BSEE-2016-xxx (Draft)

October 25, 2016

Probabilistic Risk Assessment Procedures Guide for Offshore Applications (DRAFT)

Acknowledgments

The purpose of this Guide is to focus on analyses to support the kinds of decisions that need to be made in risk management of offshore facilities; but development of this Guide has benefited substantially from numerous earlier developments aimed at other industries, including NASA's "Probabilistic Risk Assessment Procedures Guide for NASA Managers and Practitioners," NASA's Systems Engineering Handbook, the ANS/IEEE PRA Procedures Guide, and others. Many of the contributors to those earlier documents are, in some sense, indirect contributors to this Guide as well.

Forward Work

This initial draft of the Guide has selected sections omitted and yet to be identified topics for forward work. The authors acknowledge that there are unique concerns and topics to the Offshore industry for which this Guide is being developed. Therefore, it is the intent of the authors to perform some additional work in the Offshore industry in order to identify and obtain experience on what is currently missing or needed in this Guide.

Contents

List of Figures

List of Tables

Acronyms and Abbreviations

Acronym	Description
ASME	American Society of Mechanical Engineering
BAST	Best Available and Safest Technology
BE	Basic Event
BOP	Blow Out Preventer
CCF	Common Cause Failure
DIM	Differential Importance Measure
DSMCS	Dependence-Suspect Minimal Cut Sets
ESD	Event Sequence Diagram
ET	Event Tree
FMEA	Failure Modes and Effects Analysis
FMECA	Failure Modes and Effects Criticality Analysis
F-N	Frequency Exceedance
F-V	Fussell-Vesely
FT	Fault Tree
HAZOP	Hazard and Operability Study
HAZID	Hazard Identification study
IBOP	Internal BOP
IE	Initiating Event
L x C	Likelihood and Consequence
LLP	Lowest Level Practicable
LMRP	Lower Marine Riser Package
MIT	Massachusetts Institute of Technology
MLD	Master Logic Diagram
MLE	Maximum Likelihood Estimate
MODU	Mobile Offshore Drilling Unit
MTBF	Mean Time Between Failure
MTTF	Mean Time To Failure
MTTR	Mean Time To Repair (Restore)
NRC	US Nuclear Regulatory Commission
PRA	Probabilistic Risk Assessment
PW	Prevention Worth
QRA	Quantitative Risk Assessment
RAW	Risk Achievement Worth
RISC	Risk-Informed Safety Case
ROV	Remotely Operated Vehicle
RRW	Risk Reduction Worth

1. Introduction

1.1 Purpose of This Guide

This Guide is intended to assist in the development of probabilistic risk assessment (PRA) of offshore drilling facilities, in order to support decision-making by Bureau of Safety and Environmental Enforcement (BSEE) and by the industry. This Guide is not a policy document, nor does it establish regulatory requirements; it discusses particular modeling techniques that have been found to be useful in a range of applications to decision-making about complex and high-hazard facilities. In order to motivate the approach taken in the remainder of this Guide, this section discusses what kinds of analysis support what kinds of decisions.

The point of departure for development of this guide is the NASA PRA Procedures Guide [1-1], which was itself derived from earlier PRA procedural guidance; but the present development has been informed by numerous other developments from within NASA, as well as work done for the Department of Energy and the Nuclear Regulatory Commission.

1.2 Risk and Risk Management

Partly because of the broad variety of contexts in which the concepts are applied, different definitions of risk continue to appear in the literature. Most of them are generally consistent with the idea that "risk is uncertainty about the future, viewed through the lens of a value structure (i.e., focusing on outcomes that would be considered "adverse")."

In the context of making decisions about complex, high-hazard systems, "risk" is usefully conceived as a set of triplets: failure scenarios, likelihoods of those scenarios, and their actual consequences [1-2]. There are good reasons to focus on these elements rather than focusing on simpler, higher-level quantities such as "expected consequences." Risk management involves prevention of (reduction of the frequency of) adverse scenarios (scenarios having undesirable consequences), and promotion of favorable scenarios (scenarios with favorable, or at least benign, outcomes). This requires understanding the elements of adverse scenarios so that they can be prevented, and understanding the elements of successful scenarios so that they can be promoted.

Even if the decision problem is simply to decide whether a facility is deemed "adequately safe," the level of assurance (the decision-maker's confidence) derivable from understanding scenarios far exceeds the level of assurance derivable from an abstract summary of expected consequences.

1.3 Scope of This Guide

Figure 1-1 (after [1-3]) illustrates a general process for safety analysis. The leftmost portion of the figure begins the process with recognition of the decision being supported, and an assessment of what technical results are needed to support that decision. The central portion of the figure notionally suggests a range of techniques for safety analysis, going from "qualitative" techniques (thought processes such as a Hazard and Operability Study (HAZOP), which identify accident potential) to "quantitative" techniques (modeling processes that generate and quantify scenarios, frequencies, and consequences).

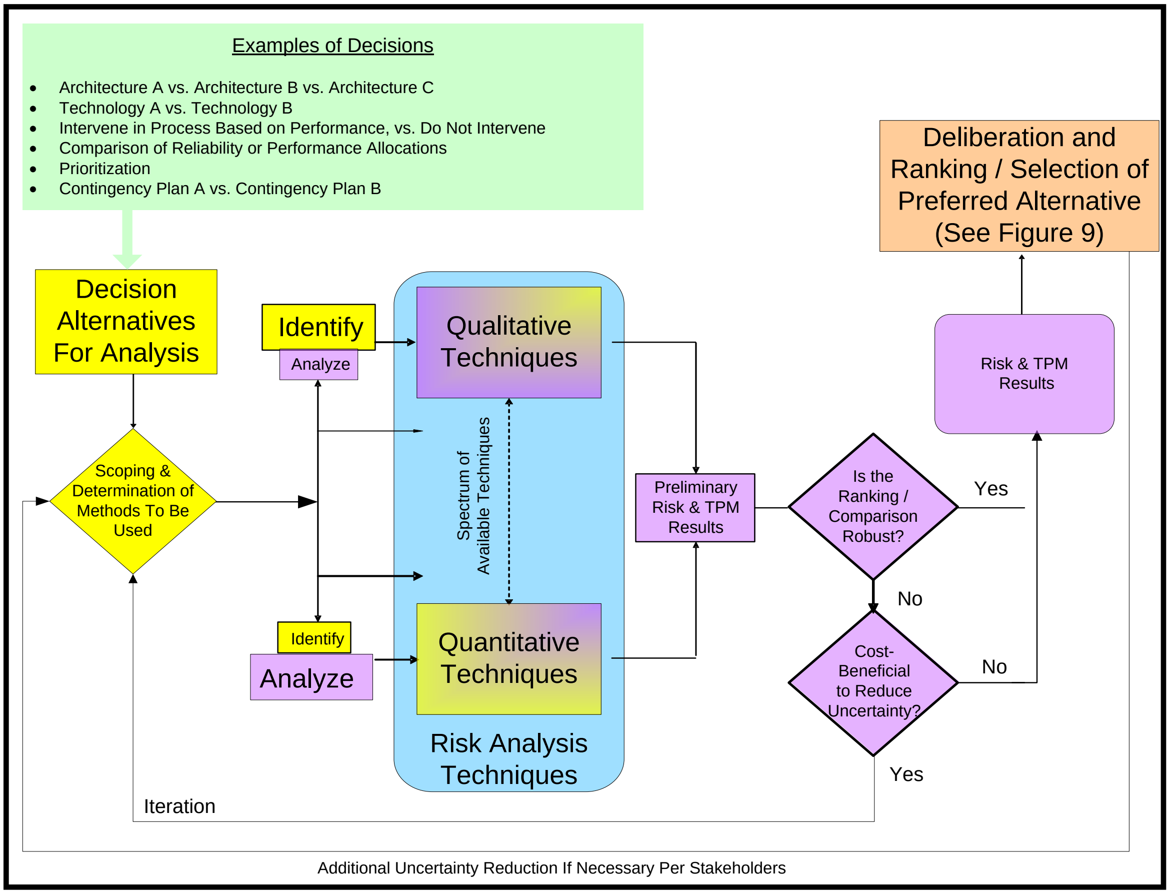
Examples of Decisions
Architecture A vs. Architecture B vs. Architecture C
Technology A vs. Technology B
Intervene in Process Based on Performance, vs. Do Not Intervene
Comparison of Reliability or Performance Allocations
Prioritization
Contingency Plan A vs. Contingency Plan B
Decision Alternatives For Analysis
Scoping & Determination of Methods To Be Used
Identify
Analyze
Qualitative Techniques
Spectrum of Available Techniques
Identify
Analyze
Quantitative Techniques
Risk Analysis Techniques
Preliminary Risk & TPM Results
Is the Ranking / Comparison Robust?
Yes
No
Cost-Beneficial to Reduce Uncertainty?
No
Yes
Risk & TPM Results
Deliberation and Ranking / Selection of Preferred Alternative (See Figure 9)
Iteration
Additional Uncertainty Reduction If Necessary Per Stakeholders

Figure 1- 1. Process for Safety Analysis

Different situations will call for a different mix of techniques. It is not always clear *a priori* what techniques are appropriate in a given situation; correspondingly, in the rightmost portion of the figure, the current state of knowledge (after analysis done to date) is assessed to determine whether there is a need to loop back and do more analysis (or get more information) in order to support the current decision.

Broadly speaking, “quantitative” techniques such as fault tree analysis are techniques that lead to (a) an explicit scenario set, (b) quantification of the likelihoods of those scenarios, and (c) analysis of the consequences of those scenarios (in short, analysis of the “triplets” discussed above). Calling the other techniques “qualitative” does not mean that they are applied absolutely without regard to probability; in fact, it is extremely difficult to absolutely decouple safety thinking from probability. Rather, the term “qualitative” is shorthand for “thought processes that help us to identify accident potential, without explicitly generating and quantifying a comprehensive scenario set.”

This Guide is focused on the “quantitative” end of the above-described analysis spectrum, using selected “qualitative” techniques as a front end to the quantitative analysis, in order to help us think appropriately about what we need to analyze in more detail.

1.4 Probabilistic Risk Assessment

1.4.1 What Are We Going to Get Out of a PRA, and How Would We Use It?

Based on modeling scenarios, frequencies, and consequences, PRA quantifies "risk metrics." The term "risk metric" refers to probabilistic performance measures that might appear in a decision model: such things as the frequency or probability of adverse consequences of a specific magnitude, or perhaps expected consequences. Figures of merit such as "system failure probability" can be used as risk metrics, but the phrase "risk metric" ordinarily suggests a higher-level, more consequence-oriented figure of merit, such as "spills of a certain magnitude."

In order to support resource allocation from a risk point of view (for a licensee or regulator), it is necessary to evaluate a comprehensive set of scenarios. The set of scenarios may need to include events that are more severe than those considered during design, and more success paths than were explicitly factored into the design. Additionally, system performance must be evaluated realistically. In order to support resource allocation decisions, the point is not usually to establish a bound on system capability or reliability, but rather to *quantify* capability and reliability (to characterize them realistically). In other words, risk-informed resource allocation requires identification and realistic quantification of all risk-significant scenarios, where "risk-significant" depends on the context of the evaluation.

In all but the simplest cases, decision support requires that uncertainty be addressed. Because risk analysis frequently needs to address severe outcomes of complex scenarios, and because these scenarios are too infrequent for us to be able to calibrate our models from experience, uncertainties may be highly significant. These uncertainties need to be reflected in the decision model, not only because they may influence the decision, but also because it is important to understand which of the uncertainties that strongly affect the decision outcome are potentially reducible through testing or research.

PRA is needed (and the effort is justified) when decisions need to be made that involve high stakes in a complex situation, as in a high-hazard mission with critical functions being performed by complex systems. Intelligent resource allocation depends on a good risk model; even programmatic research decisions need to be informed by a state-of-knowledge risk model. (Allocating resources to research programs needs to be informed by insight into which uncertainties' resolution offers the greatest payback.) Developing a comprehensive scenario set to provide decision makers with the best informed picture of threats and mitigation opportunities is a special challenge, and systematic methods are needed for development and quantification of such a model. Those methods are the subject of this guide.

1.4.2 Use of PRA in the Formulation of a Risk-Informed Safety Case (RISC)

The above discussion has been carried out with emphasis on the role of PRA in assessing system adequacy, especially with regard to selection of design features. This sort of application began even before "safety goals" were widely discussed. Increasingly, risk managers need to argue that system designs satisfy explicit risk thresholds; nowadays, even if there is no absolute regulatory or policy requirement, the promulgation of safety goals and thresholds creates an expectation that goals and thresholds will be addressed in the course of safety-related decision-making. This creates an issue for PRA, because in general, it is impossible to "prove" that the level of risk associated with a complex, real-world system is below a given decision threshold.

Partly because PRA results cannot be "proven," a "Risk-Informed Safety Case" (RISC) [1-4] is desirable. The RISC marshals evidence (tests, analysis, operating experience) and commitments to adhere to specific manufacturing and operating practices in order to assure that PRA assumptions, including the performance and reliability parameters credited in the PRA, are fulfilled. Among the commitments needed to justify confidence in the safety of the system is a commitment to analyze operating experience on an ongoing basis, including "near misses," in order to improve operations, improve the risk models, and build additional confidence in the models' completeness. This is not the same as "proving" that the PRA results are correct, but it is the best proxy for safety that can be obtained.

These matters are discussed further in the following sections of this Guide. The present discussion is simply to motivate the emphases placed in Section 2's treatments of the risk analysis techniques.

1.4.3 Characterization of Safety Margin

For purposes of making safety decisions (whether we need to modify our design or operating practices, whether our system risk is Lowest Level Practicable (LLP), whether we have reasonable assurance of adequate protection), it is useful to analyze system performance in terms of "margin," and moreover to do this in a "risk-informed" way. What attributes does a model need, in order to support a risk-informed assessment of "margin?" What do we mean by "risk-informed?"

1.4.4 Background on "Risk-Informed"

The phrase "risk-informed" originated in US Nuclear Regulatory Commission (NRC) practice. The NRC Web Site [5] offers the following definitions related to "Risk-Informed:"

- Risk-Informed Decision-Making: An approach to regulatory decision-making, in which insights from probabilistic risk assessment are considered with other engineering insights.
- Risk-Informed Regulation: An approach to regulation taken by the NRC, which incorporates an assessment of safety significance or relative risk. This approach ensures that the regulatory burden imposed by an individual regulation or process is appropriate to its importance in protecting the health and safety of the public and the environment. For additional details, see Risk Assessment in Regulation and the Fact Sheet on Nuclear Reactor Risk.

One important consideration is whether a comprehensive scenario set is modelled with a view to quantitative analysis of decision alternatives, as opposed to pass-fail compliance with prescriptive requirements derived from surrogates formulated by engineering judgment (such as large-break Loss of Coolant Accident (LOCA), an early focus of Atomic Energy Commission thinking about the regulation of light-water reactors in the US). If you are not modeling a scenario set in a way that supports saying what's important and what isn't, you're not being risk-informed.

The phrase "risk-informed" is now widely used to describe a certain thought process. It appears to have originated with NRC Chairman Jackson in the early or mid-1990s. During the 1980s and early 90s, many, many papers were being written on the subject of "risk-*based*" regulation (emphasis added). The context of those papers was deciding whether regulatory burden could legitimately be reduced (or, in principle, whether it needed to be tightened up), based on risk model results. Often, risk model results suggest that burden can be reduced; but then, as now, there was a lot of opposition to reducing burden significantly, based on PRA as the primary justification. For the traditionalists, "risk-*based*" was a non-starter. Enter Chairman Jackson: our decision-making will not be risk-*based*, but it will be risk-*informed*, meaning that we will use risk information as one of several inputs to a decision process, other inputs being things like “defense in depth” and “safety margins,” and addressing a broad range of issues of diverse kinds, and not just compliance with regulations.

The concept of “margin” has evolved in recent years. Originally, the general idea was that a system’s capacity to withstand expected loads should be designed with some leeway, recognizing that things may be a bit worse than anticipated, and this excess capacity could be specified either in terms of a point value of extra capacity, or a point value of a safety factor. Recent years [1-6, 1-7] have seen increased appreciation of the usefulness of viewing margin probabilistically, as summarized in a fairly recent doctoral thesis [1-7]:

- Safety margin is the difference between a characteristic value of the capacity and a characteristic value of the load.
- While [this measure] provides a first approximation of functional reliability, ranking different systems on safety margins alone can lead to erroneous results. The knowledge of the distance from failure in terms of safety margins is not sufficient to evaluate the risk of a system; *the breadth of the uncertain distribution* [emphasis added] is the other important part of the assessment.

The “breadth of the uncertain distribution” is suggested notionally in the figure below:

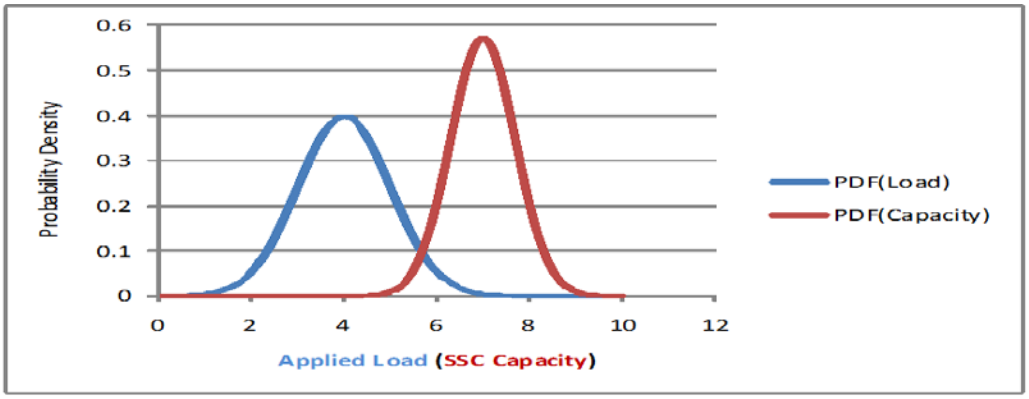

Figure 1- 2. Breadth of Uncertain Distribution

This figure shows an uncertain applied load (such as a pressure) together with the uncertain “capacity” of a component to survive that load (in this case, the pressure-retaining capability). What matters in the risk analysis is whether the pressure will exceed the actual pressure-retaining capability, and the point of the figure is that if both of these are uncertain, a naïve idea of “margin” such as the distance between the two modes is inadequate. We need to understand the probability that load exceeds capacity.

Even this load-capacity idea is oversimplified for some purposes, because it is stated above as if the two can be evaluated independently. In some cases, they cannot. Consider a notional example in which high pressure and low pressure-retaining capability are related due to high temperature; in such a case, a calculation based on the simple figure above would underestimate failure probability. In such a case, we must resort to a more simulation-based approach to risk analysis; this is discussed in Section 2.

Safety margin characterization is “risk-informed” if it is based on the following:

- An issue space is formulated, implicitly defining a class of scenarios to be analyzed probabilistically and the figures of merit[1] to be evaluated probabilistically, “margin” then to be analyzed in terms of those figures of merit in those scenarios.
    - Aleatory[2] variables are identified and assigned appropriate distributions.
    - The state of knowledge within that issue space is delineated in terms of state-of-knowledge probability distributions on uncertain variables, or perhaps probability bounds analysis.[3]
- The scenario set is analyzed in sufficient detail (with sufficient coverage of the issue space) to
    - Characterize margin in the relevant figures of merit, including the comparison of absolute margin with variability and uncertainty;
    - Understand the significance of variability and uncertainty separately;
    - Understand the probability of “failure” (the probabilistic weight of scenarios having zero or negative margin) at least semi quantitatively;
    - Understand the main drivers (particular conditions under which margin is high or low), pointing to
        - Failure modes or initial conditions, control of which would increase margin,
        - Information that needs to be obtained in order to reduce uncertainty.

This definition does not address whether the analysis is good or poor: only whether it is structured to culminate in a probabilistic characterization of “margin” in a given issue space.

1.4.5 Summary

The essence of “risk-informed” is to create a basis for resource allocation (by licensee and by regulator) that does the best job we know how to do, consistent with our state of knowledge and institutional constraints (such as limitations on the kinds of analysis we can afford). In order to be risk-informed, the analysis must be geared to supporting conclusions about which scenarios are more important than others, and how much more important, and how beneficial (or how justifiable) it would be to add preventive or mitigative measures beyond what’s already there. Modeling to support risk-informed decision-making will

[1] Typically, these will be performance metrics in terms of which system success and system failure can be defined.

[2] “Aleatory” uncertainty refers to the variability in outcomes from one trial to the next: the outcome of a roll of honest dice is uncertain, and this uncertainty is aleatory. The term “aleatory” is contrasted with “epistemic,” which refers to limitations of our state of knowledge. If we are not sure what fraction of the time a given coin will yield “heads,” this is a kind of uncertainty that we could, in principle, reduce by carrying out experiments; this kind of uncertainty is “epistemic.” These concepts are discussed in Section 2.2.

[3] “Probability bounds analysis” [1-8] is the name given to an approach to propagating uncertainty that works with intervals (upper and lower bounds) on the values of the uncertain variables, rather than sampling from explicit probability density functions of those variables.

tend to have the following attributes:

- It will comprehensively analyze representative scenarios within the slice of event space that is probabilistically significant for the decision.
- It will make little or no use of bounding (worst-case) arguments, and will instead strive for "realism" embedded within an honest treatment of uncertainty.
- It will comprehensively analyze the variability (the aleatory uncertainty) in scenario outcomes.
- It will methodically analyze the implications for the decision of the limitations of the current state of knowledge.

This, in a nutshell, is what we can get out of PRA, and how we use it. The methods and tools discussed in this guide are aimed at accomplishing these things.

1.5References

1-1 Probabilistic Risk Assessment Procedures Guide for NASA Managers and Practitioners, NASA, *NASA/SP-2011-3421*, Washington, DC. 2011.

1-2 Kaplan and Garrick, "*On the Quantitative Definition of Risk.*" Journal Risk Analysis, Vol. 1, Mar. 1981.

1-3 NASA Systems Engineering Handbook, NASA/SP-2007-6105 Rev. 1 (NASA, 2007) [downloadable from the NASA web site].

1-4 NASA System Safety Handbook Vol. 1, System Safety Framework and Concepts for Implementation, NASA/SP-2010-580 Version 1.0 (NASA, 2011) [downloadable from the NASA web site].

1-5 See the Glossary on the NRC web site, http://www.nrc.gov/reading-rm/basic-ref/glossary.html

1-6 Shaw L., Shooman M. L., and Schatz R., Time-Dependent Stress-Strength Models for Non-Electrical and Electrical Systems, Proceedings of Annual Reliability and Maintainability Symposium, pp. 186-197, IEEE Press, New York, 1973.

1-7 L. Pagani, "On the Quantification of Safety Margins" (MIT, 2004).

1-8 Verification and Validation in Scientific Computing, W. L. Oberkampf and C. J. Roy (Cambridge University Press, 2010).

2. Risk Analysis Techniques

2.1 Qualitative Risk Assessment Techniques

As noted in Section 1, the term "qualitative risk assessment techniques" is here taken to mean "thought processes that help us to identify accident potential, without explicitly generating and quantifying a comprehensive scenario set." The terms "quantitative" and "qualitative" are not perfect descriptors of the distinction that we are trying to make; refraining from using numbers in fault-tree analysis does not eliminate its capability to generate, and even notionally rank, a comprehensive scenario set (based on the order of the minimal cut sets. By "qualitative," we mean techniques such as HAZOP, which entail a great deal of thought, but do not typically involve explicit construction of a risk-model representation of a facility.

There are multiple reasons to consider qualitative techniques:

- Sometimes, qualitative techniques are adequate, by themselves, to support the current decision (for example, "design evaluation" can be a category of decision).
- In practice, quantitative techniques need to start out with the insights provided by the qualitative methods: for example, identification and grouping of initiating events, and the development of event tree structure, need to be informed by insights from techniques such as HAZOP.

For the latter reason, discussion of some qualitative techniques is provided in Section 2.2. The present section mentions a few qualitative techniques and suggests how to decide when they are sufficient. However, it is not the purpose of this section to provide detailed procedural guidance on those techniques as stand-alone applications. First, such a discussion would be beyond the scope of this guide; second, abundant material of that kind already exists elsewhere.

Accordingly, the following subsections will mention selected tools with a view to showing how they address the above two considerations (when they suffice for decision-making, and how they fit into quantitative modeling).

2.1.1. Comparison of Selected Qualitative Risk Assessment Techniques

2.1.1.1. Hazard and Operability (HAZOP) Analysis

HAZOPs are performed in a group setting where a facilitator leads a technically diverse group of experts through an exercise to identify hazards related to equipment or operations of a given system in a given operating mode. The design intent in each operating mode needs to have been specified in sufficient detail to support a sensible discussion of system behavior: in particular, nominal values need to have been specified for all important system parameters. The HAZOP discussion is then cued to analyze the system considering "deviations" of key parameters in one node at a time, based on applying "guide words" (e.g., "high," "low") to each parameter (e.g., "flow") characterizing each node (e.g. high flow in node 32, low flow in node 32, etc.). For each such deviation, the group brainstorms possible causes and possible consequences of each cause, and then may consider other factors relevant to the decision context, including possible recommendations for design changes. This discussion implicitly addresses classes of scenarios, identifying them in terms of physical behaviors, many of which could be caused by any of several different component states (good or failed), and some of which could arise even if no components are nominally “failed.” A notional example of part of a HAZOP table is shown below.

Table 2-1. Example of Hazard and Operability Analysis

HAZOP of Drilling Rig’s Mud System					
Node	**Deviation**	**Cause**	**Consequence**	**Mitigation**	**Risk Ranking**
1	Low mud weight	Improper materials	Potential underbalance condition leading to well kick	Proper vendor selection for materials Inspection of materials before use Verification of analysis Training	Likelihood - 3 Medium - 3
			…		
		Incorrect mud weight analysis			
		Human error			
2	High mud weight	….		….	….
….	…	…..	….	….	….

2.1.1.2. Failure Modes and Effects Analysis (FMEA)

FMEA is a component based technique that breaks down a system into mechanical and electrical components and postulates how each component can fail and the effect the failure has locally and on the overall system. The result is in tabular form and documents each component, the ways it can fail, and the effects on the system. Analyses also typically include how the failure can be detected and mitigations in place to prevent or lessen negative effects. An FMEA may be extended to become a Failure Modes, Effects, and Criticality Analysis (FMECA) by adding in an evaluation of the likelihood and consequences. An example of part of an FMEA is shown in Table 2-2.

Table 2-2. Example of Failure Modes and Effects Analysis

Component				Effects				
System	Component	Identifier	Failure Mode	Local Effect	Next Level Effect	System Level Effect	Detection Method	Mitigation
BOP	Pipe Ram	Upper Pipe Ram	Fails to close	Pipe ram will not close	If a well kick occurs, the annulus will not be sealed off	If a well kick occurs, formation fluid will rise past the BOP and potentially reach the drill floor	If a well kick occurs return mud flow will be high and pressure will not rise in the BOP	Redundant pipe rams and the blind shear ram
BOP	Pipe Ram	Upper Pipe Ram	Fails to seal against drill pipe	….	….	….	….	….
….	….	….	….					

2.1.2 Other Decision Aids

2.1.2.1 Bowtie Diagrams

Bowtie analysis results in a graphical representation of a class of scenarios that helps decision-makers reason appropriately.

- The middle of the bowtie represents a hazardous condition that results when control of a facility is lost (for example, an explosion in a particular location).
- The left hand side develops causes that can lead to the hazardous condition and the controls in place to prevent its occurrence.
- The controls (including physical barriers) are placed between the cause and hazard showing the failures that must occur for the hazard to occur.
- The right-hand side portrays scenarios ensuing from the occurrence of the hazardous condition, culminating in consequences on the far right. The scenarios on the right are specified in terms of the functions (including physical barriers) that limit or mitigate the consequences potentially resulting from the hazard.
- Each complete left-to-right path through a bowtie is a representation of a hazardous scenario to be considered. An example of part of a bowtie is shown in Figure 2-1.

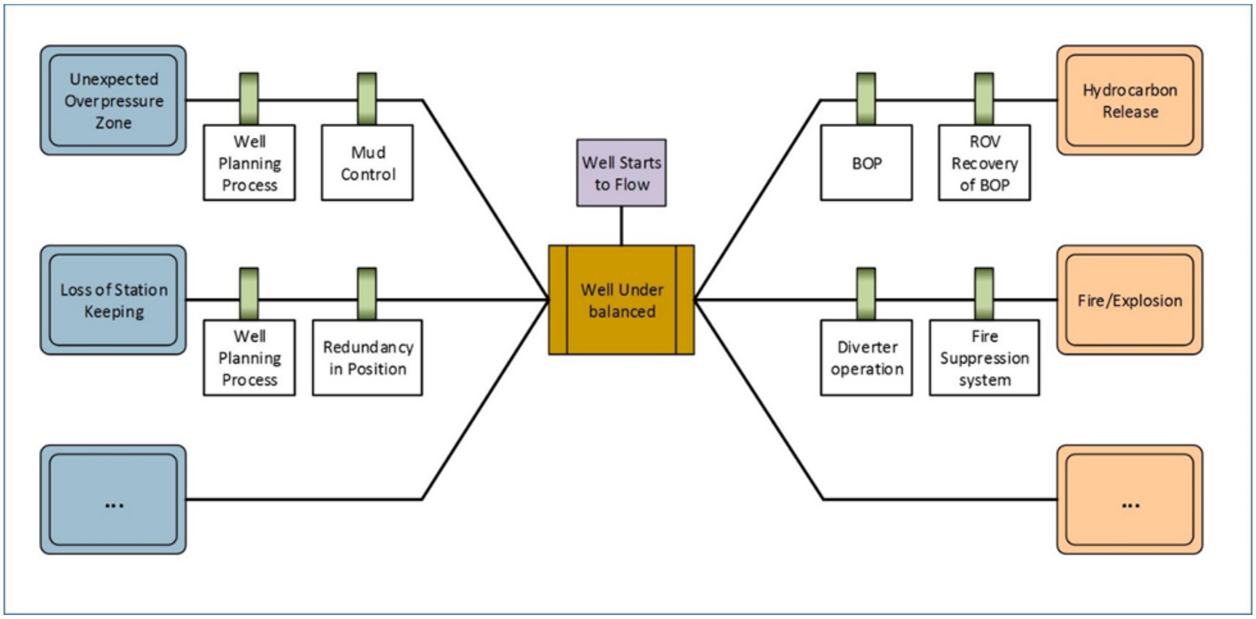

Figure 2- 1. Example of Bowtie Analysis Diagram

2.1.2.2 Risk Matrices

A "risk matrix" is commonly used to communicate perspective on the significance of particular "risks" (scenarios, or classes of scenarios having something in common). Notionally, the matrix elements correspond to discrete categories of "frequency" and "consequence" as illustrated below; individual scenarios are assigned a likelihood and consequence level and placed on the picture, as illustrated by the numbered circular symbols. Symbol #1 refers to a "risk" having both low consequences and low likelihood. Its placement in the "green" region is a way of saying that it is no real threat; either the threat is inherently minimal, or a previously identified real threat has been successfully controlled by prevention or mitigation. Symbol #3 refers to a risk having high likelihood and high consequences, and its placement in the "red" region is a way of saying that this one needs attention; it may be a showstopper. Symbol #2 refers to a risk that is in between.

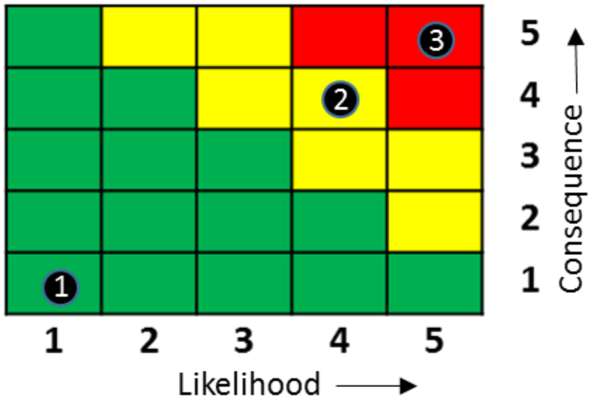

Figure 2-2. Typical Qualitative Risk Matrix

2.1.3 Strengths and Weaknesses of Qualitative Risk Assessment Techniques

The qualitative risk assessment techniques described above, as well as others not mentioned here, provide a systematic approach to evaluating risk, albeit with different focuses for each method. HAZOP, FMEA, and other methods in this general category promote completeness, which is perhaps the single most critical issue affecting the performance and application of risk modeling. FMEA promotes completeness by considering (in principle) "all" failure modes for "all" components; HAZOP does this by considering (in principle) "all" physical deviations in "all" nodes. If design intent is properly specified, then anything that would be considered an accident must represent a deviation from design intent; so, by considering "all" deviations, HAZOP creates at least an opportunity for the group to identify any accident potential that is reasonably foreseeable in the context of any given deviation. Moreover, if a similar system has some operating history, we may have a sense of the likelihood of the deviations identified, and we may have at least some perspective on their consequences.

However, neither FMEA nor HAZOP is particularly well suited to identification or ranking of scenarios that involve multiple failures, or combinations of failure events with off-normal initial conditions. This is not a fault of the thought processes involved; rather, it is because for a system of even moderate complexity, it is impractical for humans to evaluate multiple-failure scenarios without constructing an explicit scenario model and processing that model by computer. Except for very simple systems, it is difficult to determine manually whether a system is single-failure-proof. In principle, FMEA tries to capture the cascading implications of each postulated single failure, but in practice, it is difficult to propagate such implications through the system without some sort of computer aid.

Moreover, it turns out that on-the-fly assessments of probability are unreliable, and it is correspondingly difficult to estimate the likelihood of even moderately complex scenarios. The first large-scale quantitative risk analysis, the Reactor Safety Study, indicated that risk from light water reactors was dominated not by the sort of postulated pipe rupture event that had dominated safety thought for generations, but rather by events initiated by much more mundane, almost every day, deviations that are less severe but still challenge safety functions, and need to be dealt with appropriately. The high relative *frequency* of those challenges means that the *reliability of the mitigating systems* must be correspondingly high.

2.1.4 When Should PRA Be Considered?

This Guide is not a policy document, nor is it meant to prescribe to facility operators precisely when they need PRA. The present subsection is meant to help management decide what sort of analysis result they need, based on what kind of decision is being made, and what sorts of risks may be in play.

PRA is essentially a high-end risk analysis for supporting certain kinds of decisions. Quite generally, the choice of techniques in a given decision analysis needs to be made in light of the stakes associated with the decision, the complexity involved in analyzing the possible outcomes, the uncertainties, the diversity of stakeholders involved, and perhaps other considerations. By "stakes," we mean the magnitude of the consequences of accidents: fatalities or health effects, adverse environmental effects, significant expense, perhaps other adverse effects on the operating corporation. But high stakes alone may not justify the formulation of a detailed risk model. Selection of a particular course of action may need to be based on strong evidence of low accident likelihood, but if we can get that evidence without a risk model, then we don't need the model.

As an example: if operation of a facility could result in severe safety or environmental consequences and involves new technology or new environments, quantitative risk assessment such as PRA should be considered, because that situation involves high stakes, uncertainty, and (potentially) complexity. Generalizing from that example, questions such as the following can be used to help determine when a PRA should be considered:

- Is the facility design complex?
- Could the consequence of failure of the facility or operation result in higher human or environmental consequences than similar facilities or operations?
- Does the location of the facility or operation magnify the potential consequences of failure? For example, is the location in an area that is fragile, or contains a vulnerable population?
- Have the potential recovery or mitigation measures for the proposed facility or operation been proven in similar environmental situations?
- Has the facility or equipment been used in the proposed type of operation before?
    - How much experience has been gained?
    - What are the outcomes of the use?
- Is the proposed facility or operation in a new or extremely challenging environment?
- Are there any new hazards associated with the facility or operation when compared to facilities or operations performing similar jobs?
- If the facility or operation is being applied in a similar environment with similar consequences to existing facilities or operations, are there any new aspects such as material, equipment layouts, types of equipment, or positioning systems that are untested?

More generally: going back to Figure 1-1 of Section 1: in a situation with high stakes, complexity, uncertainty, and so on, it is unlikely that a qualitative model result will pass the “robustness” test in the decision diamond on the right of the figure. Correspondingly, the analysts will be directed to loop back through the figure, and choose tools that furnish the results needed to get past the “robustness” test.

2.2 Quantitative Scenario Modeling: Probabilistic Risk Assessment

When the decision has been made that the qualitative techniques do not offer the assurance necessary to make a decision, quantitative techniques (i.e. Probabilistic Risk Assessment) should be considered. The PRA ultimately presents a set of scenarios, frequencies, and associated consequences, developed in such a way as to inform decisions regarding the allocation of resources to accident prevention or mitigation. The implication of the analysis could be a change in design or operational practice, or could be a finding that the design is acceptable as is. Decision support in general requires quantification of uncertainty, and this is understood to be part of modeling and quantification.

For simplicity, the discussion below will be carried out as if the process of PRA model development were a once-through process. But this is not the case. In general, a significant amount of iteration will occur. The process of development is steered by whether the model is adequate for purposes of the decision being supported. Examples of consistency checks include:

- Have we included events that are probabilistically significant relative to the current bottom line (or to other events that we have decided to include)?
- Have we parsed events sufficiently to be able to quantify them accurately?

- Are events parsed down to a level at which we can reasonably treat them as if they were independent?

2.2.1 Elements of a PRA

This subsection discusses the elements of a PRA logic model. Major elements of the logic are introduced and briefly described; each is then illustrated with respect to simplified examples. The examples emphasize the logic-based (event tree/fault tree) modeling approach; however, some of the concepts described in this section are also applicable to other modeling approaches such as simulation as discussed in Section 2.3.

A scenario contains an Initiating Event (IE) and, usually, one or more pivotal events leading to an end state. As modeled in most PRAs, an IE is a perturbation that requires some kind of response from the crew or one or more systems. Note that for an IE to occur, there may need to be associated enabling event(s) that exist (e.g., for a fire IE to occur, there would need to be combustible material present). The pivotal events in a scenario include successes or failures of responses to the IE, or possibly the occurrence or non-occurrence of external conditions or key phenomena. Then, the scenario end state(s) are defined according to the decisions being supported by the analysis, in terms of the kind and severity of consequences, ranging from completely successful outcomes to losses of various kinds.

The first major step in logic model development is to determine the boundaries of the analysis. First, based on the goals of the analysis and decisions to be made, what end state(s) are of interest? Examples include:

- Loss of life or injury to personnel;
- Damage to the environment;
- Damage to, or loss of, equipment or property (including facilities and public properties); and
- Unexpected or collateral damage.

Determination of which end states will be analyzed will in turn determine the IEs and critical functions that must be included in the analysis.

In addition to the end state(s), the boundaries of the analysis, in many cases, would define what a successful end state would be. For instance, if the end state of interest was an uncontrolled release of hydrocarbons to the environment during exploration drilling, the success state may be defined in different ways depending on the goals of the analysis. If the goal is to evaluate a the likelihood of an accident, the success end state may be defined as successful control of the well by the Blow Out Preventer (BOP). If the goal is to evaluate the likelihood of a release as a function of the magnitude of release, considerations beyond the BOP must be taken into account such as ROV intervention, well capping, etc. and success becomes killing the well rather than successful isolation of the well by the BOP alone.

The next step involves determining what perturbations to the process, or IEs, present a challenge that could lead to the end state(s) of interest. There may be many IEs, some of which may be grouped together because the response is the same or very similar (e.g. a well kick due to surge while tripping in and a well kick due to an unexpected overpressure zone), or the IEs may have different responses (e.g. a well kick due to an unexpected overpressure zone and an inadvertent LMRP disconnect). Determination of the IEs will further determine the critical functions necessary to achieve a successful end state through development of event sequence diagrams (ESDs) that detail the response to the initiating event. Identification of IEs and ESDs / critical functions are discussed in more detail in Sections 2.2.2 and 2.2.3 respectively.

Step 3 is building the event tree(s) that develop specific accident sequences leading to the end state(s) of interest and is used, in conjunction with fault trees to quantify the frequency of each end state. One event tree is developed for each IE or group of IEs. The graphical event tree starts with the IE which is followed by a number of pivotal events determined through the accident progression / critical function assessment in Step 2. Each of the pivotal events have a potential success or failure path (although in some cases more than a binary state is possible), and are usually ordered as a time sequence of the response to the IE. A detailed discussion of how event trees are built and function is found in Section 2.2.4.

Pivotal event development is Step 4. This involves development of fault trees for each of the pivotal events in the event trees. The fault trees are models that start with a "Top Event" that is a failure or condition, and develop ways in which that event can happen, expressed in terms of "basic events." There can be many basic events (the lowest level in the fault tree) and very many combinations of basic events that can cause the Top Event. The top event model may comprise systems, human actions, environmental conditions, etc. The basic event level is where the probabilistic data are used for quantification. Fault tree development and quantification is discussed in Section 2.2.5.

With the development of the event trees and supporting fault trees, the logic model is completed. Quantification requires the development of data to populate the logic model and is discussed in Section 2.2.6.

2.2.2 Initiating Event Development

One of the first modeling issues that must be resolved in performing a PRA is the identification of accident scenarios that are related to the analysis goals. This modeling of "what can go wrong?" follows the systematic identification of accident initial causes, called initiating events, grouping of individual causes into like categories, and subsequent quantification of their likelihood. Initiating events may lead directly to undesirable outcomes or more typically require additional failures, equipment and/or human, prior to reaching a negative consequence.

The identification of initiators can come from a variety of techniques, including those discussed in Section 2.1 (e.g. HAZOPs, HAZIDs, etc.). Precursor events may also suggest the types and frequencies of applicable upsets. PRA analysts traditionally deduce initiating events through the development of a Master Logic Diagrams (MLD).

The MLD is a deductive method analogous to a fault tree. The top event of an MLD is a type of challenge to facility safety. The top levels are defined by functional events and/or external events (e.g. environment), and successive levels of the MLD are developed until the effect of the failure or event is the same as the block it feeds in to. The goal is not only to support identification of a comprehensive set of IEs, but also to group them according to the challenges that they pose (the responses that are required as a result of their occurrences). IEs that are completely equivalent in the challenges that they pose, including their effects on subsequent pivotal events, are equivalent in the risk model.

A useful starting point for identification of IEs in a MLD is a specification of "normal" operation in terms of (a) the nominal values of a suitably chosen set of physical variables and (b) the envelope in this variable space outside of which an IE would be deemed to have occurred. An example of this could be the mud return flow. There is an expected value for the mud return flow, and a deviation (increase) by a certain amount would not be "normal" and indicates a well kick may have occurred. A comprehensive set of process deviations can then be identified, and causes for each of these can be addressed in a systematic way.

Figure 2-3 shows an example of a MLD that might be used to identify initiating events (not exhaustive) related to upsets resulting in an uncontrolled release of hydrocarbons during normal drilling operations.

For this example, the end state of an uncontrolled hydrocarbon release during normal drilling operations is the top block. The next level down would be functions that either normally are barriers that prevent the end state from occurring like mud control, or functions that could induce a perturbation that could lead to the end state such as tripping in/out. In addition, a category related to environmental conditions is included since they may be somewhat unpredictable and induce conditions leading to the end state. Environment in this case is generalized and is broken down into geological conditions and weather/sea conditions. The level of the MLD in blue represents IEs that are challenges to the end state.

