Safety-Informed Design: Using Subgraph Analysis to Elicit Hazardous Emergent Failure Behavior in Complex Systems

Matthew G. McIntire, Christopher Hoyle, Irem Y. Tumer (Department of Mechanical, Industrial and Manufacturing Engineering, Oregon State University), David C. Jensen (Mechanical Engineering, University of Arkansas)

Corresponding Author:

Christopher Hoyle

418 Rogers Hall

Oregon State University

Corvallis, OR 97331-6001

Subgraph-Based Hazard Identification

Pages: 23

Tables: 0

Figures: 4

Title:

Safety-Informed Design: Using Subgraph Analysis to Elicit Hazardous Emergent Failure Behavior in Complex Systems

Abstract:

Identifying failure paths and potentially hazardous scenarios resulting from component faults and interactions is a challenge in the early design process. The inherent complexity present in large engineered systems leads to non-obvious emergent behavior, which may result in unforeseen hazards. Current hazard analysis techniques focus on single hazards (fault trees), single faults (event trees), or lists of known hazards in the domain (hazard identification). Early in the design of a complex system, engineers may represent their system as a functional model. A function failure reasoning tool can then exhaustively simulate qualitative failure scenarios. Some scenarios can be identified as hazardous by hazard rules specified by the engineer, but the goal is to identify scenarios representing unknown hazards. The incidences of specific subgraphs in graph representations of known hazardous scenarios are used to train a classifier to distinguish hazard from non-hazard. The algorithm identifies the scenario most likely to be hazardous, and presents it to the engineer. After viewing the scenario and judging its safety, the engineer may have insight to produce additional hazard rules. The collaborative process of strategic presentation of scenarios by the computer and human judgment will identify previously unknown hazards. The feasibility of this methodology has been tested on a relatively simple functional model of an electrical power system with positive results. Related work applying function failure reasoning to a team of robotic rovers will provide data from a more complex system.

1 Introduction

Complex engineered systems such as aerospace platforms and power generation facilities exhibit complex forms of failure. While some hazards may be identified and accounted for during design time, others remain unknown until the system is fully operational. These safety-critical systems do undergo rigorous testing and validation to assure safe operation, and are designed to be inherently robust and do regularly operate with degraded components. However, highly publicized, costly, and sometimes fatal accidents still occur, usually preceded by multiple seemingly innocuous events that compound and cascade across subsystems. The recent grounding of the Boeing 787 line, estimated to cost $5 billion, the immeasurable economic, environmental, and human cost of the Deep Water Horizon disaster and the space shuttle Columbia accident all demonstrate the unacceptably high cost of addressing complex failures and safety too late. For this reason, a growing field of research has been exploring how to move safety and risk analysis into the early design stage to achieve safe system design [1].

Specifically, while designing the space shuttle, NASA engineers identified ice falling from the external fuel tank as a hazard to the orbiter, and mitigated it by applying foam to the tank. They impact-tested the heat shield material for small chunks of ice and other debris, and found the risk due to falling ice after the foam installation was acceptable [2]. However, they did not consider falling foam as a potential hazard until it was observed to occur during missions. By that time, engineers were not thinking about the exact parameters of the earlier impact tests, only that they resulted in acceptable risk. Thus, a new interaction between systems resulted in an unforeseen hazard. The lack of action after the potential hazard was identified is not the focus of this paper. Instead, the methodology outlined here can be used to systematically identify unforeseen potential hazards during the design phase.

Early in the design of a complex system, engineers may represent their system as a functional model. A function failure reasoning tool can then exhaustively simulate qualitative failure scenarios. Some scenarios can be identified as hazardous by hazard rules specified by the engineer, but the goal is to identify scenarios representing unknown hazards.

The incidences of specific subgraphs in graph representations of known hazardous scenarios are used to train a classifier to distinguish hazard from non-hazard. The algorithm identifies the scenario most likely to be hazardous, and presents it to the engineer. After viewing the scenario and judging its safety, the engineer may have insight to produce additional hazard rules. The collaborative process of strategic presentation of scenarios by the computer and human judgment will identify previously unknown hazards.

The feasibility of this methodology has been tested on a relatively simple functional model of an electrical power system with positive results. Related work applying function failure reasoning to a team of robotic rovers will provide data from a more complex system.

2 Background

2.1 Design Stage Analysis of Failure and Safety

Design is fundamentally a decision-centric process [3] and the criteria used to evaluate different concept solutions provide a basis for making those decisions. While other design aspects such as performance, manufacturability, and sustainability can be design objectives in the early design stage, for safety-critical systems the focus must at some point be upon: 1) Risk and reliability analysis, and 2) Hazard (safety) analysis.

Reliability is a property of a system and represents the tendency for a system to not fail (be available). Traditional approaches for calculating reliability, such as aggregate failure rates [4] or component property distributions [5], are data-driven and require a well-defined design to provide

meaningful results. Examples include top-down approaches like Fault Tree Analysis [6] and Hazard and Operability Analysis [7] and bottom-up approaches like Failure Modes and Effects Analysis [8] and Probabilistic Risk Assessment [9].

