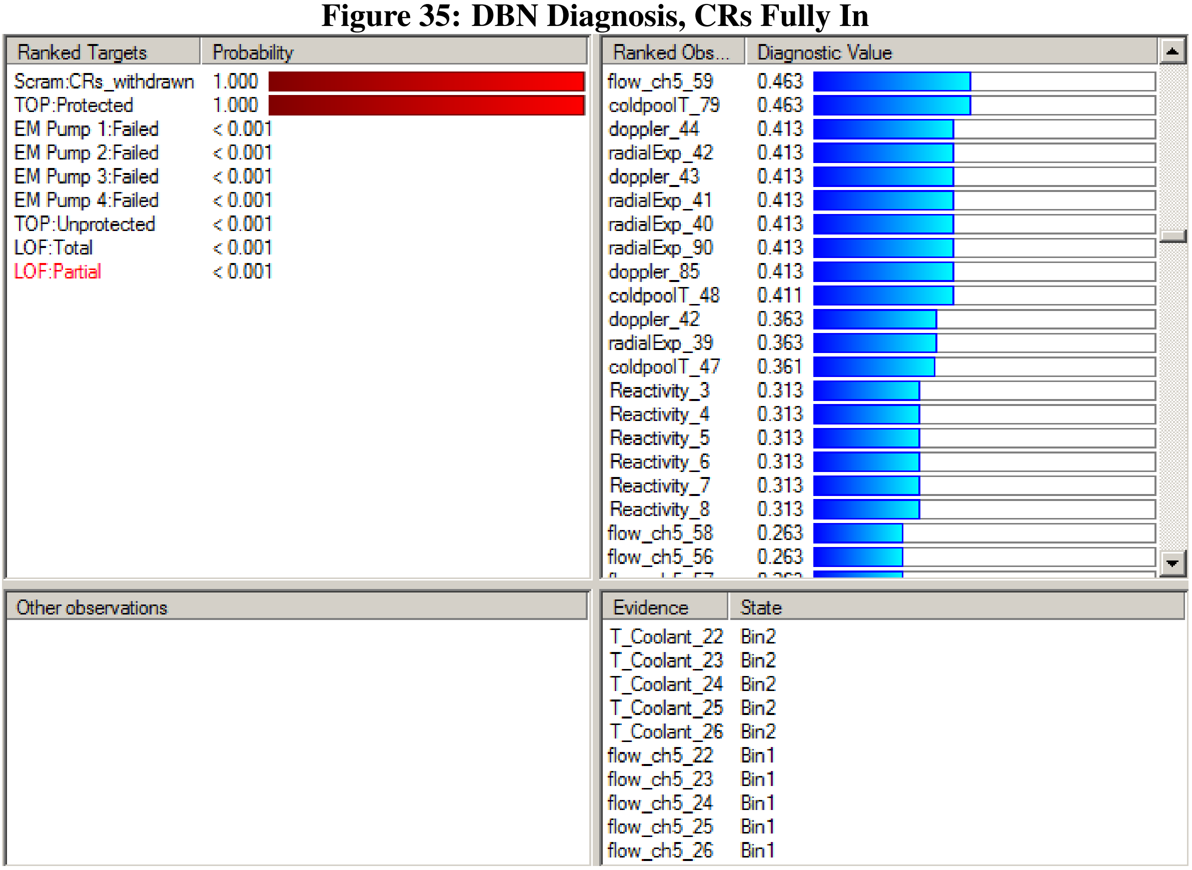

Figure 36: DBN Diagnosis, Low Flow and High Channel Coolant Temperature

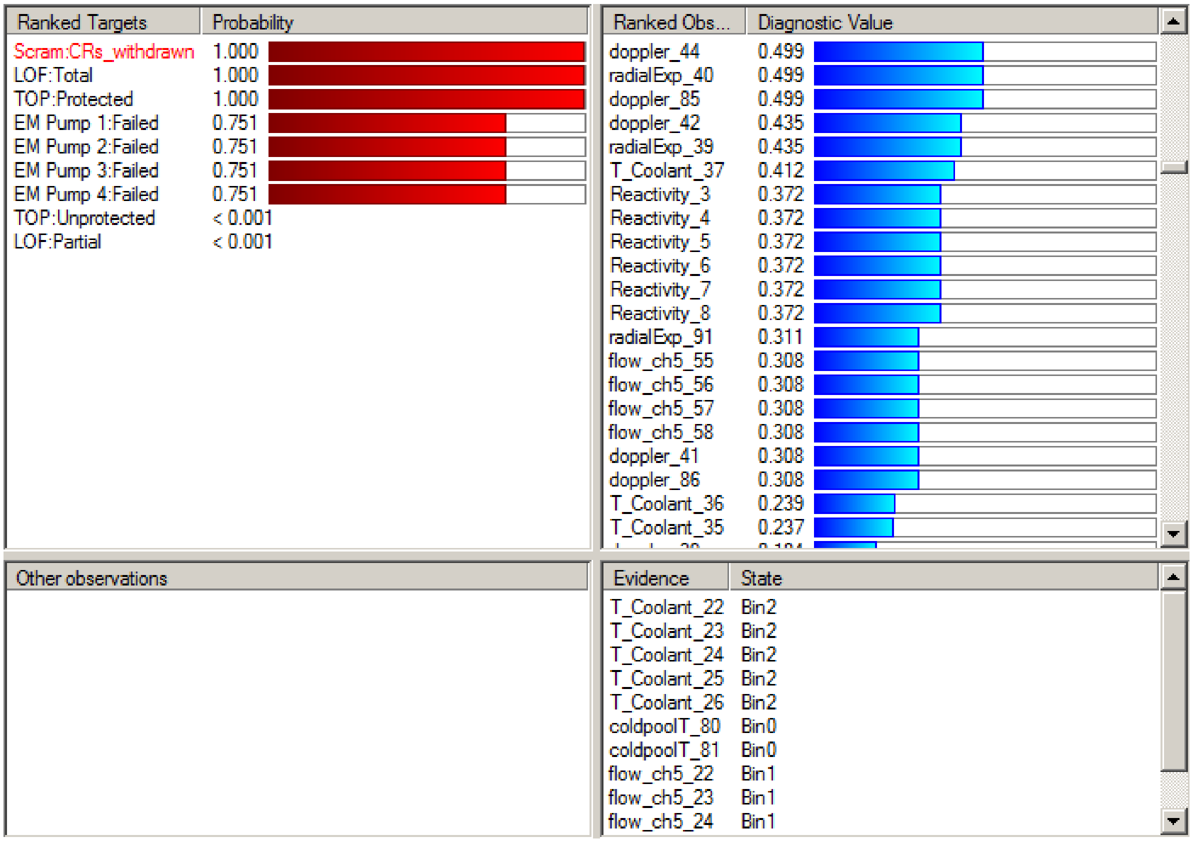

Figure 37: DBN Diagnosis, Additional Low Cold Pool Temperature

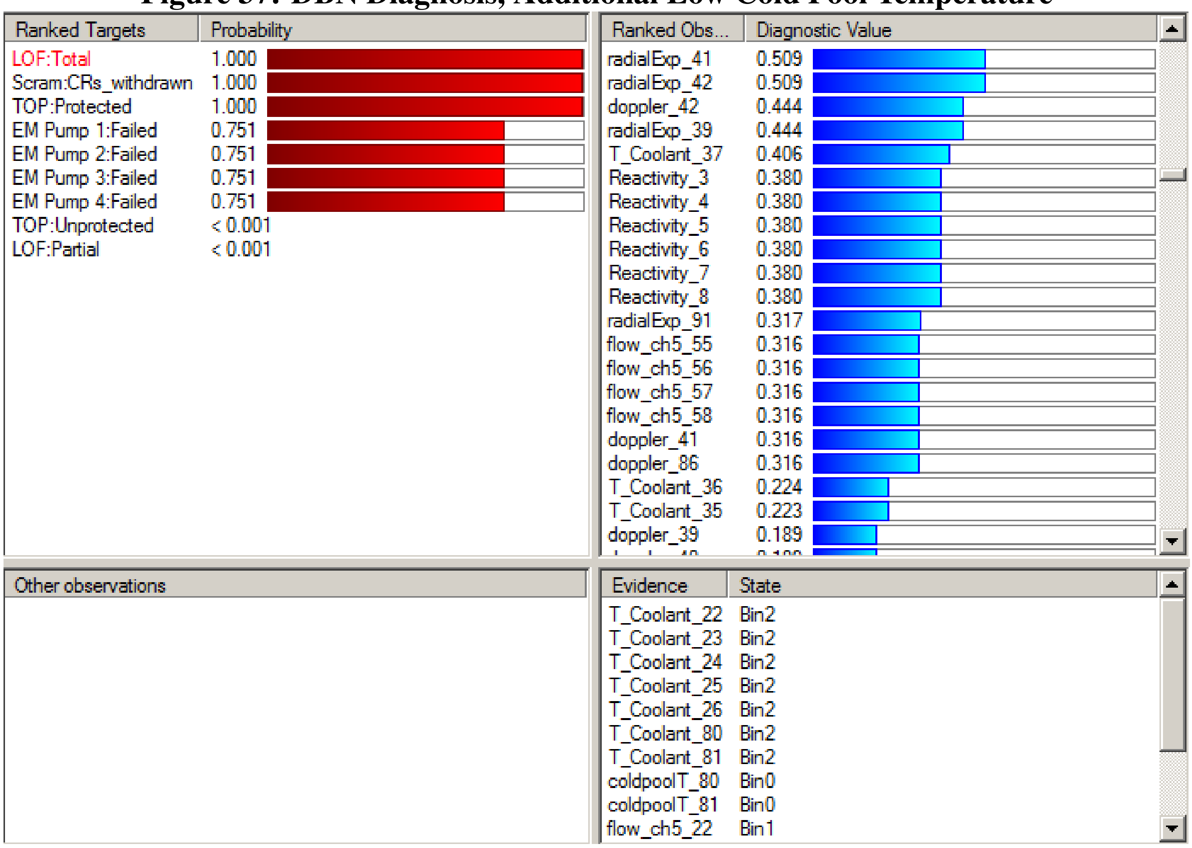

Figure 38: DBN Diagnosis, Additional Channel Coolant Temperature

5.3 DBN Time Step Effects

In order to test the effect of the number of time steps chosen for the ALADDIN data processing (as discussed in Chapter 3), the model with 96 time steps was compared to models with 190, 268, and 373 time steps. It is important to remember that for each different time step scheme different samples may be used from the overall set of accident data, which may affect binning. In all cases, flow and cold pool temperature are set to Bin 1 for time between 150,000s and the end of the simulation. This represents a loss of flow scenario with or without TOP. This spans 4 time steps in the model with 96 time steps and 16 in the model with 373, or approximately 4% of all evidence for each variable in both models. In the case of 96 time steps (Figure 39), the model believes with 0.987 probability that a protected TOP is occurring. The model recognizes channel coolant temperature later in the accident as the next most useful piece of evidence for updating the diagnosis.

With 190 time steps, the model infers the prior probabilities (Figure 40) and no additional piece of evidence is particularly useful for determining a diagnosis. With 268 time steps, the model believes (Figure 41) with a 0.021 probability that a protected TOP is occurring. This is just above the prior probability of 0.015. It recognizes various reactivities as the next most influential pieces of evidence. With 373 time steps, the model believes (Figure 42) with 1.000 probability that both a total LOF and protected TOP are occurring. It recognizes various reactivities as the most influential next pieces of evidence in changing its diagnosis. These discontinuities in results suggest that there is sensitivity to the discretization of the data which must be investigated further.