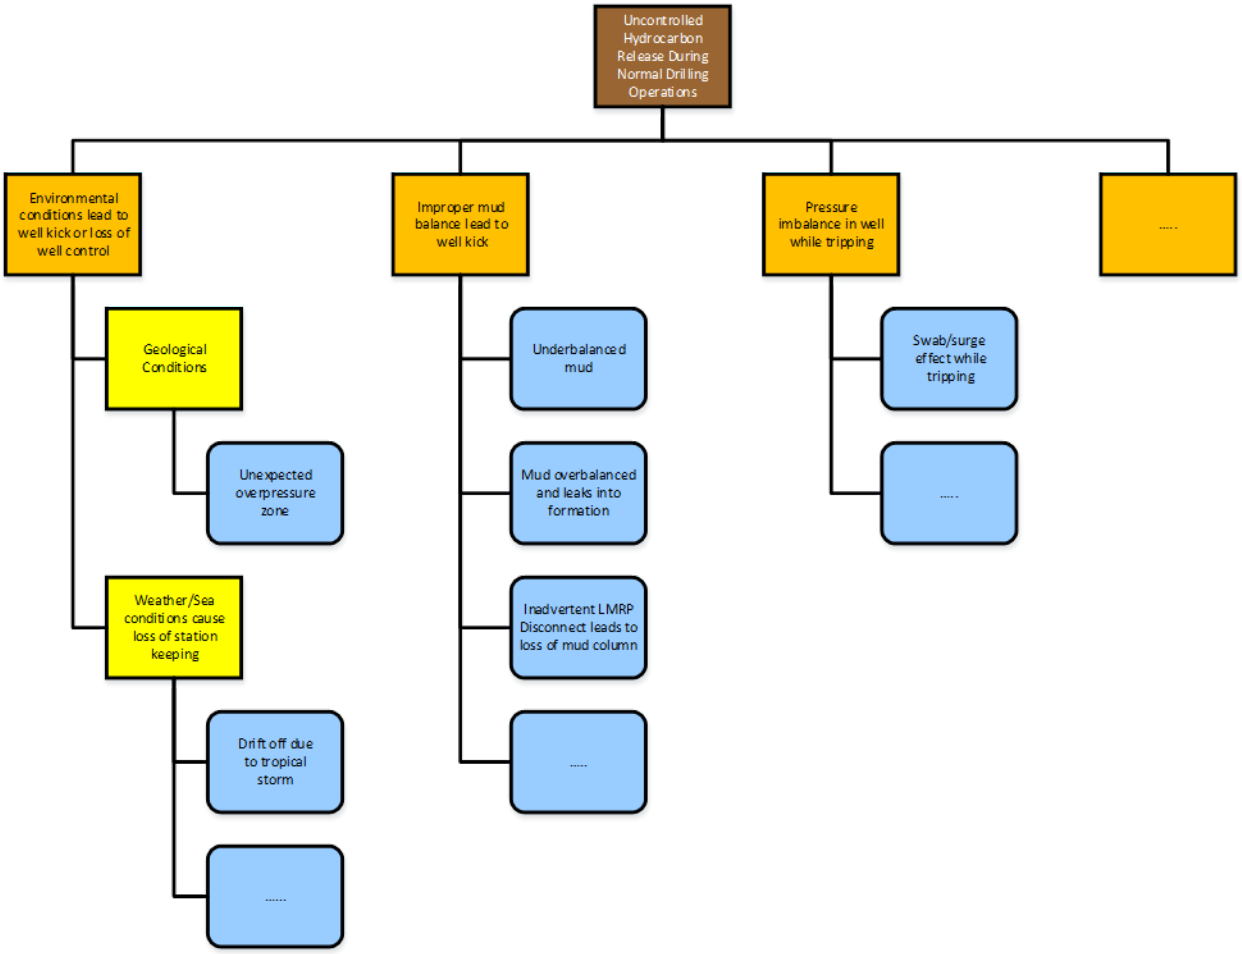
Uncontrolled Hydrocarbon Release During Normal Drilling Operations
Environmental conditions lead to well kick or loss of well control
Improper mud balance lead to well kick
Pressure imbalance in well while tripping
.....
Geological Conditions
Unexpected overpressure zone
Weather/Sea conditions cause loss of station keeping
Drift off due to tropical storm
......
Underbalanced mud
Mud overbalanced and leaks into formation
Inadvertent LMRP Disconnect leads to loss of mud column
......
Swab/surge effect while tripping
.....

Figure 2-3. Notional Master Logic Diagram Related to Candidate Initiating Events

Once an exhaustive list of IEs has been identified, the frequency of the IEs may be quantified. Note that IEs are developed as frequencies because they are on a per time or per mission basis. Some IEs may be singular events, such as the frequency of a tropical storm in a particular geographical area and of a specific magnitude that could result in station-keeping challenges. Historical data are available and might be applicable to quantification of the frequency of this IE. Other IEs may be more complex and require further development before quantification can occur. IEs like Inadvertent Disconnect of the LMRP may require a fault tree to establish the causes or enabling events, which may then be quantified in the fault tree to estimate the IE frequency. Occasionally, some IEs may be conditional. For instance, severe environmental conditions resulting in a drift-off condition may be seasonal and geographically dependent. A severe environment may be more likely in some months, e.g. June-September, in the Gulf of Mexico due to the potential formation of tropical cyclones. In this case, a temporary exploration operation may consider the time that is planned for the well, or in the case of a production platform, different initiators with different seasonal frequencies may be used to account for the IE dependencies.

Quantification of an IE frequency is often done using a Bayesian approach, where operational data are evaluated to determine the initiator frequency, including the uncertainty on the frequency (this approach is described in Sections 2.2.5.10 and 2.2.6).

2.2.3 Event Sequence Diagrams

Once an exhaustive set of IEs has been established, accident progression, with the elements shown in Figure 2-4, can be modeled using an Event Sequence Diagram (ESD) and/or its derivative, an event tree. Both are used in PRAs to provide organized displays of sequences of system failures or successes, and human errors or successes that can lead to specific end states. A traditional accident progression analysis begins with an ESD, refines it, and then transforms it into an event tree format. An ESD starts with the premise that some IE has occurred and then maps out what could occur in the future if particular systems (or humans) fail or succeed in responding appropriately to the IE. The ESD shows event sequences (or pathways) leading to different end states. ESDs are a very useful step in developing logic models: ESDs permit the complex relationships among IEs and subsequent responses to be displayed more readily and understandably than do event-tree models.

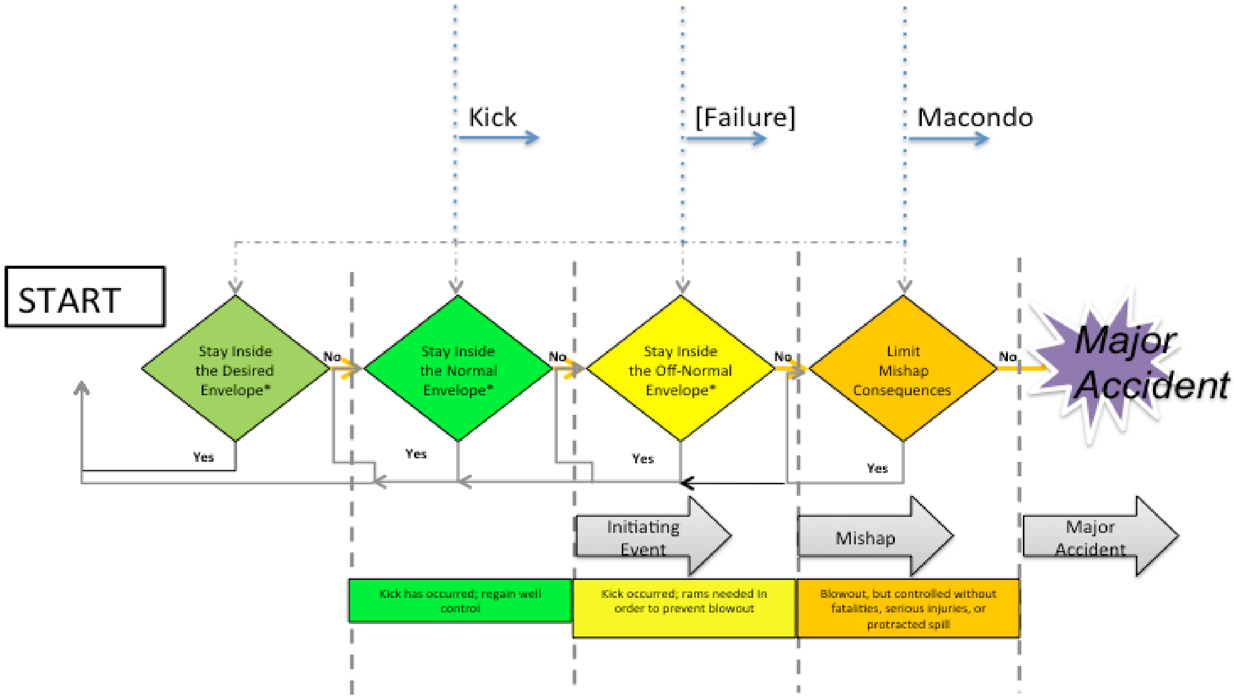
Kick
[Failure]
Macondo
START
Stay Inside the Desired Envelope*
No
Yes
Stay Inside the Normal Envelope*
No
Yes
Stay Inside the Off-Normal Envelope*
No
Yes
Limit Mishap Consequences
No
Yes
Major Accident
Initiating Event
Mishap
Major Accident
Kick has occurred; regain well control
Kick occurred; rams needed in order to prevent blowout
Blowout, but controlled without fatalities, serious injuries, or protracted spill

Figure 2-4. The Elements of an Accident Scenario

In principle, one ESD is developed for each IE; however, responses to nominally different IEs in terms of system controls or mitigations may be very similar, or even the same. In such a case, a single ESD may be used to represent the accident progression for a set of IEs. The objective is to illustrate all distinct paths from the IE to the end states.

An important attribute of an ESD is its ability to describe and document assumptions used in event trees. An ESD can be very detailed, depicting all sequences considered by the PRA analyst. When simplifying assumptions are used to facilitate event tree construction or quantification, the ESD may furnish a basis for demonstrating why such assumptions are conservative, or probabilistically justified.

Figure 2-5 depicts a simple ESD and its symbols. The Figure 2-5 ESD begins with an IE that perturbs the function being modeled from a stable state. The initial response to this perturbation is provided by System A, and if System A compensates for the IE, a successful end state results.

If System A should fail, Systems B and C together can compensate for the IE. According to Figure 2-5, a successful end state ensues if Systems B and C start and operate satisfactorily.

Failure of System B to start or operate results in End State 1. If System B is successful and System C fails to start or operate properly, successful crew intervention can still provide some mitigation for the accident and will result in a different end state (End State 2). If the crew efforts are unsuccessful, End State 1 results.

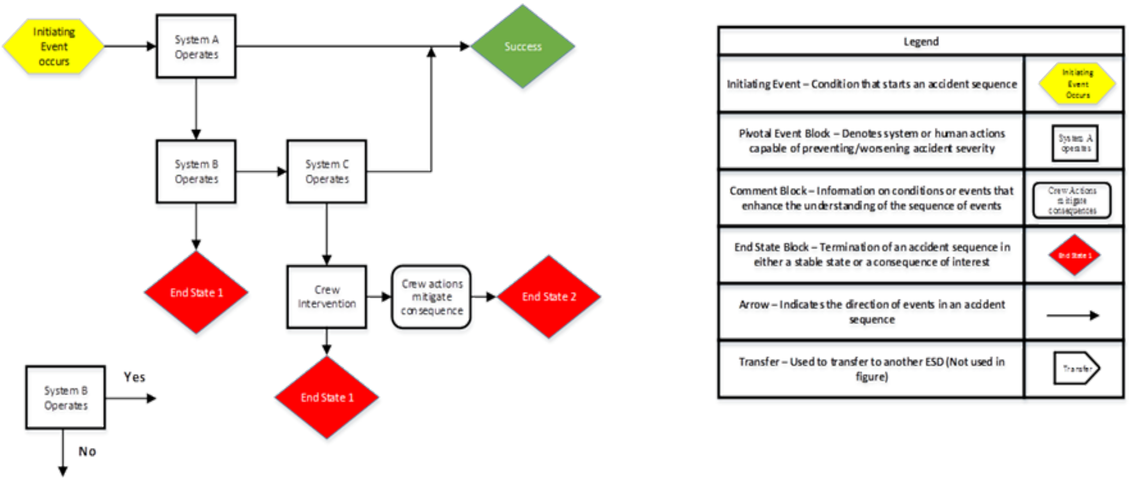

Legend	
Initiating Event – Condition that starts an accident sequence	Initiating Event Occurs
Pivotal Event Block – Denotes system or human actions capable of preventing/worsening accident severity	System A operates
Comment Block – Information on conditions or events that enhance the understanding of the sequence of events	Crew Actions mitigate consequences
End State Block – Termination of an accident sequence in either a stable state or a consequence of interest	End State 1
Arrow – Indicates the direction of events in an accident sequence	→
Transfer – Used to transfer to another ESD (Not used in figure)	Transfer

Figure 2- 5. Notional Event Sequence Diagram

Figure 2-6 is a more complex (but still relatively simple) ESD developed to evaluate accident sequences related to a well kick from drilling into an unexpected overpressure zone that results in an environmental release. Five different end state designations are used, Well Shut In and 4 different Environmental Release end states. The Well Shut In end state represents paths that result in no or insignificant environmental release and the condition of the well is stable (i.e. no hydrocarbon flow) and controlled. The Environmental Release end states represent paths where mitigating events have failed to prevent the accident from progressing resulting in a release to the environment. A different end state designation is used for each path depending on the mitigation used for gaining control of the well (e.g. ROV intervention, Well Capping, etc.). Even for a given path, the magnitude of the release can vary; for example, the magnitude of the release would depend on the flow rate of the well and the time it is flowing to the environment. If a relief well is needed to stop the flow, it will lead to a much larger release than if an ROV can intervene and stop the flow early on in the accident.

Figure 2-7 illustrates the process of ESD development. Since an ESD is success oriented, the process begins (Step 1) by identifying the anticipated response to the IE, in this case a well kick due to an unexpected overpressure zone, out to a successful outcome. For this example, the anticipated response is to first properly detect the kick before it reaches the BOP. If that occurs, then mitigating actions can take place to shut in the well. The first responses are shown as comments, stopping the rotation of the drill pipe and positioning the drill string. The mud pumps are then stopped to increase the time before the kick reaches the BOP. These first two actions are listed as comments because they affect other steps and will be accounted for in them. The next pivotal event is closing the annular preventer and opening the choke line. This isolates the well except for the path up the drill string if the drill string or another tubular is present in the BOP. With the annular successfully working the next question would be if the path through the drill string is isolated. The success of this event is conditional on whether the drill string has a float valve present or an IBOP is necessary. These conditions may be represented in a fault tree tied to the event as discussed in Section 2.2.5.8. The well is monitored for flow/pressure and if isolation is successful and no flow is present, the well is controlled, and a well kill program may be initiated. The well kill process starting with using the pipe rams and opening the annular preventer is not shown here to keep the diagram simple.

The next step (Step 2 in Figure 2-7) in developing the ESD is to consider what happens when failures occur. On the first block after the IE, if the kick is not detected prior to formation fluid reaching the BOP no other barriers exist to prevent the fluid already past the BOP from reaching the rig. A comment was placed in the ESD to show that the diverter may be used, but the diverter is for personnel safety and not for preventing environmental release. The purpose of the ESD is to estimate environmental release, which the diverter does not mitigate, so it is commented for later use if personnel safety is also analyzed. Put differently: the comment is there as a reminder, in case we need it later.

Once it is determined that that the rig will be impacted, the mitigating actions are assumed to start with an emergency disconnect from the well. This action sets in motions mitigation by operating the casing shear ram and then the blind shear ram. Successful operation of the blind shear ram is all that would be required to seal the well, as the casing shear ram is assumed to not provide an effective sealing surface. The casing shear ram is operated first however, in case there is any tubular present (e.g. casing) in the BOP that would prevent closure of the blind shear ram. In developing the sequence of events, it was noted that the casing shear may not be effective for all tubulars, and if some specific types of tubulars such as drill collar or tool joints are present, the casing shear will not be able to perform its function. A block was added to model this possibility, since, if nonshearable pipe were present, the actions of the casing and blind shear rams would be guaranteed to fail, and therefore would not need to be questioned. Success of the casing shear ram block is noted as being conditional that shearable tubular is in the BOP. If the BOP is free of any tubulars, as previously mentioned, all that is needed for successful well isolation is the blind shear ram. If the blind shear ram works, the end state is that the well is controlled, but a limited release has occurred.

The second failure path is in response to a failure of the annular. In this case, the next step would be to close the pipe rams. If the drill string is present in the BOP, the pipe rams will close around the pipe and seal the annulus. There are a minimum of 2 pipe rams available; however, there may be more, and only one has to be successful to shut down flow from the well through the annulus. If the pipe rams fail to shut in the well around the drill pipe due to a system failure or possibly a tubular that is outside of the design of the pipe ram (e.g. drill collar), formation fluid will travel up the annulus and an emergency disconnect is assumed to be the response. If the pipe rams are successful, then the drill string path is questioned to determine whether that path is isolated or not.

The last pivotal event in the top line is the pivotal event for the isolation of the drill string. In this case, a float valve may or may not be present in the string, and an Internal BOP (IBOP) may or may not be in place on the rig floor. If either/both of those are present and fail, or there is no equipment to prevent a release through the drill string, formation fluid can reach the rig. In this case, the flow will be significantly less than it would be if the annulus were open to flow, so an emergency disconnect is not questioned, but the BOP shear rams are.

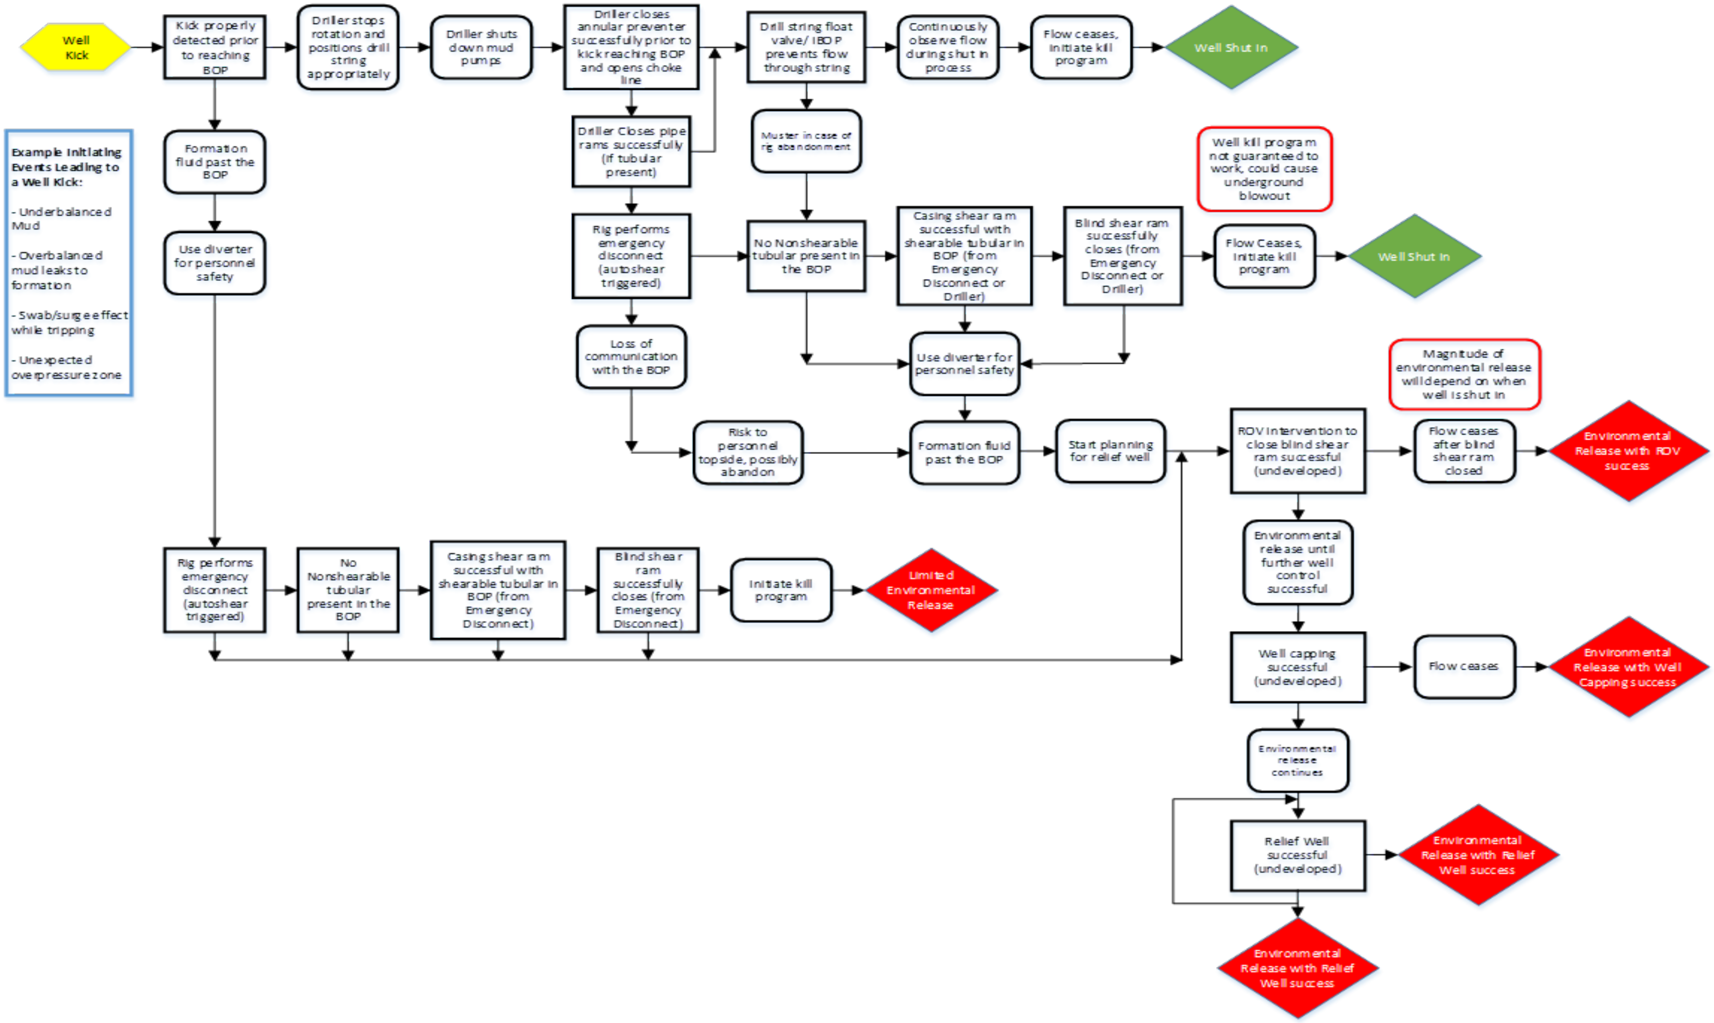
Well Kick
Kick properly detected prior to reaching BOP
Driller stops rotation and positions drill string appropriately
Driller shuts down mud pumps
Driller closes annular preventer successfully prior to kick reaching BOP and opens choke line
Drill string float valve/ IBOP prevents flow through string
Continuously observe flow during shut in process
Flow ceases, initiate kill program
Well Shut In
Example Initiating Events Leading to a Well Kick:
- Underbalanced Mud
- Overbalanced mud leaks to formation
- Swab/surge effect while tripping
- Unexpected overpressure zone
Formation fluid past the BOP
Use diverter for personnel safety
Driller Closes pipe rams successfully (if tubular present)
Muster in case of rig abandonment
Rig performs emergency disconnect (autoshear triggered)
No Nonshearable tubular present in the BOP
Casing shear ram successful with shearable tubular in BOP (from Emergency Disconnect or Driller)
Blind shear ram successfully closes (from Emergency Disconnect or Driller)
Flow Ceases, Initiate kill program
Well Shut In
Well kill program not guaranteed to work, could cause underground blowout
Loss of communication with the BOP
Use diverter for personnel safety
Magnitude of environmental release will depend on when well is shut in
Risk to personnel topside, possibly abandon
Formation fluid past the BOP
Start planning for relief well
ROV Intervention to close blind shear ram successful (undeveloped)
Flow ceases after blind shear ram closed
Environmental Release with ROV success
Environmental release until further well control successful
Rig performs emergency disconnect (autoshear triggered)
No Nonshearable tubular present in the BOP
Casing shear ram successful with shearable tubular in BOP (from Emergency Disconnect)
Blind shear ram successfully closes (from Emergency Disconnect)
Initiate kill program
Limited Environmental Release
Well capping successful (undeveloped)
Flow ceases
Environmental Release with Well Capping success
Environmental release continues
Relief Well successful (undeveloped)
Environmental Release with Relief Well success
Environmental Release with Relief Well success

Figure 2- 6. Event Sequence Diagram for a Well Kick from an Unexpected Overpressure Zone

Step 3 in Figure 2-7 further develops the ESD out to an environmental release through the accident management stage. For simplicity, the events considered are undeveloped and only shown as a single block each. If an unshearable tubular is present or the shear rams are unsuccessful (down path) then a release will occur. The next possibility for mitigation would be attempting to manipulate the BOP with an ROV. If this is successful a release will have occurred, but the magnitude will be somewhat limited due to the relatively short response time. Failure of the ROV (because of environment, BOP condition, ROV failure, etc.) will lead to the next available option which would be well capping. The development continues until all paths lead to a stable state or a release which can vary in magnitude with the last option available being a relief well.

The example ESD developed in Figure 2-7 analyzes end states corresponding to environmental release given that a well kick from an unexpected overpressure zone has occurred. A well kick may have different causes such as those shown in Figure 2-3. The ESD can provide a common response in terms of events for similar initiators; however, the probabilities in the ESD may be conditional based on what the initiator is. For instance, if the trip is caused by the swab/surge effect, is the probability of a nonshearable tubular in the BOP the same as it is for drilling into an unexpected overpressure zone? When quantifying using event trees, these conditions must be accounted for, if we are to accurately estimate the probability of the consequences of interest.

There also may be other initiating events that could lead to an environmental release that may have different event sequences. For instance, an inadvertent LMRP disconnect would not have the detection, annular, or drill string blocks on the top row in Figure 2-7 since the loss of communication with the BOP after the LMRP disconnect negates any actions by the driller. This may therefore require an ESD to be developed specifically to address those scenarios.

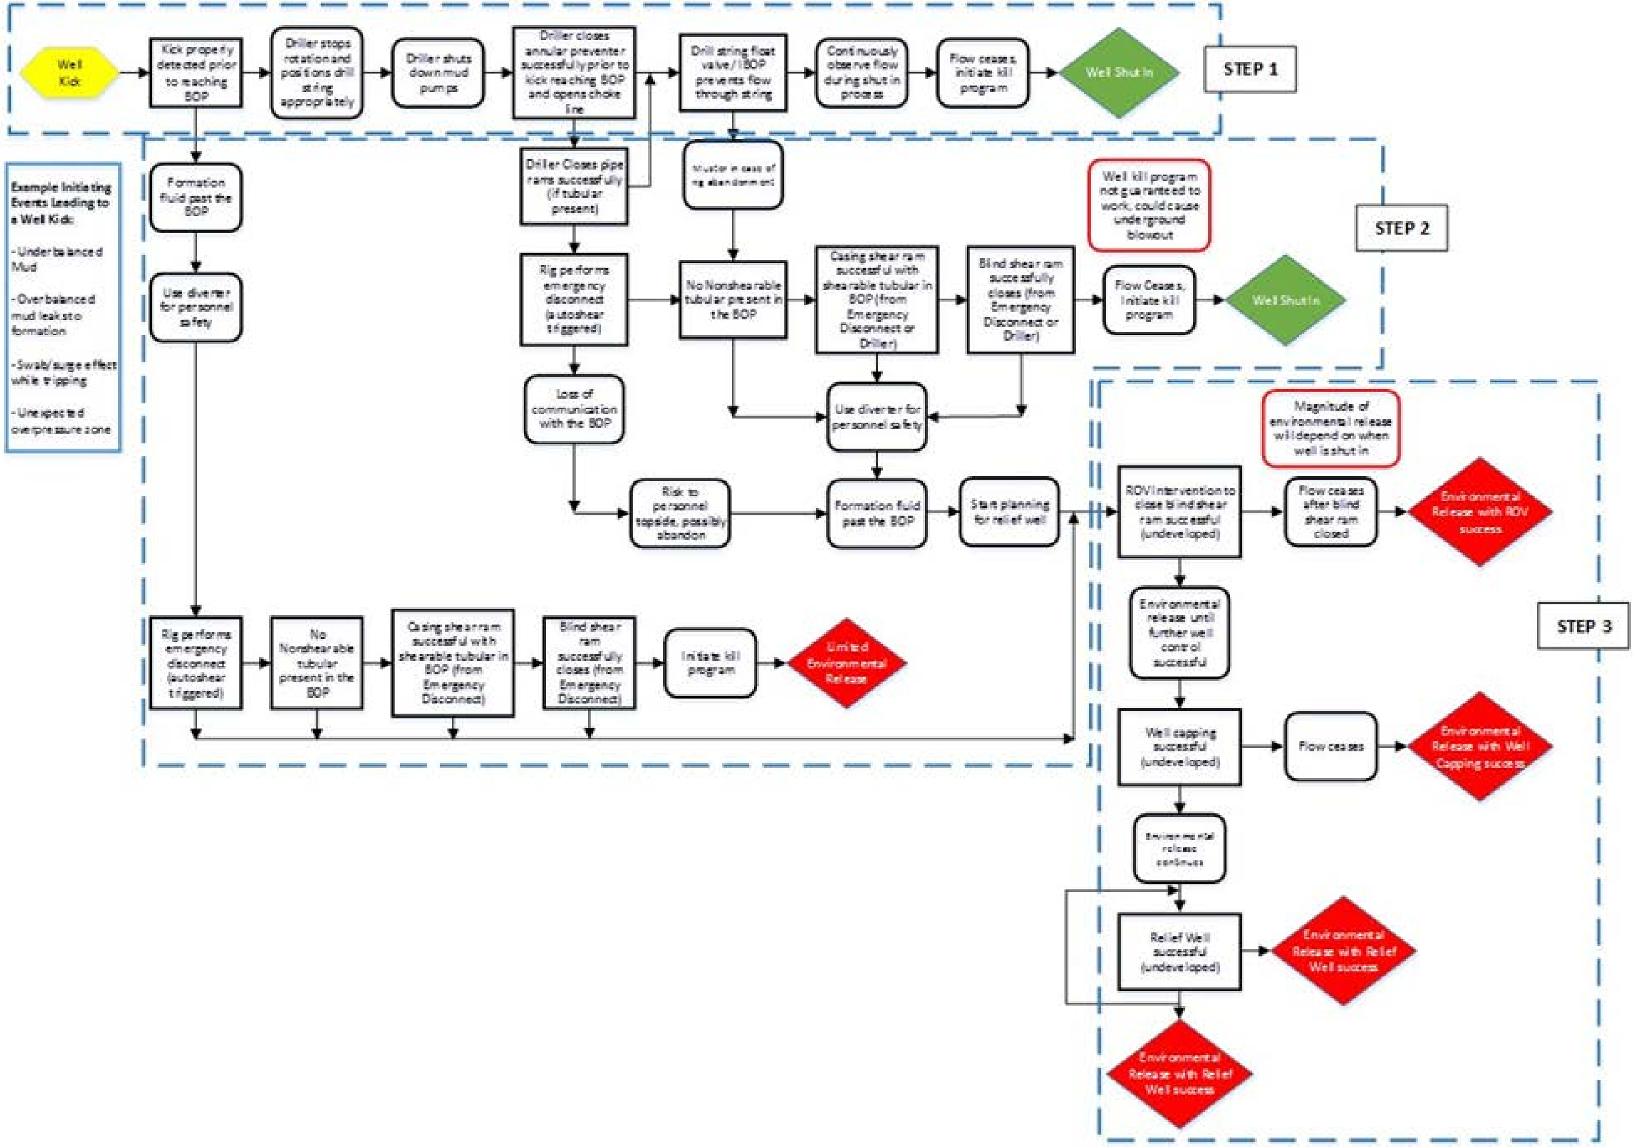
Well Kick
Kick properly detected prior to reaching BOP
Driller stops rotation and positions drill string appropriately
Driller shuts down mud pumps
Driller closes annular preventer successfully prior to kick reaching BOP and opens choke line
Drill string float valve / IBOP prevents flow through string
Continuously observe flow during shut in process
Flow ceases, initiate kill program
Well Shut In
STEP 1
Example Initiating Events Leading to a Well Kick:
- Underbalanced Mud
- Overbalanced mud leaks to formation
- Swab/surge effect while tripping
- Unexpected overpressure zone
Formation fluid past the BOP
Use diverter for personnel safety
Driller Closes pipe rams successfully (if tubular present)
Muster in case of rig abandonment
Rig performs emergency disconnect (autoshear triggered)
No Nonshearable tubular present in the BOP
Casing shear ram successful with shearable tubular in BOP (from Emergency Disconnect or Driller)
Blind shear ram successfully closes (from Emergency Disconnect or Driller)
Well kill program not guaranteed to work, could cause underground blowout
Flow Ceases, Initiate kill program
Well Shut In
STEP 2
Loss of communication with the BOP
Use diverter for personnel safety
Risk to personnel topside, possibly abandon
Formation fluid past the BOP
Start planning for relief well
Rig performs emergency disconnect (autoshear triggered)
No Nonshearable tubular present in the BOP
Casing shear ram successful with shearable tubular in BOP (from Emergency Disconnect)
Blind shear ram successfully closes (from Emergency Disconnect)
Initiate kill program
Limited Environmental Release
Magnitude of environmental release will depend on when well is shut in
ROV Intervention to close blind shear ram successful (undeveloped)
Flow ceases after blind shear ram closed
Environmental Release with ROV success
STEP 3
Environmental release until further well control successful
Well capping successful (undeveloped)
Flow ceases
Environmental Release with Well Capping success
Environmental release continues
Relief Well successful (undeveloped)
Environmental Release with Relief Well success
Environmental Release with Relief Well success

Figure 2- 7. Event Sequence Diagram Development Steps

2.2.4 Event Trees

Once the accident progression paths are understood, the next step is to build event trees for scenario quantification. An Event Tree is a graphic that displays scenarios potentially resulting from a specific IE (or a group of functionally similar IEs. Event trees are derivable from ESDs, but event trees are one step closer than ESDs to generation and quantification of scenarios. An event tree distills the pivotal event scenario definitions from the ESD and presents this information in a tree structure that is used to help classify scenarios according to their consequences and perform a quantification of the scenarios. The headings of the event tree are the IE which is the starting point, the pivotal events showing success or failure of mitigating/ aggravating events, and lastly the end state to bin the consequence of each scenario. Each individual path through the event tree is a sequence. The event tree pivotal events are linked to fault trees, and the pivotal event name should match the corresponding fault tree top event description. This is because fault trees are tied to the pivotal events and are based on potential failures for that event. An example event tree based on the ESD in Figure 2-5 is shown below.

The simple example in Figure 2-8 shows 5 sequences on the right hand side of the event tree with 3 different end states:

SUCCESS;
ENDSTATE-1; and
ENDSTATE-2.

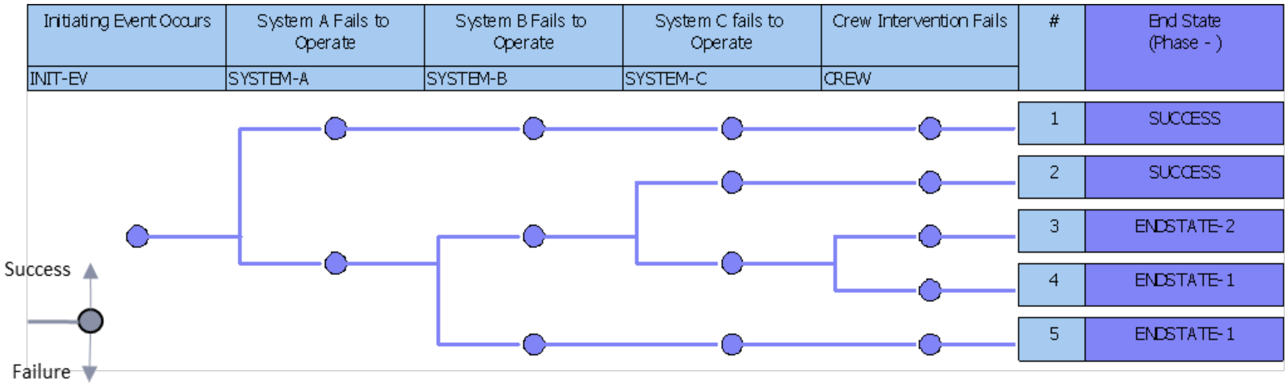

Figure 2- 8. Example Event Tree

Each sequence resulting in the end state represents a combination of the IE and success and or failures of the pivotal events. For instance, consider sequence 3 in Figure 2-9 below.

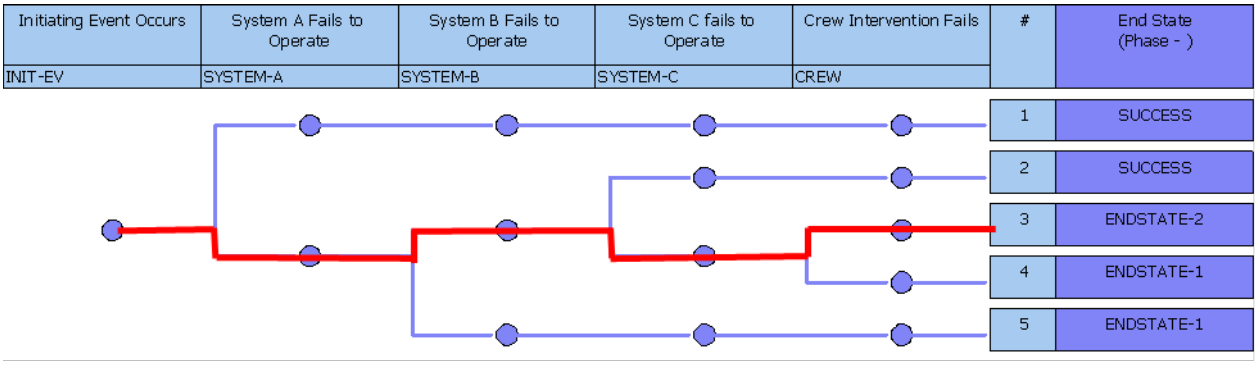

Figure 2- 9. Example Event Tree Sequence

Sequence 3 starts with the IE and results in ENDSTATE-2. The resulting sequence is a combination of successes and failures of the pivotal events along the path, yielding the expression:

INIT-EV * SYSTEM-A * \SYSTEM-B * SYSTEM-C * \CREW

In the above expression, INIT-EV represents the frequency of the IE, SYSTEM-A and SYSTEM-C represent the probabilities of failure of systems A and C respectively, as indicated by the downward step at each of those pivotal events in the event tree. \SYSTEM-B and \CREW represent NOT failure of SYSTEM-B and NOT failure of CREW (i.e. success), indicated by the upward step in the event tree.

Quantification will be discussed in more detail later. However, for purposes of illustration, the frequency of this event sequence can be quantified assuming IE and pivotal event values of:

INIT-EV = 0.10 events per unit time,
SYSTEM-A = 0.02 (failure probability of A given INIT-EV),
SYSTEM-B = 0.03 (failure probability of B given INIT-EV and failure of A),
SYSTEM-C= 0.03 (failure probability of C, given INIT-EV, failure of A, and success of B),
CREW= 0.05 (failure probability of CREW, given INIT-RV, failure of A, success of B, failure of C).

This yields the following as the frequency of occurrence of END-STATE-2 (in events per unit time):

0.1* 0.02 * (1-0.03) * 0.03 * (1-0.05) = 5.53E-5

From Figure 2-10, it can be seen that not all pivotal events are questioned in every sequence. Sequence 5 in Figure 2-9 does not question SYSTEM-C or CREW, because once SYSTEM-A and SYSTEM-B have failed, SYSTEM-C and CREW can no longer affect the end state. Dependences like this are typically accounted for when the event tree is developed, so that the resulting sequences are the minimal sets of pivotal events that must occur for that end state to occur. The expression for sequence 5 then becomes:

INIT-EV * SYSTEM-A * SYSTEM-B.

Substituting the values from above yields:

0.1 * 0.02 * 0.03 = 6.0E-5

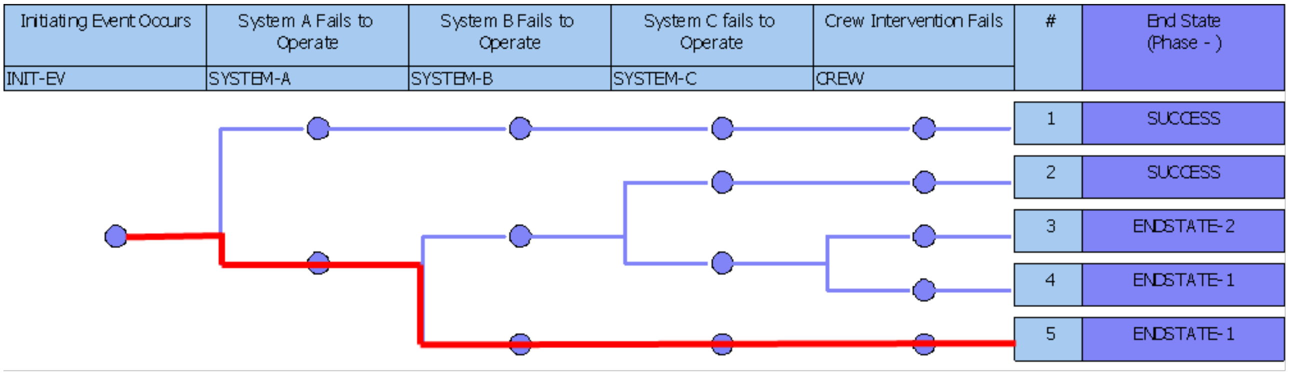

Figure 2-10. Example Event Tree Sequence where not all Pivotal Events are Questioned

When the pivotal events are replaced with fault trees, as discussed in Section 2.2.5, it becomes possible to express the event sequences in more detail, namely, in terms of basic events (e.g., component failures) rather than pivotal event names (failures of systems or entire functions). Depending on the size of the fault tree, each event tree sequence can result in many “cut sets” or unique contributors that can cause system failure, since pivotal events such as SYSTEM-A may have many different ways to fail: pump failures, valve failures, leaks, etc.

2.2.4.1 Event Tree Development

Developing an event tree usually begins with the ESD. From the ESD in Figure 2-5, 4 pivotal events were shown:

System A Operates,
System B Operates,
System C Operates, and
Crew Intervention

The event tree in Figure 2-11 is the start at mapping out the ESD paths from the ESD from the IE to the end states. Event tree development generally follows the time sequence of events from the ESD. In Figure 2.5, the initial response after the IE is the System A status, so it logically is the second event in the event tree (converted to failure). System B Fails to Operate is questioned if System A fails in the ESD, and if System B is successful, System C Fails to Operate is questioned. System B is listed after System A on the event tree because the status of System A must be known before System B is questioned. Similarly, System C is listed after System B because the status of System B must be known. Lastly, Crew Intervention is only questioned if System C is failed, so it must be listed after System C.

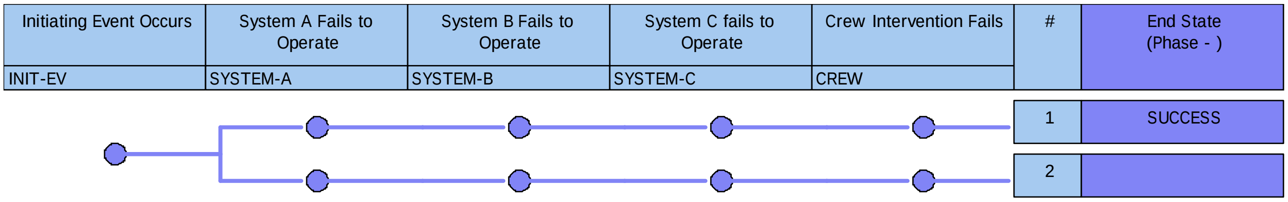

Figure 2-11. Step 1 in Building Example Event Tree

With the top line of the event tree laid out, the next step is to develop the branches for the pivotal events. From the ESD, System A operation directly follows the IE, so it must have a success and failure path as shown in the event tree. If System A is successful in the ESD, then the end state is success. The translation to the event tree in Figure 2.9 shows that none of the other pivotal events after System A Operates needs to have a downward branch for failure since they do not impact the scenario. The end state of the first sequence is labeled SUCCESS as in the ESD as shown in Figure 2-11.

Once System A has failed (down path on the event tree), the status of System B is questioned and therefore needs success and failure paths. The failure path of System B leads directly to a negative consequence labeled as ENDSTATE-1 (sequence #4 listed to the left of the End State column) since no other mitigation options are available with both System A and System B failing. If System B is successful, then the status of System C is questioned, with its success (up path) leading to a SUCCESS end state, as shown in Figure 2-12 in sequence 2.