Early work to move reliability assessment into the conceptual design stage focused on qualitative descriptions to describe the nature of faults in the conceptual design perspective [10], and how those faults affect the performance of other components in the system [11, 12,13]. Quantitative methods use descriptions of fault probability to provide a risk assessment at the early design stage [14, 15, 16, 17]. In order to provide an assessment at the concept stage, failure was viewed in terms of its effect on function [14, 16, 18, 19, 20].

Others have explored the design stage by reasoning about failures based on the mapping between components, functions, and nominal and off-nominal behavior [11, 12, 15, 18, 21, 22, 23, 24, 25, 26, 27, 28, 29]. A common element to each of these different methods for risk analysis is the use of a conceptual system representation for identifying the system-level impact of faults.

While these methods are appropriate for reliability analysis, they cannot provide an assurance of safety (i.e., hazard analysis). Safety is viewed as an emergent property of a system [30]. The functional approaches above assume but do not specify the safety of functions. For example, the functional model of a chemical reactor design would include high-level functions like ``store'' and ``mix''. However, safety functions like ``ensure no loss of human life'' are not captured explicitly in the function structure. To assure safety, other top-down approaches have been developed. A systems theoretic approach has been developed to identify means of reaching unsafe system states [30,31]. However, identifying fault propagation paths from component behaviors to system state has not yet been achieved.

The Systems Theoretic Process Analysis method (STPA) [30, 31] is an example of a top-down approach that attempts to assure safe system development. The core concept of using STPA is

identifying hazards and creating designs as control structures to mitigate those hazards. For analysis with this method, the potential for the hazard occurs when the safety constraints designed into this control structure are violated through a specific list of failures. Another method for safety analysis is the hazard and operability study (HAZOP) [7] which is based on modeling the interaction flow between components and recognizing a hazard if components deviate from the operation that was intended for the component during the design. The system-level impact of these failures is the identified hazard. Determining the probability of these failures is not possible because of the complex and unknown interaction behavior. Instead, work using this method has used an inverse approach by specifying the probability that the failure could be mitigated [30].

2.2 Functional Failure Reasoning

Risk (and safety) analysis has the greatest impact on system design when it can be incorporated into the early design-stage and be used as a decision making tool. In this capacity, safety becomes an attribute of the design and can be used in both architecture and component selection. The challenge of risk assessment at this design stage is the lack of refined system information. A fault is an undesired behavior in a component or set of components that can lead to losses in system functionality. When these losses occur, the system experiences some kind of hazard and can fail to prevent itself from being in an unsafe state. This simple model of failure and safety forms the basis of this research.

Traditional methods of failure and risk analysis rely on statistical failure data and apply methods in which expert knowledge of the system is needed to determine the impact and path of fault propagation; hence, such methods are implemented at the validation stage of design, where specific component failure probabilities and probable fault propagation paths can be defined. To achieve the benefits of early risk-based decision making, several methods for failure analysis using an abstract functional approach have been developed, including the use of historic failure rates associated with

component types or functions to identify risk [14, 16], and behavioral approaches to determine the potential impact of failures [11, 25, 32].

The *functional* approach enables a high degree of failure characterization. In particular, the Function Failure Identification and Propagation (FFIP) framework is one of the methods that use a behavioral approach for assessing the functional impact of faults in the early design stages [33]. The result of using an FFIP-based analysis of a design is an evaluation of the state of each function in the system in response to a simulated failure scenario. In previous work these results have been used to evaluate the consequences of different fault scenarios for a system design and for assessing the state of the system due to functional loss [1, 12, 22, 23, 24, 25, 34, 35, 36].

To date, the goal of the FFIP analysis approach has been to demonstrate that it is possible to identify failure propagation paths by mapping component failure states to function 'health'. While the fundamentals of FFIP have shown great promise, the value to the complex system design process has not been demonstrated. We demonstrate a new important use of the data generated by an FFIP analysis: *to help identify unforeseen hazardous scenarios.*

3 Unknown Hazard Identification Methodology

Figure 1 is a visualization of the entire methodology presented here. The engineer is of central importance, as they will create the initial functional model, specify rules for identifying hazardous scenarios, analyze individual scenarios and judge their hazard potential, and finally act on that judgment by modifying their previous input.

3.1 Functional Modeling

In order to use function failure logic, the engineer must first specify a system functional model. The level of abstraction used to create the model will determine the precision of the identified hazards. We will focus on a component-level abstraction. At this level, the engineer needs to study each system component, and specify every function that it fulfills, using a functional basis as specified in [37]. All

flows of mass, energy, and information (signals) within the system need to be accounted for. Figure 2 shows a portion of a basic functional model, demonstrating the functional basis, that will be used as an example throughout this section.

3.2 Failure Propagation

At this point the engineer must consider every state that each function can attain. From the FFIP framework [38, 39], we consider that each function state may be categorized as one of four health states: Healthy, Degraded, Lost, and No Flow.