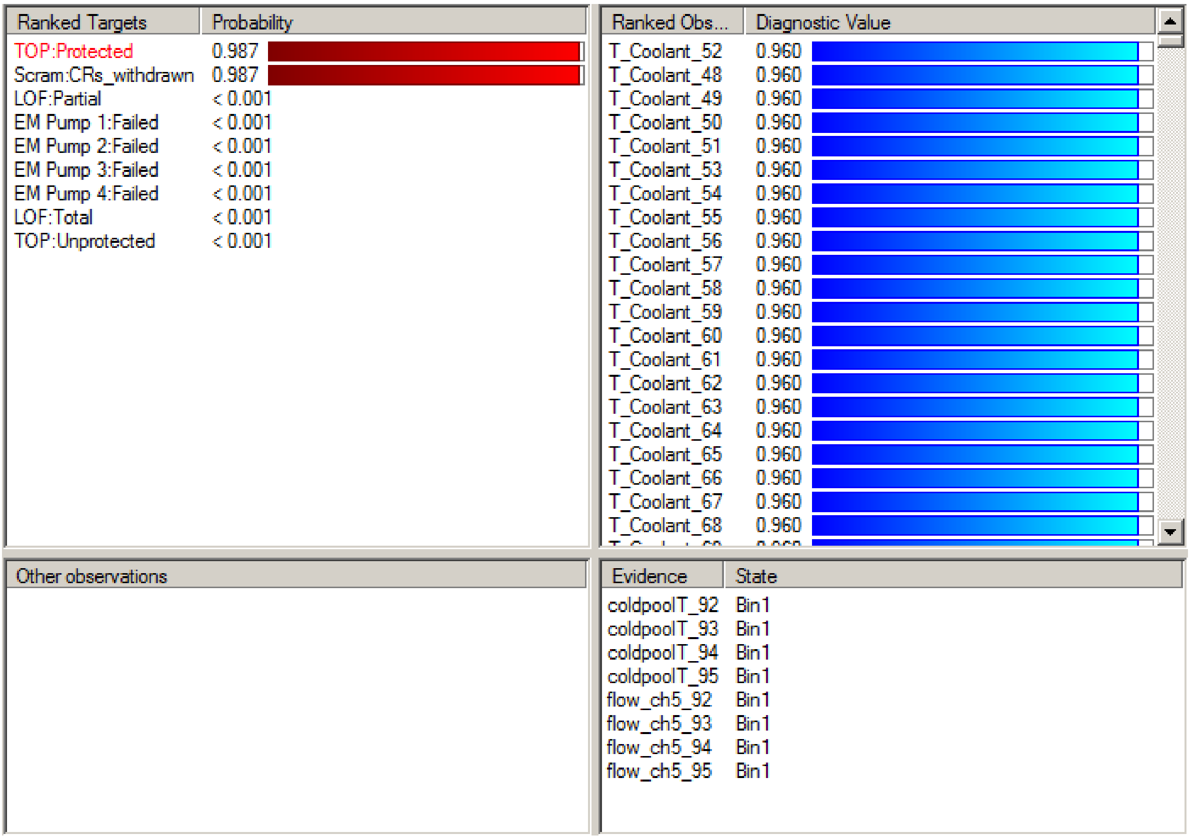

Figure 39: DBN Diagnosis, 96 Time Steps

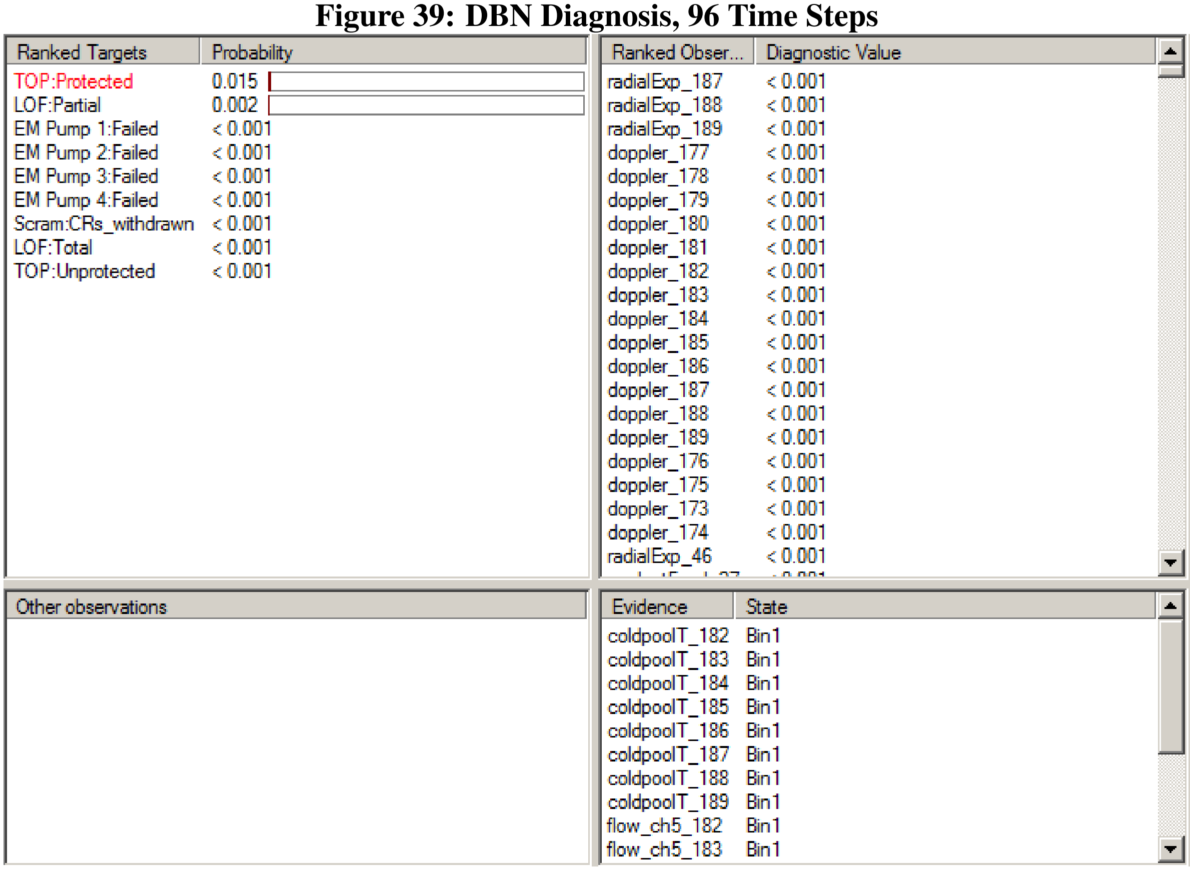

Figure 40: DBN Diagnosis, 190 Time Steps

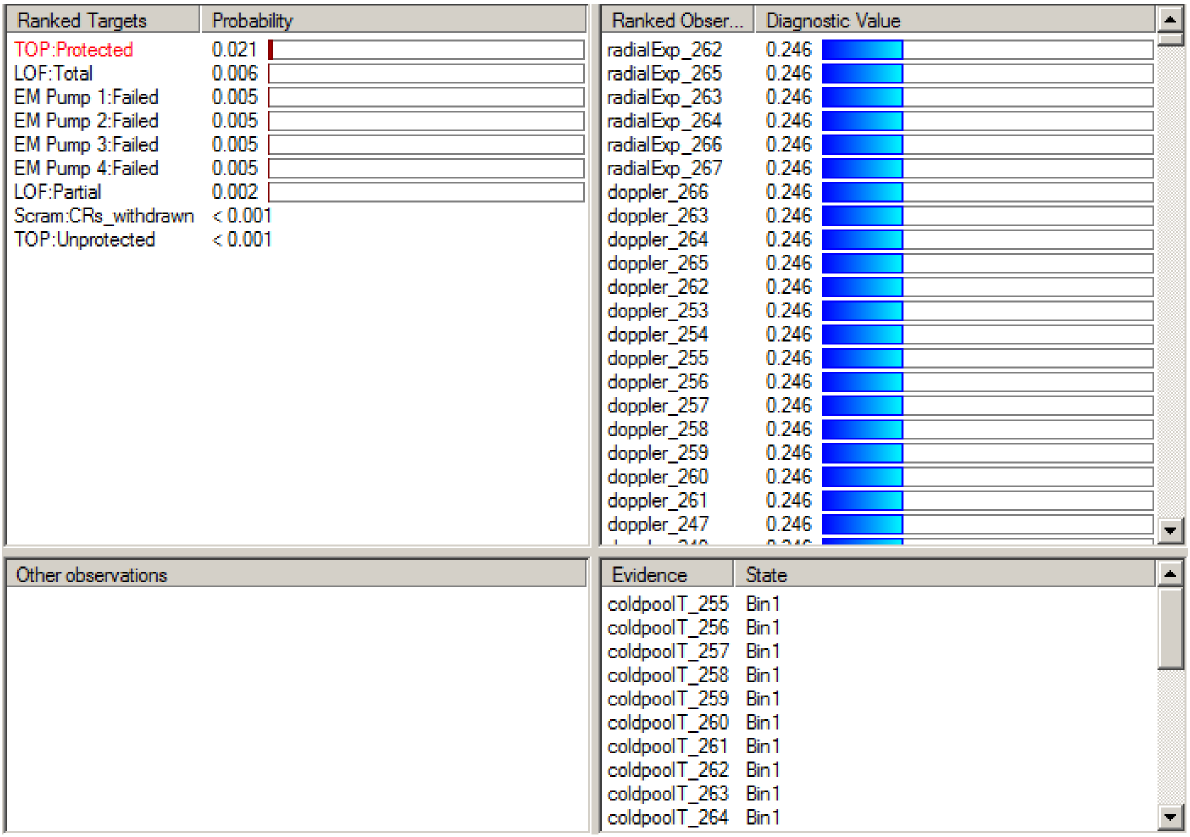

Figure 41: DBN Diagnosis, 268 Time Steps

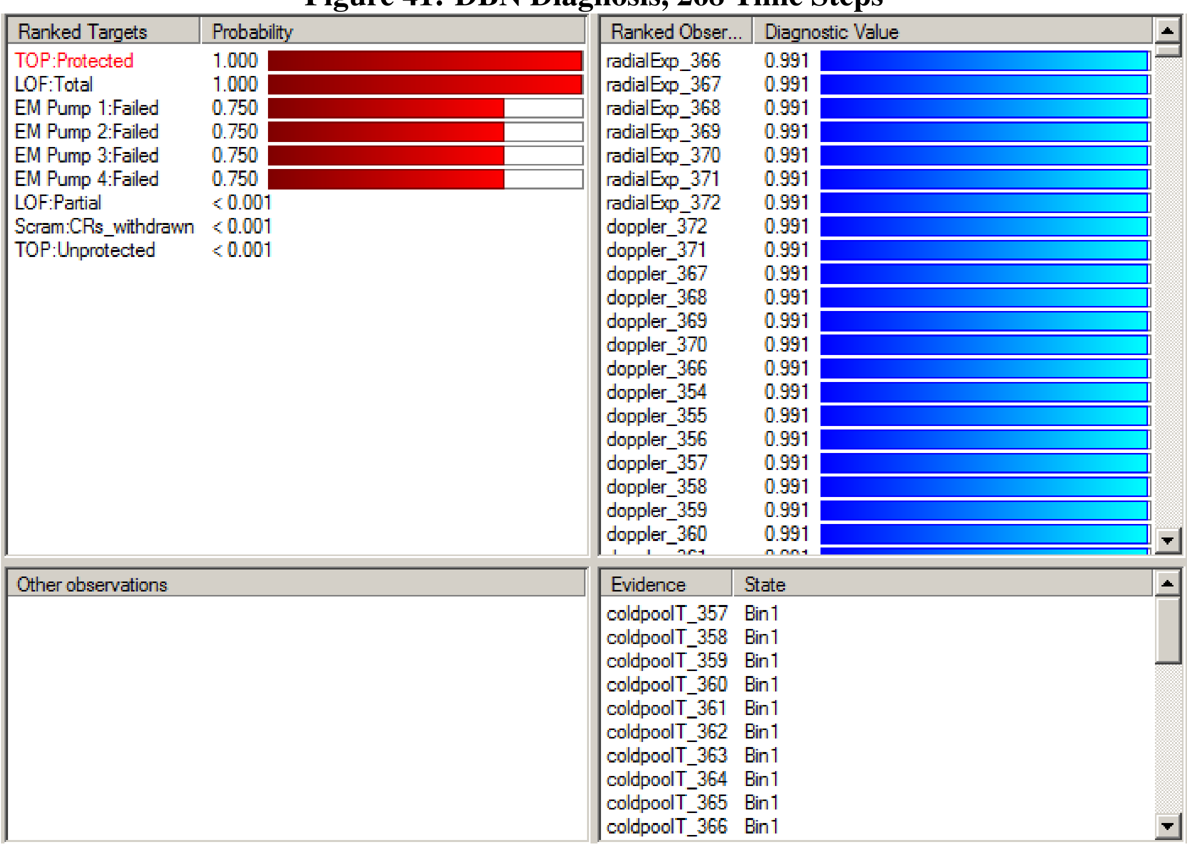

Figure 42: DBN Diagnosis, 373 Time Steps

6 Additional DBN Model Analysis

Future research will focus on validation of the proof-of-concept models through two activities. First, we will validate the predictions of the model against a test set of data generated from additional ADAPT-SAS4A simulations of the accident scenarios. We will use the test data to quantify the diagnostic accuracy of the DBNs and to identify the DBNs with highest performance. Secondly, we will attempt to compare the highest-performing DBNs against existing emergency procedures and operator expectations for the two accident scenarios (TOP and LOF) to

- Verify that the DBN can correctly diagnose the accidents based on the plant conditions in the procedures
- Verify that the level of detail in the DBN matches operator expectations

The results of this modeling and validation activity will provide the necessary foundation for an integrated probabilistic-deterministic framework for nuclear risk assessment, and will set the foundation for application of that framework to change the state of the art in accident modeling in nuclear power and beyond.