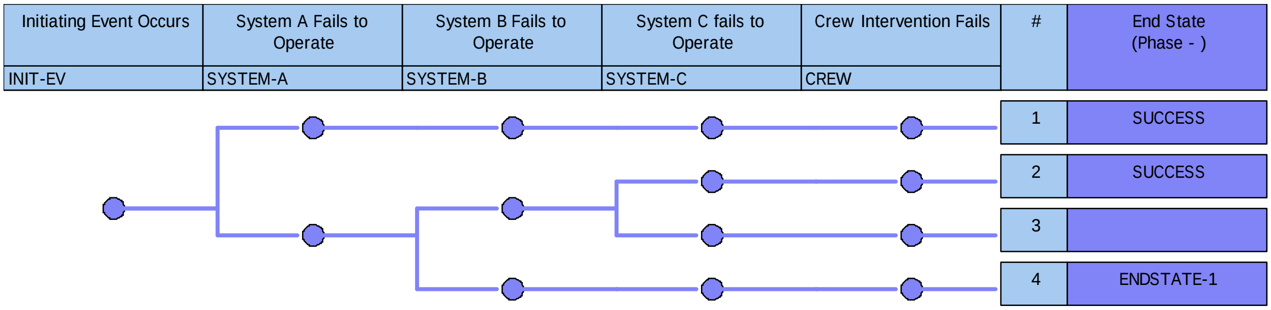

Figure 2-12. Step 2 in Building Example Event Tree

The last step in building the example event tree is to fill out the scenarios if System C fails after System A fails and System B is successful. The Crew Intervention event lessens the impact of the consequence per the ESD, and therefore the success path of the Crew Intervention event leads to a second, lesser, negative consequence labeled as ENDSTATE-2 in Figure 2.8. The failure of Crew Intervention leads to the same end state and if Systems A and B or A and C had failed, ENDSTATE-1. The final event tree is shown in Figure 2-8.

The resulting sequences of events for each of the end state in the example (Figure 2.8) are:

SUCCESS:
Sequence 1: Initiating Event Occurs * /System A Fails to Operate
Sequence 2: Initiating Event Occurs * /Systems B fails to operate and / System C fails to operate

ENDSTATE-1:
Sequence 4: Initiating Event Occurs * System A fails to operate * /System B fails to operate * System C fails to operate * Crew intervention fails
Sequence 5: Initiating Event Occurs * System A fails to operate * System B fails to operate

ENDSTATE-2:
Sequence 3: Initiating Event Occurs * System A fails * /System B fails to operate * System C fails to operate * /Crew intervention fails

The objective is to develop a tractable model for the important paths leading from the IE to the end states. Generally, risk quantification is achieved by developing fault tree models for the pivotal events in an event tree. This linking between an event tree and fault trees permits a Boolean equation to be derived for each event sequence. Event sequence quantification occurs when reliability data are used to numerically evaluate the corresponding Boolean equation.

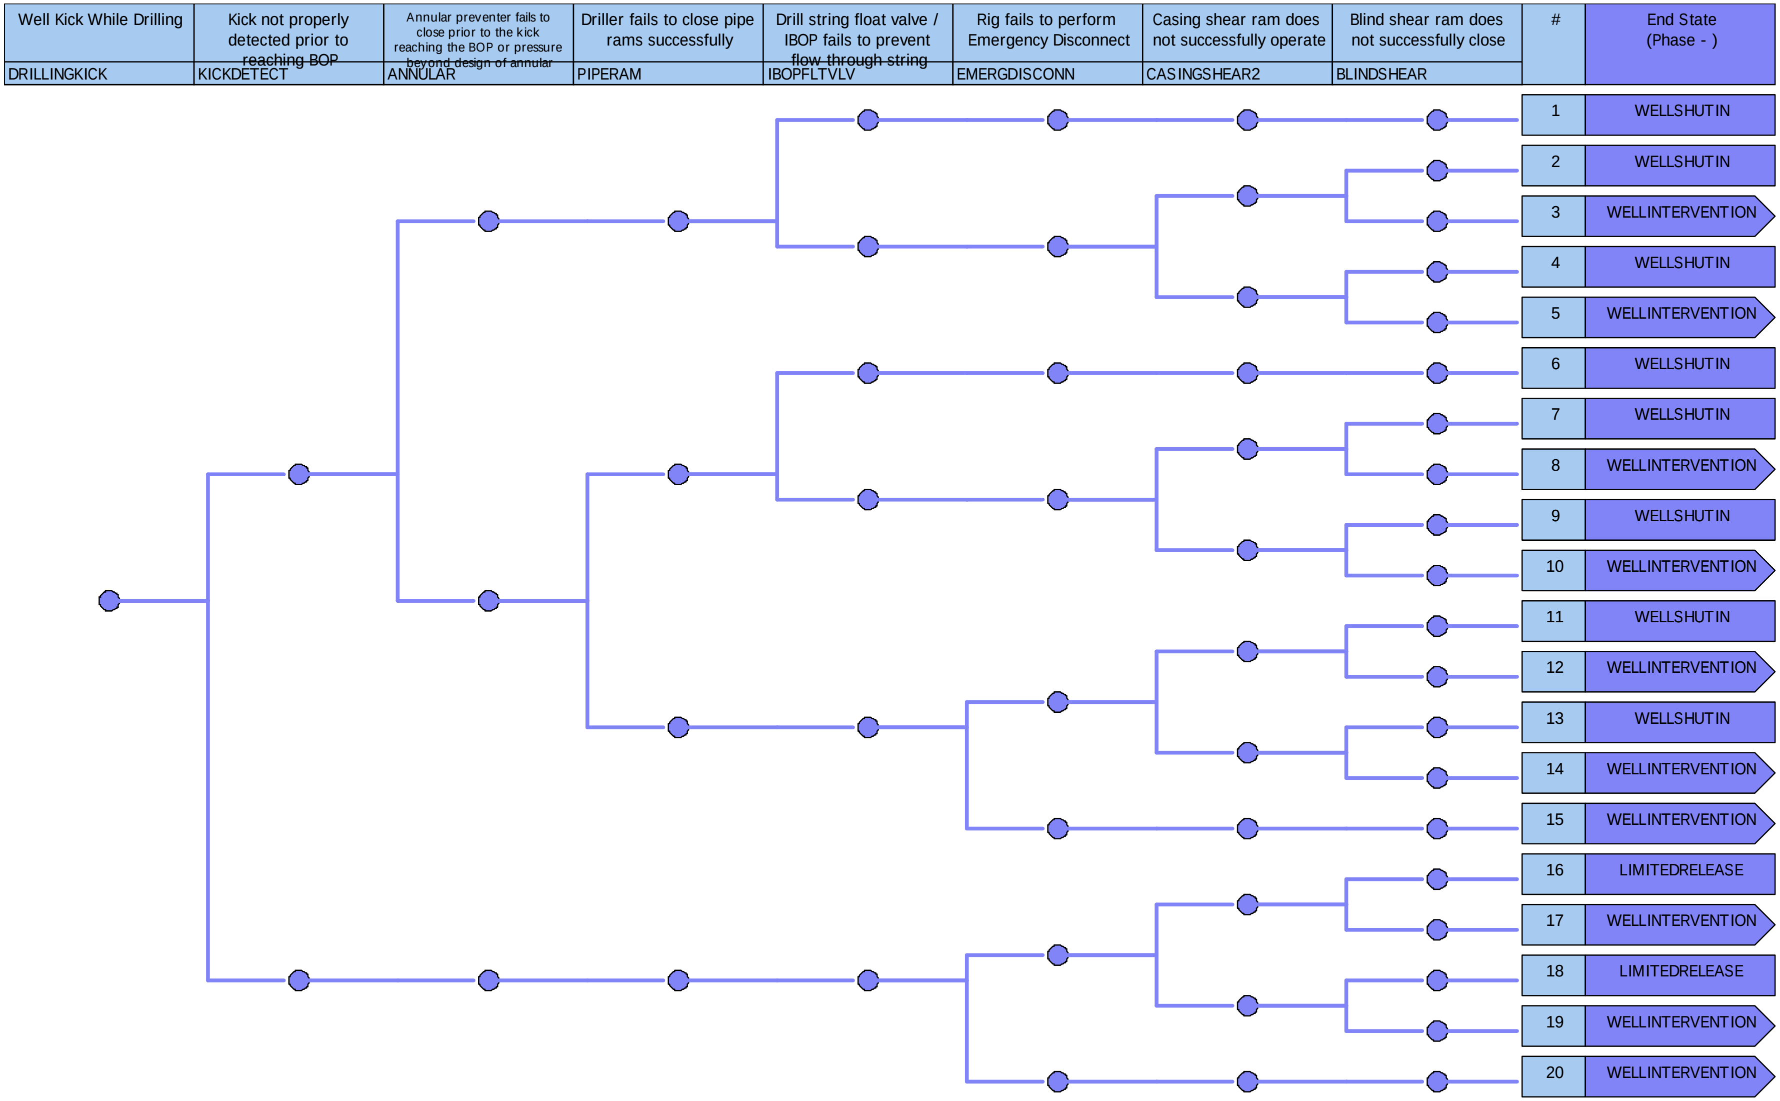

Figure 2-13. Event Tree Structure for Well Kick from an Unexpected Overpressure Zone

2.2.4.2 Event Tree Transfers

Figure 2-13 is a more complicated event tree corresponding to the Figure 2-6 ESD to estimate the risk of an environmental release in response to an unexpected overpressure zone. In this example, there are many different scenarios modeled. Going back to the original identification of initiating events and development of ESDs in Sections 2.2.2 and 2.2.3, it was noted that some IEs may have the same sequence of events and can use the same ESD/event tree, albeit perhaps with different conditional probabilities possibly assigned to the pivotal events. Other IEs may have completely different scenarios that are modeled, or may be partially the same. For initiators with common sequences of events, the common sequences may be best developed in a standalone event tree and used as an event tree transfer. Listed on the right side of Figure 2-13 under the end states is a transfer condition WELLINTERVENTION shown as the symbol in Figure 2-14.

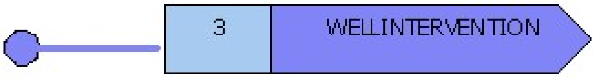

Figure 2-14. Event Tree Transfer

The WELLINTERVENTION end state is a transfer to a second event tree. A transfer is used generally in the case when there may be common elements to multiple IEs or when an event tree gets very large and has distinct and different sequential processes. The event tree in Figure 2-6 is based on a well kick for an unexpected overpressure zone. Other IEs, such as an inadvertent LMRP disconnect, may have different initial responses (no driller intervention if communication with the BOP is lost), but the responses of ROV intervention, well capping, and relief well may be the same from the ESD in Figure 2-6. In the case of Figure 2-6, it was decided to use a standalone event tree for the ROV through relief well part of the ESD, in order to manage the size of the event tree and because other IEs may have that part in common. Accordingly, the WELLINTERVENTION event tree is shown in Figure 2-15.

The first event, LARGERELEASE is simply the entry point from the previous event tree, and does not show up in event sequences. The sequences of events from the previous event tree will continue in the WELLINTERVENTION event tree to the end state where the sequences will contain all IE and pivotal event information from the initial event tree and the transfer event tree.

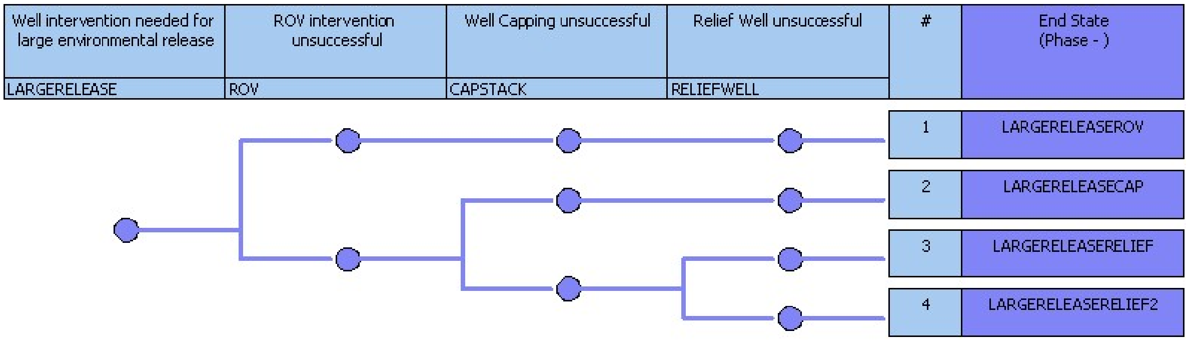

Figure 2-15. WELLINTERVENTION Event Tree

In Figure 2-13, the end states are WELLSHUTIN which would represent SUCCESS on the simple example where formation fluid has stopped flowing, the well is controlled, and a well kill program may be put in place. This end state is shown in Sequences 1, 2, 4, 6, 7, 9, 11, and 13 in Figure 2-13. The Sequences 16 and 18 end states are shown as LIMITEDRELEASE and are the scenarios where some formation fluid has risen above the BOP and will reach the environment, but the blind shear ram has sealed the well to prevent further release. The remaining sequences have their end state listed on the WELLINTERVENTION event tree. Each end state in WELLINTERVENTION is different because the failures are time dependent in that an ROV intervention may happen earlier than well capping while the relief well is much longer term. Therefore, the duration of the releases will vary based on the type of successful intervention, and these are separated into distinct end states as shown in Figure 2-15. Ultimately, if no other well intervention techniques are successful, a relief well will be needed. This is separated into 2 paths, with the success path being the initial attempt at the relief well kills the well.

The end states discussed do separate out, to some degree, the magnitude of release as a measure of consequences. It may be desired to provide more deterministic estimates showing probabilistically the magnitude of release expected (i.e. barrels of oil released as a function of probability. This can be done through the use of simulation as described in Section 2.3.

2.2.4.3 Multibranch Event Trees

The pivotal events in the previous described examples are all binary in that they only have paths related to success or failure of the event. In some cases there may be multiple states or conditions that an event may be in, each with a different probability. Related to the ESD in Figure 2-6, there are different types of tubulars that may be present in the BOP when the shear rams are called on to work depending on which operation is being performed (e.g. drilling, running casing, etc.). Since the casing shear and blind shear have different shearing capability, and there are some tubulars that are nonshearable, the type of tubular present in the BOP at the time of shearing is important to model for model accuracy. This may be done in the fault trees linked to the event tree (discussed in Section 2.2.5), but fault trees addressing multiple possible configurations or conditions can be complicated. Another way to address multiple conditions is in event tree structure using a multi branch node in the event tree as shown in Figure 2-16.

In Figure 2-16, the ESD in Figure 2-7 has been broken up into three operational conditions, no tubular in the BOP, drill pipe in the BOP, and casing in the BOP, and the first node in the event tree has three branches with each representing one of the conditions. Each branch of the event tree for this event can be assigned a probability of that event being true, and together they would add up to 1.0 since one of the conditions has to be true if the branches cover all possible conditions. Using this approach logically changes the downstream paths if the event tree is fully developed. For instance, if there is no tubular in the BOP (top branch of first node) when a shut in is required (top branch), then the status of the casing shear ram does not need to be questioned since it is not needed as it is assumed the blind shear ram provides the sealing function for the well.

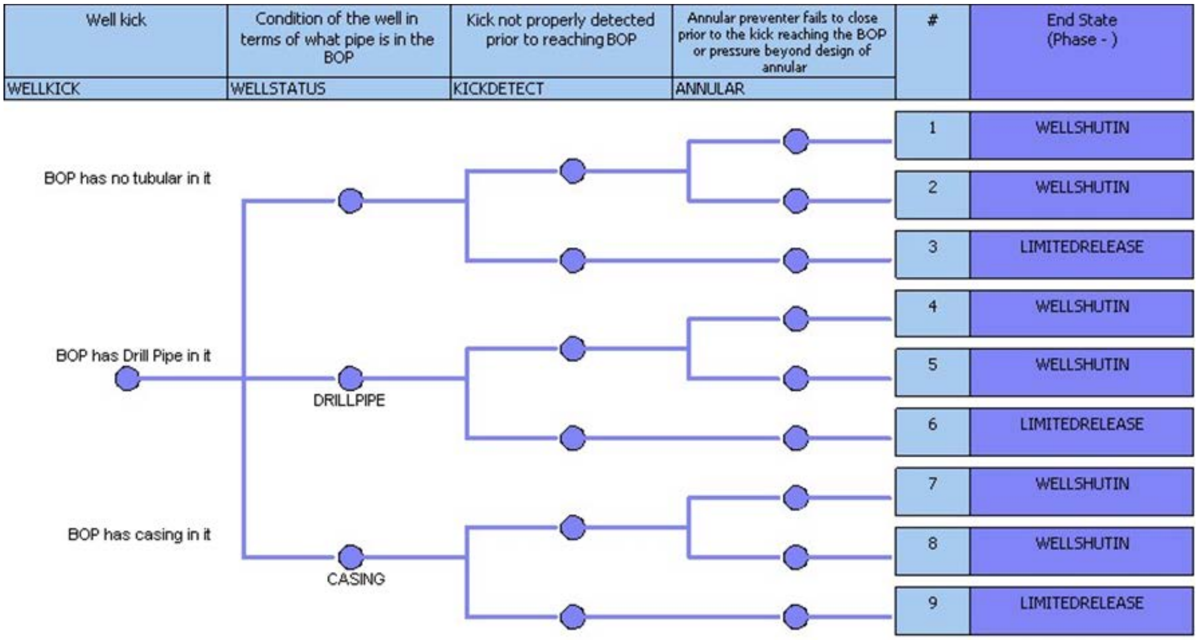
Well kick
WELLKICK
Condition of the well in terms of what pipe is in the BOP
WELLSTATUS
Kick not properly detected prior to reaching BOP
KICKDETECT
Annular preventer fails to close prior to the kick reaching the BOP or pressure beyond design of annular
ANNULAR
BOP has no tubular in it
BOP has Drill Pipe in it
DRILLPIPE
BOP has casing in it
CASING

#	End State (Phase -)
1	WELLSHUTIN
2	WELLSHUTIN
3	LIMITEDRELEASE
4	WELLSHUTIN
5	WELLSHUTIN
6	LIMITEDRELEASE
7	WELLSHUTIN
8	WELLSHUTIN
9	LIMITEDRELEASE

Figure 2- 16. Multibranch Node in an Event Tree

2.2.5 Fault Tree Modeling

In many problems of practical interest, we cannot estimate pivotal event frequencies actuarially, even if they occurred independently, because they do not occur often enough to permit useful statistical analysis; and in general, pivotal events are not independent, so even if we could quantify their frequencies (or probabilities) actuarially, we could not straightforwardly combine those results to obtain a sequence frequency. Therefore, we need to model such events synthetically:[4] that is, we need to express functional failures in terms of system failures, system failures in terms of component failures, and component failures in terms of their causes, all modeled in sufficient detail that we can begin to quantify the lowest-level elements of the model, and then work our way back up to a synthetic estimate of pivotal event probability (or conditional frequency), conditional on its role in each scenario of interest, and finally quantify the top event frequencies themselves. For PRAs, this is typically done using fault trees. More information, beyond what is in this guide, can be found in [2-1].

2.2.5.1 System Success Criteria

Prior to development of the pivotal event fault trees, success criteria are needed to define satisfactory performance in terms of the function included in the event tree. System success

[4] The term "synthetic" is used here to refer to modeling a complex event in terms of its contributors. For example, we "synthesize" an estimate of the probability of a complex event by combining estimates of the probabilities of its contributors. "Synthetic" is contrasted with "actuarial;" when we estimate a frequency actuarially, we use data on the occurrence of the actual event. For example, we would not normally synthesize an estimate of flood frequency at a specific location solely by trying to model the spectrum of rainstorms; we would also look at the historical record of flood frequency and severity. Of course, even if there have been a statistically significant number of storms, we are not certain that future conditions will match the historical record.

criteria impose operating requirements on the systems needed to successfully perform a particular function, and the duration that function is needed determines the system operating time. Once the success criteria for a function has been established, top event fault tree logic is established from the Boolean complement of the success criteria (e.g., at least 1 of 2 pipe rams must fail to close and seal around the drill pipe when demanded). Success criteria should be clearly defined. All assumptions and supporting information used to define the success criteria should be listed in the documentation (i.e., what is considered to constitute system success needs to be explicitly stated). Some examples of success criteria are:

- The blind shear ram must close and shut in the well on demand.
- At least 1 of 2 pipe rams must fail to close and seal around the drill pipe when demanded.
- At least 4 of 6 thrusters must operate to maintain station keeping under calm (specified) environmental conditions.
- At least 5 of 6 thrusters must operate to maintain station keeping under moderate (specified) environmental conditions.
- At least 6 of 6 thrusters must operate to maintain station keeping under extreme (specified) environmental conditions.

The last three examples show that success criteria may be dependent on external factors and may need to be discretely modeled. In addition, again referring to the last three examples, the conditions may require specific thrusters to be available, e.g. the 4 out of 6 case may require 2 forward and 2 aft thrusters.

2.2.5.2 Modeling Pivotal Events

Pivotal events must be modeled in sufficient detail to support valid quantification of scenarios. As a practical matter, the model must reach a level of detail at which data are available to support quantification of the model's parameters. Additionally, much of the time, pivotal events are not independent of each other, or of the IEs; the modeling of pivotal events must be carried out in such a way that these dependencies are captured properly. For example, pivotal events corresponding to system failure may have some important underlying causes in common (e.g. support systems). If the purposes of the PRA are to be served—if such underlying causes are to be identified and addressed—it is imperative to capture such dependencies in the scenario model. If pivotal events were known to be independent of each other, so that their probabilities could be combined multiplicatively, there would be less reason to analyze them in detail. Because pivotal events often share shared dependencies, their modeling in some detail is important.

Complex pivotal events can frequently be modeled using fault trees. A fault tree is a picture of a set of logical relationships between more complex (more aggregated) events such as system-level failures, and more basic (less aggregated) events such as component-level failures. Fault tree modeling is applicable not only to modeling of hardware failures, but also other complex event types as well, including descriptions of the circumstances surrounding software response and crew actions.

The mapping of scenarios into logic representations leans heavily on engineering analysis: physical simulation of system behavior in specified conditions, determination of time available for crew actions, determination of the severity of the consequences associated with scenarios. Behind every logic model is another body of modeling whose results are distilled into the logical relationships pictured in the scenario model. Assignment of system states into "success" or "failure" depends on such modeling, as does classification of scenarios into consequence categories.

Functionally, a fault tree is a deductive logic model where a top event, usually a system failure, is postulated, and reverse paths are developed to gradually link this top event with all subsystems, components, software errors, or human actions (in order of decreasing generality) that can contribute to the top event, down to those whose basic probability of failure (or success) is known and can be directly used for quantification. Graphically, a fault tree at its simplest consists of blocks (e.g., rectangles or circles) containing descriptions of failure modes and binary logic gates (e.g., union or intersection) that logically link basic failures through intermediate level failures to the top event. Figure 2-17 depicts a very simple fault tree structure.

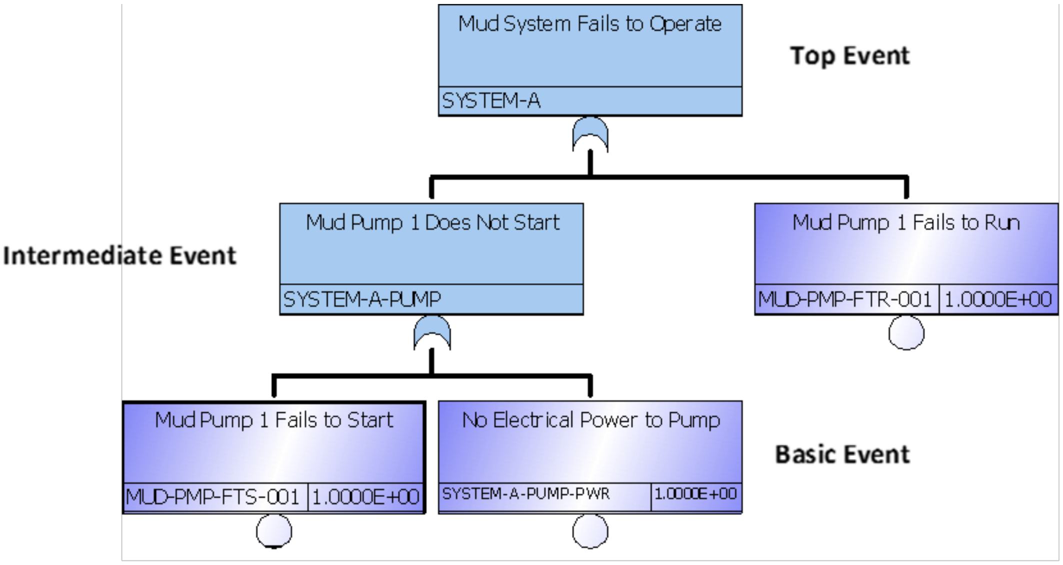

Figure 2- 17. Typical Fault Tree Structure and Symbols

Fault trees are constructed to define all significant failure combinations, called cutsets that lead to the top event. The result of a Boolean reduction of the fault tree results in combinations of failures that are the minimum set(s) required to result in the top event and are called minimal cut sets.

Ultimately, fault trees are graphical representations of Boolean expressions representing the minimal cut sets. For the fault tree in Figure 2-17, there are 3 minimal cut sets:

MUD-PMP-FTR-001;
MUD-PMP-FTS-001; and
SYSTEM-A-PUMP-PWR.

The corresponding Boolean equation for the fault tree is:

$$SYSTEM\text{-}A = MUD\text{-}PMP\text{-}FTR\text{-}001 \cup MUD\text{-}PMP\text{-}FTS\text{-}001 \cup SYSTEM\text{-}A\text{-}PUMP\text{-}PWR$$

More detail on minimal cut sets and the Boolean reduction that are the results of the fault tree are shown in Section 3.

2.2.5.3 Fault Tree Considerations

Developing a fault tree requires several considerations including:

- Identifying the objective and scope of the analysis;
- Determination of the level of detail; and
- Setting ground rules and naming conventions.

The objective and scope of the fault tree, in the context of a PRA analysis, is normally defined when the event sequences are being developed by constructing the ESDs/event trees. The critical systems/events required to respond to an initiating event are assessed by the processes in previous Sections, and incorporated as pivotal events in the event tree(s). These pivotal events become top events for the fault trees and should be worded in specific language as to highlight the failure mode of the event being analyzed based on the success criteria.

Simply labeling the top event as “Event A Fails” is generally inadequate as Event A may have different failure modes, and the objective of the analysis may only require specific ones be modeled. If extraneous failures are included in the analysis that do not contribute to the analysis objective, the results of the analysis will be erroneous. For instance, the Emergency Disconnect on a MODU has several functions including separation of the LMRP from the BOP and triggering the autoshear function on the BOP. The separation from the well is performed in an emergency situation for personnel safety to allow the MODU to move clear of the well. The intent of the autoshear function is to seal the well and prevent a hydrocarbon release. From the ESD developed in Figure 2-6, the objective of the analysis is to estimate the probability of a hydrocarbon release, so when developing the top event, only the contributors to failure of the autoshear function of the Emergency Disconnect needs to be included and the top event should be worded with that failure mode.

Defining the scope of the analysis includes understanding the initial configuration/operation of the system being analyzed. The initial state of the system will describe which components are active, which are in a standby state, and external conditions (e.g. if failure of the BOP blind shear ram is being analyzed, it is important to identify what operation is being performed, like running casing). The initial state of the components will determine the applicable failure modes for those components. A pump that is active may fail to operate while a pump in standby may fail to start or fail to operate. In cases where a component is in standby, the analysis may need to account for human error if manual activation is required for the system to start.

The level of detail on the causes resulting in the top event for a PRA analysis should be based on the level at which data is available, the objective of the analysis, and the interdependencies between systems and operations. Data analysis is discussed in Section 2.6 in detail, but generally data can be found at the major component level (e.g. pumps, valves, electrical busses, etc.) from a variety of sources. Going beyond the level that data is available will result in an unquantifiable fault tree. The objectives of the analysis must also be considered in determining the resolution of the fault tree. For a fault tree analysis on a BOP, the analysis could be performed at the yellow/blue pod level, or the analysis desires more detail down to the hydraulic component level to account for cross connect ability, for instance, that level may be modeled. Interdependencies such as cross connect capability may be a reason by itself to drive the analysis to the component level of detail.

The last consideration is setting up modeling ground rules in order to ensure consistency across the PRA. Establishing a naming convention for fault tree gates and particularly for basic events is necessary to be able to easily read cut sets and results from the analysis. Cut set Naming schemes for the basic events may include:

- The operation being performed (e.g. drilling)
- The system the component belongs to (BOP)
- The subsystem the component belongs to (e.g. Yellow pod)
- The component (e.g. shuttle valve)
- The failure mode (e.g. Fails to transfer)
- A unique identifier for the valve (usually from a drawing, e.g. SV01)

There is usually a character limit to the size of the basic event name, so abbreviations must be used for the above items such as BOP for blow out preventer, YPO for yellow pod, etc. As a minimum, the system, component, failure mode, and a unique identifier should be used when the naming scheme is developed. The overall naming scheme typically has a form like:

XXX-YYY-ZZZ-DDDDD

Where XXX corresponds to the system, YYY is the component, etc. The abbreviations for each are developed before modeling begins with the exception of the unique identifiers. The failure modes for active components should correspond to active failures and not a failed condition. For example, for a valve that is initially open and fails when commanded to close, the best way to express the failure mode is “fails to close” rather than “fails closed.” In the “fails closed” case it is not clear what the initial condition of the valve is, was the valve open and did not close when commanded, or was the valve initially closed and failed that way when commanded to open? Using the active word “to” in “fails to close” implies the valve is initially open.

A well thought out naming scheme for basic events is essential to avoid duplication of names for different events in different fault trees, particularly if multiple analysts are involved. If duplication does exist, the results produced for those events could be erroneous. Examples of typical naming conventions for failure modes and components is provided in Appendix A.

There can be some special events that are adapted to the naming scheme used for components or they may have their own separate scheme. For example, environmental conditions do not have a system or unique identifier associated with them and may therefore have a separate naming scheme developed for just those types of events.

Gate naming schemes may be more freeform since gates are not shown in the results. A consistent naming scheme for gates is advisable however to ensure that each gate is named uniquely and avoid having gates with different logic and the same name in different fault trees.

2.2.5.4 Fault Tree Symbols

Starting with the top event, the fault tree is developed by deductively determining the cause of the previous fault, continually approaching finer resolution until the limit of resolution is reached. In this fashion the fault tree is developed from the system end point backward to the failure source. The limit of resolution is reached when fault tree development below a gate consists only of basic events (i.e., faults that consist of component failures, faults that are not to be further developed, phenomenological events, support system faults that are developed in separate fault trees, software errors, or human actions). The logic of the fault tree is represented by symbols used for fault tree gates and basic events. The most common types of gates are shown in Figure 2-18 with a description of the logic for each. Appendix B gives a detailed explanation of how each gate is used and quantified. Other gate types such as "NOR" or "INHIBIT" exist but are rarely used.

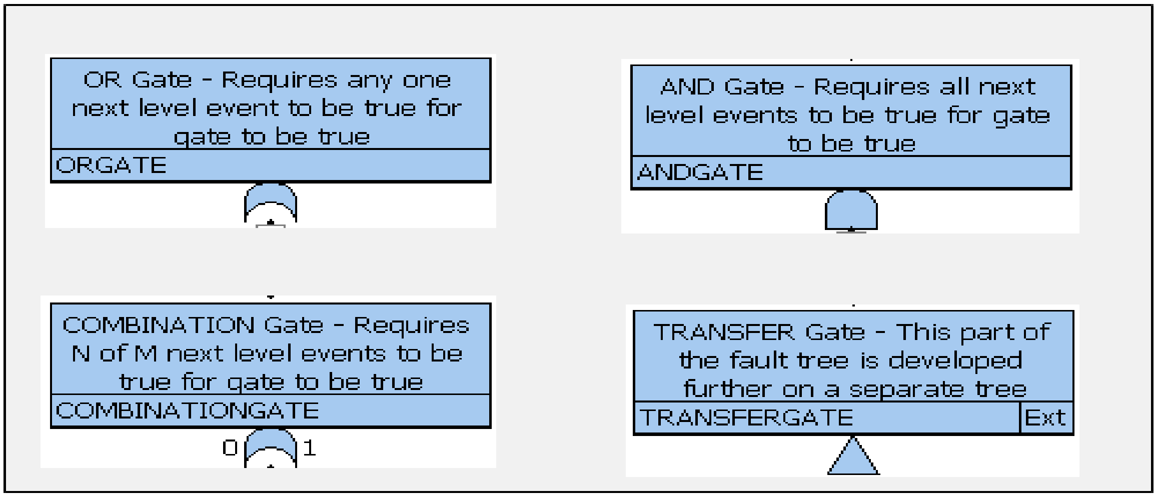
OR Gate - Requires any one next level event to be true for gate to be true
ORGATE
AND Gate - Requires all next level events to be true for gate to be true
ANDGATE
COMBINATION Gate - Requires N of M next level events to be true for gate to be true
COMBINATIONGATE
0
1
TRANSFER Gate - This part of the fault tree is developed further on a separate tree
TRANSFERGATE
Ext

Figure 2-18. Commonly Used Fault Tree Gates

The most common basic event types are shown in Figure 2-19 with a description of what they represent.

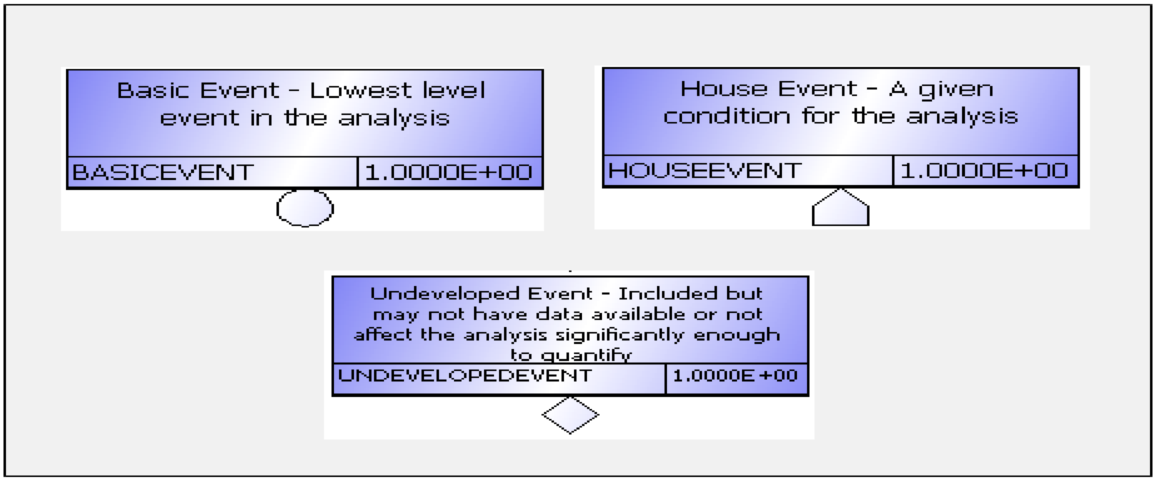
Basic Event - Lowest level event in the analysis
BASICEVENT
1.0000E+00
House Event - A given condition for the analysis
HOUSEEVENT
1.0000E+00
Undeveloped Event - Included but may not have data available or not affect the analysis significantly enough to quantify
UNDEVELOPEDEVENT
1.0000E+00

Figure 2-19. Commonly Used Basic Event Types

House events are often used in fault tree analysis as switches to turn logic on and off or represent a condition. If used as a switch, their probability is usually quantified as unity or zero, they require no reliability data input. House events are also frequently used to simulate conditional dependencies.

2.2.5.5 Simple Fault Tree Example

Using the steps described above, and going back to the example event tree provided in Figure 2-13, a simplified example of fault tree construction is developed for the “ANNULAR” event (2nd event after the initiator in Figure 2-13) which represents the failure of the annular preventer to block the annulus through the BOP. A simplified drawing of a BOP is shown in Figure 2-20.

The top event of the fault tree, as stated in the event tree, is “Annular preventer fails to close prior to the kick reaching the BOP or pressure is beyond the design of the annular.” The wording of the top event implies the initial condition is that the annular preventer is open and for success of this event, the annular preventer must close and prevent flow past the BOP. The simplified *example* diagram in Figure 2-20 shows that both the blue and yellow pods, used for control, are connected to the annular preventer. Either one is adequate to close the annular preventer and, in this simplified example, it is assumed that a crosstie exists. To switch pods, a manual action by the Driller is required.

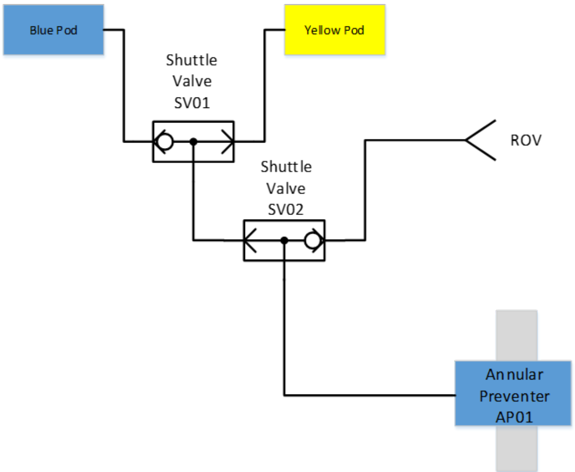

Figure 2- 20. Simple BOP Schematic

To develop the next level down in the fault tree, the design and operation is reviewed. In this example, one of the pods must provide hydraulic fluid to the preventer, the preventer itself must close, and the pressure must stay below the design pressure of the annular.

The failure of the annular preventer (BOP-CYL-FTC-AP01) is a singular event so is included under an “OR” gate as shown in Figure 2-21. For the purposes of this example, the basic event related to the pressure of the annular has been left as an undeveloped event and is also included under the OR gate (ANNULAROVERPRESSURE). The failure of the pods includes multiple events and combinations of events that must fail to satisfy the top event. Therefore, an intermediate “AND” gate is needed to develop this event (BOTHPODSFAIL) further. The left input (YELLOWPODFAILS) to the AND gate is an intermediate gate for the operating yellow pod, while the blue pod (BLUEPODFAILS) is the standby pod and addressed on the right hand side of the AND gate. For convenience, the portions of the blue and yellow pods (gate names – BLUEPODCOMMON, YELLOWPODCOMMON), have been made transfer events to allow these portions of the fault tree to be used with the blind shear, pipe, and casing shear rams. The transfer for each is also shown in Figure 2-21. For each pod, an OR gate is used with the inputs broken down into the pods themselves and the hydraulic paths from the pods. From Figure 2-18, the yellow path is aligned to the annular and the shuttle valves are in position to permit flow, so the only applicable failure mode considered is external valve leakage. Since the flow passes through both valves, both are included (BOP-SHV-LKG-SV01, BOP-SHV-LKG-SV02).

The blue pod side needs to be treated differently because it is in a standby state. Because the pods are manually selected, a basic event for the Driller failing to select the blue pod after the yellow pod fails is added (BOP-HUM-ERR-XTIEPODS). On the hydraulic path intermediate event, a basic event for the crosstie shuttle valve failing to transfer to the correct position is added (BOP-

SHV-FTT-SV01). From Figure 2-21, it should be noted that several events are included on both the yellow and blue pods, including the 2 shuttle valve external leakages and the common cause failure of both the yellow and blue pods. In a fault tree, events or gates may be used in multiple areas and when the fault tree is solved, the cut sets produced will be reduced and will not contain any duplicates.

The result of solving the ANNULAR fault tree in Figure 2-21 is shown in Table 2-3. Using the logic of the fault tree, the inputs are reduced to the “minimal cut sets” that result in the top event. The minimal cut sets are those failures or combinations of failures that if any of the basic events were not true, the top event would not be true. In Table 2-3, the first 5 cut sets are all single basic events that would result in the top event, while cut sets 6, 7, and 8 are double failures. When the basic events are assigned values, a ranked listing is produced.

Table 2-3. ANNULAR Fault Tree Minimal Cut Sets

#	Cut Set	Description
1	BOP-SHV-LKG-SV01	Crosstie shuttle valve external leakage
2	BOP-SHV-LKG-SV02	ROV shuttle valve external leakage
3	BOP-CYL-FTC-AP01	Annular preventer fails to seal
4	BOP-POD-YLBL-CCF	Common cause failure of blue and yellow pods
5	ANNULAROVERPRESS	Well pressure over the design limit of annular
6	BOP-POD-FTR-BLUE	Blue pod (standby) fails to run
	BOP-POD-FTR-YELLOW	Yellow pod (operating) fails to run
7	BOP-POD-FTR-YELLOW	Yellow pod (operating) fails to run
	BOP-SHV-FTT-SV01	Crosstie shuttle valve fails to transfer to blue pod
8	BOP-HUM-ERR-XTIEPODS	Driller fails to select blue pod after yellow pod failure
	BOP-POD-FTR-YELLOW	Yellow pod (operating) fails to run

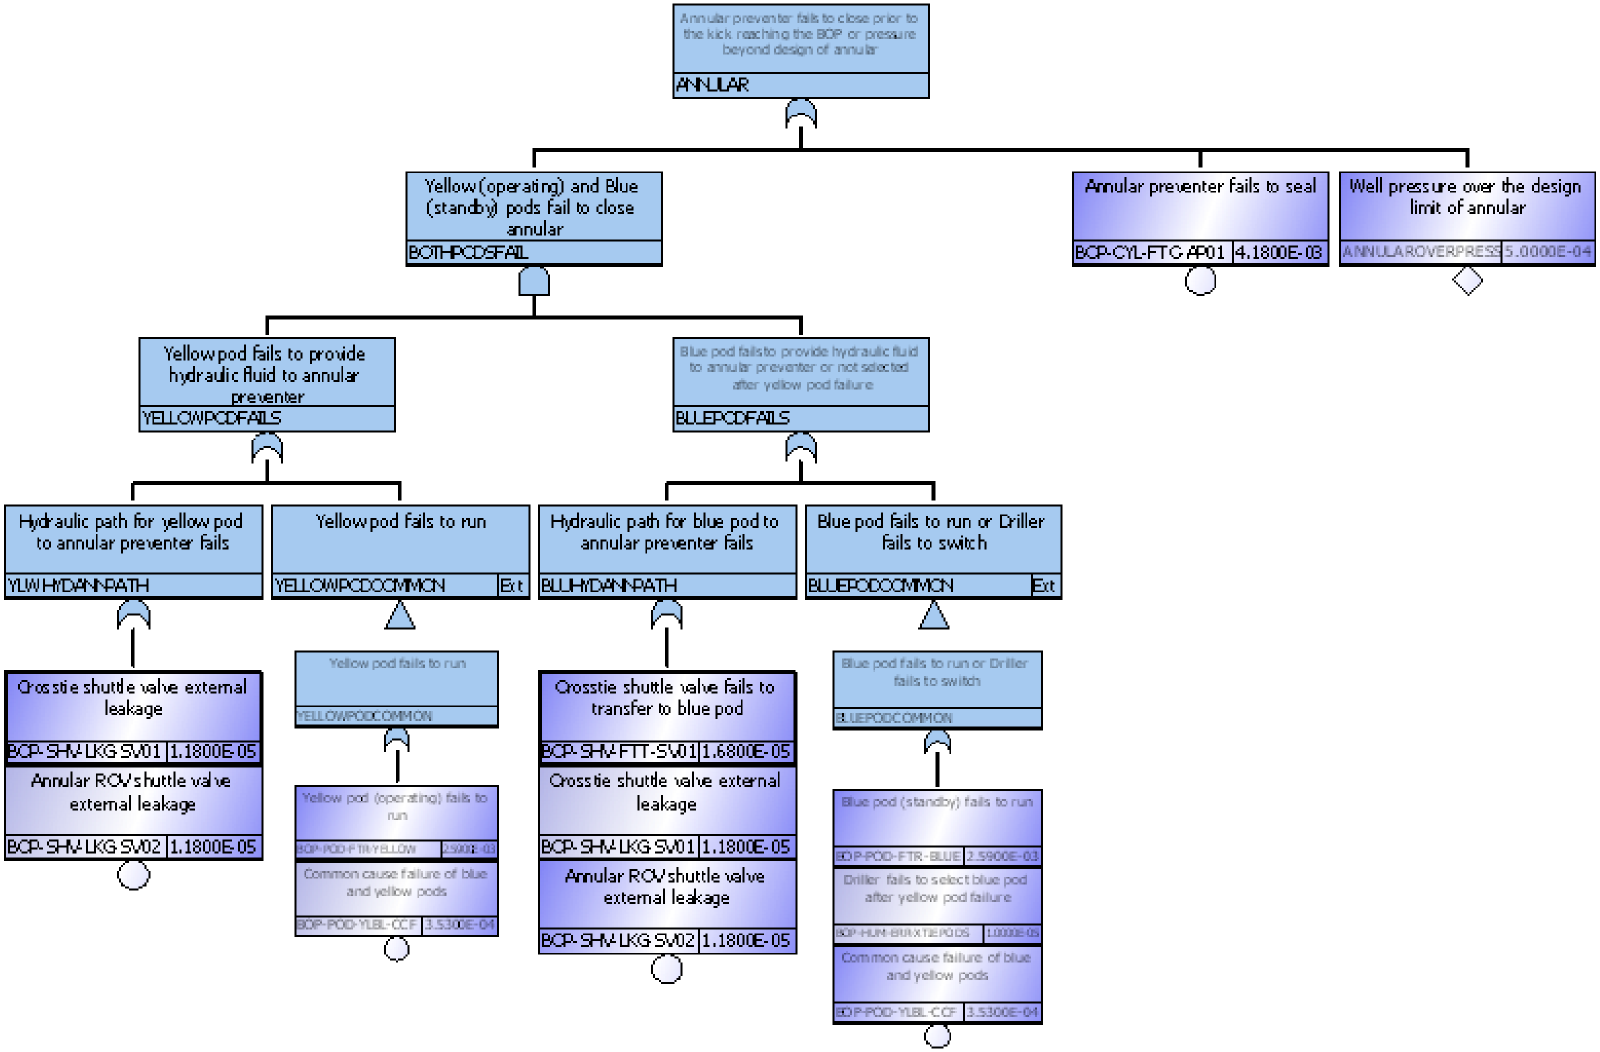
Annular preventer fails to close prior to the kick reaching the BOP or pressure beyond design of annular
ANNULAR
Yellow (operating) and Blue (standby) pods fail to close annular
BOTHPODSFAIL
Annular preventer fails to seal
BCP-CYL-FTC AP01
4.1800E-03
Well pressure over the design limit of annular
ANNULAROVERPRESS
5.0000E-04
Yellow pod fails to provide hydraulic fluid to annular preventer
YELLOWPODFAILS
Blue pod fails to provide hydraulic fluid to annular preventer or not selected after yellow pod failure
BLUEPODFAILS
Hydraulic path for yellow pod to annular preventer fails
YLWHYDANNPATH
Yellow pod fails to run
YELLOWPODCOMMON
Ext
Hydraulic path for blue pod to annular preventer fails
BLUHYDANNPATH
Blue pod fails to run or Driller fails to switch
BLUEPODCOMMON
Ext
Crosstie shuttle valve external leakage
BCP-SHV-LKG SV01
1.1800E-05
Annular ROV shuttle valve external leakage
BCP-SHV-LKG SV02
1.1800E-05
Yellow pod fails to run
YELLOWPODCOMMON
Yellow pod (operating) fails to run
BOP-POD-FTR-YELLOW
2.5900E-03
Common cause failure of blue and yellow pods
BOP-POD-YLBL-CCF
3.5300E-04
Crosstie shuttle valve fails to transfer to blue pod
BCP-SHV-FTT-SV01
1.6800E-05
Crosstie shuttle valve external leakage
BCP-SHV-LKG SV01
1.1800E-05
Annular ROV shuttle valve external leakage
BCP-SHV-LKG SV02
1.1800E-05
Blue pod fails to run or Driller fails to switch
BLUEPODCOMMON
Blue pod (standby) fails to run
BOP-POD-FTR-BLUE
2.5900E-03
Driller fails to select blue pod after yellow pod failure
BOP-HUM-BUR-XTIE PODS
1.0000E-05
Common cause failure of blue and yellow pods
BOP-POD-YLBL-CCF
3.5300E-04

Figure 2- 21. Basic Fault Tree

2.2.5.6 Modeling Support Systems

Often a function being modeled will have dependencies on other systems. From the simple example in Figure 2-17, the lower right hand basic event is related to electric power that energizes the pump. The electric power system is a separate system and may provide support for many systems. Since a fault tree is typically detailed down to the component level where data exists, the fault tree in Figure 2-17 would normally have the electric power feeding the pump detailed down to the boundary of the analysis, which on an offshore rig would be the diesel generator(s) (and maybe the fuel and air intake systems) as shown in the simplified fault tree in Figure 2-22.