The engineer must develop logic relating each function to its input and output flows. The questions to answer include: a) what effect does each flow have on the connected function state? And b), what effect does each function state have on each connected flow? In Figure 2, for example, if the *Store Electrical Energy* function is lost, the connected energy flow will be eliminated, resulting in the *Inhibit Electrical Energy* function transitioning to the No Flow state. This may be modeled using software such as the Stateflow toolbox in Matlab Simulink.

Next, faults are simulated. A Matlab script creates failure scenarios by triggering one or more faults in the model and running the behavioral model until a steady or stable state is reached. This includes every possible fault, one at a time, every pair of function-faults, and so on, until the number of coincident faults becomes either highly improbable, or the total computation time becomes intractable. A large matrix of data results from this step, containing the end health state of each function in each failure scenario. See [25] for more details of this approach.

Some sets of faults result in identical scenarios; duplicates are combined in the dataset. For example, once again looking to Figure 2, a failure in *Store Electrical Energy* might have the exact same end state as a simultaneous failure in both *Store Electrical Energy* and *Actuate Relay*. Also, some functions may be identified by the engineer whose states have no imaginable effect on the safety of the system. These may be removed from the data.

3.3 Initial Hazard Identification

When an engineer creates a component-level functional model of a system, they (as an expert) should be able to identify at least some of the critical functions or sets of functions that upon degradation or loss will result in a hazardous failure. This knowledge may come from experience, historical data, intuition, or some previously utilized hazard identification technique. For example, they might judge that any scenario based on the functional model from Figure 2 wherein simultaneously the *Store Electrical Energy* function is Nominal and the *Inhibit Electrical Energy* function is Lost is a hazard. Applying these rules to the complete set of failure scenarios reveals a subset of scenarios representing known hazards.

3.4 Graph Representation

In [39], Jensen et al proposed using latent class analysis (LCA) to group failure scenarios by functional effect similarities. This approach was initially attempted to train a classifier to identify unknown hazards, but after performing various validation tests, it was found that it performed little better than randomly guessing at scenarios. Instead, we require a method that incorporates the topology of the functional model, rather than treating the system as a list of independent functions.

Each failure scenario must then be represented as a graph. We begin by creating a graph representation of the functional model used in the FFIP process; Figure 2 is already represented as such. Each node represents a function labeled by its type and each directed edge represents a flow from one function to another labeled by its type. The labels are derived from the functional basis [37]. Final results from the proposed method may vary depending upon the model abstraction level used.

Next, the graph is repeatedly modified to represent the functional state at the end of each failure scenario. We relabel each node to indicate the end health state of the represented function (Nominal, Degraded, Lost, or No Flow), and relegate the function type to a new first-degree node

connected by an edge labeled Type. The example from Figure 2 has been modified to represent a partial failure scenario in Figure 3.

3.5 Subgraph Analysis

In order to create a classifier that distinguishes between hazardous and safe failure scenarios, hazard indicators must be identified. The frequency of occurrence of various motifs or subgraphs within each graph serves as such indicators. The goal is to identify subgraphs which occur with a different frequency in graphs representing hazardous scenarios versus graphs representing safe scenarios. To identify subgraphs, we used the software package Subdue: Graph Based Knowledge Discovery from Washington State University. It finds subgraphs most prevalent in a set of graphs by way of compressing the graph data [40].

In order for Subdue to identify subgraphs containing failed functions, which might be helpful in identifying new hazards, we run it on the set of graphs representing the known hazards. However, because even in the known presence of hazards most functions still finish in a nominal state, we first trim excess nominal functions from the graphs. This is done to focus the subgraph identification on those functions critical to identifying hazards. Any nominal node that is not adjacent to a failed node is removed. This is done for each known hazardous scenario.

3.6 Naive Bayes Classifier

In order to estimate which unknown scenario is most likely to be hazardous, we calculate the unknown scenario most likely to be classified as hazardous, given a naive Bayes classifier constructed from the frequency of subgraph occurrence in the failure scenarios. We use a naïve Bayes classifier due to its simplicity of implementation and because the classifier model fits our problem well.

Each subgraph *i* of *n* subgraphs becomes a feature in the naive Bayes classifier. The simplest measure to use for the classifier is the number of times *x* the subgraph occurs in the graph representation of a scenario. A distribution is approximated for this frequency of occurrence of each

subgraph in the known hazardous scenarios $p(x_i|hazard)$. This is repeated for the unknown scenarios $p(x_i|unknown)$ and the known safe scenarios $p(x_i|safe)$. Note that initially there may not be any known safe scenarios, and so $p(x_i|safe)$ will be zero for all *x*.

Under the naive independence assumptions, the Bayes classifier has the following form

$$p(C_k) = \frac{p(C_k)\prod_{i=1}^{n} p(x_i|C_k)}{\prod_{i=1}^{n} p(x_i)} \tag{1}$$

In other words, the probability of an event with features ***x*** belonging to class *k* is the probability of any event belonging to class *k* times the product of the conditional probabilities of each x_i given class *k* divided by the product of the total probabilities of each x_i.