Additional nodes in the DBN that may increase the utility of the model have been identified. The first is a target node that reflects the state of the fuel, in order to identify which pieces of evidence suggest an accident sequence that will lead to fuel damage. It is expected that this target would follow cladding thickness closely, similarly to how differential pressure follows flow. An additional accident type node may be added for loss of operating heat removal (LOHR), which would be closely related to the states of targets BOP and DRACS. Both of these changes, in addition to experimentation with more refined binning processes, will lead to a more powerful diagnosis tool.

To provide insight into which plant parameters are most important, the authors will expand the use of Kullback-Leibler (KL) divergence [14]. Formally, KL divergence measures the distance between two probability distributions (e.g., between two DBN models). In probability theory, KL divergence is used to measure the amount of information lost when Q is used to approximate P. In a general probability application, P could be defined as the true distribution of a variable and Q could be defined as a theoretical model of the data. For application to the current problem, KL divergence is used to compare the base DBN model with a DBN with one plant parameter removed. Essentially, the KL divergence calculates the information lost when an arc is removed from the model [15].

$$\sum_{i \in p} P(i) log \left(\frac{P(i)}{Q(i)} \right) \tag{1}$$

In calculating the KL divergence of an arc in the DBN, $P(i)$ is the model with the arc being measured while $Q(i)$ is the model without the arc being measured. 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 to that found in [16]. Joint KL divergence calculations are conducted over all the target nodes for each observation node. In calculating the joint KL divergence, we treated each combination of possible target node states as a single state in a joint target node that collected all targets into a single node.

7 Conclusion

In this report a notional small sodium-cooled fast reactor system was presented with a set of accidents, including earthquake-induced transient overpower, control rod removal transient overpower, loss of primary coolant flow, and loss of operating and shutdown heat removal. The reactor's response to 7,188 unique accident scenarios was examined, and then processed for inclusion in a Dynamic Bayesian Network for the purpose of diagnosing an accident by the state of various reactor parameters. Particular emphasis is paid to those parameters which are monitored and displayed in the control room, as these are the information that would be available to operators in the event of such accidents. As the procedure reaches a higher level of refinement, it can be used for a variety of purposes:

- Inform SAMGs and accident response in general,
- Assess the challenges of accident diagnosis to examine the acceptability of reduced staffing levels,
- Inform operator training as an exploration tool, demonstrating the relationships between accident conditions and plant parameters
- Inform instrument design by identifying those monitored plant parameters that are essential for proper diagnosis of specific accidents

Previous reports on SMART SAMG development have focused on accident management insights for the sodium fast reactor. While accident management is still the overriding goal for the SNL component of DOE's advanced reactor PRA initiative, it was apparent to the authors that the diagnostic capability of the DBN must mature further before significant accident management insights could be gained. Thus, SNL efforts in FY15 were focused on developing the ALADDIN data translator which allows for rapid prototyping of accident diagnostic DBNs. Using the ALADDIN data translator to expertly determine the selection of the diagnostic binning values, along with rapid iteration of the DBN structure, should dramatically increase the diagnostic capabilities of the DBN.

The DBN approach for post-processing accident data presented in this report demonstrated limited diagnostic capabilities due to the initial n-ary discretization. The three bins produced by the n-ary discretization procedure for this analysis produced certain conditional probabilities that were consolidated in one bin and did not vary with accident time. Even with these limitations, inferences were still achievable, including the diagnosis of:

- Protected transient overpower inferred from early hot core outlet temperature and low flow indications and
- Loss of flow inferred through further incorporation of low inlet coolant temperatures further into the accident.