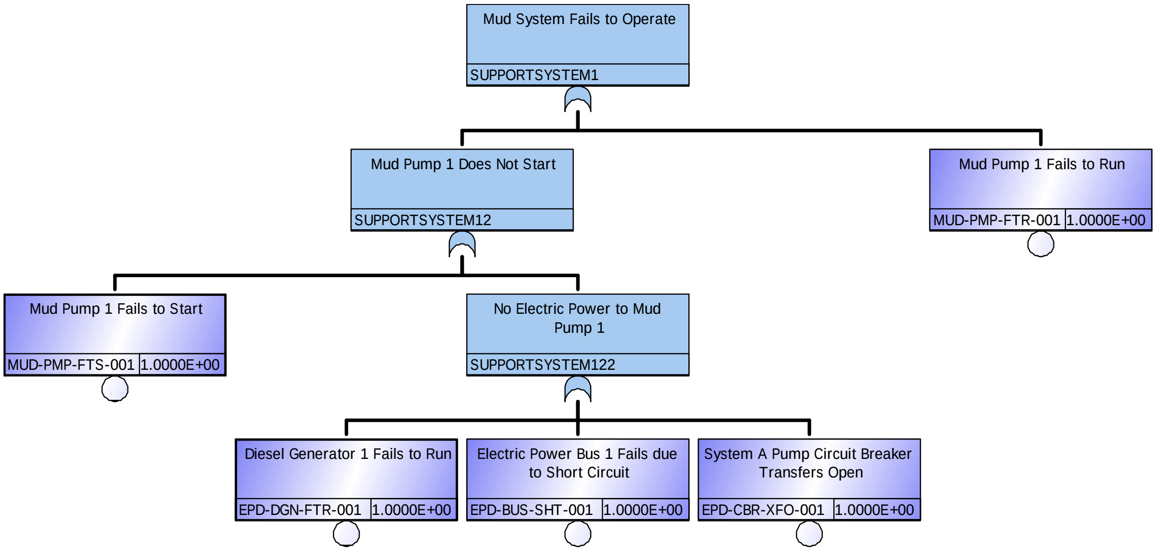
Mud System Fails to Operate
SUPPORTSYSTEM1
Mud Pump 1 Does Not Start
SUPPORTSYSTEM12
Mud Pump 1 Fails to Run
MUD-PMP-FTR-001
1.0000E+00
Mud Pump 1 Fails to Start
MUD-PMP-FTS-001
1.0000E+00
No Electric Power to Mud Pump 1
SUPPORTSYSTEM122
Diesel Generator 1 Fails to Run
EPD-DGN-FTR-001
1.0000E+00
Electric Power Bus 1 Fails due to Short Circuit
EPD-BUS-SHT-001
1.0000E+00
System A Pump Circuit Breaker Transfers Open
EPD-CBR-XFO-001
1.0000E+00

Figure 2-22. Simple Fault Tree with Support System Modeled

From Figure 22, the electric power support is shown to be composed of the diesel generator, power bus, and the circuit breaker that feeds mud pump 1. In Figure 2-22, the diesel generator and electric power bus likely would feed multiple items, and should be separated so they can be modeled only once and then used where necessary. Figure 2-23 shows the proper way to model this situation. The common parts of the electric power system have been separated under a separate OR gate and generically labeled as power from diesel generator1. If these components are needed in another fault tree the OR gate can be made a transfer event as shown in Figure 2-24, and that transfer event can be used wherever needed.

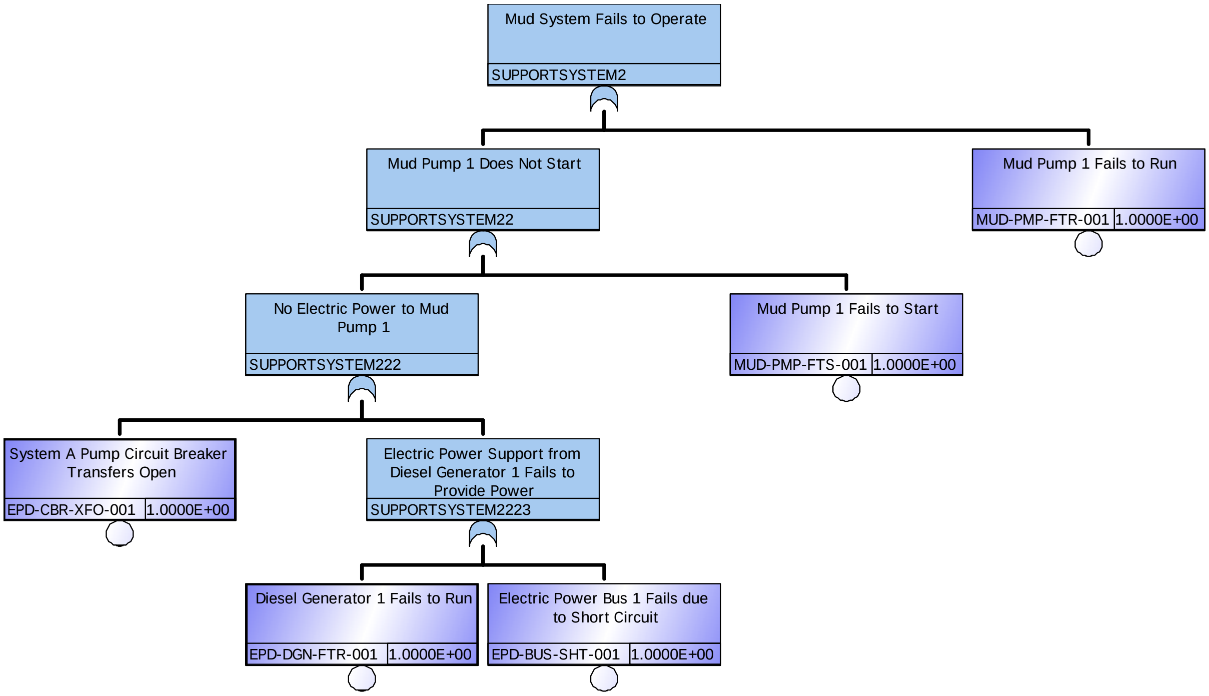
Mud System Fails to Operate
SUPPORTSYSTEM2
Mud Pump 1 Does Not Start
SUPPORTSYSTEM22
Mud Pump 1 Fails to Run
MUD-PMP-FTR-001
1.0000E+00
No Electric Power to Mud Pump 1
SUPPORTSYSTEM222
Mud Pump 1 Fails to Start
MUD-PMP-FTS-001
1.0000E+00
System A Pump Circuit Breaker Transfers Open
EPD-CBR-XFO-001
1.0000E+00
Electric Power Support from Diesel Generator 1 Fails to Provide Power
SUPPORTSYSTEM2223
Diesel Generator 1 Fails to Run
EPD-DGN-FTR-001
1.0000E+00
Electric Power Bus 1 Fails due to Short Circuit
EPD-BUS-SHT-001
1.0000E+00

Figure 2-23. Support System (Diesel Generator 1 and the power bus) Modeled so it can be used in Multiple Fault Trees

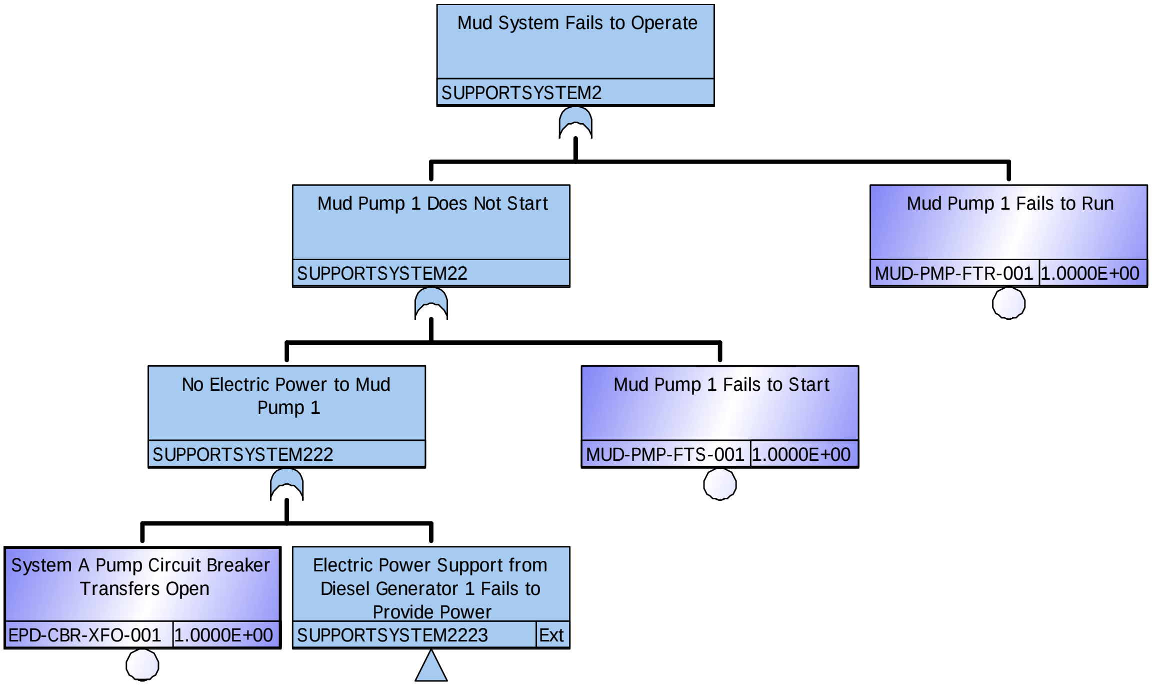
Mud System Fails to Operate
SUPPORTSYSTEM2
Mud Pump 1 Does Not Start
SUPPORTSYSTEM22
Mud Pump 1 Fails to Run
MUD-PMP-FTR-001
1.0000E+00
No Electric Power to Mud Pump 1
SUPPORTSYSTEM222
Mud Pump 1 Fails to Start
MUD-PMP-FTS-001
1.0000E+00
System A Pump Circuit Breaker Transfers Open
EPD-CBR-XFO-001
1.0000E+00
Electric Power Support from Diesel Generator 1 Fails to Provide Power
SUPPORTSYSTEM2223
Ext

Figure 2- 24. Power from Diesel Generator 1 Modeled as a Transfer that will be used in Multiple Fault Trees

2.2.5.6.1 Dependency Matrices

Because system dependencies may get complex, PRA analysts may map the system relationships out prior to starting a fault tree to ensure all dependencies are properly accounted for. A method for doing that has typically been developing a system dependency matrix. Figure 2-25 shows a simple block diagram of typical systems on a drill ship. The support systems (electric power, cooling, etc.) are those that do not directly contribute to the accomplishment of a primary task such as drilling, and frontline systems are those that are used to accomplish a primary task such as drilling (mud, drawworks, etc.)

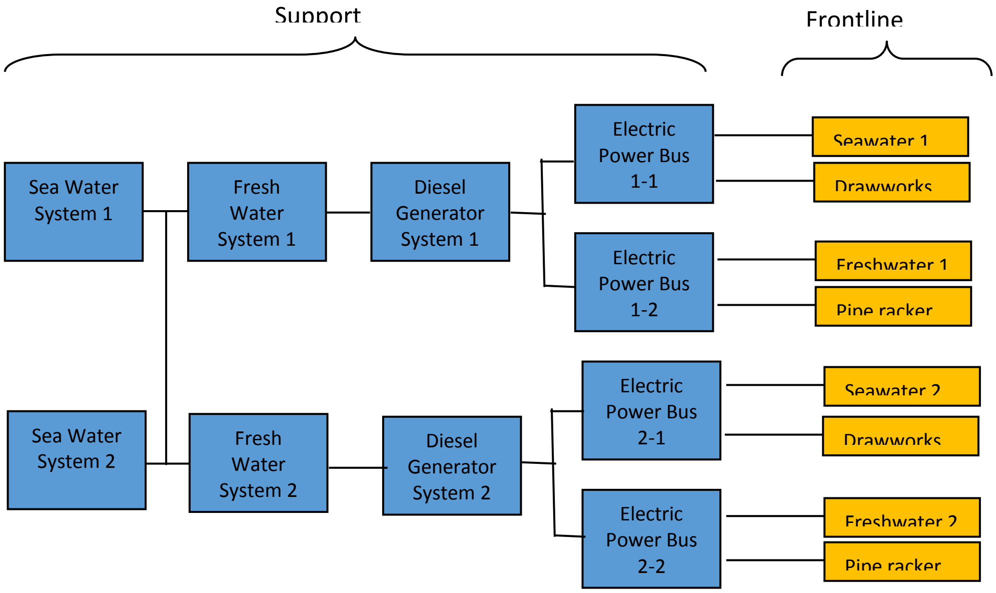

Figure 2- 25. Simplified System Block Diagram

Knowing the relationships in Figure 2-25, a dependency matrix may be constructed as shown in Figure 2-26. Using the dependency matrix assists the PRA analyst in ensuring the correct support system dependencies are modeled. The added notes should detail any special situations such as crossties.

When modeling the support systems, occasionally there are “loops” in the systems. For instance, from Figure 2-26, the seawater system supports the fresh water system, which in turn supports the diesel generator. The diesel generator, however, powers the sea water system and the fresh water system. From a modeling perspective, support systems should be modeled in the fault tree the way they support the frontline system being modeled. For a fault tree of the drawworks, the first support system put in, based on this example, would be electric power, the supporting busses and the diesel generator. The fresh water system would be next as a support to the diesel generator, and finally the sea water system supporting the fresh water system. In this case, the diesel generator, the electric bus (1-1) are already in the model, so there is no need to create a “loop” in the fault tree. The one missing element would be electric power bus 1-2 which powers the fresh water system, but not the drawworks. The specific electric bus (1-2) would have to be included separately with the fresh water system for this example.

	Sea Water System 1	Sea Water System 2	Fresh Water System 1	Fresh Water System 2	Diesel Generator System 1	Diesel Generator System 2	Electric Power Bus 1-1	Electric Power Bus 1-1	Electric Power Bus 2-1	Electric Power Bus 2-2	Drawworks	Pipe Racker	…….
Sea Water System 1	■		A	A									
Sea Water System 2		■	A	A									
Fresh Water System 1			■		X								
Fresh Water System 2				■		X							
Diesel Generator System 1					■		X	X					
Diesel Generator System 2						■			X	X			
Electric Power Bus 1-1	X						■				B		
Electric Power Bus 1-2			X					■				C	
Electric Power Bus 2-1		X							■		B		
Electric Power Bus 2-2				X						■		C	
……..													

Figure 2- 26. Example Dependency Matrix

Notes:

A – Seawater Systems 1 and 2 are normally separated can be crosstied so either System 1 or 2 can be used to cool one or both freshwater systems.

B – The Drawworks can be used from either electric power bus 1-1 or 2-1.

C – The Pipe Racker can be used from either electric power bus 1-2 or 2-2.

2.2.5.7 Modeling Common Cause

For complex systems with redundancy, common cause failure of like components can be a major risk contributor. The specifics on how common cause is evaluated and quantified is shown in Appendix D. This section discusses the options on how it should be represented in the fault tree model as basic events.

Going back to Figure 2-21, the basic fault tree example of the annular preventer, common cause was modeled for the blue/yellow pods designated by the basic event name ending in CCF. In this example, the basic event BOP-POD-YLBL-CCF is included under the intermediate gates for both the yellow pod and the blue pod. Since the yellow and blue pods are under an AND gate, and both are require to fail for the top event to be true, this one basic event satisfies that condition. The result after solving the fault tree is shown in Table 2-3, where the fourth cut set is a single common cause event. Repeating a common cause basic event wherever the effect is appropriate, in this case with the yellow and blue pods, is the best method of modeling common cause because it maintains the relationship of the basic event to the intermediate events. When fault tree transfers are needed, this can be important to ensure accuracy in the model. For systems that have three or more redundant components /system s (e.g. dynamic positioning thrusters), this will lead to multiple common cause events under each thruster as shown in Figure 2-27 (shown as stacked basic events for simplicity). In Figure 2-27, the DPS thruster 1 has been filled out with all common cause terms and the appropriate ones involving thruster 1 have also been included for the other three thrusters.

For highly redundant systems, PRA analysts sometimes include only a global common cause term, one that accounts for all components failing. This is done for simplicity and can be a good approximation as the global common cause terms are often the dominant contributor for common cause failure. If this approximation is considered appropriate, the single basic event can be included under each system as previously described, or as a single basic event at the same fault tree level as the AND or N of M gate modeling the redundancy.

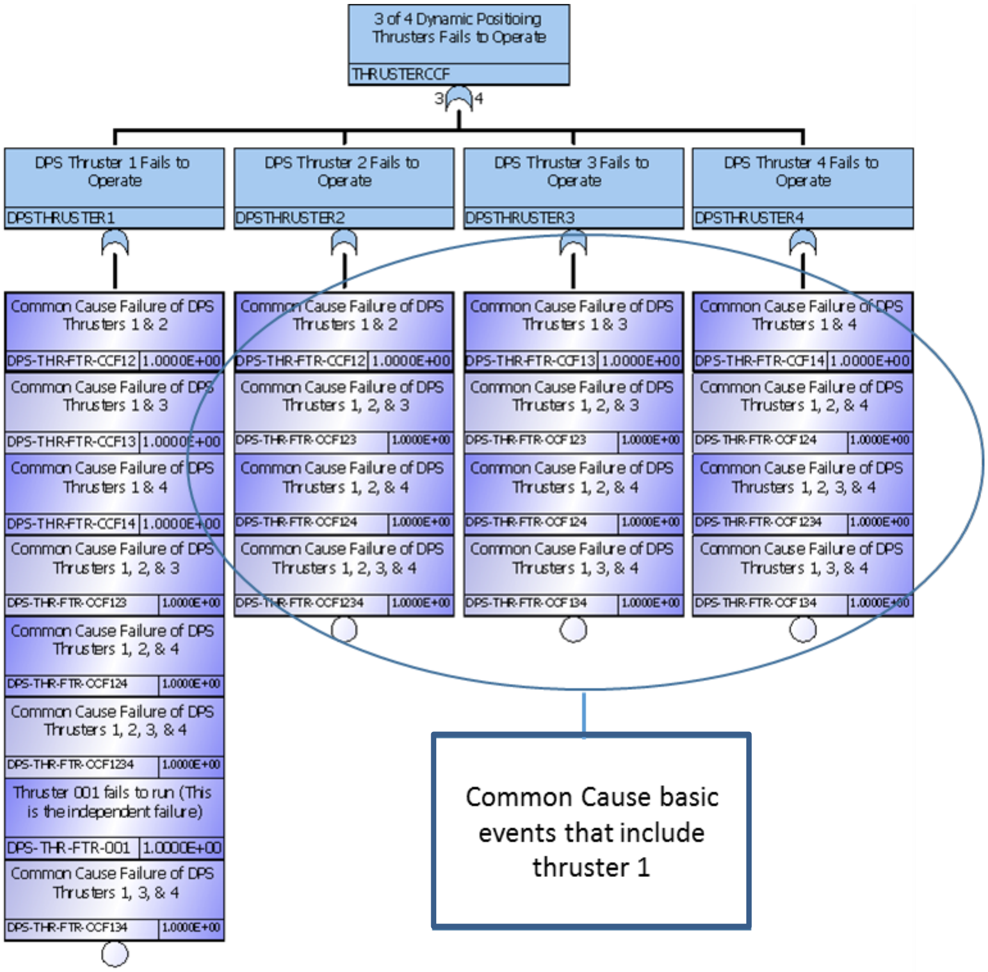
3 of 4 Dynamic Positioing Thrusters Fails to Operate
THRUSTERCCF
3
4
DPS Thruster 1 Fails to Operate
DPSTHRUSTER1
DPS Thruster 2 Fails to Operate
DPSTHRUSTER2
DPS Thruster 3 Fails to Operate
DPSTHRUSTER3
DPS Thruster 4 Fails to Operate
DPSTHRUSTER4
Common Cause Failure of DPS Thrusters 1 & 2
DPS-THR-FTR-CCF12 1.0000E+00
Common Cause Failure of DPS Thrusters 1 & 3
DPS-THR-FTR-CCF13 1.0000E+00
Common Cause Failure of DPS Thrusters 1 & 4
DPS-THR-FTR-CCF14 1.0000E+00
Common Cause Failure of DPS Thrusters 1, 2, & 3
DPS-THR-FTR-CCF123 1.0000E+00
Common Cause Failure of DPS Thrusters 1, 2, & 4
DPS-THR-FTR-CCF124 1.0000E+00
Common Cause Failure of DPS Thrusters 1, 2, 3, & 4
DPS-THR-FTR-CCF1234 1.0000E+00
Thruster 001 fails to run (This is the independent failure)
DPS-THR-FTR-001 1.0000E+00
Common Cause Failure of DPS Thrusters 1, 3, & 4
DPS-THR-FTR-CCF134 1.0000E+00
Common Cause Failure of DPS Thrusters 1 & 2
DPS-THR-FTR-CCF12 1.0000E+00
Common Cause Failure of DPS Thrusters 1, 2, & 3
DPS-THR-FTR-CCF123 1.0000E+00
Common Cause Failure of DPS Thrusters 1, 2, & 4
DPS-THR-FTR-CCF124 1.0000E+00
Common Cause Failure of DPS Thrusters 1, 2, 3, & 4
DPS-THR-FTR-CCF1234 1.0000E+00
Common Cause Failure of DPS Thrusters 1 & 3
DPS-THR-FTR-CCF13 1.0000E+00
Common Cause Failure of DPS Thrusters 1, 2, & 3
DPS-THR-FTR-CCF123 1.0000E+00
Common Cause Failure of DPS Thrusters 1, 2, & 4
DPS-THR-FTR-CCF124 1.0000E+00
Common Cause Failure of DPS Thrusters 1, 3, & 4
DPS-THR-FTR-CCF134 1.0000E+00
Common Cause Failure of DPS Thrusters 1 & 4
DPS-THR-FTR-CCF14 1.0000E+00
Common Cause Failure of DPS Thrusters 1, 2, & 4
DPS-THR-FTR-CCF124 1.0000E+00
Common Cause Failure of DPS Thrusters 1, 2, 3, & 4
DPS-THR-FTR-CCF1234 1.0000E+00
Common Cause Failure of DPS Thrusters 1, 3, & 4
DPS-THR-FTR-CCF134 1.0000E+00
Common Cause basic events that include thruster 1

Figure 2- 27. Common Cause Modeling for a 3 of 4 System

2.2.5.8 Modeling Conditionality

The house event, shown in Figure 2-19, is used to show whether a particular condition that affects the analysis is present or not. This is often used as a switch by the analyst to turn a condition on (set the event probability to 1.0) or off (set the event probability to 0.0) and see what the affect is on the results. In other cases there may be a condition that exists a fraction of the time, and the logic that satisfies the top event changes depending on whether the condition is present or not. A case like this was assumed for the IBOPFLTVLV event in Figure 2-13. This top event considers the failure to isolate the path up the drill string, and it was assumed that a float valve may or may not be present when the kick occurs. The simplified fault tree for this event is shown in Figure 2-28.

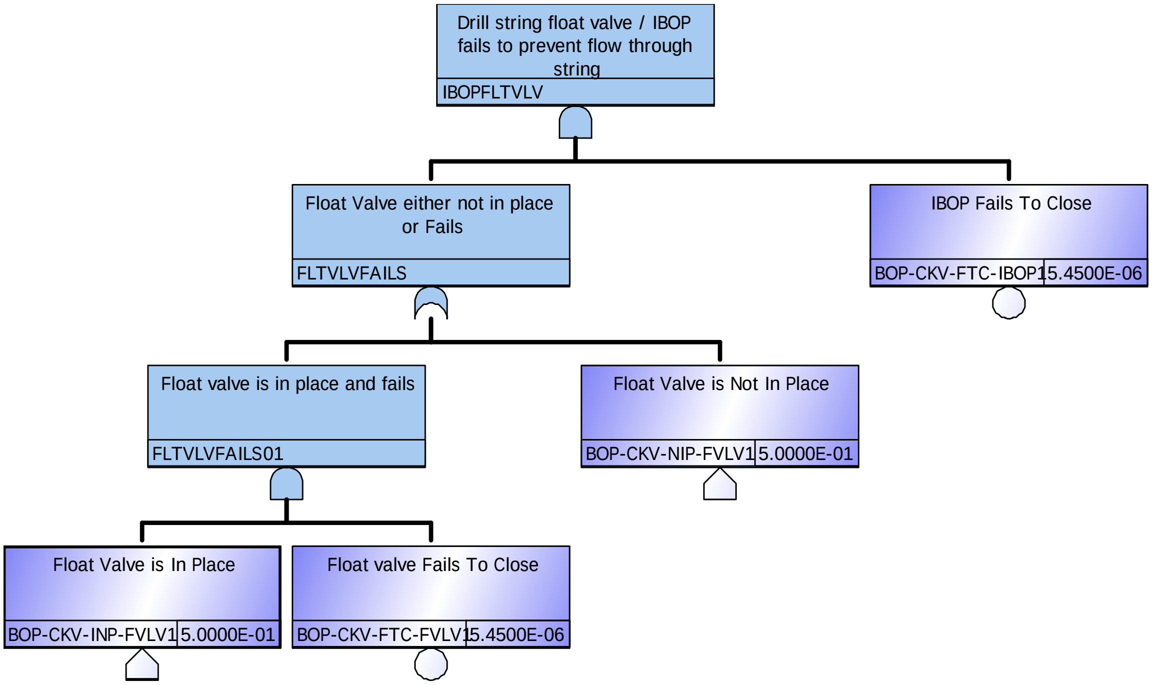
Drill string float valve / IBOP fails to prevent flow through string
IBOPFLTVLV
Float Valve either not in place or Fails
FLTVLVFAILS
IBOP Fails To Close
BOP-CKV-FTC-IBOP1
5.4500E-06
Float valve is in place and fails
FLTVLVFAILS01
Float Valve is Not In Place
BOP-CKV-NIP-FVLV1
5.0000E-01
Float Valve is In Place
BOP-CKV-INP-FVLV1
5.0000E-01
Float valve Fails To Close
BOP-CKV-FTC-FVLV1
5.4500E-06

Figure 2- 28. Modeling Conditionality in a Fault Tree

As can be seen in Figure 2-28, there is an intermediate event (OR gate) for the float valve that models the two conditions that may exist, the float valve is in place or it is not. For the condition that it is not in place, the resulting cut set is that the IBOP fails and for the fraction of time that there is no float valve in place, the top event will be true. For the time when the float valve is in place, an AND gate is used because the top event being true also requires the float valve to fail. The two cut sets are shown in Table 2-4.

Table 2-4. Cut Sets for Figure 2-28 Fault Tree

Cut set	Basic Event Name	Basic Event Description
1	BOP-CKV-FTC-IBOP1	IBOP Fails To Close
	BOP-CKV-NIP-FVLV1	Float Valve is Not In Place
2	BOP-CKV-FTC-FVLV1	Float valve Fails To Close
	BOP-CKV-FTC-IBOP1	IBOP Fails To Close
	BOP-CKV-INP-FVLV1	Float Valve is In Place

When modeling these type of conditions, the analyst needs to ensure that the dependence is maintained. In this example, the condition of the float valve being present or not is the only two possibilities, so the probabilities must add up to 1.0 when combined.

2.2.5.9 Modeling Maintenance

To Be Added

2.2.5.10 Modeling Initiating Events

To Be Added

2.2.5.11 Linking Fault Trees and Event Trees

Once the event trees and their associated fault trees have been developed and linked as shown in Figure 2-29, the qualitative part of the PRA model is completed. Fault trees and event trees are said to be "linked" when the fault trees for pivotal events and the event trees containing those pivotal events are tied together properly in the software being used to evaluate the accident sequence cut sets. The scenarios are formed from the basic events and fault tree logic combined with the event tree sequences and end states. The model can now be evaluated qualitatively to review individual scenarios. Using the event tree model previously developed in Section 2.2.4 in Figure 2-13, simplified fault trees such as the one in Figure 2-21 for the annular preventer were developed. Table 2-4 shows a sample of the output from the model in terms of cut sets.

Each cut set has the initiating event, DRILLINGKICK, followed by other basic events whose combined occurrence leads to the end state. Typically all the events shown are failure basic events, however, as shown in some of the cut sets (e.g. 1,3), success events are shown there, /ROV, /CAPSTACK. Success events are usually ignored because the values of failure are so small. With the very small failure values, the success approximates 1.0 so does not affect the results. In the cases of the cut sets in Table 2-4 with the success terms, failure events like ROV and CAPSTACK are assumed to be much larger, and therefore success terms must be accounted for to provide accurate results.

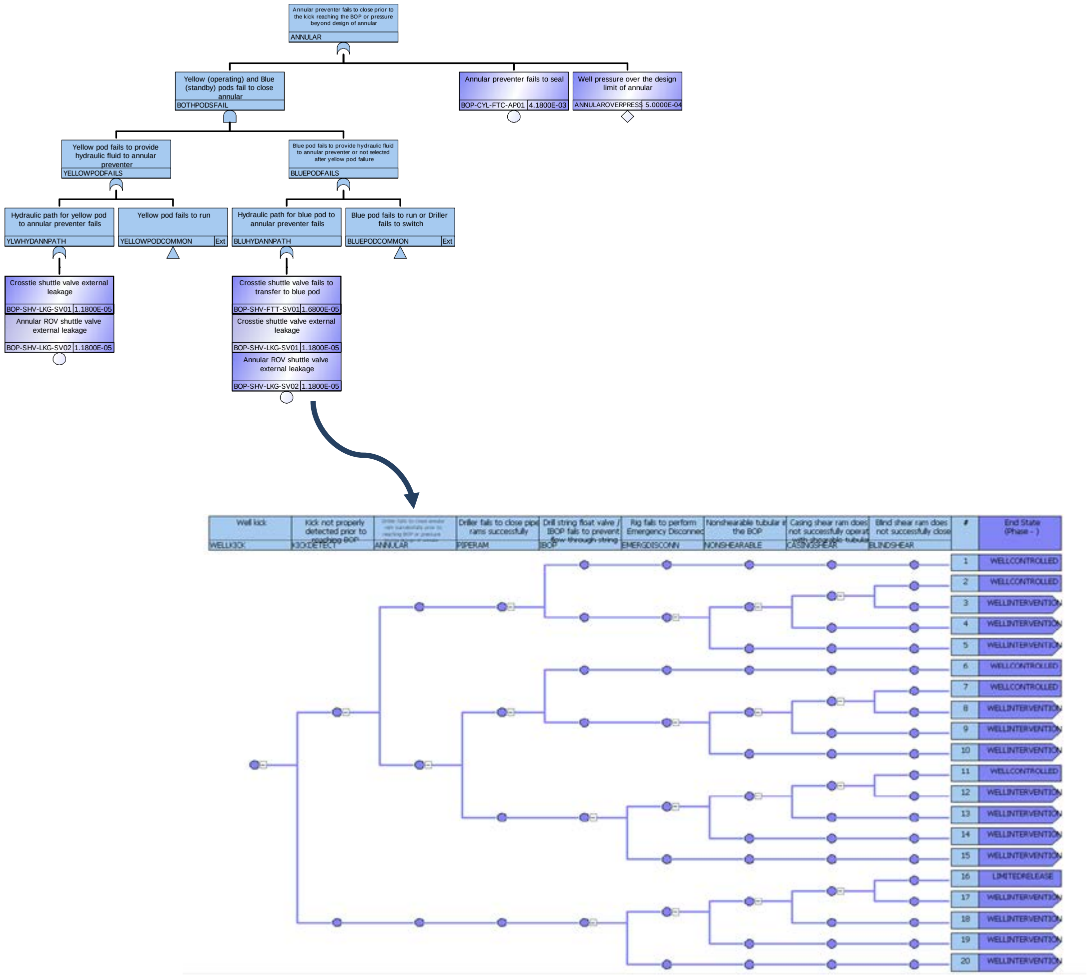
Annular preventer fails to close prior to the kick reaching the BOP or pressure beyond design of annular
ANNULAR
Yellow (operating) and Blue (standby) pods fail to close annular
BOTHPODSFAIL
Annular preventer fails to seal
BOP-CYL-FTC-AP01
4.1800E-03
Well pressure over the design limit of annular
ANNULAROVERPRESS
5.0000E-04
Yellow pod fails to provide hydraulic fluid to annular preventer
YELLOWPODFAILS
Blue pod fails to provide hydraulic fluid to annular preventer or not selected after yellow pod failure
BLUEPODFAILS
Hydraulic path for yellow pod to annular preventer fails
YLWHYDANNPATH
Yellow pod fails to run
YELLOWPODCOMMON
Ext
Hydraulic path for blue pod to annular preventer fails
BLUHYDANNPATH
Blue pod fails to run or Driller fails to switch
BLUEPODCOMMON
Ext
Crosstie shuttle valve external leakage
BOP-SHV-LKG-SV01
1.1800E-05
Annular ROV shuttle valve external leakage
BOP-SHV-LKG-SV02
1.1800E-05
Crosstie shuttle valve fails to transfer to blue pod
BOP-SHV-FTT-SV01
1.6800E-05
Crosstie shuttle valve external leakage
BOP-SHV-LKG-SV01
1.1800E-05
Annular ROV shuttle valve external leakage
BOP-SHV-LKG-SV02
1.1800E-05
Well kick
WELLKICK
Kick not properly detected prior to reaching BOP
KICKDETECT
ANNULAR
Driller fails to close pipe rams successfully
PIPERAM
Drill string float valve / IBOP fails to prevent flow through string
IBOP
Rig fails to perform Emergency Disconnect
EMERGDISCONN
Nonshearable tubular in the BOP
NONSHEARABLE
Casing shear ram does not successfully operate with shearable tubular
CASINGSHEAR
Blind shear ram does not successfully close
BLINDSHEAR
#
End State (Phase -)
1 WELLCONTROLLED
2 WELLCONTROLLED
3 WELLINTERVENTION
4 WELLINTERVENTION
5 WELLINTERVENTION
6 WELLCONTROLLED
7 WELLCONTROLLED
8 WELLINTERVENTION
9 WELLINTERVENTION
10 WELLINTERVENTION
11 WELLCONTROLLED
12 WELLINTERVENTION
13 WELLINTERVENTION
14 WELLINTERVENTION
15 WELLINTERVENTION
16 LIMITEDRELEASE
17 WELLINTERVENTION
18 WELLINTERVENTION
19 WELLINTERVENTION
20 WELLINTERVENTION

Figure 2- 29. Fault Tree Linked to Event Tree

Tracing a single cut set through the event tree shows how linking the fault trees through the event tree accounts for support system dependencies that may be common through different top events. Cut set 1 from Table 2-5 is a relatively simple combination of events, the initiator – DRILLINGKICK, a common cause failure of both the blue and yellow pods - BOP-POD-YLBL-CCF, and success of the ROV - /ROV. The end state is listed as LARGERELEASEROV, and the sequence number is 14-1, which in this case is related to sequence 14 on the initiating event tree (from Figure 2-13) and the 1 is the sequence from the transfer tree (from Figure 2-15). The path is traced out in Figure 2-30.

When cut set 1 (Sequence 14-1) is traced through the event tree, the path shows that the annular preventer, the pipe rams, the casing shear ram, and the blind shear ram have all failed (they are all on the down paths). The basic event representing common cause failure of both the blue and yellow pods - BOP-POD-YLBL-CCF, is found in each of the fault trees for the failed events, and therefore has caused all of those BOP functions to fail. The proper modeling of dependencies like this allow large integrated models of complex systems to sort through the integrated system and reduce the failures to the minimal combinations of interest.

Table 2-5. Sample Cut Sets from Linked Fault Tree/Event Tree Model

#	Cut set	Description
1	DRILLING : 14-1	End State LARGERELEASEROV
	DRILLINGKICK	Well Kick While Drilling
	BOP-POD-YLBL-CCF	Common cause failure of blue and yellow pods
	/ROV	ROV intervention unsuccessful
2	DRILLING : 16	End State LIMITEDRELEASE
	DRILLINGKICK	Well Kick While Drilling
	DRL-HUM-ERR-001	Kick not properly detected
3	DRILLING : 14-2	End State LARGERELEASECAP
	DRILLINGKICK	Well Kick While Drilling
	BOP-POD-YLBL-CCF	Common cause failure of blue and yellow pods
	/CAPSTACK	Well Capping unsuccessful
	ROV-FTR-001	ROV intervention unsuccessful
4	DRILLING : 14-3	End State LARGERELEASERELIEF
	DRILLINGKICK	Well Kick While Drilling
	BOP-POD-YLBL-CCF	Common cause failure of blue and yellow pods
	CAP-LKG-001	Well capping unsuccessful
	/RELIEFWELL	Relief Well unsuccessful
	ROV-FTR-001	ROV intervention unsuccessful
5	DRILLING : 14-1	End State LARGERELEASEROV
	DRILLINGKICK	Well Kick While Drilling
	BOP-POD-FTR-BLUE	Blue pod (standby) fails to run
	BOP-POD-FTR-YELLOW	Yellow pod (operating) fails to run
	/ROV	ROV intervention unsuccessful
6	DRILLING : 14-2	End State LARGERELEASECAP
	DRILLINGKICK	Well Kick While Drilling
	BOP-POD-FTR-BLUE	Blue pod (standby) fails to run
	BOP-POD-FTR-YELLOW	Yellow pod (operating) fails to run
	/CAPSTACK	Well Capping unsuccessful
	ROV-FTR-001	ROV intervention unsuccessful
7	DRILLING : 14-4	End State LARGERELEASERELIEF2
	DRILLINGKICK	Well Kick While Drilling
	BOP-POD-YLBL-CCF	Common cause failure of blue and yellow pods
	CAP-LKG-001	Well capping unsuccessful
	REL-WELL-LKG-001	Relief well not successful on first attempt
	ROV-FTR-001	ROV intervention unsuccessful
8	DRILLING : 14-3	End State LARGERELEASERELIEF
	DRILLINGKICK	Well Kick While Drilling
	BOP-POD-FTR-BLUE	Blue pod (standby) fails to run
	BOP-POD-FTR-YELLOW	Yellow pod (operating) fails to run
	CAP-LKG-001	Well capping unsuccessful
	/RELIEFWELL	Relief Well unsuccessful
	ROV-FTR-001	ROV intervention unsuccessful

9	DRILLING : 19-1 DRILLINGKICK BOP-POD-YLBL-CCF DRL-HUM-ERR-001 /ROV	End State LARGERELEASEROV Well Kick While Drilling Common cause failure of blue and yellow pods Kick not properly detected ROV intervention unsuccessful
10	DRILLING : 15-1 DRILLINGKICK BOP-POD-YLBL-CCF EDI-HUM-ERR-001 /ROV	End State LARGERELEASEROV Well Kick While Drilling Common cause failure of blue and yellow pods emergency disconnect fails ROV intervention unsuccessful

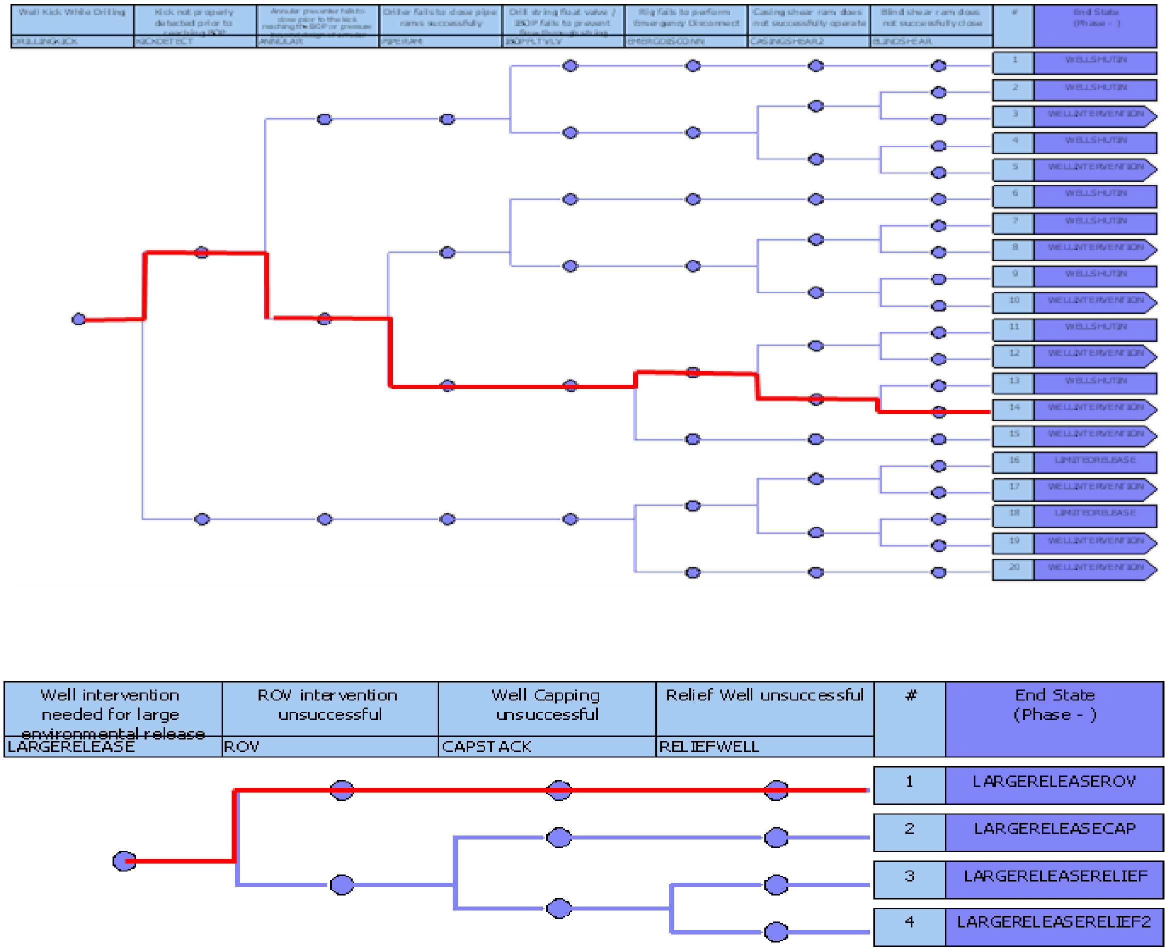

Figure 2- 30. Sequence 14-1 for Cut Set 1

The Boolean reduction of cut sets prevents other combinations of events from being produced. For instance, if the annular preventer fails due to the common cause failure of the yellow and blue pods, it will not appear in the results coupled with a failure of the blind shear ram to close or one of the shuttle valve failures. This is because the common cause failure of the pods has already guaranteed the failure of the blind shear ram, and the failure of the ram or shuttle valve is inconsequential if the common cause failure has occurred.