In this case, we are only interested in the relative likelihoods of each scenario, represented by their respective *x* values, belonging to the hazard class. Thus we can reduce (1) to

$$p(hazard|\boldsymbol{x}) \propto \prod_{i=1}^{n} \frac{p(x_i|hazard)}{p(x_i|hazard) + p(x_i|unknown) + p(x_i|safe)} \tag{2}$$

where the probability of a scenario represented by ***x*** belonging to the hazard class is proportional to the product of the ratios of each conditional probability of occurrence of x_i given a hazard classification and the sum of the conditional probabilities given each class.

3.7 Iterative Hazard Identification

The scenario of unknown safety with the highest likelihood calculated by Eq. (2) is estimated to be the scenario most likely to be hazardous. It is then presented to the engineer, who will study its functional state. A graphical representation of the functional model will be displayed on screen, with the health state of the functions and flows clearly indicated. If the model is too large to display all at once, fully nominal and or fully failed sections of the model may be collapsed into blocks related to a

higher level system function. The engineer must judge the safety of a given scenario. If the engineer judges the scenario as safe, then it will be reclassified as known safe. The conditional probabilities of subgraph frequencies for unknown and known safe classes are recalculated, and the likelihood of each scenario is updated. This is represented in Figure 1 by the engineer making the No decision.

If, however, the engineer classifies the scenario as hazardous, they can create a new hazard identification rule that will not only account for the scenario at hand, but potentially others within the set of simulated failures. This necessitates rerunning the subgraph analysis. Once again, this causes the appropriate conditional distributions and all likelihoods to be updated, and a new scenario to be presented. This is represented in Figure 1 by the engineer making the Yes decision.

A third option is required should the engineer judge that the hazard potential of the scenario depends upon some factor not included in the system representation. At this point, they can go back to the functional model and incorporate new functions and connections as needed. While the failure scenarios and clustering results are being updated, the scenario at hand will be temporarily classified as safe, so that the engineer may continue to judge scenarios with the algorithm penalizing those which are similar to the current one.

This process is continued until one of a number of stopping conditions is met. First, a minimum likelihood value may be established, below which scenarios will no longer be justifiably similar enough to known hazards to be considered. Second, a predetermined consecutive number of safe judgments may be deemed sufficient, especially if the identification of new hazards has been shown to decrease approximately exponentially. Third, the resources allocated to identify new unknown hazards have been exhausted.

4 Electrical Power System Case Study

We utilized data from an electrical power system (EPS) functional model and accompanying fault propagations from [39] to test our methodology. The system includes four load types: pump, fan, light, and generic DC load; it includes the power supply to those loads, including battery and invertor; and it includes the control of that power by circuit protection, sensors, software, and relays. The system includes redundancy in load satisfaction (two pumps, two fans, etc.), redundancy in power supply, and a fully interconnected control system. Each component has its own failure modes to populate a list of failure scenarios, except the software control, which is assumed infallible. The behavioral model and function-fault logic were written in Matlab Simulink and simulated exhaustively under single and double fault conditions for previous work (see [39] and references therein) The data consists of 3508 unique failure scenarios represented by the states of 58 functions. A block diagram of the EPS is shown in Figure 4.

In order to demonstrate the power of the unknown hazard identification method, we assumed that any scenario wherein Fan 1 and Pump Sensor 1 were simultaneously failed was a known hazard. Correspondingly, we assumed that any scenario wherein Fan 2 and Pump Sensor 2 were simultaneously failed was also a hazard, albeit an unknown one. Thus there were 16 known hazardous scenarios and 16 unknown hazardous scenarios. We tested how many iterations of our method were required to identify one of the unknown hazards.

First, we represented the functional model as a graph, with each component-as-a-function given a single word label from the functional basis such as Store, Sense, Inhibit, or Actuate. This effectively hid the unknown (test) hazards among many similar failure scenarios.

Next, following the method as laid out in Sections 3.4 and 3.5, we used Subdue to identify 40 representative isomorphically unique subgraphs from the graphs of the 16 known hazardous scenarios. Continuing with the iterative process in Section 3.6, we counted the frequency of each subgraph in each failure scenario, and fitted probability distributions to the occurrence of each subgraph in each class, the

class of 16 known hazards, and the class of 3492 unknown scenarios. After inspecting the histograms of the subgraphs, we decided to fit a mixture of two normal distributions to each to represent the distributions parametrically, due to their obviously bimodal nature.

This was implemented in a Python script which made an external call to run Subdue. We used the graph-tool package to handle subgraph isomorphisms and frequency counts, and wrote our own implementation of the Naive Bayes classifier.

We then set up a `while` loop which identified the most-likely-to-be-hazardous scenario and judged it as safe, repeating the process until one of the test hazards appeared. In our test, one of the test hazards appeared as the most likely hazard on the 11th iteration.