One open question that remains from the current analysis relates to the ability of the DBN to diagnose unprotected accidents. Namely: how much contrary evidence is needed before the DBN is able to overcome the low probability prior probability associated with unprotected events? Because many of the interesting accident management decisions for sodium reactors exist in the beyond design basis unprotected accident regime, this question must be answered before the ultimate usefulness of the DBN SMART SAMG methodology is truly realized.

References

[1] K. M. Groth, M. R. Denman, J. N. Cardoni, and 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, 22–27 June 2014.

[2] K. M. Groth, M. R. Denman, J. N. Cardoni, and T. A. Wheeler, "Proof of principle framework for developing dynamic risk-informed severe accident management guidelines," Sandia National Laboratories, Albuquerque, NM, Tech. Rep. SAND2013-8324, September 2013.

[3] M. R. Denman, K. M. Groth, J. Cardoni, and T. Wheeler, "Advance liquid metal reactor Discrete Dynamic Event Tree/Bayesian network analysis and incident management guidelines (risk management for sodium fast reactors)," Sandia National Laboratories, Albuquerque, NM, SAND2015-2484, 2015.

[4] L. J. Koch, *Experimental Breeder Reactor II: An Integrated Experimental Fast Reactor Nuclear Power Station*, 2003.

[5] "Clinch river breeder reactor project preliminary safety analysis report," Project Management Corporation, Tech. Rep., 1976.

[6] "An overview of fftf contributions to liquid metal reactor safety," Westinghouse Hanford Company, Richland, Washington, Tech. Rep. WHC-SA–0973, 1990.

[7] "Preapplication safety evaluation report for the power reactor innovative small module (prism) liquid-metal reactor," United States Nuclear Regulatory Commission, Washington, DC, Tech. Rep. NUREG-1368, 1994.

[8] 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.

[9] Argonne National Laboratory, "The SAS4A/SASSYS-1 Safety Analysis Code System," Argonne, IL, ANL/NE-12/4, 2011.

[10] T. Sumner, A. Moisseytsev, and T. Y. C. Wei, "Afr-100 safety analyses," in *Proceedings of ICAPP '12*. Argonne National Laboratory, 2012.

[11] M. C. Darling, T. B. Jones, G. F. Luger, and K. M. Groth, "Aladdin: The automatic loader of accident data for dynamic inferencing networks," Sandia National Laboratories, Albuquerque, NM, Tech. Rep. SAND2015-DRAFT, 2015.

[12] "Probabilistic risk assessment of the advanced liquid metal reactor," General Electric, Oakland, California, Tech. Rep. GEFR-00873, 1994.

[13] "Prism preliminary safety information document," General Electric, San Jose, California, Tech. Rep. GEFR-00793, 1987.

[14] 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.

[15] J. R. Vergara and P. A. Estévez, “A review of feature selection methods based on mutual information,” *Neural Computing and Applications*, vol. 24, no. 1, pp. 175–186, 2014.

[16] J. R. Koiter, “Visualizing inference in Bayesian networks,” Master’s thesis, Delft University of Technology, 2006.

[17] “Severe accident management programmes for nuclear power plants,” International Atomic Energy Agency, Vienna, Austria, Tech. Rep. NS-G-2.15, 2009.

[18] D. Clayton and W. P. Poore, “Fukushima daiichi - a case study for bwr instrumentation and control system performance during a severe accident,” Oak Ridge National Laboratory, Oak Ridge, Tennessee, Tech. Rep. ORNL/TM-2013/154, 2013.

[19] N. Todreas and M. Kazimi, *Nuclear Systems: Thermal Hydraulic Fundamentals, Second Edition*, 2012, no. v. 1.

[20] Y. Yang and G. I. Webb, “A comparative study of discretization methods for naive-bayes classifiers,” in *Proceedings of PKAW*, vol. 2002. Citeseer, 2002.

[21] A. Darwiche, “A differential approach to inference in bayesian networks,” *Journal of the ACM (JACM)*, vol. 50, no. 3, pp. 280–305, 2003.

DISTRIBUTION:

1	MS 0748	Katrina Groth, 6231
1	MS 0748	Matthew Denman, 6231
1	MS 0784	Mitch McCrory, 6231
1	MS 0784	Randy Gauntt, 6232
1	MS 0899	Technical Library, 9536 (electronic copy)

Sandia National Laboratories