The final step in building the PRA model is populating the fault trees with data. The next several sections discuss the various methods needed to establish the most credible data for quantification.

2.2.6 Quantification / Allocation

For purposes of this section, we assume that implementation of the processes described above has developed a scenario set (a collection of "scenarios" leading to some consequence); we need to model their frequencies (or their conditional probabilities) and perhaps their consequences, and to carry out these activities with appropriate regard for uncertainty. Those tasks are the subjects of the present section.

In order to claim a low level of risk, it is necessary either to be able to argue some kind of inherent safety (the situation is safe because of physical laws), or to develop a lot of evidence and argument regarding the performance of our engineered systems. In the latter case, the key claims are:

- For the given design, built and maintained according to specified engineering standards, and for a stated body of operating procedures and conventions, we have identified the scenarios that lead to the undesirable consequences, and the severity of their consequences.
- We know how to quantify the scenarios;
    - We know what probabilities to assign to basic events, and to combinations of basic events;
        - We have analyzed what levels of reliability performance are achievable, and we know (modulo some uncertainty) what it takes to achieve them.
        - We have analyzed the potential for linkages between the occurrence of various combinations of basic events, and have factored this into our quantification (with uncertainty).
- We commit to the measures needed to attain (and maintain, and assure, and demonstrate on an ongoing basis) the levels of reliability performance credited in the analysis.

Event probabilities are not constants of nature. To assign a low failure probability to an engineered system or component is to take credit for an engineering accomplishment, and the results of our analysis are conditioned on that credit. Our reason for reviewing operating experience is not that past performance guarantees future performance; we review operating experience (1) in order to understand how past engineering investments have panned out in past performance, (2) as a sanity check of the numbers that we put into our analysis, and (3) as a guide to the insurance activities to which we need to commit if our risk estimates are to come true.

2.2.6.1 Quantification of individual scenarios

2.2.6.1.1 Quantifying the probability / frequency of individual basic events:

Historically, many analyses have implicitly treated PRA input numbers as if they were objectively significant: uncertain, to be sure, but having some objective (albeit unknown) value, analogous to the value of a physical constant, in the sense that you can look it up. In some cases [2-2, 2-3], it is recognized that basic event probabilities may be, in effect, influenced by operating conditions at the subject facility.

But in some applications, there seems to be a tacit supposition that a PRA result is an attribute of a facility. This interpretation is inappropriate. A more complete interpretation is the following.

Assuming that

- the scenario model is structurally complete (!),
    - addresses all initiating events, accounts for dependencies of all types, CCF, etc.,
- the data we have used to quantify our basic events are relevant,

  - derived from components appropriately similar to ours, similar service conditions, similar maintenance practices, similar component ages,
- and that we will follow operational practices similar to those to which the performance data pertain,

we can reasonably hope to achieve performance comparable to the performance achieved by the facilities represented in the data base.

In most cases, the basic event quantifications are not guarantees of future performance: they are simply PRA inputs, implicitly tied to a commitment of sorts (to the regulator) to an investment in achieving the level of reliability performance claimed. We will return to this topic in Section 3 of this guide.

On that understanding, we will now discuss a range of techniques for basic event quantification: the values for frequency (or, as appropriate, probability) of such things as component failures, initiating events, and human failures.

The figure below is to be interpreted in the context of an existing scenario model that we wish to quantify. That model should have been developed down to a level of detail at which the basic events are largely independent, and we can obtain some data bearing on how often the basic events occur. This discussion focuses on events that we try to quantify from experience. Common-cause failures and failure events that are driven by scenario-dependent conditions are treated in other sections.

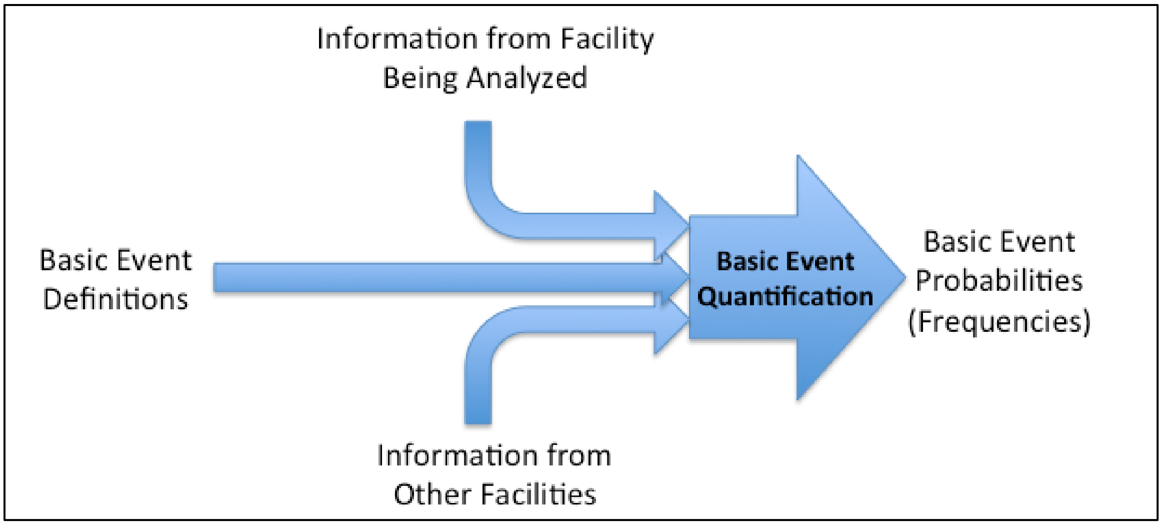
Information from Facility Being Analyzed
Basic Event Definitions
Basic Event Quantification
Basic Event Probabilities (Frequencies)
Information from Other Facilities

Figure 2- 31. Sources of Information for Quantification of Basic Event Likelihood

The Scenario Context of Basic Events

Normally, the scenario model is developed in such a way that each "scenario" can be expressed in a narrative way. For example:

> "Initiating event IE1 occurred; the intended response was for Valve A to open, but Valve

> A failed to open; in that circumstance, Pump C was required to operate, and to continue to operate for at least 6 hours, but it was unavailable at the time. As a result of this chain of events, the top event occurred."

This sort of narrative tells us what the basic events mean. We need to know how often IE1 occurs (the characteristic number of failures in a given time interval), the fraction of demands in which Valve A fails to open after IE1 occurs, and the fraction of time that Pump C is "unavailable" (under repair, for example) in the facility states within which IE1 can occur. These are only examples, but serve to suggest the kinds of information that we would hope to obtain from operating experience.

2.2.6.1.2 Estimating Parameters in Models for Basic Event Probability (Frequency)

The two main phases of developing a PRA database are:

- Information Collection and Classification
- Parameter Estimation

Typical quantities of interest are:

- Initiating Event (IE) Frequencies
- Component Failure Frequencies
- Component Test and Maintenance Unavailability
- Common Cause Failure (CCF) Probabilities
- Human Error Rates
- Software Failure Probabilities

Developing a PRA database of parameter estimates involves the following steps:

- Model-Data Correlation (identification of the data needed to correspond to the level of detail in the PRA models, determination of component boundaries, failure modes, and parameters to be estimated, e.g., failure rates, MTTR)

- Data Collection (determination of what is needed, such as failure and success data to estimate a failure rate, and where to get it, i.e., identification of data sources, and collection and classification of the data)

- Parameter Estimation (use of statistical methods to develop uncertainty distributions for the model parameters)

- Documentation (how parameter uncertainty distributions were estimated, data sources used, and assumptions made)

Typical PRA parameters, and the underlying probability models, are summarized in the Table below. We do not simply examine experience and directly obtain a number (for a probability or frequency) that we can use in our scenario quantifier; we *model* the probability or frequency in terms of underlying parameters, which we seek to learn from experience. Typically, there is epistemic uncertainty about the values of these underlying parameters, and carrying this uncertainty through the quantification can be important. Parameters for which there is epistemic uncertainty are shown in bold in the table.

Note: Model parameters for which there is epistemic uncertainty are shown in bold in the center column of the table. Data needed to estimate those parameters are listed in the right-hand column. Other model parameters (such as "mission time") are determined by the application.

Table 2-6. Typical Probability (or Frequency) Models in PRAs and their Parameters

Basic Event Type	Commonly-Used Models of Basic Event Probability	Data Required In Order to Quantify Models
Initiating event	Poisson model for probability of seeing k events in time t: $Pr(k) = e^{-\lambda t}\frac{(\lambda t)^k}{k!}$ where t: Mission time **λ: frequency**	Number of events k in time t
Component fails on demand	Constant probability of failure on demand, or **q**	Number of failure events k in total number of demands N
Standby component fails in time, or component changes state between tests (faults revealed on functional test only)	Constant standby failure rate $Q = 1 - \frac{1 - e^{-\lambda_s T_s}}{\lambda_s T_s}$ T_S: Time between tests **$\boldsymbol{\lambda}_S$: Standby failure rate**	Number of events k in total time in standby T
Component in operation fails to run, or component changes state during mission (state of component continuously monitored)	Constant failure rate $U = 1 - e^{-\boldsymbol{\lambda}_0 T_m} \approx \boldsymbol{\lambda}_0 T_m$ T_m: Mission time **$\boldsymbol{\lambda}_0$: Operating failure rate** Approximation is adequate when $\boldsymbol{\lambda}_0 T_m \ll 1$	Number of events k in total exposure time T (total time standby component is operating, or time the component is on line)
Component unavailable due to test	$Q = \frac{T_{TD}}{T_S}$ T_{TD} : Test duration (only in the	Average test duration (T_{TD}) and time between tests (T_S)

	case of no override signal) T_S: Time between tests	
Component unavailable due to corrective maintenance (fault revealed only at periodic test, or preventative maintenance performed at regular intervals)	$Q = \frac{T_U}{T_T}$ T_U: Total time unavailable while in maintenance (out of service) T_T: Total operating time	Total time out of service due to maintenance acts while system is operational, T_u, and total operating time T_T.
Component unavailable due to unscheduled maintenance (continuously monitored components)	$Q = \frac{\mu T_R}{1 + \mu T_R}$ T_R: Average time of a maintenance outage ["Repair time"]. **μ: Maintenance rate**	Number of maintenance acts r in time T (to estimate μ)
Standby component that is never tested. Assumed constant failure rate.	$Q = 1 - e^{-\lambda_m T_p}$ T_p : Exposure time to failure **λ_m : Standby failure rate**.	Number of failures r, in T units of (standby) time
Common-Cause Failure Probability (Refer to Appendix D)	$\boldsymbol{\alpha}$□through $\boldsymbol{\alpha_m}$, where *m* is the redundancy level	n_1 through n_m where n_k is the number of CCF events involving k components

The Table also shows the data needed to estimate the various parameters. The type of data needed varies depending on the type of event, and on how its frequency or probability is modeled. For example, probabilities typically require event counts (e.g., Number of Failures), and exposure or "Success Data" (e.g., Total Operating Time). Other parameters may require only one type of data, such as Maintenance/Repair Duration for mean repair time distribution, and counts of multiple failures in the case of CCF parameter estimates.

Sources of Information

Ideally, parameters of PRA models of a specific system should be estimated based on operational data of that system. As previously discussed, even past performance of "that system" does not guarantee future performance, for several reasons; but data from "that system" must be among the most relevant data available, unless something fundamental has recently changed.

If system-specific data of adequate quantity, quality, or availability are lacking, the analysis has to rely on other sources and types of information. In such cases, surrogate data, generic information, or expert judgment are used directly or in combination with (limited) system-specific data. It bears repeating that in submittals to regulators, the submitter is accountable for the treatment on which the conclusions are based.

Parameter Estimation Methods

Bayesian methods are widely used in PRA, while classical estimation has found only limited and restricted use in PRA. Accordingly, this section describes only the Bayesian approach to parameter estimation.

Bayesian estimation incorporates information beyond that contained in the data sample; this is part of what makes Bayesian inference different from classical estimation. In practice, Bayesian estimation comprises two main steps. The first step involves using previous information to develop a prior distribution for the parameters of a basic event model, such as a failure rate. The second step of Bayesian estimation involves using additional or new data (e.g., recent performance history) to update the prior distribution, yielding a "posterior" distribution for the parameters of that basic event model. This step is often referred to as "Bayesian updating" of the prior distribution. This process is illustrated in Appendix F.

For PRA applications, determining the prior distribution is usually based on generic data, and the new or additional data usually involve system-specific test or operating data. The resulting posterior distribution would then be the system-specific distribution of the parameter. If system-specific data do not exist, the applicability of other data or information would need to be evaluated and used. Refer to Appendices E and G.

Within the standard approach, one formulates explicit state-of-knowledge probability distributions about uncertain variables, both epistemic and aleatory. If these uncertain variables are model inputs, and one has distributions for them, one can infer the distribution of the model output(s), or at least the "parameter uncertainty" portion (which, in principle, ought to be evaluated together with other uncertainties) during quantification. Given a proper understanding of the uncertainties that affect the analysis, one can proceed to apply the standard machinery of decision-making under uncertainty.

Within the standard Bayesian approach, information is gathered about epistemically uncertain variables (or hypotheses regarding which we are uncertain), including formulation of a “prior” distribution on the values of those variables (or the probabilities of the various hypotheses being true); those distributions are then “updated” as new information becomes available, and one’s state of knowledge is improved (sometimes). Bayes’ so-called “theorem” states that

$$p(H_i \mid E) = p(H_i) * \frac{p(E \mid H_i)}{p(E)}, \qquad (1)$$

where

- H_i represents a hypothesis whose probability is to be updated with new evidence,
- $p(H_i\ )$ is the prior probability of H_i,
- E represents a new piece of evidence,
- $p(x|y)$ is the conditional probability of x given y,
- p(E), the prior probability of the observed evidence, can be written as

$$p(E) = \sum_i p(E \mid H_i) p(H_i) \ . \qquad (2)$$

Hereafter, this is referred to as the “update rule.” The update rule says that the conditional posterior probability of hypothesis H_i, given new evidence E, is equal to the prior probability of hypothesis H_i, multiplied by the conditional probability of observing E if H_i is true, divided by the total prior probability of observing E, calculated as shown in (2). In essence, new evidence that favors hypothesis H_i more than it favors hypothesis H_j (i.e., $p(E|H_i\ ) > p(E|H_j\ )$) tends to increase the posterior probability of hypothesis H_i relative to the posterior probability of H_j. In accordance with the update rule, new evidence causes the probabilities of the competing hypotheses to shift towards the implications of the new evidence.

The above paragraph has been worded as if the hypotheses were discrete, but it also applies if the hypotheses are understood to refer to different possible values of a continuous variable. In the latter case, the quantity on the left is understood to be a posterior probability density function of that variable.

The form of the update rule follows easily enough starting with the identity

$$p(a) p(b \mid a) = p(b) p(a \mid b) \text{, for any a, b,} \qquad (3)$$

dividing through by p(a), and identifying b with H_i and a with E. The identity, in turn, is easily understood with reference to a Venn diagram:

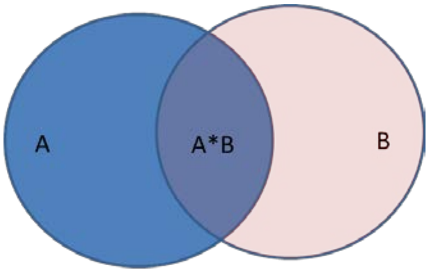

P(A|B)=p(A*B)/p(B)
P(B|A)=p(A*B)/p(A)

Figure 2- 32. Venn Diagram

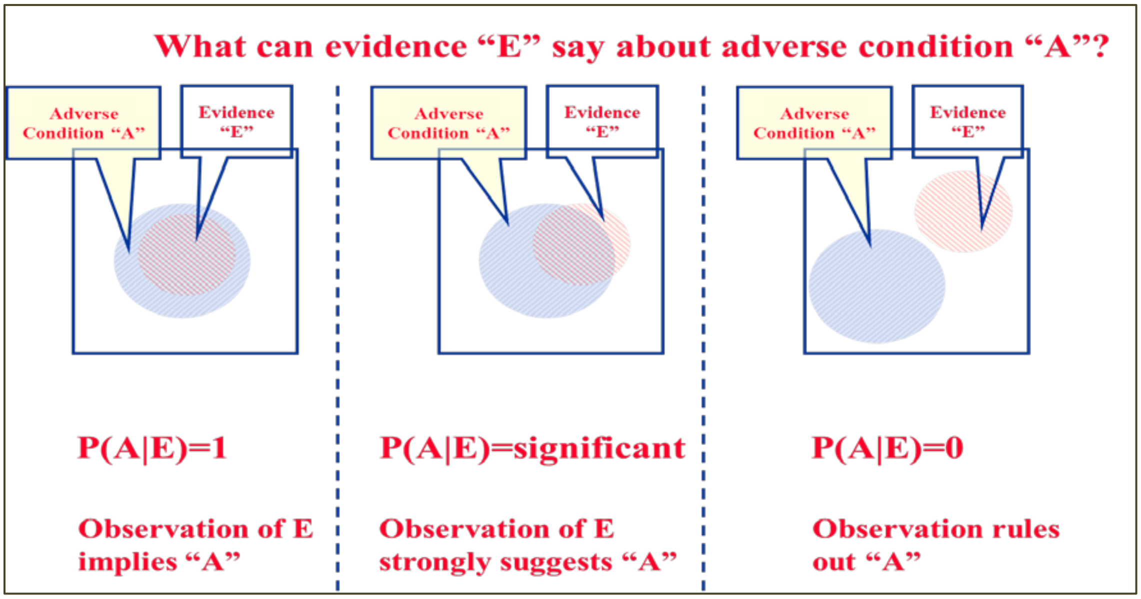

What can evidence "E" say about adverse condition "A"?

P(A|E)=1

Observation of E implies "A"

P(A|E)=significant

Observation of E strongly suggests "A"

P(A|E)=0

Observation rules out "A"

Figure 2- 33. Venn Diagram Illustration

Examination of a few cases may serve to aid intuition. Suppose our hypothesis (H_i, in the earlier notation) is that an adverse condition is present in a particular system ("A" for "adverse condition is present" in the above figure), and we have gathered evidence E to help determine whether A is true. In the above figure, the Venn diagram on the right illustrates the situation in which A and E do not overlap, so p(E|A) is zero, and the Bayes update rule will yield p(A|E)=0. The Venn diagram on the left illustrates a situation in which we see evidence E only if A is true, and putting the indicated numbers into the update rule will yield p(A|E)=1. The case in between – partial overlap of A and E – is where the practical applications lie.

Among the important properties of the update rule is that as new evidence is gathered, the process can be iterated; for a given collective body of evidence and a given starting prior and a given likelihood function, the same conclusion will be reached, regardless of how the evidence is parsed and applied in subsets. An illustration of this is shown in Appendix F.

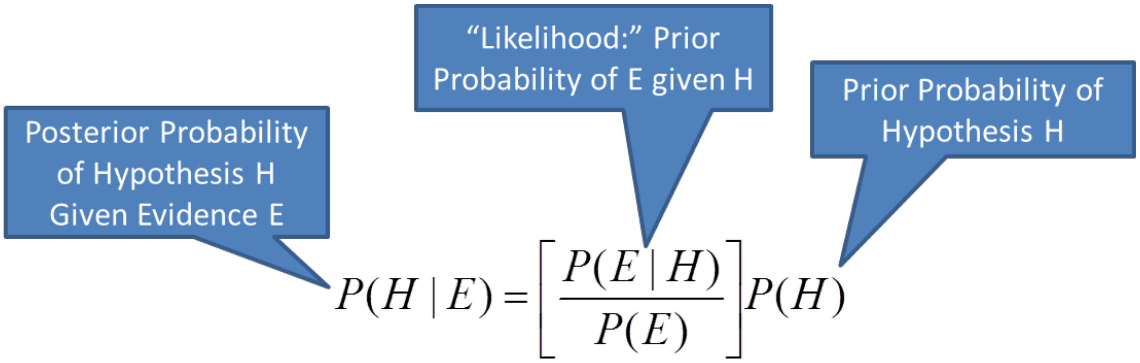
Posterior Probability of Hypothesis H Given Evidence E
"Likelihood:" Prior Probability of E given H
Prior Probability of Hypothesis H

$$P(H \mid E) = \left[\frac{P(E \mid H)}{P(E)} \right] P(H)$$

Figure 2- 34. The Update Rule

The preceding statement calls to mind a much stronger claim advanced by Bayesians: that all rational individuals will reach the same conclusion from a given body of evidence.

The current state of knowledge depends not only on the evidence, but also the prior distributions of the variables, and the form of the likelihood function: how the evidence is interpreted in the context of the current application. There is a vast literature on formulating the prior, but unfortunately, some of it shortchanges the topic of the likelihood function. We will return to the likelihood function later.

Prior Distributions

Prior distributions can be specified in different forms depending on the type and source of information as well as the nature of the random variable of interest. Functional forms widely used in PRA of engineered systems include:

- Parametric (gamma, lognormal, beta):
  - Gamma or lognormal for rates of events (time-based reliability models)
  - Beta or truncated lognormal for event probabilities per demand
- Numerical (histogram, CDF values/percentiles)
  - Applicable to both time-based and demand-based reliability parameters.

Among the parametric forms, a number of probability distributions are extensively used in risk studies as prior and posterior distributions. These are:

- Lognormal (μ, σ)

$$\pi(x) = \frac{1}{\sqrt{2\pi}\sigma x} e^{-\frac{1}{2}\left(\frac{\ln x - \mu}{\sigma}\right)^2}, 0 < x < \infty,$$

where μ and σ are the parameters of the distribution. The lognormal distribution can be truncated (truncated lognormal) so that the random variable is constrained to be less than a specified upper bound; if this sort of truncation is applied, then the distribution needs to be renormalized.

- Gamma(α, β)

$$\pi(x) = \frac{x^{\alpha-1}\beta^{\alpha}}{\Gamma(\alpha)}e^{-\beta x} \quad 0 \le x < \infty$$

where a and b are the parameters of the distribution.

- Beta(α, β)

$$\pi(x) = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)}x^{\alpha-1}(1-x)^{\beta-1} \quad 0 \le x \le 1$$

where α and β are the parameters of the distribution.

Information content of prior distributions can be based on:

- Previous system-specific estimates
- Generic, based on actual data from other (similar) systems
- Generic estimates from reliability sources
- Expert judgment (see discussion in Appendix H)
- Ignorance (i.e. lack of applicable data).

In the above list, the first four situations lead to prior distributions that may reflect considerable uncertainty about the parameters, but nevertheless assign higher probability to some values than to others. In those cases, application of situation-specific information through the update process is supposed to drive the posterior distribution to where it needs to be (or perhaps merely to reduce the uncertainty spread in that distribution). In situations where essentially no *a priori* information exists, attempts are made to formulate a prior reflecting this ignorance. A common approach to this is using a prior distribution that is uniform (constant) over the interval of interest. *Unfortunately, despite generations of work on how best to formulate such a prior, choice of prior distribution remains a research topic. If the current decision is sensitive to the tails of the posterior distribution, extra attention to this issue is warranted.*

Selection of the Likelihood Function

The form of the likelihood function depends on the nature of the assumed Model of the World representing the way the new data/information is generated:

For data generated from a Poisson Process (e.g., counts of failures during operation), the Poisson distribution is the proper likelihood function:

$$\Pr(k|T,\lambda) = \frac{(\lambda T)^k}{k!}e^{-\lambda T},$$

which gives the probability of observing k events (e.g., number of failures of a component) in T units of time (e.g., cumulative operating time of the component), given that the rate of occurrence of the event (failure rate) is λ. The Maximum Likelihood Estimate (MLE) of λ is

$$\lambda_{MLE} = \frac{k}{T}.$$

It is also possible to combine data from several independent Poisson processes, each having the same rate. This applies to the case where data are collected on different but identical units of equipment to estimate their common failure rate. The failure counting process for each unit is assumed to be a Poisson process. In particular, suppose that the ith Poisson process is observed for time t_i, yielding the observed count k_i. The total number of event occurrences is $k = \sum_i k_i$, where the sum is taken over all of the processes, and the exposure time is $T = \sum_i t_i$. This combined evidence can be used in the likelihood function given above.

For data generated from a Bernoulli Process (e.g., counts of failures on system demands), the Binomial distribution is the proper likelihood function:

$$\Pr(k|N, q) = \binom{N}{k} q^k (1-q)^{N-k}$$

which gives the probability of observing k events (e.g., number of failures of a component) in N trials of a component (e.g., total number of tests of the component), given that the probability of failure per trial (failure on demand probability) is q. The MLE of q is:

$$q_{MLE} = \frac{k}{N}$$

Analogously to the Poisson processes discussed above, data from independent trials that are known to be exchangeable (they are known to be determined by the same q, because they are the same component or identical components operated similarly) can be pooled: the failures can be summed, $k = \sum_i k_i$, and the demands can be summed, $N = \sum_i n_i$, and the results used in the binomial likelihood formula given above.

These cases are simple ones but are widely used. Likelihood functions for parameters of physical models are discussed in Appendix F.

In some cases, resort must be had to a process that relies on experts to furnish input. This is discussed in Appendix H.

Development of the Posterior Distribution

Using the update rule in its continuous form, the prior probability distribution of a continuous unknown quantity, $\Pr_O(x)$, can be updated to incorporate new evidence E as follows:

$$\Pr(x|E) = \frac{L(E|x)Pr_0(x)}{\int L(E|x)Pr_0(x)dx}$$

where Pr(x|E) is the posterior or updated probability distribution of the unknown quantity x given evidence E (occurrence of event E), and L(E|x) is the likelihood function (i.e., probability of the evidence E, assuming that the value of the unknown quantity is x). Illustrative combinations of prior and likelihood functions as well as the form of the resulting posterior distributions are listed in Table 2-7.

Table 2-7. Typical Prior and Likelihood Functions Used in PRAs

Functional Form of Prior	Functional Form of the Likelihood	Resulting Functional Form of the Posterior
Lognormal	Poisson	Numerical
Gamma	Poisson	Gamma
Beta	Binomial	Beta
Truncated Lognormal	Binomial	Numerical

For certain cases in the above table, the posterior has the same functional form as the prior. This occurs when there is a certain similarity between the functional form of the likelihood and that of the prior. For example, the Beta prior is proportional to powers of q multiplying powers of 1-q, as is the binomial distribution used as the likelihood; and as a result, the posterior is likewise a product of powers of q and powers of (1-q). In such a case, the update can be done analytically (in closed form). When a combination of prior and likelihood has this property, the prior is said to be "conjugate" to the likelihood. In the case of non-conjugate priors - for example, the case of "lognormal * Poisson => numerical" - resort must be had to numerical integration.

Two commonly used conjugate distributions are listed in Table 2-8. The formulas used to calculate the mean of the resultant posterior in terms of the parameters of prior and likelihood functions are provided.

Table 2-8. Common Conjugate Priors Used in Reliability Data Analysis

Functional Form of Prior Distribution, Mean Value	Functional Form of Likelihood	Posterior Distribution (same as prior)	Mean of Posterior
Beta (α,β), $\overline{x_{prior}} = \frac{\alpha}{\alpha + \beta}$	Binomial (k, N)	Beta	$\overline{x_{posterior}} = \frac{\alpha + k}{\alpha + \beta + N}$
Gamma (α,β); $\overline{x_{prior}} = \frac{\alpha}{\beta}$	Poisson (k, T)	Gamma	$\overline{x_{posterior}} = \frac{\alpha + k}{\beta + T}$

In the case of the conjugate priors listed in the above table, because we can compute the prior and posterior means in closed form, we can see how new data cause the mean to shift.

In principle, a prior distribution should reflect a state of knowledge, not a choice made to avoid the need for numerical integration. This has always been true, but is even more emphatically true in light of the very real improvements in computational capability in recent generations. It used to be argued that given a halfway reasonable prior, updates with new data would eventually drive posterior distributions to where they need to be; but in practical applications, where there is not always a surfeit of new data, this ideal is not always realized.

Developing Prior Distributions from Multiple Sources of Generic Information

Typically, generic information can be categorized into two types:

- Type 1 Failure data from operational experience with other similar but not identical components, or from identical components under different operating conditions. This information is typically in the form of failure and success data collected from the performance of similar equipment in various systems. The data in this case are assumed to come from a "non-homogenous" population.

- Type 2 Failure rate estimates or distributions contained in various industry compendia, such as several of the databases discussed earlier. Estimates from expert judgment elicitations would be included in this category. Type 2 data are either in the form of point estimates (or "best estimates"), or a range of values centered about a "best estimate." Ranges of the best estimate can be expressed in terms of low, high, and recommended values, or as continuous probability distributions.

When multiple sources of generic data are available, then it is likely that we are dealing with a non-homogeneous population. In these cases, the data cannot be pooled, and the reliability parameter of interest (e.g., failure rate) will have an inherent variability. The probability distribution representing this variability is known as a population variability distribution of the reliability parameter of interest. Refer to Appendix G for a discussion of this point.

2.2.6.2 Quantifying the Scenario Set

2.2.6.2.1 "Point" Estimates

Reference is frequently made to "point estimates" of important quantities such as top event probability. The term "point estimate" refers to the use of specific numbers for the inputs to the calculation, without immediate regard to uncertainty or variability. (Propagation of uncertainty is discussed below.) Quantification of point estimates is discussed here not because point estimates should be used uncritically in decision-making, but rather because the point estimate is a first step towards a more complete model quantification, and as such, plays a role even in scenario generation (a point discussed further in Section 3.2) by helping determine the truncation cutoff to be used and allowing a sanity check of the results.

Given the minimal cut sets for a particular top event, presented in sum-of-products form, we can obtain a point estimate of top event probability (or frequency) by:

- Quantifying each cut set

- multiplying the basic event point estimates of probability or frequency, as appropriate), or
- if the events in the cut set are related in some way, doing a side calculation to quantify the probability of the conjunction, and
- summing over cut set results to obtain the estimate for the top event.

This result, called the "rare-event approximation," is (in some respects) easy to calculate, and provides a rigorous upper bound on top event probability: not that the point estimate is guaranteed to be an upper bound on top event probability, but that the exact calculation of the "point estimate" would be less than or equal to this estimate. This follow because, for any two events X and Y (which could be basic events, or cut sets, or complex functional failures), we have

$$P(X+Y) = P(X) + P(Y) - P(X*Y).$$

So

$$P(X) + P(Y) \geq P(X+Y).$$

The term P(X*Y) is small in many applications (hence the "rare event" nomenclature), and in those cases, neglect of it is reasonable; moreover, it is often straightforward to check on the magnitude of its effect. In other words: Summing the cut set probabilities overestimates the top event, but is frequently reasonable.

A somewhat more involved calculation is the "min cut upper bound," obtained as

P(TOP) ~ 1- sum product (1-p(xi))

2.2.6.2.2 Propagating uncertainty through the scenario set

Given a way of calculating a point estimate for any setting of the basic event model parameters, we can propagate parameter uncertainty through the model in the same way that we can propagate parameter uncertainty through any model to obtain an uncertainty distribution on its output: we can sample from the joint distribution of the inputs, compute the result of that sample, iterate until the result is deemed to have converged to the point where key metrics can be evaluated (mean, median, mode, key percentiles).

Special cases:

Epistemic Coupling

The above statement referred to sampling from the **joint** distribution of all of the variables. If some of the variables are correlated epistemically, it is necessary to reflect this in the calculation. Consider the example of two essentially identical valves in series that are required to close under a certain challenge. They are of common manufacture and are assumed to see the same operating conditions, including test and maintenance practices. Arguably they should have the same failure probability. Since they are in series, and are required to close, failure of this function entails failure of both valves; so the top event probability will contain a contribution that is proportional to p(X*X), where "X*X" means "failure of both identical valves." Treating this as if it were equal to p(X)*p(X) underestimates the contribution for possibly several reasons. Temporarily setting aside the issue of common cause failure, there is an epistemic issue: in general, for any quantity Z,

$$< Z^2 > \geq < Z >^2,$$

so failure to acknowledge this epistemic coupling tends by itself to underestimate the result.

2.3 Simulation

2.3.1 Phenomenological Modeling

To Be Added

2.3.2 Discrete Event Simulation

Section 2.2 describes logic modeling: that is, modeling techniques using event trees and fault trees. By their nature, logic modeling techniques discretize the scenario descriptions, and thereby introduce approximations that may be severe. For example, as illustrated earlier, event tree end states are typically rather broadly specified, such as “large” and “limited” breaches of containment. It is possible to improve very significantly on these approximations using discrete-event simulation (e.g. estimate the number of deaths or barrels of oil spilled).

Discrete event simulation modeling is similar to developing an ESD as discussed in Section 2.2. Time ordered events are developed and decision blocks are used with probabilities that direct the flow of the simulation. In addition, events can be used to simulate variables such as well flow rates and recovery times. The model is run by performing numerous replications (i.e., thousands or more) using Monte Carlo sampling to obtain the probabilities or values at each decision event in the model, and the outcome of each replication is recorded. Obtaining a sufficient number of replications is important to ensure that all desired events get sampled and all reasonable paths in the model are exercised. For example, if a decision block has a probability of 0.01 (1 in 100), running 100 replications would, on average, only go down that path once. A single data point on a path would fail to show the range of outcomes for that path.

In the example discussed below, the specified outputs are the duration and magnitude of a release.

An example of a discrete event simulation model was developed for the ESD shown in Figure 2-4; the model is shown in Figure 2-37. The model was developed to mimic the ESD with the comment blocks included. Additional blocks are included to assign flow rates and record the results.

The data in the discrete event model was derived from the fault tree model. Some events in the discrete event model use the same data used by the fault trees. Other events have dependencies (i.e., conditional probabilities) that must be accounted for. For example, the fault trees ANNULAR and PIPERAM from the event tree in Figure 2-11 both contain the failure of both the yellow and blue pods. In the event tree software, this dependency is automatically accounted for in the cut set minimization process. These dependencies could be explicitly modeled in the discrete event model or the model can be simplified by manually deriving the dependencies from the fault tree results and then using them as inputs in the discrete event model.
The fault tree results for ANNULAR are shown in Table 2-9, while the fault tree results for PIPERAM are shown in Table 2-10. The cut sets highlighted in yellow (e.g., common cause failure

of the yellow and blue pods) are the same in each result, which means they cause both top events to be true. Since the PIPERAM event is questioned after the ANNULAR event, the failures in ANNULAR that would fail both need to be determined. From the yellow highlighted cut sets, about 7.0% of the failures are due to a common dependency. In other words, if the ANNULAR event is true, 7.0% of the time the PIPERAM event will occur for the same causes. In the PIPERAM event, the yellow highlighted cut sets account for approximately 96.5% of the total failure probability, which leaves 3.5% of the total failure probability or about a 1.3E-5 probability of failure due to independent causes. To determine the total dependent failure probability of the PIPERAM event for the discrete event simulation model, the 7% (0.07) is added to the 1.3E-5, which is still approximately 0.07.

A similar exercise is performed with the casing shear and blind shear rams to develop those probabilities. This is done for both the failure and success paths as applicable. On the success path, only the independent failures may cause a top event to be true if the preceding event is successful and has some shared dependency.

Table 2-9. ANNULAR Fault Tree Results

#	Prob/Freq	Total %	Cut Set	Description
Total	5.059E-3	100	Displaying 8 Cut Sets. (8 Original)	
1	4.180E-3	82.62		
	4.180E-3		BOP-CYL-FTC-AP01	Annular preventer fails to seal
2	5.000E-4	9.88		
	5.000E-4		ANNULAROVERPRESS	Well pressure over the design limit of annular
3	3.530E-4	6.98		
	3.530E-4		BOP-POD-YLBL-CCF	Common cause failure of blue and yellow pods
4	1.180E-5	0.23		
	1.180E-5		BOP-SHV-LKG-SV01	Crosstie shuttle valve external leakage
5	1.180E-5	0.23		
	1.180E-5		BOP-SHV-LKG-SV02	Annular ROV shuttle valve external leakage
6	6.708E-6	0.13		
	2.590E-3		BOP-POD-FTR-BLUE	Blue pod (standby) fails to run
	2.590E-3		BOP-POD-FTR-YELLOW	Yellow pod (operating) fails to run
7	4.351E-8	< 0.01		
	2.590E-3		BOP-POD-FTR-YELLOW	Yellow pod (operating) fails to run
	1.680E-5		BOP-SHV-FTT-SV01	Crosstie shuttle valve fails to transfer to blue pod
8	2.590E-8	< 0.01		
	1.000E-5		BOP-HUM-ERR-XTIEPODS	Driller fails to select blue pod after yellow pod failure
	2.590E-3		BOP-POD-FTR-YELLOW	Yellow pod (operating) fails to run

Table 2-10. PIPERAM Fault Tree Results (Top 10 Cut Sets)

#	Prob/Freq	Total %	Cut Set	Description
Total	3.728E-4	100	Displaying 20 Cut Sets. (20 Original)	
1	3.530E-4	94.68		
	3.530E-4		BOP-POD-YLBL-CCF	Common cause failure of blue and yellow pods
2	1.310E-5	3.51		
	1.310E-5		BOP-CYL-JAM-PRAM12	Common cause failure of upper and lower pipe rams
3	6.708E-6	1.80		
	2.590E-3		BOP-POD-FTR-BLUE	Blue pod (standby) fails to run
	2.590E-3		BOP-POD-FTR-YELLOW	Yellow pod (operating) fails to run
4	2.590E-8	< 0.01		
	1.000E-5		BOP-HUM-ERR-XTIEPODS	Driller fails to select blue pod after yellow pod failure
	2.590E-3		BOP-POD-FTR-YELLOW	Yellow pod (operating) fails to run
5	8.538E-9	< 0.01		
	9.240E-5		BOP-LPR-FTC-PRAM01	Lower pipe ram jams
	9.240E-5		BOP-UPR-FTC-PRAM02	Upper pipe ram jams
6	1.090E-9	< 0.01		
	9.240E-5		BOP-LPR-FTC-PRAM01	Lower pipe ram jams
	1.180E-5		BOP-SHV-LKG-SV09	Crosstie shuttle valve external leakage upper Pipe Ram
7	1.090E-9	< 0.01		
	1.180E-5		BOP-SHV-LKG-SV08	ROV Shuttle valve external leakage lower Pipe Ram
	9.240E-5		BOP-UPR-FTC-PRAM02	Upper pipe ram jams
8	1.090E-9	< 0.01		
	1.180E-5		BOP-SHV-LKG-SV07	Crosstie shuttle valve external leakage lower Pipe Ram
	9.240E-5		BOP-UPR-FTC-PRAM02	Upper pipe ram jams
9	1.090E-9	< 0.01		
	9.240E-5		BOP-LPR-FTC-PRAM01	Lower pipe ram jams
	1.180E-5		BOP-SHV-LKG-SV10	ROV Shuttle valve external leakage upper Pipe Ram
10	1.392E-10	< 0.01		
	1.180E-5		BOP-SHV-LKG-SV08	ROV Shuttle valve external leakage lower Pipe Ram
	1.180E-5		BOP-SHV-LKG-SV09	Crosstie shuttle valve external leakage upper Pipe Ram

Other distributions for well flow rates and timing of recovery events such as when an ROV successfully manipulates the BOP may be put in as histograms as shown in Figures 2-35 and 2-36.

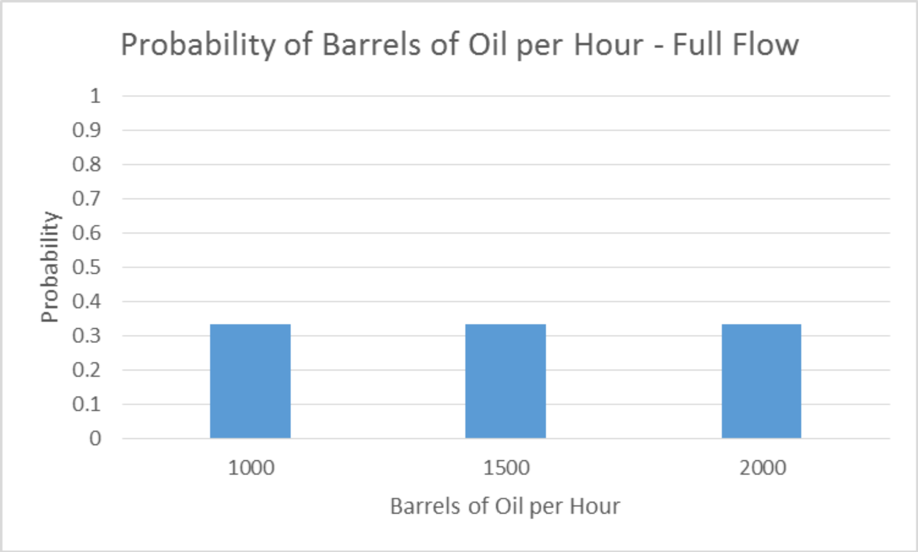

Figure 2- 35. Example Histogram of Probability vs. Well Flow Rate

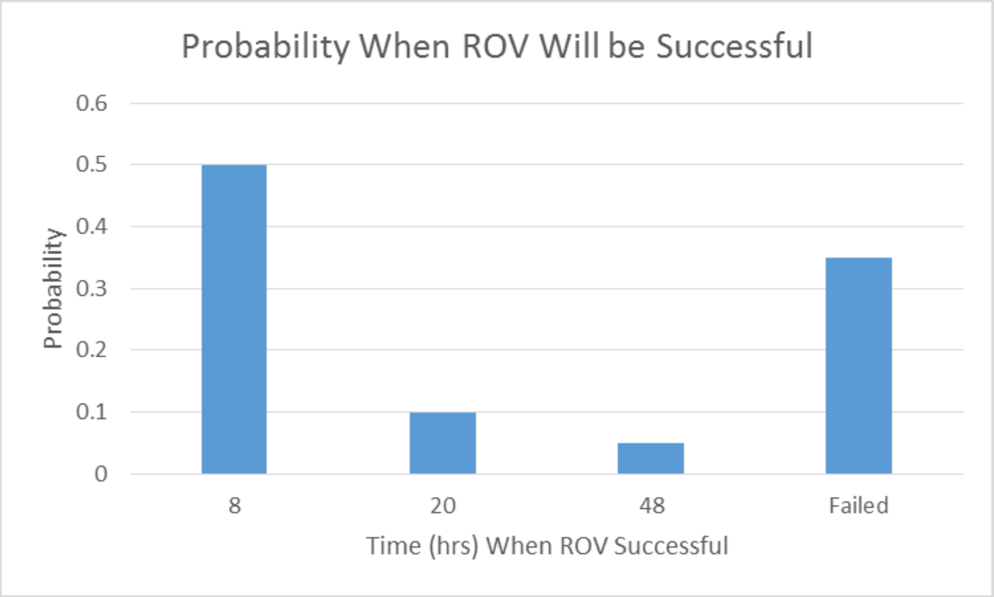

Figure 2- 36. Example Probability of Success vs. Time for ROV

Results from the discrete event simulation, shown in Table 2-11, are a sample of results from model simulation replications showing the end event and parameters of interest, in this case the time to shut in the well and the number of barrels of oil released. The replication column shows on which model simulation run the event occurred. The results can be manipulated to develop different kinds of useful products such as frequency of exceedance (F-N) curves discussed in Section III.