In order to show the significance of this result, we performed a statistical test to determine if identifying the hazard in 11 iterations is likely to occur randomly. We used the negative hypergeometric distribution implemented in the `tolerance` package in R. The negative hypergeometric is the appropriate distribution to use when sampling from a finite population (e.g. population of scenarios) without replacement in which each drawn sample can be classified into two mutually exclusive categories, such as hazard/no hazard. To calculate the probability of at least one test scenario being drawn randomly from the set without replacement after eleven draws we used the R command `pnhyper(11,16,3508,1)`, which returns a probability of 4.9 percent. Thus it is highly unlikely that our results using the proposed method occurred because of random chance, there is less than a 5 percent chance of finding the hazard in eleven iterations using random sampling. To put this distribution into perspective, its mean is 219 draws (and its median is 149 draws) meaning that on average it would take 219 draws (versus 11) to identify the hazard through random sampling. Thus, we conclude with 95 percent confidence that our result of finding a test scenario on the eleventh iteration is significant (i.e. our method is significantly different than random sampling).

5 Method Assumptions

While the method is general in nature, there are a few assumptions we must make due to the human-computer interaction. In order for the method to be useful, we must assume that a subset of the failure scenarios implied by the functional model specification represent hazards, and that they are recognizable as hazards by the engineer analyzing the system. And furthermore, that a subset of those scenarios recognizable as hazards is not identified by hazard patterns specified by the engineer *a priori*. We then assume that each additional rule created by the engineer during the human-machine collaborative process to identify more hazardous failure scenarios is useful to the engineer in order to mitigate risk. We believe these assumptions to be plausible, but they should be further tested.

We also make the assumption that types of hazards (groups of hazards identified by rules) are inherently rare among the failure scenarios. Therefore, we attempt to identify as many as possible by presenting the engineer with the scenarios most likely to represent a hazard. This assumption is true for many high safety systems which have evolved over time or are generally well understood by the engineering community; however, newer, more innovative highly complex systems may not meet this assumption.

Alternatively, we could have viewed hazards as more commonplace. Under this assumption, we would present the engineer with the scenario whose safety estimate is the most uncertain. We would be attempting to reduce the total uncertainty in a measure of system safety. In this case, the engineer would not be presented with the most likely to be hazardous scenarios; those would be assumed hazardous.

Finally, we assume that any potential hazards reachable purely through nominal operation of the various functions have already been identified and mitigated, or require a different type of model to identify. See, for example, the functional resonance accident model [41].

6 Conclusions and Future Work

In this paper, we have described a new method for eliciting and identifying unknown hazards in a complex system described by a functional model. We suggest using subgraphs of graph representations of known hazardous scenarios to build a classifier capable of distinguishing hazard from non-hazard. We used a naive Bayes classifier as a simple first attempt. The classifier is updated by the expert judgments made by an engineer, thus providing an innovative man-machine classification system. We have shown that this method is superior to a simple random selection of scenarios.

We plan to test this method on a slightly larger system currently being modeled to validate the function failure identification and propagation method. This will involve a swarm of autonomous rovers. We plan to once again validate our method by defining a list of hazards, removing some, then using the method to see if they reappear quickly. We intend to use a variety of hazards in our test, rather than the single hazard presented here to demonstrate the methodology.

We will also study the further application of subgraph analyses on those failure scenarios identified as hazardous, in an attempt to present the engineer with types or groups of faults that often result in hazards, or common failure paths through the model that result in hazards. Many challenges remain, though, including testing the method with engineers familiar enough with a complex system to fully test its effectiveness.

7 Acknowledgment

This research was supported by the NSF award CMMIs 1363349 and 1363509. Any opinions or findings of this work are the responsibility of the authors, and do not necessarily reflect the views of the sponsors or collaborators.

Author Biographies:

Matthew G. McIntire

Matthew McIntire is a PhD student at Oregon State University in Mechanical Engineering Design. He has studied large-scale optimization under uncertainty, and functional modeling of complex systems for early-stage design risk analysis. He received his BS in Engineering and Applied Science with Mission Applications from Seattle Pacific University in 2008, and worked as a missionary-engineer with Students International for three years in Guatemala.

Christopher J Hoyle

Dr. Christopher Hoyle is currently Assistant Professor in the area of Design in the Mechanical Engineering Department at Oregon State University. His current research interests are focused upon decision making in engineering design, with emphasis on the early design phase. His areas of expertise are uncertainty propagation methodologies, Bayesian statistics and modeling, stochastic consumer choice modeling, optimization and design automation. He is coauthor of the book Decision-Based Design: Integrating Consumer Preferences into Engineering Design. He received his PhD from Northwestern University in Mechanical Engineering in 2009 and his Master's degree in Mechanical Engineering from Purdue University in 1994.

Irem Y. Tumer

Irem Y. Tumer is a Professor and Associate Dean for Research and Economic Development in the College of Engineering at Oregon State University. She was previously a Research Scientist and Group Lead in the Complex Systems Design and Engineering group in the Intelligent Systems Division at NASA Ames. She received her PhD in mechanical engineering from UT Austin in 1998. Her expertise touches on systems engineering, model-based design, risk-based design, system analysis and optimization, function-based

design, and integrated systems health management. Her research focuses on the overall problem of designing highly complex and integrated systems with reduced risk of failures.

David C. Jensen

David C. Jensen is an Assistant Professor in the Department of Mechanical Engineering at the University of Arkansas. He attained PhD, MS, and BS degrees in mechanical engineering at Oregon State University. He also leads the research effort for the Complex Adaptive Engineered Systems Research Laboratory. He has worked extensively in modeling, simulating, and validating complex engineered systems. His research has been supported by awards through NSF, NASA, the Air Force Office of Scientific Research, and DARPA. Dr. Jensen's teaching and research are centered on design and mechanics, complex system design, and risk and safety in complex system design.

List of Figures and their captions:

Figure 1: The Iterative Hazard Identification Process

Figure 2: A partial functional model

Figure 3: A single failure scenario

Figure 4: A block diagram of the EPS

[1] Papakonstantinou, N., Jensen, D., Sierla, S., and Tumer, I. Y., 2011. “Capturing interactions and emergent failure behavior in complex engineered systems and multiple scales". In Proceedings of the ASME Design Engineering Technical Conferences; Computers in Engineering Conference.

[2] Columbia Accident Investigation Board, 2003. Columbia accident investigation board report, volume 1. Tech. rep., aug.

[3] Ullman, D. G., 2003. The mechanical design process. McGraw-Hill.

[4] Carter, A., 1986. Mechanical reliability, Vol. 1. Macmillan London.

[5] Bain, L., and Engelhardt, M., 1991. Statistical analysis of reliability and life-testing models: theory and methods, Vol. 115. CRC.

[6] Vesely, W. E., Goldberg, F. F., Roberts, N., and Haasi, D. F., 1981. The Fault Tree Handbook. US Nuclear Regulatory Commission, NUREG0492, Washington, D.C.

[7] Redmill, F., Chudleigh, M., and Catmur, J., 1999. System safety: HAZOP and Software HAZOP. Wiley.

[8] MIL-STD-1629A. Procedures for Performing Failure Mode, Effects, and Criticality Analysis. Department of Defense.

[9] Stamatelatos, M., and Apostolakis, G., 2002. Probabilistic Risk Assessment Procedures Guide for NASA Managers and Practitioners v 1.1. NASA, Safety and Mission Assurance, Washington, D.C.

[10] Wang, K., and Jin, Y., 2002. “An analytical approach to function design". In 14th Int. Conference on Design Theory and Methodology IDETC/CIE2002, pp. 449-459.

[11] Huang, Z., and Jin, Y., 2008. “Conceptual Stress and Conceptual Strength for Functional Design-for-Reliability". In Proceedings of the ASME Design Engineering Technical Conferences; International Design Theory and Methodology Conference.

[12] Kurtoglu, T., and Tumer, I. Y., 2008. “A graph-based fault identification and propagation framework for functional design of complex systems". Journal of Mechanical Design, 130(5).

[13] Smith, J., and Clarkson, P. J., 2005. “Design concept modelling to improve reliability.". Journal of Engineering Design, 16(5), pp. 473-492.

[14] Grantham-Lough, K., Stone, R. B., and Tumer, I. Y., 2009. “The risk in early design method". Journal of Engineering Design, 20(2), pp. 144-173.

[15] Hata, T., Kobayashi, N., Kimura, F., and Suzuki, H., 2000. “Representation of functional relations among parts and its application to product failure reasoning". International Journal for Manufacturing Science and Production, 3(2/4), pp. 77-84.

[16] Stone, R. B., Tumer, I. Y., and VanWie, M., 2005. “The Function Failure Design Method". Journal of Mechanical Design, 14, pp. 25-33.

[17] Tumer, I. Y., and Stone, R. B., 2003. “Mapping Function to Failure during High-Risk Component Development". Research in Engineering Design, 14, pp. 25-33.

[18] Clarkson, P., Simons, C., and Eckert, C., 2004. “Predicting change propagation in complex design". Journal of Mechanical Design, 126, p. 788.

[19] Grantham-Lough, K., Stone, R. B., and Tumer, I. Y., 2008. “Implementation Procedures for the Risk in Early Design (RED) Method". Journal of Industrial and Systems Engineering, 2(2), pp. 126-143.

[20] Stone, R. B., Tumer, I. Y., and Stock, M. E., 2006. “Linking product functionality to historical failures to improve failure analysis in design". Research in Engineering Design, 16(2), pp. 96-108.

[21] Clausing, D., 1994. Quality function deployment. MIT Press, Cambridge, MA.

[22] Jensen, D., Tumer, I. Y., and Kurtoglu, T., 2008. “Modeling the propagation of failures in software-driven hardware systems to enable risk-informed design". In Proceedings of the ASME International Mechanical Engineering Congress and Exposition.