Table 2-11. Discrete Event Simulation Model Results

Replication #	End Event (what stopped flow)	Duration of Release (hrs)	Barrels of Oil Released
272	ROV Intervention	3	3,029
1703	ROV Intervention	33	55,211
5019	ROV Intervention	7	8,965
5556	Well Cap	93	183,980
6603	Well Cap	80	124,490
6725	Well Cap	92	88,314
8174	Well Cap	66	127,280
9287	Limited Release	0	741
14068	ROV Intervention	12	11,360
14287	ROV Intervention	5	4,390
15444	Well Cap	141	295,550

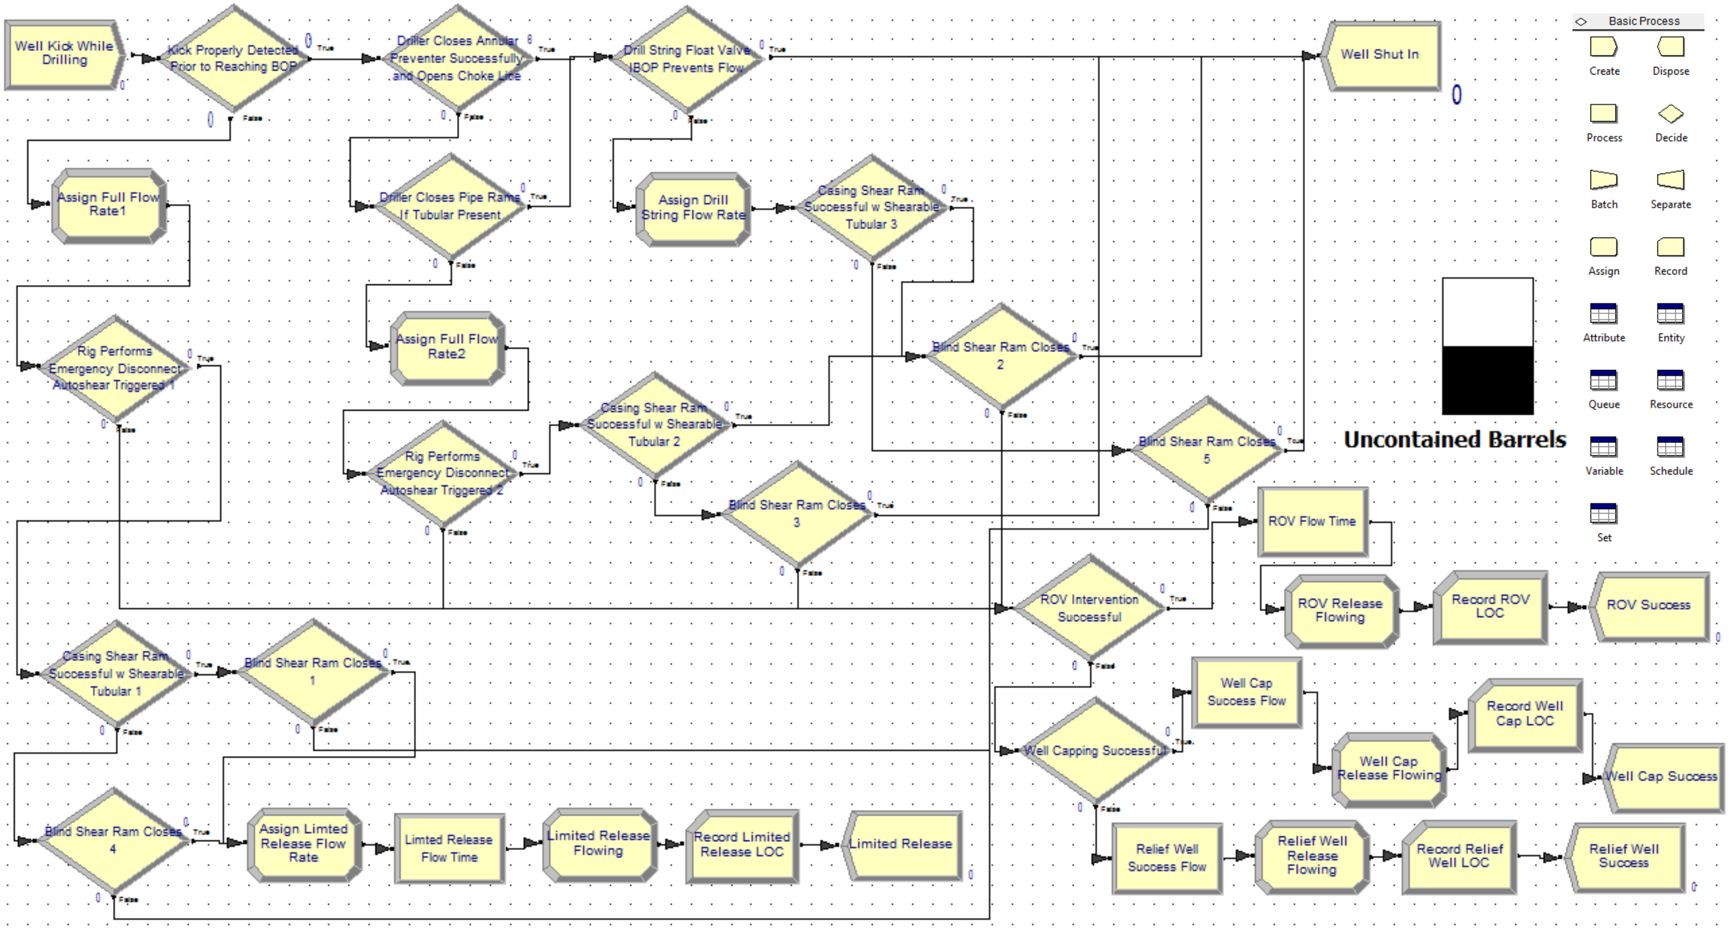

Figure 2- 37. Example Discrete Event Simulation Model

2.4 References

2-1 Fault tree Handbook with Aerospace Applications, Version 1.1, NASA, August 2002

2-2 NUREG-1816, “Independent Verification of the Mitigating Systems Performance Index (MSPI) Results for the Pilot Plants (Final Report),” USNRC, 2005

2-3 H. Hamzehee, R. W. Youngblood, et al, “Risk-Based Performance Indicators: Results of Phase 1 Development,” NUREG-1753 (USNRC, 2002).

3. Results: Presentation and Interpretation

3.1 Risk Analysis Support to a Notional Safety Case

Figure 3-1 shows a notional "claims tree:" a hierarchy of the claims that might be made in a safety case presented to a regulator. The regulation of facilities by BSEE is beyond the scope of this document, but it is nevertheless useful to organize the discussion of certain topics around a figure like this.

In many venues of application of risk models, the models are developed by parties associated with the (proposed) facilities, even though the model results are to be applied in assurance cases put before regulatory decision-makers who are accountable to different parties for different considerations (e.g., regulators may be more accountable for public safety than for facility economics).

The technical content of the present guidance is trying to be useful both to applicants and to regulators. The facility operators (and investors) need to "ensure" that the facilities are (or will be) safe; the regulator needs "assurance" that the facility is (or will be) safe. The needs of the two are distinct. The "claims tree" is aimed specifically at promoting a successful dialogue between applicants and regulators.

The premise of the figure is that a finding has to be made regarding the safety of a specific facility, and this finding needs to be based in part on an analysis. The analysis needs to address certain figures of merit (such as risk metrics) and, potentially, to show that certain other requirements are met (such as requirements on barrier availability and performance). The four major sections of the figure are:

1. Design characterization

2. Analysis of risk (and possibly other metrics), conditional on a particular baseline allocation of performance (e.g., reliability)

    The analysis satisfies certain process requirements
    The analysis provides sensitivity and uncertainty information
    A process exists for identifying, and dealing with, unresolved safety questions

3. [Optional] A process has been carried out to substantiate a claim that the facility is as safe as reasonably practicable.
Note: This figure can be specialized to refer to a more general "Best Available and Safest Technology" (BAST) rather than strictly safety. In either case, process-based arguments to support the respective conclusions are called for.

4. The performance allocation credited in the analysis is, in fact, feasible. The items considered critical, the associated levels of performance, and the activities

needed to make the risk analysis “come true” have been identified and committed to. This includes making reliability allocations come true, barrier availabilities come true, and so on, and includes a commitment to analysis of operating experience, looking for deficiencies in the model. Ordinarily, the model used for Item 2 above is also a starting point for this item.

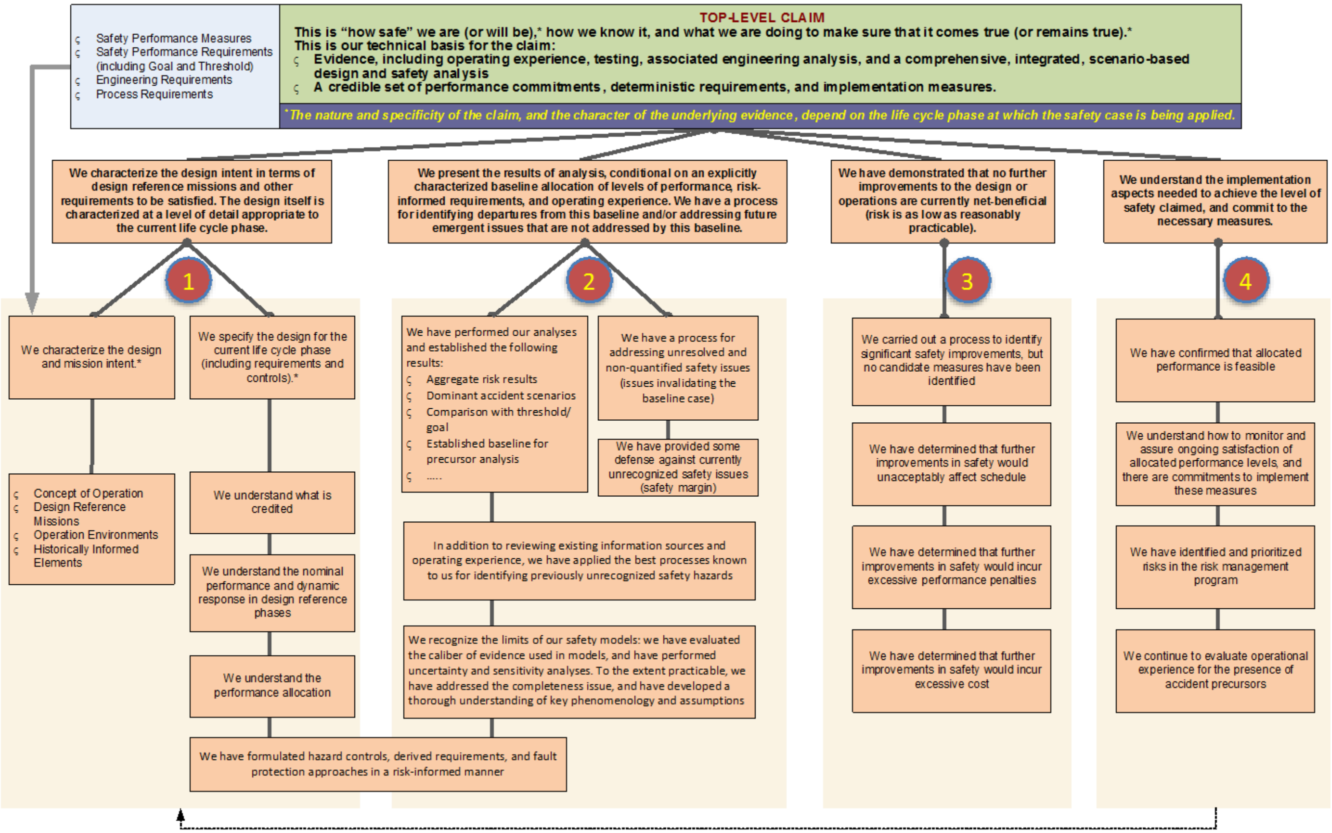

ς Safety Performance Measures
ς Safety Performance Requirements (including Goal and Threshold)
ς Engineering Requirements
ς Process Requirements

TOP-LEVEL CLAIM

This is "how safe" we are (or will be),* how we know it, and what we are doing to make sure that it comes true (or remains true).*
This is our technical basis for the claim:

ς Evidence, including operating experience, testing, associated engineering analysis, and a comprehensive, integrated, scenario-based design and safety analysis
ς A credible set of performance commitments, deterministic requirements, and implementation measures.

*The nature and specificity of the claim, and the character of the underlying evidence, depend on the life cycle phase at which the safety case is being applied.

We characterize the design intent in terms of design reference missions and other requirements to be satisfied. The design itself is characterized at a level of detail appropriate to the current life cycle phase.

We present the results of analysis, conditional on an explicitly characterized baseline allocation of levels of performance, risk-informed requirements, and operating experience. We have a process for identifying departures from this baseline and/or addressing future emergent issues that are not addressed by this baseline.

We have demonstrated that no further improvements to the design or operations are currently net-beneficial (risk is as low as reasonably practicable).

We understand the implementation aspects needed to achieve the level of safety claimed, and commit to the necessary measures.

We characterize the design and mission intent.*

ς Concept of Operation
ς Design Reference Missions
ς Operation Environments
ς Historically Informed Elements

We specify the design for the current life cycle phase (including requirements and controls).*

We understand what is credited

We understand the nominal performance and dynamic response in design reference phases

We understand the performance allocation

We have performed our analyses and established the following results:

ς Aggregate risk results
ς Dominant accident scenarios
ς Comparison with threshold/ goal
ς Established baseline for precursor analysis
ς .....

We have a process for addressing unresolved and non-quantified safety issues (issues invalidating the baseline case)

We have provided some defense against currently unrecognized safety issues (safety margin)

In addition to reviewing existing information sources and operating experience, we have applied the best processes known to us for identifying previously unrecognized safety hazards

We recognize the limits of our safety models: we have evaluated the caliber of evidence used in models, and have performed uncertainty and sensitivity analyses. To the extent practicable, we have addressed the completeness issue, and have developed a thorough understanding of key phenomenology and assumptions

We have formulated hazard controls, derived requirements, and fault protection approaches in a risk-informed manner

We carried out a process to identify significant safety improvements, but no candidate measures have been identified

We have determined that further improvements in safety would unacceptably affect schedule

We have determined that further improvements in safety would incur excessive performance penalties

We have determined that further improvements in safety would incur excessive cost

We have confirmed that allocated performance is feasible

We understand how to monitor and assure ongoing satisfaction of allocated performance levels, and there are commitments to implement these measures

We have identified and prioritized risks in the risk management program

We continue to evaluate operational experience for the presence of accident precursors

Figure 3-1. "Claims Tree"

Portions of Items 2 and 4 of the above list are within the scope of the present document. Item 2 notionally covers the safety analysis, and the results that need to be presented, including sensitivity and uncertainty information. Item 4 captures the claims that tie the analysis results to reality, including commitments to scrutinize operating experience as part of an effort aimed at model improvement. Certain aspects of Item 4 are beyond the scope of this document, but an emphasis of this document is that the numbers on which PRA results are based are engineering accomplishments, not constants of nature; submittal of a PRA to a regulator needs to be tied to commitments to make those numbers come true, including identification of model inadequacies revealed by operating experience. This is part of Item 4.

3.2 Quantifying the Model

Quantification of a PRA model is a simple process; however, steps must be taken to ensure that the output is relatively complete and accurate. The term “relatively” is used here with “complete” because PRA models can have millions of scenarios, and often many are of such a low probability that they do not need to be considered. A PRA model can take a long time to run and produce a very large amount of data, so determining the right level for quantification is an important step to ensure the necessary results are obtained while still being manageable. This applies to both the classical event tree – fault tree models and the discrete event models as discussed below.

3.2.1 Event Tree – Fault Tree Model Quantification

Quantifying an event tree – fault tree model typically involves choosing different options for performing the quantification. Three of the most important are discussed below. It should be noted that the model quantification and results discussed in this section is a simplified example for illustrating the results of PRA, and not an actual analysis of a real facility.

3.2.1.1 Truncation Cutoff

The truncation cutoff used in quantifying an event tree is used to stop the quantification of scenarios (i.e. minimal cut sets) below a user selected value. A large PRA model can have millions or tens of millions of minimal cut sets and evaluating all of them can take a long time to run the model and result in an excessively large amount output that is hard to manipulate to display results. Typically, lower probability scenarios may be many orders of magnitude lower than the overall result, and, therefore, are not significant. Most PRA software therefore implements a user-supplied truncation value, discarding cut sets whose probability is below that truncation value. In principle, this is an uncontrolled approximation, but is frequently the practical thing to do, provided that steps are undertaken to understand the effects of the truncation.

The process to determine what the best truncation cutoff is should start at a level the analyst expects would be consistent with the result (this is essentially a guess based on construction of the model), and then vary the truncation cutoff by reducing it an order of magnitude until the results at least converge to a value less with less than a percent difference between two successive quantifications.

Using the example developed in Section 2 (Figure 2-11), it is possible to see the effect the truncation limit has on the results in Table 3-1.

Table 3-1. Effect of Truncation Limit on Event Tree Quantification

Truncation Cutoff	Number of Cut sets	Overall Likelihood	End States
1.0E-4	2	4.471E-4	LARGERELEASEROV, LIMITEDRELEASE
1.0e-5	4	5.525E-4	LARGERELEASEROV, LIMITEDRELEASE, LARGERELEASECAP, LARGERELEASERELIEF
1.0E-6	6	5.590E-4	LARGERELEASEROV, LIMITEDRELEASE, LARGERELEASECAP, LARGERELEASERELIEF
1.0E-7	8	5.597E-4	LARGERELEASEROV, LIMITEDRELEASE, LARGERELEASECAP, LARGERELEASERELIEF, LARGERELEASERELIEF2
1.0E-8	17	5.599E-4	Same as 1.0E-7
1.0E-9	27	5.599E-4	Same as 1.0E-7
1.0E-10	38	5.599E-4	Same as 1.0E-7
1.0E-12	68	5.599E-4	Same as 1.0E-7
0.0	9794	5.599E-4	Same as 1.0E-7

When reviewing the results in Table 3-1, two effects can be seen as the truncation cutoff is varied. First, when the truncation value has been reduced to 1.0E-8, the overall likelihood has converged to several decimal places, and the number of cut sets is 17 at that level. When the truncation cutoff is set to 0.0, the model produces 9794 cut sets, and from the table, the majority of them are below 1.0E-12 so do not affect the overall result.

A second consideration should also be given to the end states found in the results. As seen in Table 3-1, the initial quantification of the model only resulted in two of the five end states being evaluated. If the goal of the analysis is based on the overall risk and major contributors, whether end states show up in the result may not make a difference if they do not significantly contribute, but if end states are to be evaluated separately for dominant contributors, the truncation cutoff should also be selected to get representative cut sets from each end state.

3.2.1.2 Solution Method

A second quantification option that can have a significant impact on results is the solution method as the fault tree scenarios are combined in the event tree. The success path on an event tree can be treated in different ways. Linked fault tree software typically defaults to a solution that uses a "delete term" function. The delete term function removes invalid cut sets from sequential top events with common basic events. An example of an invalid cut set would arise if, in a particular sequence expression, failure of System A is combined with success of System B, and some of the "failure" cut sets for A are inconsistent with success of System B. This condition cannot exist, so the cut set is deleted from that sequence expression.

Using "delete term" is reasonable, but it is an approximation: the success of System B is set to a probability of 1.0. Because PRA analyses usually are evaluating rare events, the top event probabilities are small most of the time, and the approximation of a success path to 1.0 is acceptable. In some cases, a top event may have a relatively large probability (> 0.01), and in this case, choosing a solution method that accounts for the proper success path probability may be required for a sufficiently accurate calculation.

3.2.1.3 Uncertainty

Quantifying the model to obtain the uncertainty distribution is similar to the truncation cutoff issue in that enough iterations of the model must be performed to ensure the mean value has converged. PRA software generally runs quickly so the number of iterations needed to converge is usually not an issue.

In many cases, the mean will actually differ from the point estimate, because in general, basic event distributions are correlated. Correlation of basic events results if the same uncertainty distribution is used for a number of basic events. Each sampling in the uncertainty calculation for the basic events that are correlated uses the same value from the common distribution. This has a tendency to increase the mean value if an AND gate is used as the sampling from the extreme ends of the distribution compound and to stretch the distribution, resulting in a mean value higher than the point estimate.[5]

3.2.2 Discrete Event Simulation Model Quantification

For analyses using discrete event simulation, as discussed in Section 2.3.2, the main option for quantification that must be considered is the number of replications (i.e. the number of passes through the model). Determining a sufficient number of replications is an iterative process that should start by reviewing the paths through the model along with the inputs to estimate the expected number of replications that result in a particular end state. For example, if a review of the model inputs shows that the output should occur with a frequency of about 1E-4, then the number of replications to get a result on that path, on average, would be 1/1E-4 or 10,000. Because the mean will not have converged with a single point (or a small number of points), for this example a reasonable starting point would be 100,000 replications, or 10 times the average to get a single result. Ideally, a sufficient number of replications can be run in order to obtain a mean within acceptable convergence bounds. If an initial number of replications is not sufficient to establish convergence, then additional replications are required. At some point, time constraints might limit the number of replications so that the desired convergence is not met. In these cases the uncertainty due to the limited number of replications should be presented. This uncertainty is often presented as 90% confidence bounds about the mean.

Table 3-2 shows output from the model shown in Section 2.3.2 based on 50,000, 100,000, and 250,000 replications.

[5] In general, $<x^2>$ is greater than or equal to $<x>^2$.

Table 3-2. Discrete Model Simulation Results For Different Numbers of Replications

			Barrels Leaked			Time to Effect (Hours)		
End Event	**Count**	**Probability**	**Min**	**Mean**	**Max**	**Min**	**Mean**	**Max**
Limited Release	16	1 in 3125	0.16	0.62	1.36	0.37	0.82	1.22
ROV	8	1 in 6250	1938.60	18796.70	48353	2.12	11.43	23.26
Well Cap	6	1 in 8333	89597	243699.50	484670	68.40	162.16	266.67
Relief Well	1	1 in 50000	4818100	4818100	4818100	3771.93	3771.93	3771.93
Replications:	**50000**							
			Barrels Leaked			**Time to Effect (Hours)**		
End Event	**Count**	**Probability**	**Min**	**Mean**	**Max**	**Min**	**Mean**	**Max**
Limited Release	25	1 in 4000	0.16	0.63	1.3616	0.28	0.81	1.22
ROV	15	1 in 6700	1938.6	18766.99	48353	2.12	12.45	24.06
Well Cap	13	1 in 7700	85475	271075.5	717960	68.36	186.02	440.56
Relief Well	1	1 in 100000	4818100	4818100	4818100	3771.93	3771.93	3771.93
Replications:	**100000**							
			Barrels Leaked			**Time to Effect (Hours)**		
End Event	**Count**	**Probability**	**Min**	**Mean**	**Max**	**Min**	**Mean**	**Max**
LimitedRelease	45	1 in 5600	0.16	0.62	1.4531	0.28	0.82	1.22
ROV	44	1 in 5700	1938.6	18645.58	100480	2.12	12.24	49.84
WellCap	33	1 in 7600	79285	302953.8	724090	68.36	201.56	440.56
ReliefWell	1	1 in 250000	4818100	4818100	4818100	3771.93	3771.93	3771.93
Replications:	**250000**							

3.3 Reviewing the Results

The PRA model is typically developed to answer a specific question or questions regarding the risk of a facility or operation, and a range of results are produced and may be reviewed at a variety of different levels (e.g. from system reliabilities to magnitudes of oil released to the environment). Common results evaluated as outputs from a PRA include:

- Total likelihood[6] of various end states;
- The relative ranking of each scenario to the total end state likelihood or total risk;
- Estimates of scenario consequences (e.g., environmental release, damage to property, number of injuries or fatalities, dollar loss, …);
- Importance measures;
- Display of uncertainties associated with various estimates; and
- System level reliabilities.

Each of these types of results are discussed in more detail in the following sections with examples based on the environmental release model developed in Section 2.

3.3.1 Overall End State Likelihood and Relative Risk Ranking

[6] In this section, the term "likelihood" is used to refer either to probability, frequency, or both.

The overall objective of performing a PRA is typically to evaluate a design or operation with respect to the risk involved. The purpose could be to ensure the design or process is acceptably safe relative to safety goals or requirements, or to understand if there are any driving vulnerabilities that can be addressed further. The first metrics assessed are usually the overall risk and a relative risk ranking of scenarios. Sample output from the simplified model developed in Section 2 is shown in Table 3-3.

Table 3-3. Sample PRA Model Output

#	Prob/Freq	Cut Set Contribution %	Cut Set	Description
Total	5.598E-4	100	Displaying 10 Cut Sets. (9794 Original)	
1	2.471E-4	44.14	DRILLING : sequence 14-1	
	1.000E+0		DRILLINGKICK	Well Kick While Drilling
	3.530E-4		BOP-POD-YLBL-CCF	Common cause failure of blue and yellow pods
	7.000E-1		/ROV	ROV intervention unsuccessful
		End State	LARGERELEASEROV	Added through Event Tree Add
2	2.000E-4	35.73	DRILLING : sequence 16	
	1.000E+0		DRILLINGKICK	Well Kick While Drilling
	2.000E-4		DRL-HUM-ERR-001	Kick not properly detected
		End State	LIMITEDRELEASE	Added through Event Tree Add
3	9.531E-5	17.03	DRILLING : sequence 14-2	
	1.000E+0		DRILLINGKICK	Well Kick While Drilling
	3.530E-4		BOP-POD-YLBL-CCF	Common cause failure of blue and yellow pods
	9.000E-1		/CAPSTACK	Well Capping unsuccessful
	3.000E-1		ROV-FTR-001	ROV intervention unsuccessful
		End State	LARGERELEASECAP	Added through Event Tree Add
4	1.006E-5	1.80	DRILLING : sequence 14-3	
	1.000E+0		DRILLINGKICK	Well Kick While Drilling
	3.530E-4		BOP-POD-YLBL-CCF	Common cause failure of blue and yellow pods
	1.000E-1		CAP-LKG-001	Well capping unsuccessful
	9.500E-1		/RELIEFWELL	Relief Well unsuccessful
	3.000E-1		ROV-FTR-001	ROV intervention unsuccessful
		End State	LARGERELEASERELIEF	Added through Event Tree Add
5	4.696E-6	0.84	DRILLING : sequence 14-1	
	1.000E+0		DRILLINGKICK	Well Kick While Drilling
	2.590E-3		BOP-POD-FTR-BLUE	Blue pod (standby) fails to run
	2.590E-3		BOP-POD-FTR-YELLOW	Yellow pod (operating) fails to run
	7.000E-1		/ROV	ROV intervention unsuccessful
		End State	LARGERELEASEROV	Added through Event Tree Add
6	1.811E-6	0.32	DRILLING : sequence 14-2	
	1.000E+0		DRILLINGKICK	Well Kick While Drilling
	2.590E-3		BOP-POD-FTR-BLUE	Blue pod (standby) fails to run
	2.590E-3		BOP-POD-FTR-YELLOW	Yellow pod (operating) fails to run
	9.000E-1		/CAPSTACK	Well Capping unsuccessful
	3.000E-1		ROV-FTR-001	ROV intervention unsuccessful
		End State	LARGERELEASECAP	Added through Event Tree Add
7	5.295E-7	0.09	DRILLING : sequence 14-4	

	1.000E+0		DRILLINGKICK	Well Kick While Drilling
	3.530E-4		BOP-POD-YLBL-CCF	Common cause failure of blue and yellow pods
	1.000E-1		CAP-LKG-001	Well capping unsuccessful
	5.000E-2		REL-WELL-LKG-001	Relief well not successful on first attempt
	3.000E-1		ROV-FTR-001	ROV intervention unsuccessful
		End State	LARGERELEASERELIEF2	Added through Event Tree Add
8	1.912E-7	0.03	DRILLING : sequence 14-3	
	1.000E+0		DRILLINGKICK	Well Kick While Drilling
	2.590E-3		BOP-POD-FTR-BLUE	Blue pod (standby) fails to run
	2.590E-3		BOP-POD-FTR-YELLOW	Yellow pod (operating) fails to run
	1.000E-1		CAP-LKG-001	Well capping unsuccessful
	9.500E-1		/RELIEFWELL	Relief Well unsuccessful
	3.000E-1		ROV-FTR-001	ROV intervention unsuccessful
		End State	LARGERELEASERELIEF	Added through Event Tree Add
9	4.942E-8	< 0.01	DRILLING : sequence 19-1	
	1.000E+0		DRILLINGKICK	Well Kick While Drilling
	3.530E-4		BOP-POD-YLBL-CCF	Common cause failure of blue and yellow pods
	2.000E-4		DRL-HUM-ERR-001	Kick not properly detected
	7.000E-1		/ROV	ROV intervention unsuccessful
		End State	LARGERELEASEROV	Added through Event Tree Add
10	2.471E-8	< 0.01	DRILLING : sequence 15-1	
	1.000E+0		DRILLINGKICK	Well Kick While Drilling
	3.530E-4		BOP-POD-YLBL-CCF	Common cause failure of blue and yellow pods
	1.000E-4		EDI-HUM-ERR-001	emergency disconnect fails
	7.000E-1		/ROV	ROV intervention unsuccessful
		End State	LARGERELEASEROV	Added through Event Tree Add

The above table shows an overall probability of approximately 5.6E-4 for having an environmental release of hydrocarbons during drilling. The table then lists the top 10 cut sets that can be individually reviewed. For this simplified model, the results show several things that could be of interest. The top 3 cut sets account for over 95 percent of the risk, and therefore would be a focus of improvement if the overall risk was considered to be too high (implying that improvement is needed). The drivers for the first 3 cut sets are a common cause failure of the blue and yellow pods on the BOP (cut sets 1 and 3), and Driller error for cut set 2. Further inspection of the top 3 cut sets shows that the cut set 2 end state is Limited Release while cut sets 1 and 3 are large release end states with potentially significantly higher consequences.

In practice, PRA results are often more "flat" than shown in the results from the simplified model used in this guide because they are stated at a finer level of resolution, so that no single cut set contributes a significant percentage to the total. Often there may be hundreds or even thousands of cut sets that make up 95 percent or more of the high level risk number. In this case, the results may be put in a spreadsheet and manipulated by grouping cut sets related to a particular end state, component, or system to develop insights.

3.3.2 Estimates of Consequences

While the discussion in Section 3.3.1 focused on the frequency of end states, also of interest is the magnitude of the end state(s). Classical PRA models using event trees and fault trees have end states that are subjectively defined, e.g. Limited Release, Large Release, etc. By manipulating the cut sets in the previous section, the frequencies of these end states may be estimated. For many applications, this approach may be acceptable as it shows that frequency of the end state and any scenarios that are major contributors. This information allows actions to be identified that may reduce risk.

In some applications the actual magnitude of end states may be needed or desired. For instance there may be governmental requirements on the expected casualty rate from a particular facility or operation. In this case the model must estimate the number of deaths for individual scenarios to develop that result. A discrete event simulation model as discussed in Section 2.3.2 is a method that can be used to perform that analysis. The output from a discrete event simulation is the results related to each replication of the model and can be very large. Table 3-4 shows a sample of typical results.

Table 3-4. Sample Results from a Discrete Event Simulation Model

Replication	End State				Release Duration (hrs)
	Limited Release	Relief Well	Well Cap	ROV	
272	0.30	0	0	0	0.5
947	0.00	4818100	0	0	3771.9
1064	0.00	0	0	18503	11.7
6603	1.36	0	0	0	1.2
6725	0.59	0	0	0	0.5
8174	0.67	0	0	0	0.9
9080	0.00	0	0	3408.7	3.3
9287	0.42	0	0	0	1.2
9882	0.00	0	0	1938.6	2.1
14287	0.30	0	0	0	0.5

The results are then manipulated to be useable by developing plots. Because the inputs are based on distributions (e.g. flow rates), the results must be binned into logical ranges to show results. Figure 3-2 shows the output from the example developed in Section 2.3.2 in terms of probability versus magnitude of release.

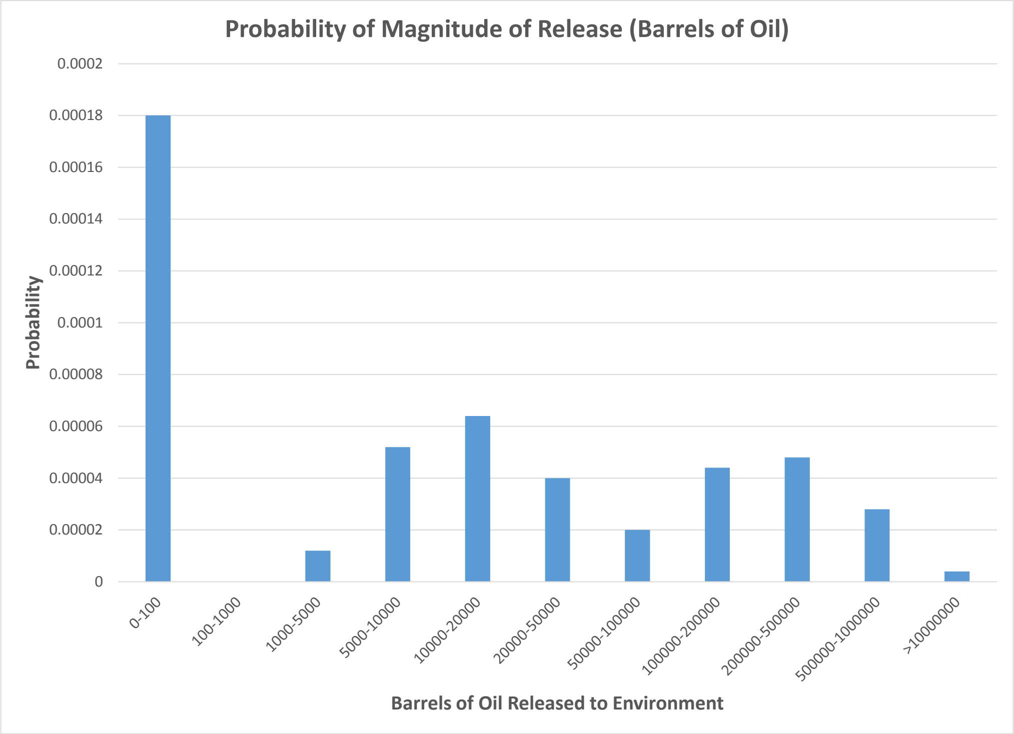

Figure 3- 2. Example Output from Discrete Event Simulation

The bins (e.g. 0-100, 100-1000) are chosen after reviewing the results to determine logical groupings.

A special type of graph called a frequency of exceedance (F-N) curve can be a valuable tool if the probability of exceeding a particular magnitude of consequence is of interest. This type of plot displays the magnitude of the end state on the x axis and the probability of exceeding it on the y axis.

Table 3-5 shows the data needed to construct an F-N curve. The magnitude is developed by the analyst subjectively based on reviewing the results and assigning the output to bins. Note that each bin has a "greater than" designation. The number of replications assigned to each bin is obtained from the results and then divided by the total number of replications, in this case 250,000. This gives the frequency of exceedance for each bin. The results is a graph as shown in Figure 3-3.

Table 3-5. Frequency of Exceedance Calculation

Magnitude (Barrels released)	Replications	Frequency of Exceedance
>0	123	0.000492
>100	78	0.000312
>1000	78	0.000312
>5000	75	0.0003
>10000	62	0.000248

>20000	46	0.000184
>50000	36	0.000144
>100000	31	0.000124
>200000	20	0.00008
>500000	8	0.000032
>10000000	1	0.000004

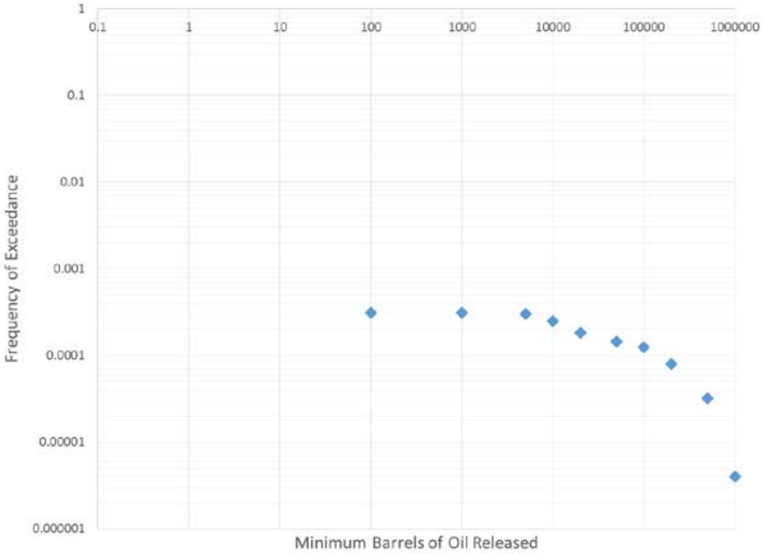

Figure 3- 3. Example Frequency of Exceedance Curve

The frequency of exceedance curves are typically plotted on a log-log scale because the data can span orders of magnitude. Looking at the curve from right to left, flat frequency of exceedance curves indicate that the failures occurring between points are having a minimal effect on the overall magnitude of the consequence. When the curve goes more or less vertical, and the frequency drops very significantly, this means that it is difficult or impossible (highly infrequent) to exceed the corresponding magnitude of release.

3.3.3 Importance Measures

Ranking of risk scenarios based on their frequencies as discussed in Section 3.3.1 provides limited insight regarding the contribution of individual events such as component failures to the total risk. Scenario ranking provides insights on importance of group of failures, not failure of individual events. An event (say, component x failure) can appear in the structure of many low frequency scenarios, yet it may be absent in the definition of the dominant risk scenarios. If the contribution of low frequency scenarios to the total risk is comparable to that of a few dominant risk scenarios, then scenario ranking will not capture the risk importance of component x. To address this issue, and to provide perspective on importance of individual events or parameters of the PRA model, several quantitative importance measures are calculated.

Once the importance measures are calculated, the events or parameters of the risk model can be ranked according to the relative value of the importance measure. This provides some insight into what is, or could be, influencing actual risk. This insight can inform risk-informed decision making (e.g., allocating resources), or point to the need for risk mitigation efforts, such as redesign of hardware components, the addition of redundancy, etc., but should not be a sole basis for decision-making.

The quantitative importance measures typically found in PRA software include:

- Fussell-Vesely (F-V)
- Risk achievement worth (RAW);
- Risk reduction worth (RRW); and
- Birnbaum.

Another measure, the "Differential" importance measure, is discussed in Appendix I, and is reflects the fractional change of a risk metric due to a particular basic event given a change in a basic event probability. This metric is not typically found in PRA software.

All of the above are formulated in failure space: they say something about sets of minimal *cut sets* that involve a specific event or model parameter. A measure based on success space (i.e. *path sets*), Prevention Worth (PW), is a single-event measure that can afford different insights from the failure-space measures.

The three most commonly used importance measures (F-V, RAW, and RRW) are discussed below with examples. Detailed information on the derivation of the failure space based importance measures are included in Appendix I. Prevention Worth, based in success space, is detailed in Appendix J. Finally, a more comprehensive way of looking at model results, "Prevention analysis," is discussed in Appendix K. Instead of looking at basic events one at a time, Prevention analysis answers the question "what *combinations* of basic events should I undertake to prevent, in order to reduce risk in the most cost-effective way?"

3.3.1.1 Fussell-Vesely Importance

The most frequently used importance measure is the F-V importance of basic events. The F-V importance of a given basic event is the fraction of overall risk contributed by the cut sets containing that basic event. This is similar to the scenario risk ranking in Section 3.3.1, but performed at a basic event level. A basic event may show up in many cut sets that are lower frequency than the top scenarios, but the summation of the lower frequency cut sets for that component may show that basic event to be a significant risk because it is included in many scenarios. Because most cut sets are made up of multiple basic events, and the cut set frequency is counted for each basic event F-V importance value, the F-V contributions summed over all basic events will normally be greater than 1.0.

Table 3-6 displays the top 5 cut sets from the example model developed in Section 2. The basic event DRILLINGKICK occurs in all of the cut sets (since it is the only initiating event used), so it has a F-V importance of 1.0. The Driller failing to detect a kick (DRL-HUM-ERR-001) is only found in cut set 2 which has a cut set value of 2.0E-4. The F-V importance for this event is calculated simply by dividing 2.0E-4 by the total risk, 5.572E-4. The result is 0.359 which is the same as the cut set value because the basic event is only included in that single cut set. The basic event ROV-FTR-001 is found in cut sets 3 and 4 with cut set values of approximately 9.53E-5 and 1.01E-5 respectively. In this case, the cut set values are added (1.05E-4) and divided by the total risk for a F-V importance of 0.189. Table 3-7 shows the F-V importance for each of the failure events in Table 3-6. The "/" basic events are not included because they represents success paths on the event tree.

Table 3-6. Top 5 Minimal Cut sets Example for Importance Measures

#	Prob/Freq	Total %	Cut Set	Description
Total	5.572E-4	100		
1	2.471E-4	44.35	DRILLING : 14-1	
	1.000E+0		DRILLINGKICK	Well Kick While Drilling
	3.530E-4		BOP-POD-YLBL-CCF	Common cause failure of blue and yellow pods
	7.000E-1		/ROV	ROV intervention unsuccessful
2	2.000E-4	35.90	DRILLING : 16	
	1.000E+0		DRILLINGKICK	Well Kick While Drilling
	2.000E-4		DRL-HUM-ERR-001	Kick not properly detected
3	9.531E-5	17.11	DRILLING : 14-2	
	1.000E+0		DRILLINGKICK	Well Kick While Drilling
	3.530E-4		BOP-POD-YLBL-CCF	Common cause failure of blue and yellow pods
	9.000E-1		/CAPSTACK	Well Capping unsuccessful
	3.000E-1		ROV-FTR-001	ROV intervention unsuccessful
4	1.006E-5	1.81	DRILLING : 14-3	
	1.000E+0		DRILLINGKICK	Well Kick While Drilling
	3.530E-4		BOP-POD-YLBL-CCF	Common cause failure of blue and yellow pods
	1.000E-1		CAP-LKG-001	Well capping unsuccessful
	9.500E-1		/RELIEFWELL	Relief Well unsuccessful
	3.000E-1		ROV-FTR-001	ROV intervention unsuccessful
5	4.696E-6	0.84	DRILLING : 14-1	
	1.000E+0		DRILLINGKICK	Well Kick While Drilling
	2.590E-3		BOP-POD-FTR-BLUE	Blue pod (standby) fails to run
	2.590E-3		BOP-POD-FTR-YELLOW	Yellow pod (operating) fails to run
	7.000E-1		/ROV	ROV intervention unsuccessful

Table 3-7. F-V Importance Calculation Example

Basic Event	Description	Cut sets with Basic Event	Total Cut set Value	F-V Importance
DRILLINGKICK	Well Kick While Drilling	1, 2, 3, 4, 5	5.572E-4	1.00E+00
BOP-POD-YLBL-CCF	Common cause failure of blue and yellow pods	1, 3, 4	3.52E-04	6.33E-01
DRL-HUM-ERR-001	Kick not properly detected	2	2.000E-4	3.59E-01
ROV-FTR-001	ROV intervention unsuccessful	3, 4	1.05E-04	1.89E-01
CAP-LKG-001	Well capping unsuccessful	4	1.006E-5	1.81E-02
BOP-POD-FTR-BLUE	Blue pod (standby) fails to run	5	4.696E-6	8.43E-03
BOP-POD-FTR-YELLOW	Yellow pod (operating) fails to run	5	4.696E-6	8.43E-03

The F-V importance is based on basic event contributions. When common cause of a component is modeled using separate common cause basic events, the F-V importance for the common cause events is treated separately from the independent failure basic event. In this case, the F-V values from the common-cause cut set and the independent-failure cut set must be added to obtain the total for that particular component. An example of this is shown in Table 3-7 with the blue and yellow pods. The common cause failure event (BOP-POD-YLBL-CCF) has a F-V of 0.633, and the blue pod (BOP-POD-FTR-BLUE) (and yellow pod) has a F-V importance of 0.00843. The total for the blue (and yellow) would be 0.641 after adding the common cause and independent failures.