[23] Jensen, D., Tumer, I. Y., and Kurtoglu, T., 2009. “Design of an Electrical Power System using a Functional Failure and Flow State Logic Reasoning Methodology". In Proceedings of the Prognostics and Health Management Society Conference.

[24] Jensen, D., Tumer, I. Y., and Kurtoglu, T., 2009. “Flow State Logic (FSL) for analysis of failure propagation in early design". In Proceedings of the ASME Design Engineering Technical Conferences; International Design Theory and Methodology Conference.

[25] Kurtoglu, T., Tumer, I. Y., and Jensen, D., 2010. “A Functional Failure Reasoning Methodology for Evaluation of Conceptual System Architectures". Research in Engineering Design, 21(4), p. 209.

[26] Padhke, M., 1989. Quality engineering using robust design. Prentice Hall, Englewood Cliffs, NJ.

[27] Sasajima, M., Kitamura, Y., Mitsuru, I., and Mizoguchi, R., 1996. “A representation language for behavior and function: Fbrl". Expert Systems with Applications, 10(3-4), pp. 471-479.

[28] Umeda, Y., Tomiyama, T., and Yoshikawa, H., 1992. “A design methodology for a self-maintenance machine based on functional redundancy". In International Conference on Design Theory and Methodology, Amer Society of Mechanical, p. 317.

[29] Umeda, Y., Tomiyama, T., Yoshikawa, H., and Shimomura, Y., 1994. “Using functional maintenance to improve fault tolerance". IEEE Expert: Intelligent Systems and Their Applications, 9(3), pp. 25-31.

[30] Leveson, N., 2011. “Engineering a safer world". MIT Press.

[31] Pereira, S., Lee, G., and Howard, J., 2006. A System-Theoretic Hazard Analysis Methodology for a Non-advocate Safety Assessment of the Ballistic Missile Defense System. Citeseer.

[32] Krus, D., and Grantham Lough, K., 2007. “Applying function-based failure propagation in conceptual design". In Proceedings of the ASME Design Engineering Technical Conferences; International Design Theory and Methodology Conference.

[33] Kurtoglu, T., Johnson, S., Barszcz, E., Johnson, J., and Robinson, P., 2008. “Integrating system health management into early design of aerospace systems using functional fault analysis". In Proc. of the International Conference on Prognostics and Heath Management, PHM'08.

[34] Coatanea, E., Nonsiri, S., Ritola, T., Tumer, I., and Jensen, D., 2011. “A framework for building dimensionless behavioral models to aid in function-based failure propagation analysis". Journal of Mechanical Design, 133, p. 121001.

[35] Sierla, S., Tumer, I., Papakonstantinou, N., Koskinen, K., and Jensen, D., 2012. “Early integration of safety to the mechatronic system design process by the functional failure identification and propagation framework". Mechatronics, p. doi:10.1016/j.mechatronics.2012.01.003.

[36] Tumer, I., and Smidts, C., 2010. “Integrated design and analysis of software-driven hardware systems". IEEE Transactions on Computers, 60, pp. 1072-1084.

[37] Hirtz, J., Stone, R., McAdams, D., Szykman, S., and Wood, K., 2002. “A Functional Basis for Engineering Design: Reconciling and Evolving Previous Efforts". Research in Engineering Design, 13, pp. 65-82.

[38] Kurtoglu, T., Tumer, I. Y., and Jensen, D. C., 2010. “A functional failure reasoning methodology for evaluation of conceptual system architectures". Research in Engineering Design, 21(4), pp. 209-234.

[39] Jensen, D. C., Bello, O., Hoyle, C., and Tumer, I. Y., 2014. “Reasoning about system-level failure behavior from large sets of function-based simulations". Artificial Intelligence for Engineering Design, Analysis and Manufacturing, 28(04), pp. 385-398.

[40] Ketkar, N. S., Holder, L. B. and Cook, D. J., 2005. “Subdue: Compression-based frequent pattern discovery in graph data”. In Proceedings of the 1st international workshop on open source data mining: frequent pattern mining implementations, pp. 71-76.

[41] Hollnagel, E. and Goteman, O., 2004. “The functional resonance accident model”. In Proceedings of Cognitive System Engineering in Process Plant, pp. 155-161.