3.3.1.2 Risk Achievement Worth (RAW)

The F-V importance shows the relative contributions of components and basic events to the overall risk, *given the probability numbers put into the model*. But it cannot be concluded that components and basic events that do not show large contributions are "unimportant." It may simply be that as a result of their low presumed failure probabilities, they do not contribute much to top event likelihood. Another way to review results is to use the RAW importance measure. The RAW basically executes a drastic sensitivity study: it assumes the basic event is failed by substituting a value of 1.0 for the basic event probability in all cut sets containing the event, and recalculating the total risk.[7] The new total risk is divided by the total risk before the substitution to establish a ratio of how much the risk would increase if the basic event was failed.

Using the sample data from Table 3-8 for the ROV intervention unsuccessful (basic event ROV-FTR-001), if a value of 1.0 is substituted for the nominal value of 3.0E-1 (in cut set 2), the new total risk estimate is 8.03E-4 and ratioing over the original estimate (5.572E-4) gives a RAW of 1.44.

Table 3-8. RAW Examples

#	Prob/Freq	Total %	Cut Set	Description
Total	5.572E-4	100		
1	2.471E-4	44.35	DRILLING : 14-1	
	1.000E+0		DRILLINGKICK	Well Kick While Drilling
	3.530E-4		BOP-POD-YLBL-CCF	Common cause failure of blue and yellow pods
	7.000E-1		/ROV	ROV intervention unsuccessful
2	2.000E-4	35.90	DRILLING : 16	
	1.000E+0		DRILLINGKICK	Well Kick While Drilling
	2.000E-4		DRL-HUM-ERR-001	Kick not properly detected
3	9.531E-5	17.11	DRILLING : 14-2	
	1.000E+0		DRILLINGKICK	Well Kick While Drilling
	3.530E-4		BOP-POD-YLBL-CCF	Common cause failure of blue and yellow pods
	9.000E-1		/CAPSTACK	Well Capping unsuccessful
	3.000E-1		ROV-FTR-001	ROV intervention unsuccessful

[7] Ideally, a value of "TRUE" is substituted in the logic model, the top event Boolean expression is re-evaluated and only then requantified.

4	1.006E-5	1.81	DRILLING : 14-3	
	1.000E+0		DRILLINGKICK	Well Kick While Drilling
	3.530E-4		BOP-POD-YLBL-CCF	Common cause failure of blue and yellow pods
	1.000E-1		CAP-LKG-001	Well capping unsuccessful
	9.500E-1		/RELIEFWELL	Relief Well unsuccessful
	3.000E-1		ROV-FTR-001	ROV intervention unsuccessful
5	4.696E-6	0.84	DRILLING : 14-1	
	1.000E+0		DRILLINGKICK	Well Kick While Drilling
	2.590E-3		BOP-POD-FTR-BLUE	Blue pod (standby) fails to run
	2.590E-3		BOP-POD-FTR-YELLOW	Yellow pod (operating) fails to run
	7.000E-1		/ROV	ROV intervention unsuccessful

The RAW is a particularly good measure for identifying single failure points in the model. An example in Table 3-8 is the driller failing to detect a kick in cut set 2. If a 1.0 is substituted for that basic event (DRL-HUM-ERR-001), the RAW is about 1800, which is the inverse of the original total risk estimate (5.572E-4).

The same caution on common cause as applied to F-V applies to the RAW. If a component has independent and common-cause failure basic events, the basic event RAW only applies to that type of failure for that basic event. Caution should also be applied to the RAW when selecting a truncation cutoff. If the truncation cutoff is set too high, some basic events may not appear in the results, and therefore have erroneous RAW values.

3.3.1.3 Risk Reduction Worth

Risk Reduction Worth is closely related to the F-V importance. Where the F-V importance shows the fractional contribution of a basic event to the total risk, the RRW is a ratio of the total risk if the basic event failure probability is set to 0.0 over the nominal total risk. An easy method for calculating the RRW is:

$$1.0 / (1.0\text{-F-V importance})$$

The resulting ratio is the factor by which the risk would be reduced if the failure probability of the basic event was set to 0.0.

3.3.2 Uncertainty

The failure data inputs to a PRA are typically distributions describing the uncertainty around each event being analyzed. These individual uncertainties are used to estimate the uncertainty around the end state(s) of interest in the PRA model. The output from PRA software is typically displayed in two ways, as a probability density curve (Figure 3-4), or a cumulative distribution (Figure 3-5). The probability density represents the relative likelihood (y-axis) for a given probability value (x-axis). The cumulative distribution shows the probability (y-axis) that the end state or event will be less than or equal to the probability value (x-axis).

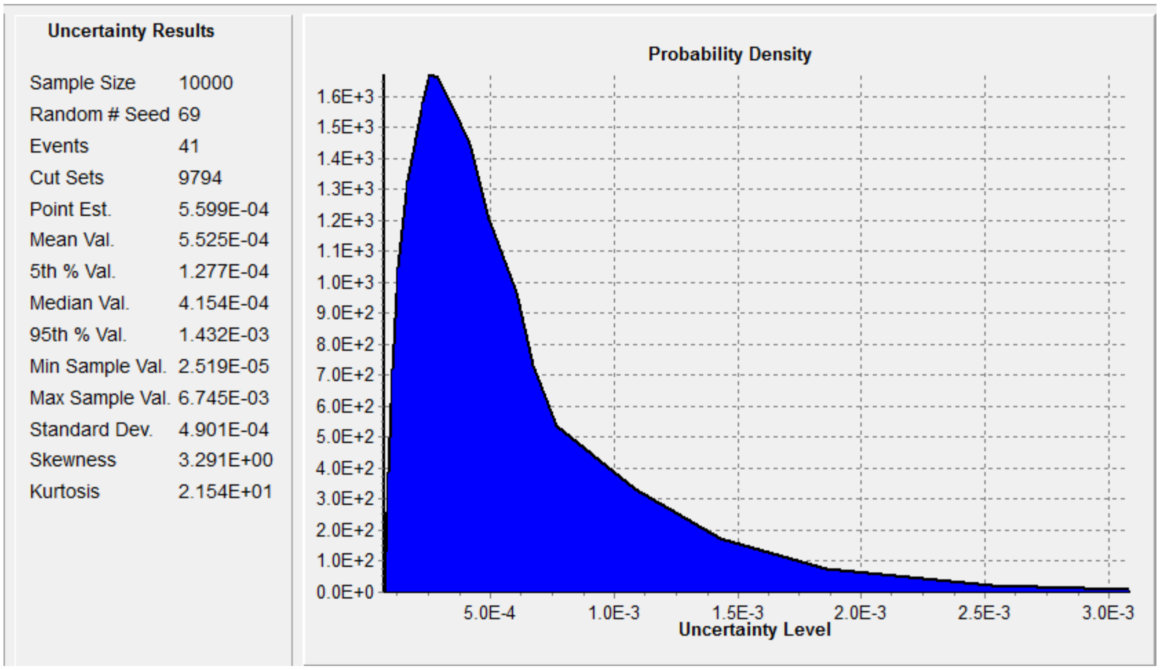

Uncertainty Results

Sample Size	10000
Random # Seed	69
Events	41
Cut Sets	9794
Point Est.	5.599E-04
Mean Val.	5.525E-04
5th % Val.	1.277E-04
Median Val.	4.154E-04
95th % Val.	1.432E-03
Min Sample Val.	2.519E-05
Max Sample Val.	6.745E-03
Standard Dev.	4.901E-04
Skewness	3.291E+00
Kurtosis	2.154E+01

Figure 3- 4. Example Probability Density Function

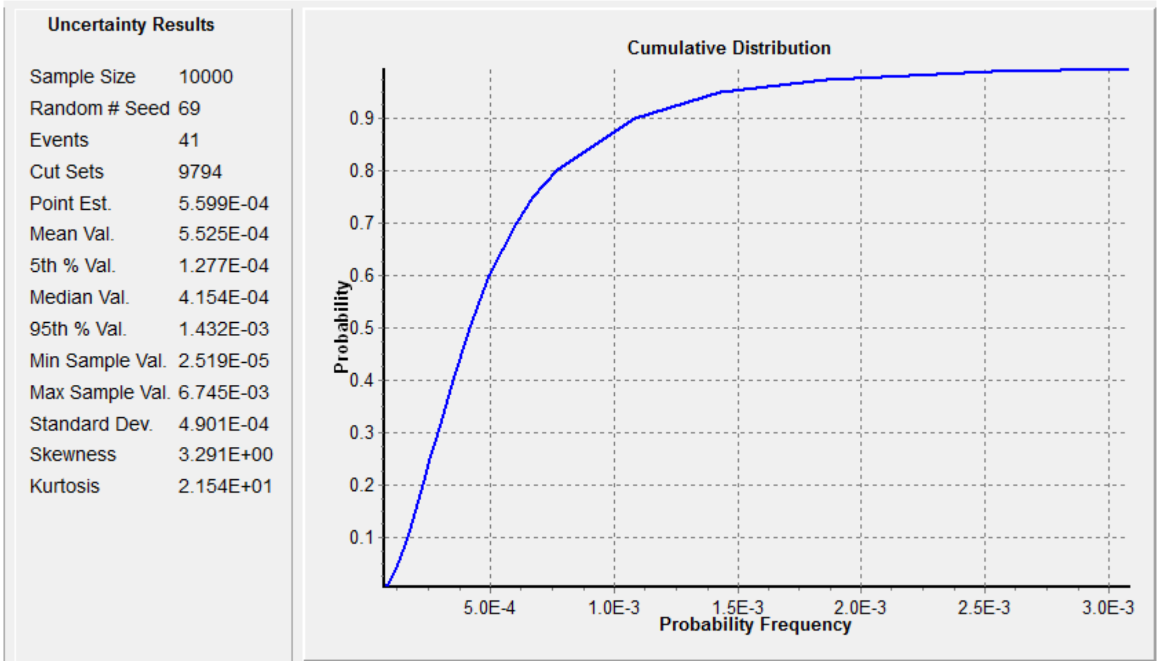

Uncertainty Results

Sample Size	10000
Random # Seed	69
Events	41
Cut Sets	9794
Point Est.	5.599E-04
Mean Val.	5.525E-04
5th % Val.	1.277E-04
Median Val.	4.154E-04
95th % Val.	1.432E-03
Min Sample Val.	2.519E-05
Max Sample Val.	6.745E-03
Standard Dev.	4.901E-04
Skewness	3.291E+00
Kurtosis	2.154E+01

Figure 3- 5. Example Cumulative Probability Distribution

The probability density and cumulative distributions are good for describing the uncertainty of a single end state or event. When one is comparing distributions, a chart like that shown in Figure 3-6 can be used that readily displays a comparison of where the mean values lie as well as the distribution around the means for each point.

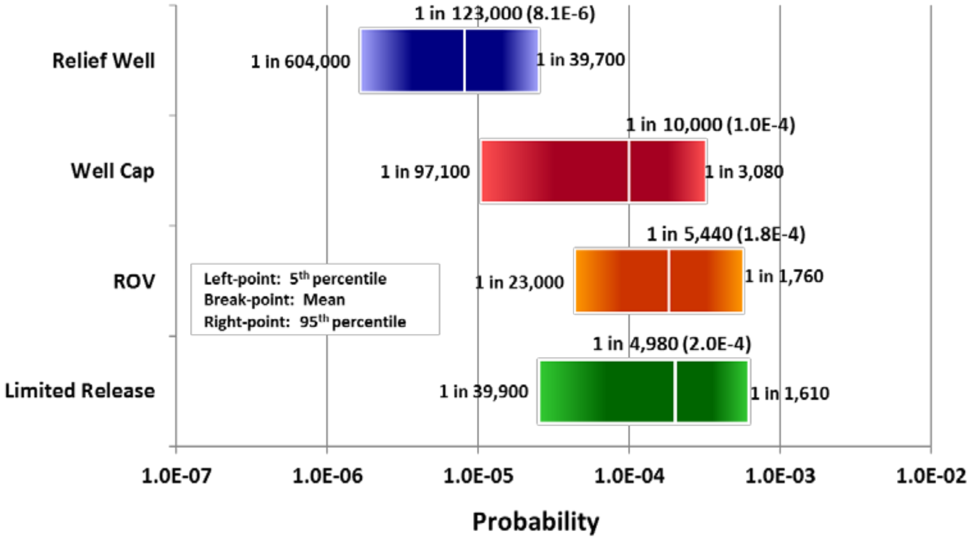

Figure 3- 6. Example Comparison of End State Distributions

3.3.5 System Level Reliability

A PRA is typically done on a facility with interest in a particular end state or states. In developing the model, many systems or functions are analyzed and may be isolated to give insights specifically for those systems and functions. For instance, in the model developed in Section 2, if the analyst wanted to review the causes and contributions to failure of the annular preventer to close, the fault tree results could be used to provide those insights. These are shown in Table 3-9 for the fault tree ANNULAR. When reviewing system or function level results from the PRA model, it is important to note the context for which the fault tree was developed. The PRA will have a specific focus and the system analysis may not include all failures of the system.

Table 3-9. Sample Fault Tree Results for ANNULAR

#	Prob/Freq	Total %	Cut Set	Description
Total	5.059E-3	100	Displaying 8 Cut Sets. (8 Original)	
1	4.180E-3 4.180E-3	82.62	BOP-CYL-FTC-AP01	Annular preventer fails to seal
2	5.000E-4 5.000E-4	9.88	ANNULAROVERPRESS	Well pressure over the design limit of annular
3	3.530E-4 3.530E-4	6.98	BOP-POD-YLBL-CCF	Common cause failure of blue and yellow pods
4	1.180E-5 1.180E-5	0.23	BOP-SHV-LKG-SV01	Crosstie shuttle valve external leakage

5	1.180E-5	0.23		
	1.180E-5		BOP-SHV-LKG-SV02	Annular ROV shuttle valve external leakage
6	6.708E-6	0.13		
	2.590E-3		BOP-POD-FTR-BLUE	Blue pod (standby) fails to run
	2.590E-3		BOP-POD-FTR-YELLOW	Yellow pod (operating) fails to run
7	4.351E-8	< 0.01		
	2.590E-3		BOP-POD-FTR-YELLOW	Yellow pod (operating) fails to run
	1.680E-5		BOP-SHV-FTT-SV01	Crosstie shuttle valve fails to transfer to blue pod
8	2.590E-8	< 0.01		
	1.000E-5		BOP-HUM-ERR-XTIEPODS	Driller fails to select blue pod after yellow pod failure
	2.590E-3		BOP-POD-FTR-YELLOW	Yellow pod (operating) fails to run

3.3.6 Sanity Checks of the Results

When the model quantification is completed and results are obtained, the analyst must perform a sanity check to ensure inputs and outputs are appropriate. On a large model small errors in inputs can make a difference in results.

3.3.6.1 Basic Event Input

A basic check on the input data should be performed at the time of first quantification. A quick review of a basic event listing will reveal any data that has not been input or has a value of 1.0 or 0.0. These should be adjusted as necessary.

3.3.6.2 Fault Tree Linking

Once cut sets are obtained from the first quantification, a review of the basic event names included should be performed to ensure that the correct fault tree results are being used and the fault trees are linked correctly. Improper fault tree linking may yield cut sets with fault tree names versus basic event names. If developed events are used, this is expected; if not, then the fault tree linking must be performed again.

3.3.6.3 Reviewing Results

With the basic event input and fault tree linking verified, the results should be reviewed to determine if other problems exist. This is a subjective review, based on the analyst's knowledge from building the model. Questions to ask are:

Does the overall risk make sense?

Do the top scenarios make sense?

Are any scenarios that were expected to be risk drivers missing?

If the answer to any of these is less than adequate, the analyst may have to trace through specific scenarios to determine why the expected result is not showing up or higher than expected. The answer may be reasonable, if not, troubleshooting is required to make the appropriate fixes to the model.

3.3.6.4 Sensitivity Studies on Assumptions

When information is lacking heavy reliance is placed on the analyst's judgment. For example, assumptions made regarding success requirements for pivotal events and for accident progression can significantly affect the PRA results. The effect of such assumptions needs to be investigated by sensitivity analyses. The results of sensitivity analyses should be reported in tabular form and it should include the base-case assumption (the basis for the nominal PRA results), the alternative assumption and its basis, and the change in the numerical results between the base case and the alternative case.

3.4 Can the model support the decision being made?

A risk model cannot be perfect; complex risk models contain too many idealizations and abstractions to be literally correct at a high level of detail, even without uncertainties; and in many cases, the uncertainties are significant as well. Is the model good enough to be used in the present decision situation? Or should we do additional work on the model? If the model's results point to one decision alternative with a high degree of confidence – and if we believe the model results – then our work is done. On the other hand, if:

- there is sufficient uncertainty about the model's results to limit our confidence in the present decision, and
- there is a way to reduce that uncertainty, and
- the decision stakes are high enough to justify the additional effort,

then more should be done. However, it is necessary first to understand gain a better understanding of what the risk model is saying.

Risk analysis must accomplish the following:

Identification of accident scenarios;-

Estimation of the likelihood of each scenario; and

Evaluation of the consequences of each scenario.

Once this is done, it is necessary to integrate the results into an assurance case, suitable for use by decision-maker(s).

1. The integration includes, among other things, development of best estimates for frequencies and consequences, development of distributions reflecting the uncertainty associated with those estimates, propagation of the uncertainties to obtain final results, and development of appropriate displays to communicate the results with their associated uncertainties. Documentation related to PRA models whose analysis results are used to make critical decisions regarding design, development, manufacturing, and operations that may impact human safety or environmental damage should be reviewed. Specific methods and procedures should be used for assessing and communicating the credibility of PRA model analysis results based on factors such as peer review, input pedigree, uncertainty analysis, results robustness, use history, qualifications of the analysts, and so on.

2. To provide focus for the presentation of results, the results should include identification of

system features that are the most important contributors to risk. Insights into relative importance of various features of the system, and the relative importance of various modeling assumptions, may be developed from uncertainty and sensitivity analyses. A discussion of these insights is required to provide the proper interpretation of the "bottom line" conclusions. Such insights should include an appreciation of the overall degree of uncertainty about the results and an understanding of which sources of uncertainty are critical to those results and which are not. In general, many of the insights gained are not strongly affected by the uncertainties. The numerical results need only be accurate enough to allow the decision maker to distinguish risk- significant elements from those of lesser importance.   The level of detail and the style of presentation of risk results depend on the risk assessment objectives. The results section must communicate the project's motivations and objectives and should be done in a way that clearly establishes the appropriateness of the generated results in meeting the risk assessment objective. For example, if the risk assessment is intended for evaluation of alternative design features as in risk-informed decision-making, the results should be presented in a structure that allows comparison of various design options according to an appropriate ranking scheme.

3. Ultimately the question must be asked: Are the results robust enough to support a decision? If not, what are the soft spots in the analysis (e.g., dominant uncertainties), and what can we do about them?

Appendix A – Example Basic Event Naming Conventions for Fault Trees

As discussed in Section 2.2.5.3, a consistent naming convention for fault tree basic events is necessary for several reasons. Ultimately, a good naming convention helps in reading and parsing model results in an efficient manner; but even more importantly, the Boolean processing function requires that a given basic event be named consistently in all of the logic models in which it occurs. Serious errors can result if this is not done correctly. If a model development is being carried out by more than one individual, enforcement of this condition and similar conditions is a management priority.

Usually the naming convention is in a form similar to:

XXX-YYY-ZZZ-DDDDD,

where XXX, YYY, etc. represent identifying attributes to the component and failure mode that may include (as previously discussed):

- The operation being performed (e.g. drilling);
- The system the component belongs to (BOP);
- The subsystem the component belongs to (e.g. Yellow pod);
- The component (e.g. shuttle valve)
- The failure mode (e.g. Fails to transfer)
- A unique identifier for the valve (usually from a drawing)(e.g. SV837)

The system, component, failure mode, and unique identifier should be included as a minimum, and other fields may be added based on the analysis and the character limitations of the PRA software being used. Table A-1 and A-2 show typical naming conventions for failure modes and components. A 3-letter identifier was used for each, but that can vary depending on the analyst's choice. The number of characters should be related to the number of items to be accounted for in the field. For instance if the operations being analyzed are; drilling, tripping, running casing, and an empty hole, an identifier for operation may only be one letter since only 4 operations are being considered. Failure modes and components typically have many variations, so allocating 3 letters allows flexibility for those items as shown in the examples in Tables A-1 and A-2. Unique identifiers from drawings or other documents may be variable, so as the last field in the name, it may be desirable to not specify the number of characters for that field.

Table A-1. Example Basic Event Naming Convention for Failure Modes

Fails to close	FTC
Fails to open	FTO
Fails to operate on demand	FOD
Fails to reseat	FRS
Fails to run	FTR
Fails to start	FTS
Fails to transfer	FTT
Jammed	JAM

Leakage	LKG
Plugged	PLG
Premature opening	PMO
Rupture	RUP
Short Circuit	SHT
Structural Failure	STR
Transfer closed	XFC
Transfer open	XFO

Table A-2. Example Basic Event Naming Convention for Components

Check Valve	CKV
Gate Valve	VLV
Safety Relief Valve	SRV
Hydraulic/Pneumatic Cylinder	CYL
Shuttle Valve	SHV
Accumulator	ACC
Reservoir	RES
Pump	PMP
Filter	FLT
Diesel Generator	DGN
Circuit Breaker	CBR
Electric Power Bus	BUS
Relay	RLY
Battery	BAT
Flow Switch	FSW
Pressure Switch	PSW
Level Switch	LSW
Manual Valve	MNV

Appendix B – Fault Tree Gate Logic and Quantification

The primary logic gates used in fault tree modeling are the OR gate, AND gate, and the COMBINATION gate shown in Figure B-1.

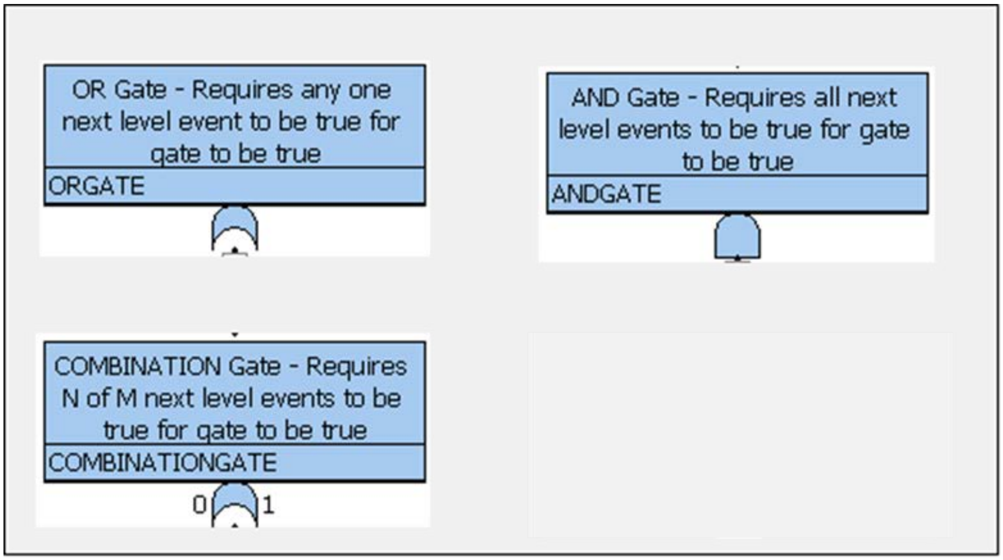
OR Gate - Requires any one next level event to be true for gate to be true
ORGATE
AND Gate - Requires all next level events to be true for gate to be true
ANDGATE
COMBINATION Gate - Requires N of M next level events to be true for gate to be true
COMBINATIONGATE
0
1

Figure B- 1. Common Fault Tree Logic Gates

B.1 OR Gates

The OR gate is used for situations where if any event under the gate is TRUE (has occurred), then the OR gate will be TRUE (or occur). For instance, if a well kick occurs, the Driller has to recognize the well kick and close the annular preventer. For a fault tree, which is developed in failure space, the top event would be “Failure to close the annular preventer after a well kick.” Since both actions have to occur for success, either one (Driller recognizing the kick or annular preventer closing) failing would result in the top event being true. The simple OR gate is shown in Figure B-2.

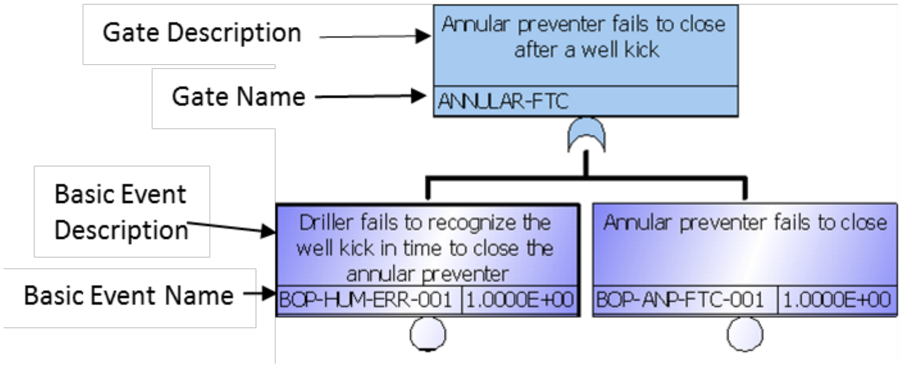
Gate Description
Gate Name
Basic Event Description
Basic Event Name
Annular preventer fails to close after a well kick
ANNULAR-FTC
Driller fails to recognize the well kick in time to close the annular preventer
BOP-HUM-ERR-001
1.0000E+00
Annular preventer fails to close
BOP-ANP-FTC-001
1.0000E+00

Figure B- 2. Simple OR Gate

The OR gate ANNULAR-FTC is true if either basic event BOP-HUM-ERR-001 or BOP-ANP-FTC-001 is true.

The output from the OR gate would result in the following cut sets:

BOP-HUM-ERR-001, and
BOP-ANL-FTC-001.

In Boolean logic the equation for the results becomes:

ANNULAR-FTC = BOP-HUM-ERR-001 ∪ BOP-ANL-FTC-001

The events in a fault tree are generally considered to be independent (with the exception of common cause, discussed in Appendix D). That is, the occurrence of one event does not affect the likelihood of another. Figure B-3 shows a representation of how the events in Figure B-2 are viewed in a Venn Diagram if they are independent. As shown in the Venn Diagram, if the events are truly independent, either one of them could occur; or, some percentage of the time, both could be true as represented by the overlap of the two events. In reality, both would never occur because if the Driller fails, the annular preventer will not have a chance to fail, even if a latent failure is present.

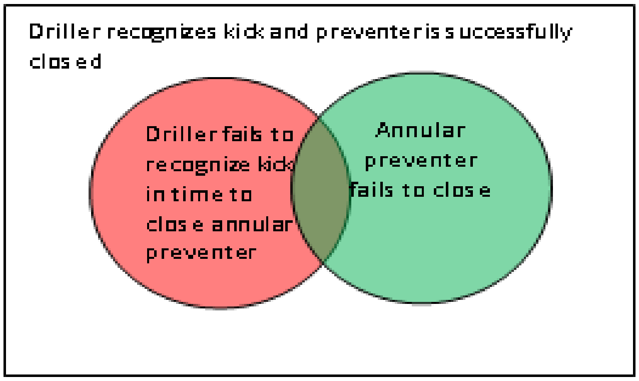
Driller recognizes kick and preventer is successfully closed
Driller fails to recognize kick in time to close annular preventer
Annular preventer fails to close

Figure B- 3. Venn Diagram for Fault Tree Independent Events

Quantification of the OR gate is performed once probabilities are assigned to the basic events. The possibility of the two events being true concurrently must be accounted for if the values of the probabilities are relatively large, i.e. great than 0.01. In order to do this, the intersection of the two events is subtracted from the total in the form:

ANNULAR-FTC = BOP-HUM-ERR-001 + BOP-ANL-FTC-001 - BOP-HUM-ERR-001 * BOP-ANL-FTC-001.

Assigning the probabilities below:

BOP-HUM-ERR-001 = 0.001
BOP-ANL-FTC-001 = 0.001

Gives the equation:

ANNULAR-FTC = 0.001 + 0.001 – 0.001 * 0.001

ANNULAR-FTC = 1.999E-3

In this case the probabilities are small, so the intersection term does not affect the answer significantly. In these cases, the rare event approximation, leaving off the last intersection term, may provide a reasonable answer. If the probabilities were significantly higher, for example 0.5, the equation becomes:

ANNULAR-FTC = 0.5 + 0.5 – 0.5 * 0.5

ANNULAR-FTC = 0.75

In this case, with large probabilities, the answer is significantly affected due to the event independence (the probability of ANNULAR-FTC would be calculated as 1 if it were not corrected for the intersection term).

OR gates may have many inputs, including other gates.

B.2 AND Gates

The AND gate is used for situations where all events under the gate must be true in order for the AND gate to be “TRUE.” For instance, if a BOP has 3 pipe rams and closure of any one would stop the well from flowing, the top event for functional failure would be “Failure to close a pipe ram after a well kick.” Since any one of the 3 pipe rams suffices for success, they all must be failed for the top event to be true. The simple AND gate for this situation is shown in Figure B-4.

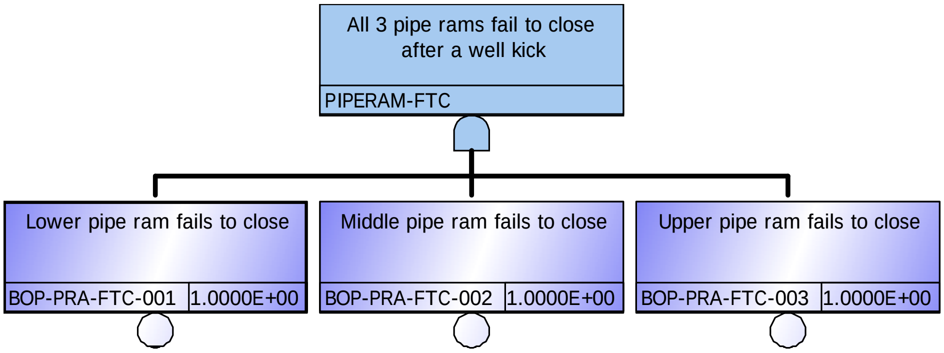
All 3 pipe rams fail to close after a well kick
PIPERAM-FTC
Lower pipe ram fails to close
BOP-PRA-FTC-001
1.0000E+00
Middle pipe ram fails to close
BOP-PRA-FTC-002
1.0000E+00
Upper pipe ram fails to close
BOP-PRA-FTC-003
1.0000E+00

Figure B- 4. Simple AND Gate

The AND gate, PIPERAM-FTC is true if basic events BOP-PRA-FTC-001 AND BOP-PRA-FTC-002 AND BOP-PRA-FTC-003 are true.

The output from the AND gate would result in the single cut set:

BOP-PRA-FTC-001* BOP-PRA-FTC-002 * BOP-PRA-FTC-003

In Boolean logic the equation for the results becomes:

PIPERAM-FTC = BOP-PRA-FTC-001 $\cap$ BOP-PRA-FTC-002 $\cap$ BOP-PRA-FTC-003

Figure B-5 shows a representation of how the events in Figure B-4 are viewed in a Venn Diagram if they are independent. In the previous case of the OR gate, the area representing failure of the top event was the total shaded area. For the AND gate, the area that satisfies the top event condition is that where all shaded areas overlap (“intersect”): label “A” in Figure B-5. Anywhere outside of area A, at least one pipe ram has not failed.

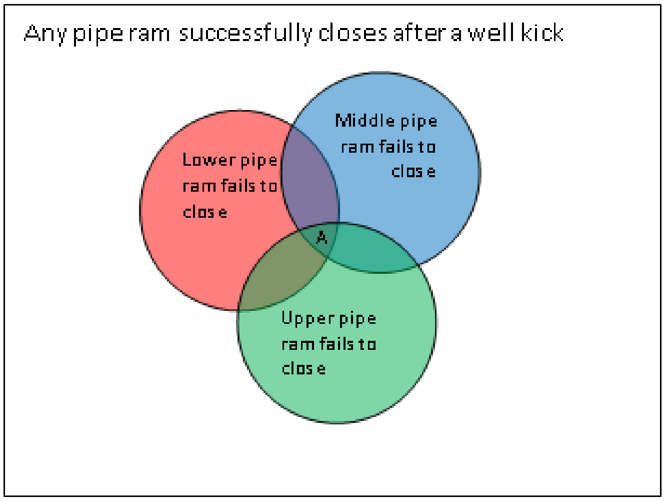
Any pipe ram successfully closes after a well kick
Middle pipe ram fails to clos e
Lower pipe ram fails to clos e
A
Upper pipe ram fails to clos e

Figure B- 5. Venn Diagram for Three Independent Events

Quantification of the AND gate is performed once probabilities are assigned to the basic events. The equation formed to calculate the probability of area "A" in Figure B-5 is:

PIPERAM-FTC = BOP-PRA-FTC-001 * BOP-PRA-FTC-002 * BOP-PRA-FTC-003

Assigning the probabilities below:

BOP-PRA-FTC-001 = 0.001
BOP-PRA-FTC-002 = 0.001
BOP-PRA-FTC-003 = 0.001

gives the equation:

PIPERAM-FTC = 0.001 * 0.001 * 0.001

ANNULAR-FTC = 1.0E-9

B.3 COMBINATION Gates

The COMBINATION gate is used for situations where M (at least 3) events are under the gate and N events (where N is at least 2 but less than M) must be true in order for the COMBINATION gate to be true or occur. For instance, if a MODU has 3 thrusters (3 is used for simplicity in the example) for position keeping and any 2 operating is enough to keep position, the top event would be "At least 2 thrusters fail and position keeping is lost." Since 2 of the 3 thrusters must be operating, if 2 of the 3 fail the top event will be true. The simple COMBINATION gate for this situation is shown in Figure B-6.

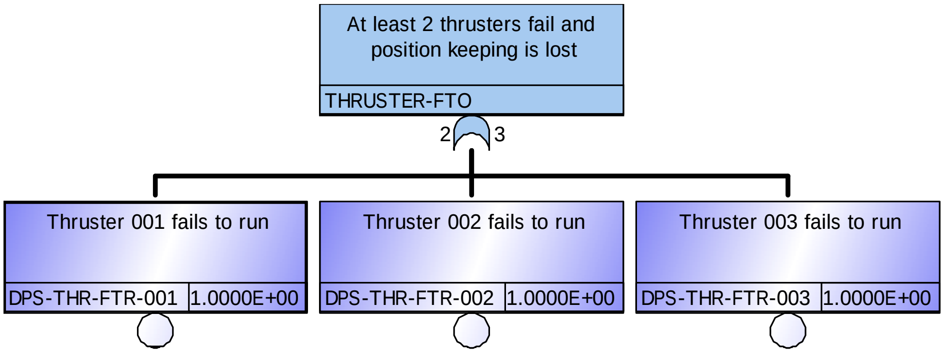
At least 2 thrusters fail and position keeping is lost
THRUSTER-FTO
2
3
Thruster 001 fails to run
DPS-THR-FTR-001
1.0000E+00
Thruster 002 fails to run
DPS-THR-FTR-002
1.0000E+00
Thruster 003 fails to run
DPS-THR-FTR-003
1.0000E+00

Figure B- 6. Simple COMBINATION Gate

The COMBINATION gate, THRUSTER-FTO is true if any two of the basic events DPS-THR-FTR-001, DPS-THR-FTR-002, and DPS-THR-FTR-003 are true.

The output from the COMBINATION gate would result in the three cut sets:

DPS-THR-FTR-001 * DPS-THR-FTR-002,
DPS-THR-FTR-002 * DPS-THR-FTR-003, and
DPS-THR-FTR-001 * DPS-THR-FTR-003.

In Boolean logic the equation for the results becomes:

THRUSTER-FTO = DPS-THR-FTR-001 ∩ DPS-THR-FTR-002 ∪ DPS-THR-FTR-001 ∩ DPS-THR-FTR-003 ∪ DPS-THR-FTR-002 ∩ DPS-THR-FTR-003

Figure B-7 shows a representation of how the events in Figure B-6 are viewed in a Venn Diagram if they are independent. For the COMBINATION gate, the area that satisfies the top event condition is that where at least 2 shaded areas overlap. These areas are labeled “A,” “B,” “C,” and “D” in Figure B-6. Areas “A,” “B,” and “C” are overlaps between 2 thrusters and represent the probabilities that each specific combination of 2 will fail. Area “D” is the overlap of all 3 thrusters and represents the probability that they all fail. This area will also satisfy the top event of *at least* 2 thrusters failing.

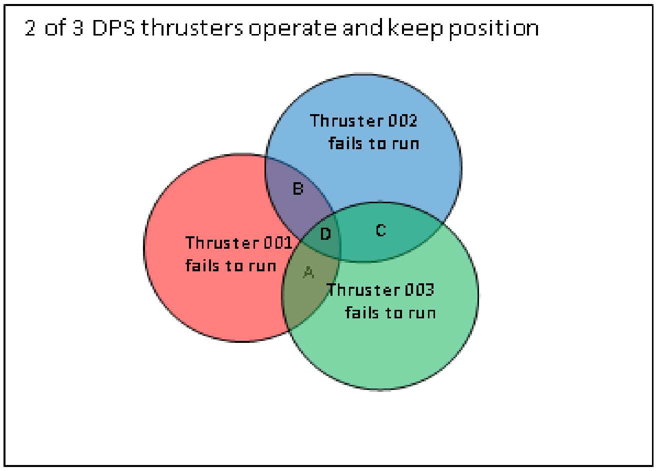
2 of 3 DPS thrusters operate and keep position
Thruster 002 fails to run
B
D
C
Thruster 001 fails to run
A
Thruster 003 fails to run

Figure B- 7. Venn Diagram for Three Independent Events

Quantification of the COMBINATION gate is performed once probabilities are assigned to the basic events. The equation formed to calculate the probability of areas “A,” “B,” ”C,” and “D” in Figure B-7 is:

> THRUSTER-FTO = (DPS-THR-FTR-001 * DPS-THR-FTR-002) + (DPS-THR-FTR-001 * DPS-THR-FTR-003) + (DPS-THR-FTR-002 * DPS-THR-FTR-003) – (2 * DPS-THR-FTR-001 * DPS-THR-FTR-002 * DPS-THR-FTR-003)

The first 3 terms in parentheses in the above equation represent the intersection of each pair of thrusters (areas A, B, and C in Figure B-7). The area D is included in each intersection term, and is therefore counted 3 times if the intersection terms are simply added. The last term in the equation is a correction to account for this over-counting. In this case, the correction is small because the probabilities are small.

Assigning the probabilities below:

DPS-THR-FTR-001 = 0.001
DPS-THR-FTR-002 = 0.001
DPS-THR-FTR-003 = 0.001

Gives the equation:

THRUSTER-FTO = 0.001 * 0.001 + 0.001 * 0.001 + 0.001 * 0.001 – 2 * 0.001 * 0.001 * 0.001

THRUSTER-FTO = 2.998E-6

Appendix C - Calculating Frequency, Reliability, and Availability Metrics

This appendix provides a highly simplified discussion of the basics of quantifying reliability, availability, and frequency of failure metrics for components and for systems. It does so using conventional Markov model graphics. This is done in order to clarify how most PRA software uses the component performance information that must be entered in order to run the program. Most PRA software does not actually use Markov models, but the standard Markov model representation is a useful reminder of what sort of thing the subject calculations actually do, whether they are based on simulation, solution of Markov models, or hand calculations.

For present purposes, it is assumed that the functions of systems and components are well defined, the failure modes of components and their effects have been identified, the rates of occurrence of these failure modes are quantified in some way, and the system configurations that would be considered "successful" have been defined. For a system having redundancy (more than one way to succeed, despite failed components), this would mean that the number of trains or divisions needed for "success" has been defined; and it is additionally recognized that for different kinds of functional demands, different definitions of "success" might apply.

At any given time, a component that is capable of doing its job is "available" (i.e., it is not out for test or maintenance) and in a "good" state. A very simplified state diagram for a single component is shown in Figure C-1.

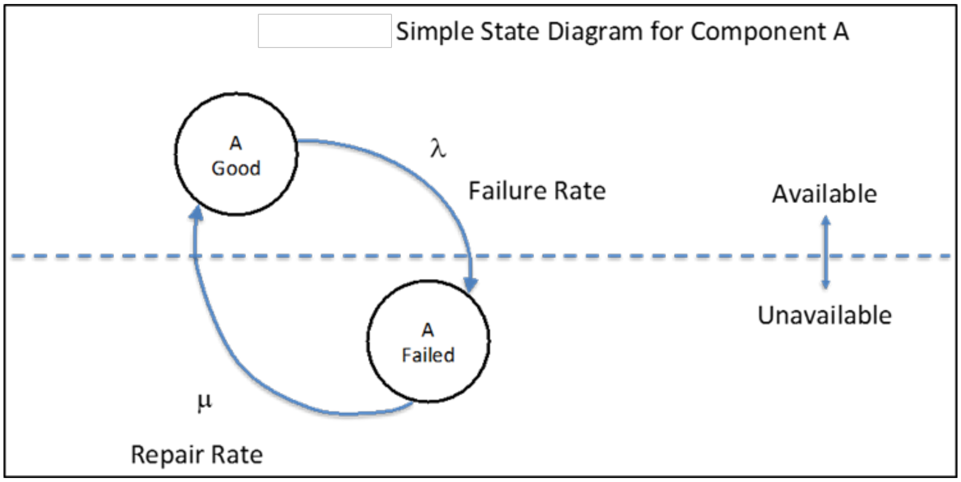

Figure C- 1. Simple State Diagram for Component A

The circles represent component states, and the arcs represent transitions between states. The arc labeled λ is a "failure:" a transition from "good" to "failed." This occurs at the failure rate λ, which has the units of "events per unit time." Analogously, the arc labeled μ corresponds to restoration of the component to "good" status, occurring at the repair rate μ. Within the model underlying this figure, a component is either "good" or "failed," and the probabilities of these states must therefore sum to 1. If we treat λ and μ as stochastic and constant in time, we can write the following equations for the time rate of change of the probabilities of "up" and "down:"

$$\frac{d(up)}{dt} = -\lambda * p(up) + \mu * p(down),$$

$$\frac{d(down)}{dt} = \lambda * p(up) - \mu * p(down),$$

$$p(up) + p(down) = 1.$$

This is an extremely simple example of a class of models called “Markov.” A distinguishing feature of these models is that what happens at any given instant depends only on the state of the system at that instant: a Markov model has no memory of what went before. The modeling of “repair” as a purely stochastic phenomenon, occurring independently of how long a component has been down, is a drastic and unrealistic approximation. But it makes this set of equations trivial to solve, and, for some purposes, is a useful starting point. We can do a better job of modeling things like this in discrete event simulation, which is discussed in Section 2.3.

One can solve the above equations straightforwardly. Typically, the initial condition is: at time 0, p(up)=1 and p(down)=0.

This simple model has the property that over sufficient time, it will converge to a condition in which

$$< p(up) >= \frac{\mu}{\mu+\lambda} \ and < p(down) >= \frac{\lambda}{\mu+\lambda}.$$

This follows from setting the time derivatives to zero, and solving for <p(up)> and <p(down)> using simple algebra. If we were reasoning intuitively, we might argue that the occupancy of the down state is given by the frequency of entering that state (λ), multiplied by the average dwell time in that state ($1/\mu$). This slight difference between this result (1μ) and the above formula results from the need to correct for the availability, discussed below.