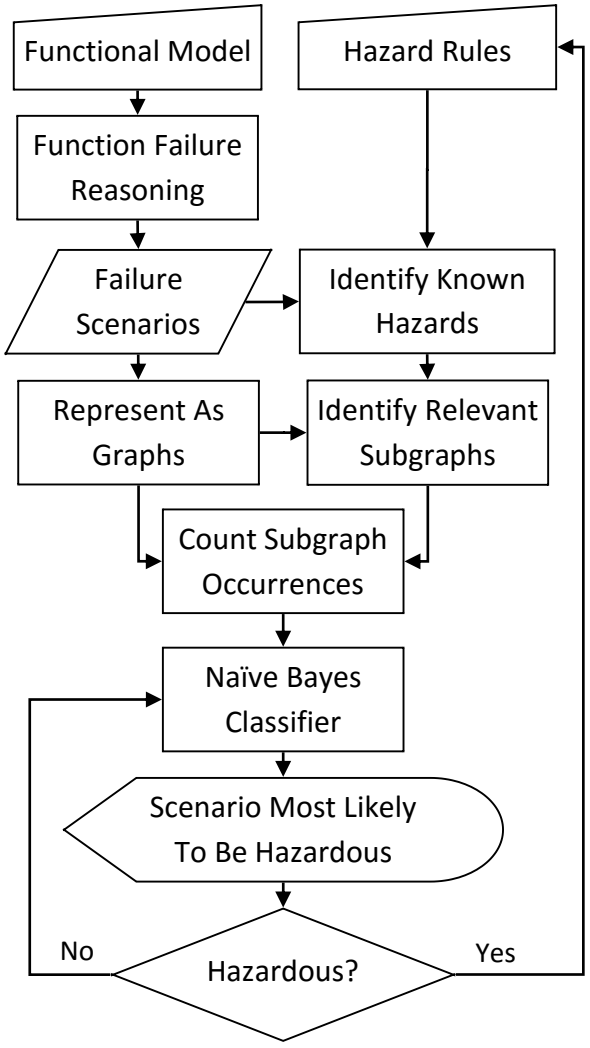

Figure 1: The Iterative Hazard Identification Process

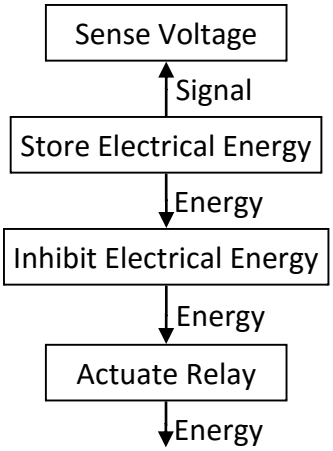

Figure 2: A partial functional model

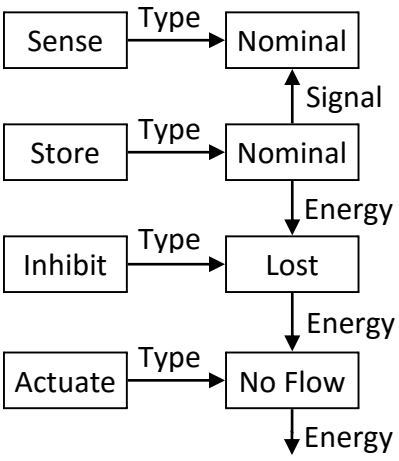

Figure 3: A single failure scenario

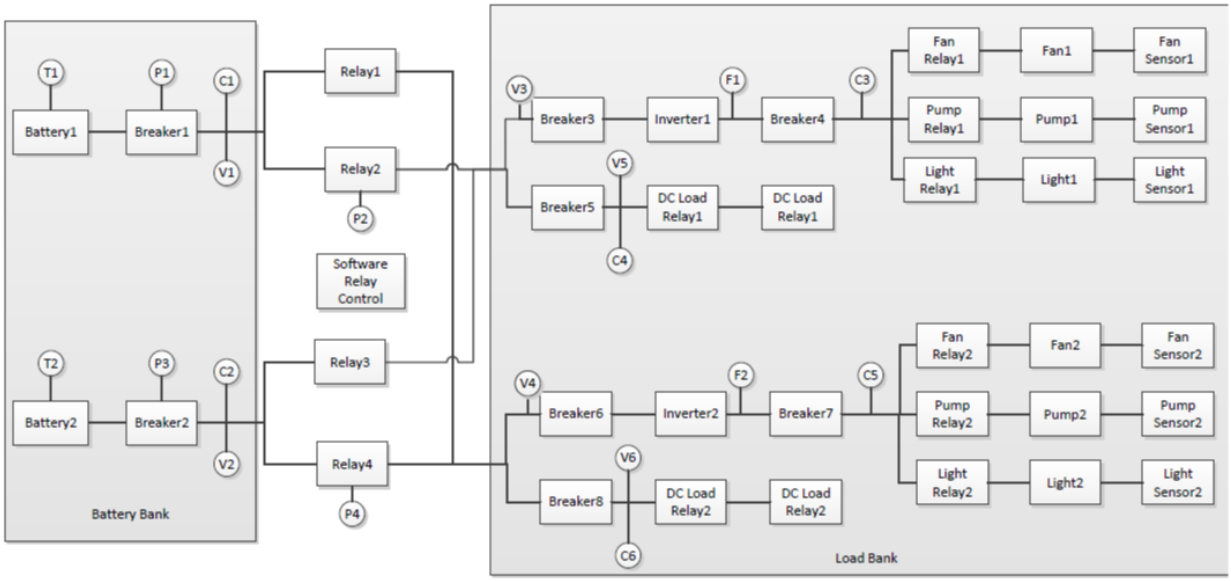

Figure 4: A block diagram of the EPS [39]