Convergence to the steady state is seen in Figure C-2, for illustrative values of λ and μ. The system evolves from its initial condition (p(up)=1) to the steady state given by the above formulas, for the values of λ and μ given on the figure.

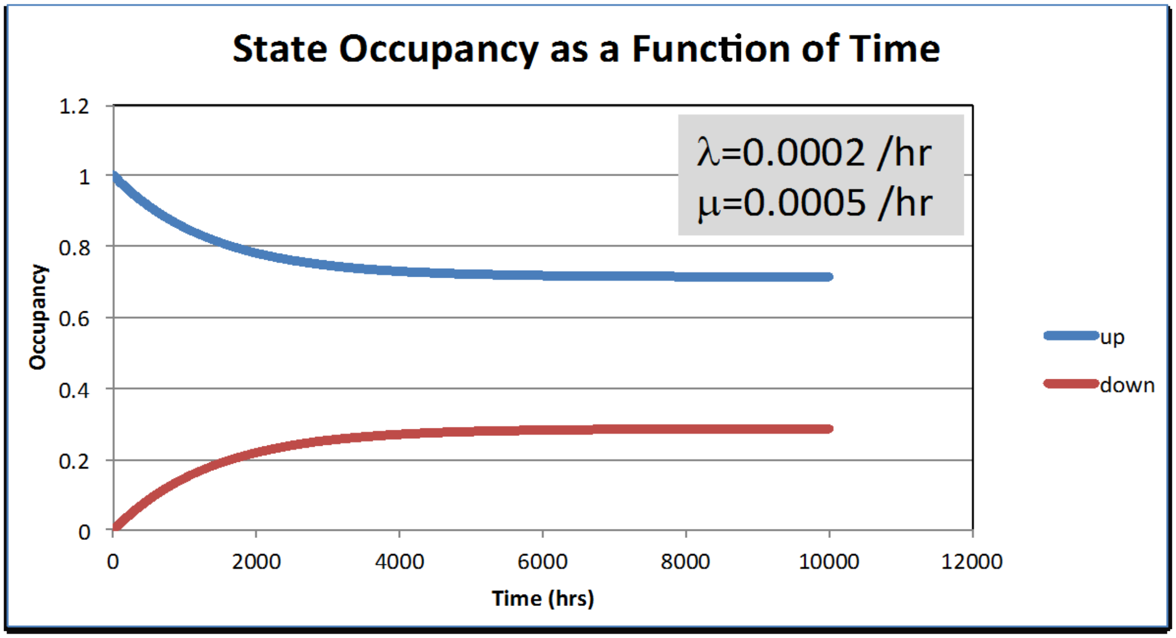

Figure C- 2. Steady State Diagram for Component A

For many components, typical failure rates are on the order of one per many thousands of hours, and typical repair rates are on the order of one per some tens of hours, or less; putting these numbers into the formulae for <p(up)> and <p(down)> yields a number close to 1 for time-averaged availability (<p(up)>), and a small number (equal to 1-availability) for time-averaged *un*availability (<p(down)>).

The rate of failure events actually experienced is not given simply by λ; the component must be "up" in order to be able to fail. The *rate* of failures actually experienced is therefore λ*p(up)*. If *p(up)* is close to unity, then equating the expected rate of failures to λ is not a severe approximation; but in checking computer calculations, the difference between λ and the observed rate of failures may be observable, if unavailabilities are on the order of a few percent, which can easily be the case.

This point generalizes: the rate at which any arc is traversed is the product of the rate associated with that arc, multiplied by the occupancy (the probability) of the state from which the arc originates.

If a component is known to be "good" at time = 0, then the probability that it is failed at time T is ~ λ*T, for T<<1/λ. *Averaged over this interval,* the probability of being in a failed state is (1/2)*λ*T.

Figure C-3 makes a slightly different point. In Figure C-1, a component was either "good" or "in repair." In some systems, failure will not immediately be detected, and we need more than a "repair rate" concept to build a model. In Figure C-3, a failed component is placed into repair only when the failure is detected, which could occur either as a result of an actual demand on the system, or as a result of a scheduled test, carried out for the very purpose of detecting a failed state.

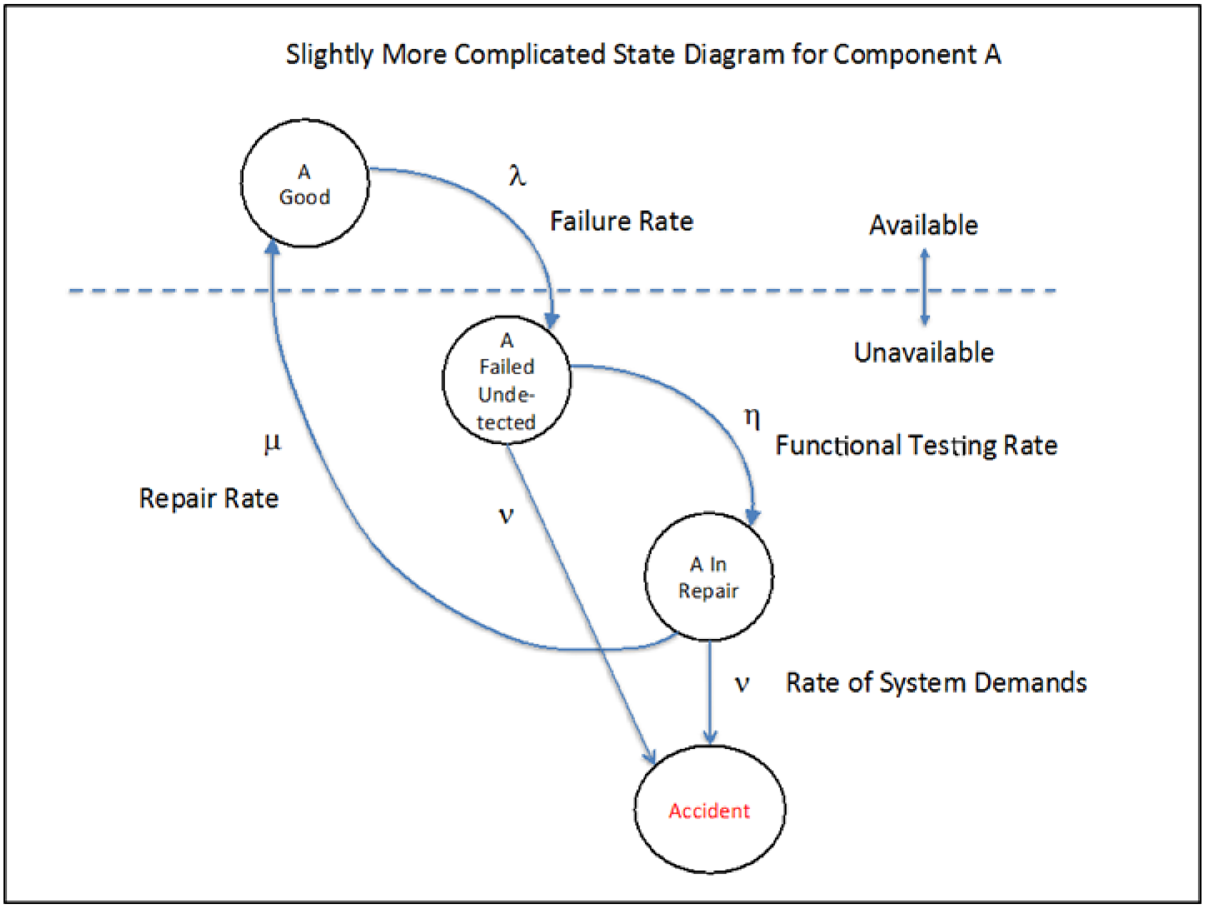

Figure C- 3. Slightly More Complicated State Diagram for Component A

These figures begin to illustrate a general principle underlying the calculation of complex “rates” (accident frequencies, functional failure frequencies, …). Figure C-4 shows a two-component system with a one-out-of-two success criterion: if the system is demanded and either component works, the system succeeds; if both components are down, the system fails. The accident rate is the rate of demands multiplied by the probability of both components being down. As modeled in Figure C-4, that latter probability depends on the underlying failure rates and repair rates; but as mentioned above, we need a way of detecting component failures (as in Figure C-3) before we initiate repair.

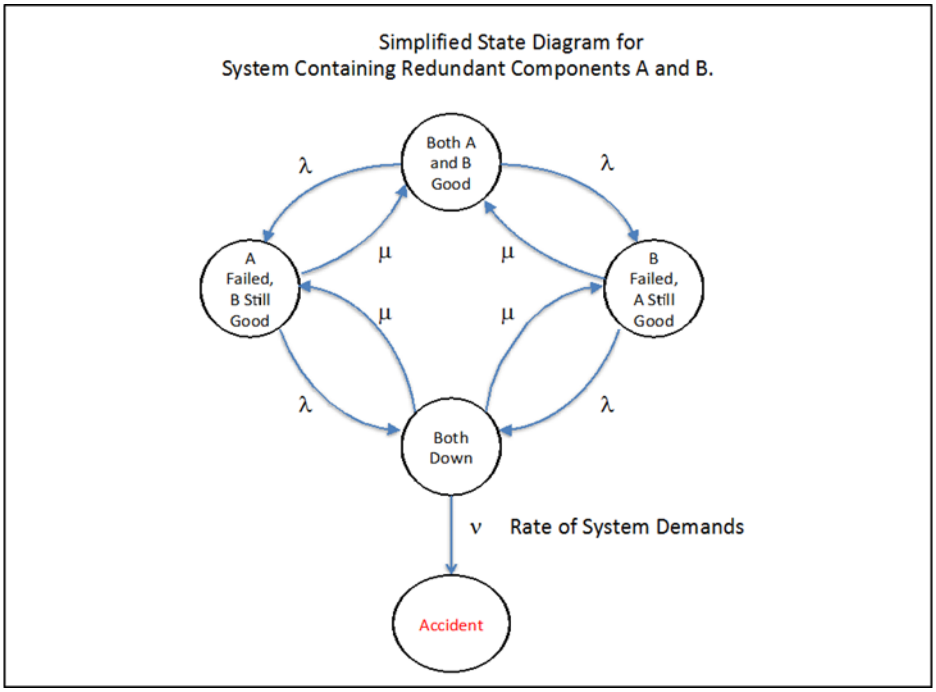

Figure C- 4. Simplified State Diagram for System Containing Redundant Components A and B

If, for some reason, we are interested in the *rate* of system failure, we obtain this by summing the rates of traversing the two arcs into the "both down" state: that is,

$\lambda *$ p(A failed, B still good) + $\lambda *$ p(B failed, A still good).

In all of this, we have assumed that

- there is no causal linkage between the failures of A and B,
- there is no causal linkage between the rate of demands and the failure rates,
- all of the rates are constant in time (even the repair rate, and even though this is unlikely to be a realistic description).

"Failure on Demand"

The Reactor Safety Study (WASH-1400) defines "Demand Probabilities" as:

> … the probability that the device will fail to operate upon demand for those components that are required to start, change state, or function at the time of the accident [*sic*]. The demand probabilities, denoted by Q_d, incorporate contributions from failure at demand, failure before demand, as well as failure to continue operation for a sufficient period of time for successful response to the need. When pertinent, the demand data Q_d can be associated with standard cyclic data or can be interpreted as a general unavailability. Human error data can also be associated with demand probabilities (i.e. per action) as discussed in the human evaluation section.

Not all communities of practice make use of all aspects of this definition. Some argue that if a component is "good" in the instant before a demand, it will (by definition) function during the demand; within this concept, “failures” are either failures in standby, or failures to run, and the occupancies of failed states are quantified accordingly (i.e., in terms of a standby failure rate or a rate of failure to run). Others argue that owing to variability in the stresses imposed by a particular demand, there is a nonzero probability that a nominally “good” component will fail upon the arrival of a demand. Still others would argue for modeling a state between “good” and “bad” (i.e., “degraded”) having a probability of failure on demand that is significant but still less than unity. This is a modeling decision to be evaluated on a case-by-case basis; the present point is that operationally, Q_d is simply the state probability that one multiplies by a “demand” arc to get the frequency of accidents or functional failures, as the case may be.

Although it may seem simple and convenient to lump all causes of component non-performance together, it is conventional to split out maintenance unavailability contributions from actual component failures, because in some applications, operational rules proscribe having multiple components out for maintenance, and the logic model needs to reflect that point: the model should not generate system failure cut sets in which everything is out for maintenance, unless that can, in fact, occur. Sometimes it is necessary to split out failure to start from failure to run, because the consequences are different, or perhaps because common cause failure considerations are different for the two failure modes, and so on.

Appendix D – Common Cause Failure

To be added

Appendix E - Sources of Failure Rate and Event Data

E.1 Background

A fundamental requirement to quantify a risk assessment model is the basic equipment failure rate data. These data comprised of numerical estimates of failure rate and event data that are used in the model and best represent the failure rate characteristics of the facility. There are several categories of failure that are included in a risk model. These include:

- loss of containment (leaking or rupture) of equipment that belong to the hydrocarbon containment envelope
- failure on demand of a component within a safeguard system when required, and
- external event rate of occurrence for events that challenge the facility to maintain critical safety and environmental integrity functions.

Ideally, parameters of PRA models of a specific system should be estimated based on operational data of that system. The next most representative data is that from the fleet of similar facilities operated by the same entity.

Often, however, the analysis has to rely on a number of sources and types of information if the quantity or availability of system-specific data are insufficient. In such cases surrogate data, generic information, or expert judgment are used directly or in combination with (limited) system-specific data. According to the nature and degree of relevance, data sources may be classified by the following types:

- Historical performance of successes and failures of an identical piece of equipment under identical environmental conditions and stresses that are being analyzed (e.g., direct operational experience).

- Historical performance of successes and failures of an identical piece of equipment under conditions other than those being analyzed (e.g., test data).

- Historical performance of successes and failures of a similar piece of equipment or similar category of equipment under conditions that may or may not be those under analysis (e.g., another program's test data, or data from handbooks or compilations). General engineering or scientific knowledge about the design, manufacture and operation of the equipment, or an expert's experience with the equipment.

E.2 Generic Data Sources

Generic data is surrogate or non-specific information related to a class of parts, components, subsystems, or systems. Most generic data sources cover hardware failure rates. All other data categories, particularly human and software failure probabilities, tend to be much more mission-specific, system-specific, or context dependent. As such, generic data either do not exist or need to be significantly modified for use in a PRA.

The international offshore industry has performed risk and reliability assessments for a variety of facilities for over 30 years. Each of these quantitative evaluations tends to increase the general collection of risk and reliability information when this information is stored or published for later use. In addition to the individual quantitative evaluations, various industry entities also manage failure data and incident reporting systems, for example the *** system. A selection of offshore industry data collection systems includes:

- Guidelines for Process Equipment Reliability Data with Data Tables
- Process Equipment Reliability Database (PERD)
- Failure Rate and Event Data for use within Risk Assessments (HSE PCAG)
- Failure Frequency Guidance Process Equipment Leak Frequency Data for use in QRA
- Lees' Loss Prevention in the Process Industries (Third Edition)
- OGP Risk Assessment Data Directory
- OIR/12
- Offshore Reliability Data
- Pipeline and Riser Loss of Containment (PARLOC) Report
- Wellmaster RMS
- Worldwide Offshore Accident Database

These data sources are presented in this guideline along with content descriptions. These sources are commonly utilized in Probabilistic Risk Assessments conducted for offshore facilities. This list is not exhaustive nor endorsed for use, but simply a compilation of frequently used sources. They are presented in alphabetical order.

Table E-1. Guidelines for Process Equipment Reliability Data with Data Tables

Name	Guidelines for Process Equipment Reliability Data with Data Tables
Sponsor/Author:	Center for Chemical Process Safety of the American Institute of Chemical Engineers
Data Types	The level three taxonomy contains 50 component types under the following groups; • Electrical Equipment • Instrumentation • Process Equipment • Protection Systems
Description	Failure rate data handbook, multi-industry sources.
Number and Type of Records:	Book 300 pages, 75 individual failure rate estimate pages.
Frequency of Update	None
Time Frame:	Prior to 1989
Data Access	Commercial publication
Notes	The PERD handbook is a compilation of data tables based on literature review of estimates from many industries and from proprietary files of previously analyzed and selected information. There is no clear relationship to analysis of individual failure events although the format resembles other handbooks which are based on estimates derived from analysis of event data from equipment populations. The intent of the data is for use in the Chemical Process Industry. Failure rate estimates are given as lower, mean and upper bound as failure rates or demand failure probabilities, by failure mode.
Reference	ISBN 0-8169-0422-7

Table E-2. Process Equipment Reliability Database (PERD)

Name	**Process Equipment Reliability Database (PERD)**
Sponsor/Author:	Center for Chemical Process Safety of the American Institute of Chemical Engineers
Data Types	Relief Devices
Description	Event failure database
Number and Type of Records:	2,000 relief valve inventory records and over 5,000 proof test event records.
Frequency of Update	0
Time Frame:	2001-2013
Data Access	Tiered membership scheme. Access to raw event data to allow statistical data analysis for contributing members.
Notes	The PERD database is based on taxonomies developed within the PERD project and is an extension of the Guidelines for Process Equipment Reliability Data with Data Tables. Relief devices were selected to collect event and test data and implement the database.
Reference	http://www.aiche.org/ccps/resources/process-equipment-reliability-database-perd

Table E-3. Failure Rate and Event Data for use within Risk Assessments (HSE PCAG)

Name	Failure Rate and Event Data for use within Risk Assessments (HSE PCAG)
Sponsor/Author:	UK Health and Safety Executive, Hazardous Installations Directorate
Data Types	Categories include: mechanical, electrical, bulk transport and moveable storage. Specific types include: vessels, reactors, valves, pumps, hoses and couplings, flanges and gaskets, Pipelines and compressors.
Description	Non-mandatory reference compiled by the agency for assisting their
Number and Type of Records:	A compilation of many of references ranging from proprietary study reports to textbooks comprising 96 pages of data tables and background information
Frequency of Update	0
Time Frame:	1972-2012
Data Access	Publication available from HSE website.
Notes	HID CI5 has an established set of failure rates that have been in use for several years in QRAs submitted for land use planning cases. The estimates "do not necessarily take account of all factors that could be relevant and significant at particular installations." However, in the absence of site specific data, the values given here may serve as a starting point for safety reports.
Reference	http://www.hse.gov.uk/landuseplanning/failure-rates.pdf

Table E-4. Failure Frequency Guidance Process Equipment Leak Frequency Data for use in QRA

Name	Failure Frequency Guidance Process Equipment Leak Frequency Data for use in QRA
Sponsor/Author	DNVGL
Data Types	Compressors Centrifugal and reciprocating Filters Flanges Heat Exchangers (Air Cooled, Plate, Shell and Tube) Pig Traps Process Piping Pumps (Centrifugal and Reciprocating) Instruments Valves (Actuated and Manual) Pressurized Process Vessels Atmospheric Storage Tanks
Description	A proprietary publication containing guidance and data on process equipment leak frequency for use in QRA. In this document, a detailed review and comparison is made between the DNV taxonomy and frequency values and the UK HSE HCRD Database taxonomy and frequency values. Additional comparisons are made to guidance developed by Flemish and Dutch governments for the same purpose.
Number and Type of Records:	The guide is 40 pages, with 20 pages of data tables presenting leak frequency by equivalent hole size for each equipment type.
Frequency of Update	Continuously
Time Frame:	2005-2012
Data Access	Proprietary publication available for purchase from DNVGL.
Notes	The leak frequency data contained in the guidance document was generated by the LEAK software, which is an application that contains a continuously updated database of leak frequency data and a structured computational capability for leak frequency calculations
Reference	https://www.dnvgl.com/services/calculate-leak-frequency-data-leak-1759

Table E-5. Lees' Loss Prevention in the Process Industries (Third Edition)

Name	Lees' Loss Prevention in the Process Industries (Third Edition)
Sponsor/Author	Texas A&M University, Department of Chemical Engineering
Data Types	Vessels and tanks Pipework Heat Exchangers Rotating Machinery Valves Instruments Process Computers Relief Systems Fire and Gas Detection Systems Fire Protection Systems Emergency Shutdown Systems Utility Systems LNG Plants Leaks Ignition Explosion following ignition Fires Explosion Transport External Events
Description	A well-known, seminal reference 3 volume text compiling the wide range of topics relevant to process safety.
Number and Type of Records:	Appendix 14 of this reference is titled Failure and Event Data. It compiles 38 pages of reference failure rate data
Frequency of Update	3rd Ed (2005), 2nd Ed (1994), 1st Ed (1979)
Time Frame:	Cited references range from 1960 – 2004
Data Access	Commercial publication
Notes	Failure rate data contained in the book are compilations of many failure rate publications from numerous industries. Failure rate estimates are reproduced from cited publications.
Reference	ISBN 0-7506-7589-3

Table E-6. OGP Risk Assessment Data Directory

Name	OGP Risk Assessment Data Directory
Sponsor/Author	International Association of Oil and Gas Producers
Data Types	Major accidents Occupational risk Land transport accident statistics Aviation transport accident statistics Water transport accident statistics Construction risk for offshore units Process release frequencies Risers & pipeline release frequencies Storage incident frequencies Blowout frequencies Mechanical lifting failures Ship/installation collisions Ignition probabilities Consequence modelling Structural risk for offshore installations Guide to finding and using reliability data for QRA Vulnerability of humans Vulnerability of plant/structure Escape, evacuation and rescue Human factors in QRA
Description	The Risk Assessment Data Directory is a series of guidance documents that provide data and information for use to improve the quality and consistency of risk assessments with readily available benchmark data. The directory includes references for common incidents analyzed in upstream production operations.
Number and Type of Records:	20 individual documents (Datasheets)
Frequency of Update	1st Ed (1997), 2nd Ed (2009)
Time Frame:	Prior to 2009
Data Access	Commercial publication
Notes	This series of documents was commissioned with the specific goal of defining generic data for use in QRAs
Reference	http://www.iogp.org/pubs

Table E-7. OIR/12

Name	OIR/12
Sponsor/Author	UK Health and Safety Executive
Data Types	Hydrocarbon release event database compiled by UK Health and Safety Executive, with periodic publications of the analysis of these data in publically available report format.
Description	Event data are required to be submitted under the Reporting of Injuries, Diseases and Dangerous Occurrences Regulations 1995 - (RIDDOR 95). OIR/12 addresses offshore hydrocarbon release events.
Number and Type of Records:	585 event records.
Frequency of Update	Continuous
Time Frame:	2001-2008, previous data deemed inconsistent with current analysis taxonomy and analysis requirements
Data Access	Publication available online
Notes	Data are analyzed by time trend, platform type, platform age, release magnitude, system, cause. Data prior to 2001 are presented with disclaimer.
Reference	http://www.hse.gov.uk/research/rrpdf/rr672.pdf

Table E-8. Offshore Reliability Data

Name	Offshore Reliability Data
Sponsor/Author	OREDA, managed, produced and distributed by Veritec, followed by Sintef/Det Norske Veritas
Data Types	Comprehensive topsides and subsea production equipment, safety equipment and limited onshore exploration and production equipment
Description	OREDA is a project organization sponsored by eight oil and gas companies with worldwide operations. OREDA's main purpose is to collect and exchange reliability data among the participating companies and act as The Forum for co-ordination and management of reliability data collection within the oil and gas industry.
Number and Type of Records:	Event records from 278 installations, 17,000 equipment items with 39,000 failure and 73,000 maintenance records. The database also includes subsea fields with over 2,000 years operating experience.
Frequency of Update	6th Ed (2015), 5th Ed (2009) 4th Ed (2002), 3rd Ed (1997), 2nd Ed (1992), 1st Ed (1984)
Time Frame:	Corresponding with updates
Data Access	Tiered membership scheme. Access to raw event data to allow statistical data analysis for contributing members. Handbook available for purchase.
Notes	All estimates in these handbooks are derived from statistical analysis of event data.
Reference	https://www.oreda.com/

Table E-9. Pipeline and Riser Loss of Containment (PARLOC) Report

Name	Pipeline and Riser Loss of Containment (PARLOC) Report
Sponsor/Author	Oil and Gas UK
Data Types	Pipeline and Riser
Description	The Pipeline and Riser Loss of Containment (PARLOC) report is a source of risk assessment data for generic loss of containment frequencies and covers pipelines and risers in the offshore oil and gas industry.
Number and Type of Records:	206 incident events, loss of containment and near miss. 10,000 km-yr pipeline, 4,000 riser-yr from the UK sector of North Sea, eastern Irish Sea, West of Shetland
Frequency of Update	1990, 1992, 1994, 1996, 2001online in 2006/07Hard copy update commenced 2013
Time Frame:	1988 - Current
Data Access	Commercial publication
Notes	Most complete and homogeneous dataset of subsea pipeline and riser incident event data
Reference	http://oilandgasuk.co.uk/parloc.cfm

Table E-10. Wellmaster RMS

Name	Wellmaster RMS
Sponsor/Author	7 member companies, managed by exprosoft
Data Types	Subsea (subsurface) equipment types, limited subsea (seabed) equipment, e.g. X-mas Trees.
Description	The world's largest database of reliability data for well and subsea equipment
Number and Type of Records:	6,000 wells/40,000 well years
Frequency of Update	Continuous
Time Frame:	1986-2016
Data Access	Online access available commercially
Notes	All estimates are derived from statistical analysis of event data using the online application
Reference	https://wellmaster.exprosoft.com

Table E-11. Worldwide Offshore Accident Database

Name	Worldwide Offshore Accident Database
Sponsor/Author	DNVGL
Data Types	Accident events from global population
Description	Event data including name, type and operation mode of the unit involved in the accident, date, geographical location, chain of events, causes and consequences, and evacuation details
Number and Type of Records:	6451 accidents occurring among 3795 operating units
Frequency of Update	Continuous
Time Frame:	1986-2016
Data Access	Purchase of data search consultancy or a database subscription. The program is a web application
Notes	Comprehensive database of offshore accident event data
Reference	https://www.dnvgl.com/services/world-offshore-accident-database-woad-1747

It is important to recognize the perspective of the risk modeler in order to establish requirements on the quality of failure rate and event data to be used in a risk model. Once a complete risk model is constructed and quantified, it is often the case that a large number of individual failure rate and event data input values do not strongly influence the overall calculated level of risk or contribute to insights provided by analyzing risk contributors. This being the case, the requirements for high fidelity and representative failure rate and event data should vary corresponding to the significance to the calculated risk results. In short, if the failure rate and event data do not significantly influence the results, then we can use lower quality estimates.

E.3 System-Specific Data Collection and Classification

System-specific data can be collected from sources such as:

- Maintenance Logs
- Test Logs
- Operation Records

In the majority of cases, system-specific data are gathered from operation and test records in their "raw" form (i.e., in the form that cannot be directly used in a statistical analysis). Even when data have already been processed (e.g., reduced to counts of failure), care must be exercised to ensure that the data reduction and processing are consistent with QRA modeling requirements, such as having a consistent failure mode classification, and correct count of the total number of tests or actual demands on the system).

In collecting and classifying hardware failure, a systematic method of classification and failure taxonomy is essential. A key element of such taxonomies is a classification of the functional state of components. One such classification system has been offered in Reference [E-1]. Using a taxonomy implies a knowledge structure used to describe a parent-child relationship (i.e., a hierarchy). Under the guidelines for evaluation of risk and reliability-related data, the taxonomy provides the structure by which data and information elements provide meaning to analysts. Within the risk and reliability community, a variety of taxonomies and associated definitions are used. ISO 14224 provides a taxonomy for collection and processing of equipment failure data in the petroleum industry.

When concerned about the physical *causes* of failures, a set of physics-based causal factors would be required. However, this low level of information is not necessary if the inference being made for a specific component or system is concerned with – in general – failures or successes. If, instead, we wished to infer the probability of failure conditional upon a specific failure mechanism, we would need to have information related to the nature of failure (e.g., the physical causal mechanisms related to specific failures).

In other words, this classification can take place via a failure modes and effects analysis, similar to the functional failure modes and effects analysis. Henley and Kumamoto [E-2] carried this idea one step further when they proposed a formal cause-consequences structure to be stored in an electronic database. In their approach, specific keywords, called modifiers, would be assigned to equipment failures. For example, modifiers for on-off operation included: close, open, on, off, stop, restart, push, pull, and switch. Alternative hierarchy related to system/component/failure modes may look like:

System
└ Component
└ Failure Mode
└ Affected Item
└ Failure Mechanism
└ Failure Cause

With regard to the intended function and in reference to a given performance criterion, a component can be in two states: *available* or *unavailable*. The unavailable state includes two distinct sub-states: *failed* and *functionally unavailable*, depending on whether the cause of the unavailability is damage to the component or lack of necessary support such as motive power. The state classification also recognizes that even when a component may be capable of performing its function (i.e., it is available), an incipient or degraded condition could exist in that component, or in a supporting component. These failure situations are termed *potentially failed* and *potentially functionally unavailable*, respectively. These concepts have proven useful in many PRA data applications.

Another aspect of reliability data classification is the identification of the failure cause. In the context of the present discussion, the cause of a failure event is a condition or combination of conditions to which a change in the state of a component can be attributed. It is recognized that the description of a failure in terms of a single cause is often too simplistic. A method of classifying causes of failure events is to progressively unravel the layers of contributing factors to identify *how* and *why* the failure occurred. The result is a chain of causal factors and symptoms.

A hierarchy of *parts or items* that make up a component is first recognized, and the functional failure mode of the component is attributed to the failure or functional unavailability of a subset of such parts or items. Next the physical sign or *mechanism* of failure (or functional unavailability) of the affected part(s) or item(s) are listed. Next the *root cause* of the failure mechanism is identified. Root cause is defined as the most basic reason or reasons for the failure mechanism, which if corrected, would prevent reoccurrence. The root cause could be any causal factor, or a combination of various types of causal factors.

E.4 References

E-1 OREDA, "Offshore and Onshore Reliability Data 6th edition," 2015

E-2 H. Kumamoto and E. J. Henley, "Probabilistic Risk Assessment and Management for Engineers and Scientists 2nd Edition," IEEE Press, Piscataway, New Jersey, 1996

Appendix F - Further Discussion of Bayesian Updating

F.1 Simple Examples

F.1.1 Updating of Prior for a Poisson Example

In this example, the goal is to estimate an hourly failure rate for a component, assuming that the failures obey a Poisson distribution. We choose a lognormal distribution for the prior, and a Poisson distribution for the likelihood model. [8] The operational data for the component category indicate 2 failures in 10,000 hours.

Since the prior distribution is lognormal, and the likelihood function is Poisson, and these two are not "conjugate," the posterior distribution must be derived numerically. The prior and posterior distributions are shown in Figure F-1, along with the "maximum likelihood estimate" (the MLE). The MLE, the value of the parameter for which the likelihood function P(E|parameter) is maximum, is 2E-4 in this case (failures / hours). Note that the probability density functions are plotted as a function of *log* frequency.

The posterior distribution is shifted from the prior distribution towards the MLE. This is typical.

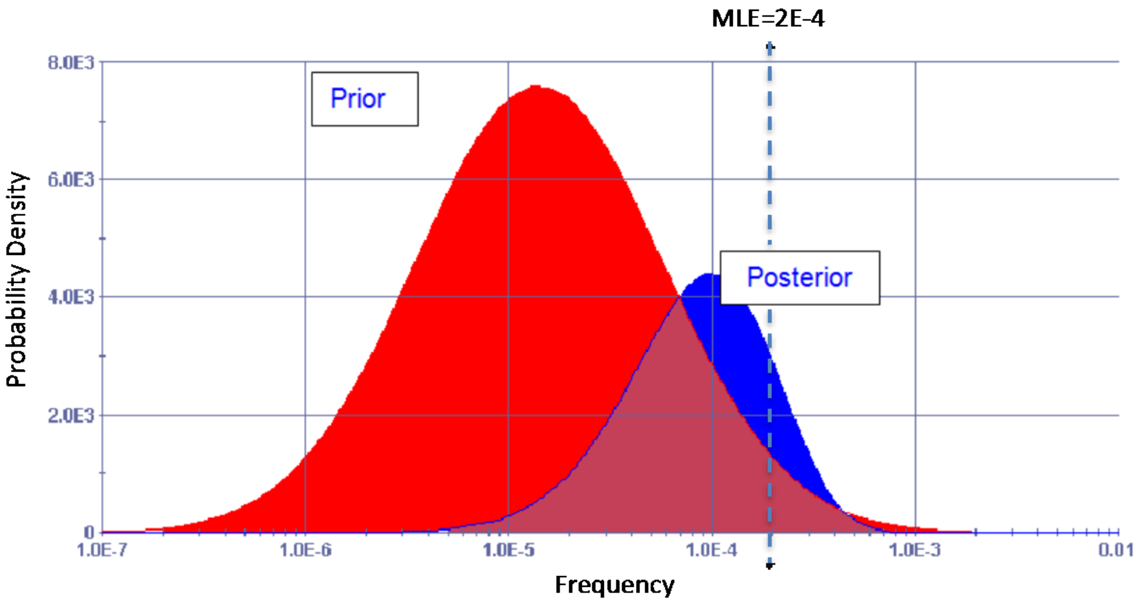

Figure F-1. The Prior Distribution Distributions for the Failure Rate Example

F.1.2 Updating Distribution of Failure-on-Demand Probability

In this example, the goal is to estimate a failure-on-demand probability. We have chosen the prior distribution of a particular component failure probability on demand to be a beta distribution with Mean = 1E-4 failures per demand, and Standard Deviation = 7E-5. The operational data for the component category are 1 failure in 2,000 demands. Our chosen likelihood model is the Binomial distribution, which is conjugate to the Beta prior. Therefore, the posterior distribution is also a Beta

[8] These distributions are discussed in Section 2.2.1.6.2.

distribution. The prior and posterior distributions are shown in Figure F-2, along with the MLE (1/2000=5E-4).

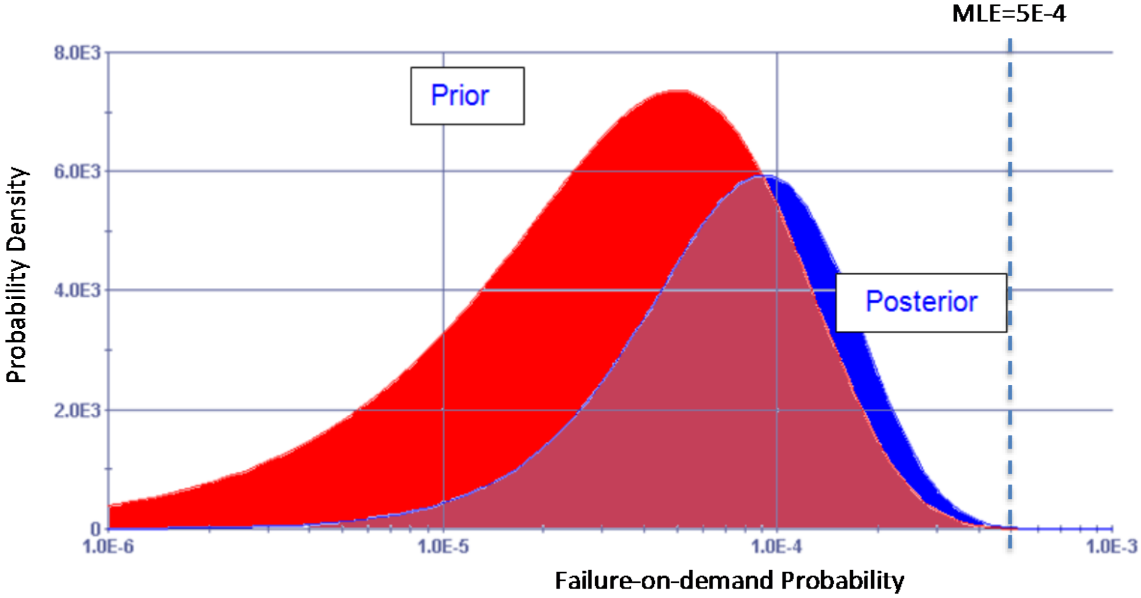

Figure F-2. The Prior and Posterior Distributions for the Failure-on Demand Example

Once again, we see the posterior distribution shifted towards the MLE. In this example, however, we also see something else: the MLE is quite unexpected, given the prior. This situation should give us pause, and is discussed below in Section F.2 of this Appendix.

F.1.3 Sequential Updating

Bayes' Theorem provides a mechanism for updating the state of knowledge when the information is accumulated in pieces. The updating process can be performed sequentially and in stages corresponding to the stages in which various pieces of information become available. If the total amount of information is equivalent to the "sum" of the pieces, then the end result (posterior distribution) is the same regardless of whether it has been obtained in stages (by applying Bayes' Theorem in steps) or in one step (by applying Bayes' Theorem to all the evidence at once).

Example: Updating Failure Rate for a Poisson Process

A component is tested for 1000 hours in one test and 4000 hours in another. During the first test the component does not fail, while in the second test one failure is observed. We are interested in an updated estimate of the component failure rate assuming a gamma prior distribution with parameters α= 1, β = 500.

Approach 1: Sequential. We first start with prior (Gamma distribution): $\Gamma(x|\alpha=1, \beta=500)$. We also use Poisson as the likelihood function: $Pr(k_1=0|T_1=1000, \lambda)$, representing the first data set (k_1 = 0 in T_1 = 1000 hours). The parameters of the resulting Gamma posterior distribution are $\alpha'=\alpha+k_1=1+0=1$, and $\beta'=\beta+T_1=500+1000=1500$ (refer to Section 2.2.6 for a general discussion of the update process).

Next, we use this posterior as the prior distribution in for a new update, using the second data set. The prior is $\Gamma'(l|\alpha'=1, \beta'=1500)$ and the likelihood is again Poisson: $Pr(k_2=1|T_2=4000, l)$. The parameters of the posterior after the second update are $\alpha''=\alpha'+k2=1+1=2$, and

$\beta''=\beta'+T_2=1500+4000=5500$. The posterior mean is given by

$$\bar{\lambda} = \frac{\alpha''}{\beta''} = \frac{2}{5500} = 3.6E - 4 \text{ failures/hour}$$

Approach 2: Use all the data at once. The total evidence on the failure history of the component in question is $k=k_1+k_2=0+1=1$, and $T=T_1 +T_2 =1000+4000=5000$. Starting with our prior distribution with parameters $\alpha = 1$, $b = 500$, the above cumulative evidence can be used in one application of Bayes' Theorem with Poisson likelihood: $\Pr(k=1|T_2 =5000, l)$. The parameters of the resulting Gamma posterior distribution are $\alpha'=\alpha+k=1+1=2$, $\beta'=\beta+T=500+5000$, and

$$\bar{\lambda} = \frac{\alpha'}{\beta'} = \frac{2}{5500} = 3.6E - 4 \text{ failures / hour,}$$

as for Approach 1. In this case, the equivalence of the two approaches is clear from the functional form: the numerator of the mean is given by the sum of the prior α and the total failures, and similarly for the denominator. But the principle holds generally. Note that the validity of the result depends on the trials all being exchangeable; this point is discussed below in Section 2 of this appendix.

F.2 Prior Distributions, Likelihood Models, and Data Applicability

It is difficult to avoid making choices in the assessment of uncertainty. This subsection discusses the kinds of things that need to be considered when those choices are made, taking as a point of departure the situation noted in the demand failure probability update illustrated above.

F.2.1 All Prior Distributions Contain Information

The amount of information that a prior distribution contains can be quantified (in various ways), and the distributions that contain the least amount of information can be found and used if so desired; but there is no such thing as a prior that contains no information. This is not necessarily a bad thing: none of the questions to which Bayesian analysis is applied, even the simplest, can be answered without some information. Classical methods avoid this need by answering different questions, which may or may not be similar enough to the questions we ask for the answers to be useful to us.

Take the simple example of determining whether a coin is fair or two-headed. If we flip a coin and it comes up heads six times in a row, we have collected some evidence against the coin being fair. But we must use our prior belief about the coin to reach a conclusion about whether we believe this coin is fair. If we just obtained the coin at the bank, where it is extremely unlikely to find a forged or misprinted coin mixed in amongst thousands of real coins, we should not be in a hurry to assume the coin is two-headed even after ten consecutive flips. On the other hand, if we found the coin on the floor of a magic shop, we should seriously consider the possibility that someone dropped a joke coin. Without prior probabilities we can't answer the question "how likely is this to be a fair coin?" All we can say is that it is sixty-four times easier for a two-headed coin to generate the data we collected than for a fair coin to generate it. A classical statistician would say "there is less than a 5% chance that a fair coin will come up heads 6 times in a row," but he will say nothing about the chance that we are holding a two-headed coin in our hands now. A blind devotee of non-informative priors will assign a prior probability 1/2 to the fair and two-headed possibilities, calculate posterior probabilities of 1/65 for fair, 64/65 for two-headed -- and probably make a lot of false accusations of two-headed coins. A reasonable Bayesian will assign a prior probability somewhere between 1/1000 and 1/1000000 to the two-headed possibility if he is at the bank (and not suspicious at all of a coin that comes up heads six times), and something closer to 1/100 if he is at a magic shop (and start seriously considering the possibility of a two-headed coin after several heads in a row come up.)

Proper choices of model form and prior should encapsulate all available information that we had about a problem before we started collecting data. This is an easy task if we have a small set of alternatives and easily quantifiable information, but it can be a very hard task for real-world problems -- often impossible to do perfectly. The better of a job we do at choosing the prior, the better our final answer will be.

The importance of choosing the correct model -- choosing the correct family of distributions to try to fit one's data to -- is often downplayed, but the success of the whole model-fitting enterprise depends on the reasonableness of this model. This is also a convenient and effective way to encapsulate information about the allowable range of the data.

For example, consider the distribution of times between failures of some component (or times between eruptions of a volcano, or some similar problem). If failure appears to be a completely random process, perhaps controlled by some external process, then a constant hazard rate, and exponential distribution of failure times, may be appropriate. If components accumulate damage through use, or pressure builds up during the interval between eruptions, then the hazard function increases with time, and the underlying distribution has a thinner-than-exponential tail. The Weibull distribution is popularly used to model component lifespans because it has a polynomial hazard function, convenient to model rapidly increasing risk of failure as the component exceeds its design lifespan. An nth degree polynomial hazard function corresponds to a distribution with tail thickness proportional to e^-(x^(n+1)).

On the other hand, consider the length of time a car sits in a parking lot. One's natural reaction to seeing the car sit there for a full day is not "surely the owner is going to be back any second now!" but rather "while originally I thought that car would only be there for a few minutes or hours, I should now entertain the possibility that it will be left here all week or all month" -- a situation modeled by a distribution with decreasing hazard function and a very long (decaying more slowly than exponential) tail, such as a lognormal or power-law distribution.

If the variable of interest takes on values only in a certain range, then it usually makes good sense to choose a prior that covers only that same range. When the data are confined to [0,1], the Beta prior is a common choice; when the data are confined to positive numbers, the gamma, Weibull, and lognormal are common choices. If there is a theoretical reason to expect the data to take a certain form, the model should usually be chosen to match that form: if you are trying to estimate the size of a 100-year flood from annual peak discharges of a river, fitting a Gumbel distribution to the data is probably a better choice than fitting some general-purpose distribution.

Modelers should be wary of choosing a prior that is at odds with the real world. In particular, remember that the normal distribution has support from -∞ to +∞, and if one models lengths or times or some other variable which is nonnegative by a normal distribution, the resulting posterior will always assign a nonzero probability to negative values (maybe very small, but still nonzero). The normal distribution is also very thin-tailed: it is extremely hard to make the posterior mean be more than a few standard deviations from the prior mean. This is a feature for some applications like modeling measurement errors that are known to be small compared to the quantity being measured; but this can be a significant flaw if a diffuse prior is desired.

A similar problem arises if one uses a lognormal distribution to model the probability of a rare event, as is commonly done in human reliability analysis. The lognormal is an excellent choice when a prior distribution spanning several orders of magnitude is needed (as when you don't know whether a rare event has probability 0.001 or 0.00001) but the lognormal has support on [0, ∞], not just [0,1], so care must be taken to handle the case when the posterior has substantial mass beyond 1. Truncation (treating all mass beyond 1 as if it were concentrated at 1) is only reasonable if very little mass is beyond 1.

F.2.2 How much information does a prior contain?

Sometimes it is easy to directly interpret how much information the prior contains relative to the data set. The simplest example is the beta-binomial model: a Beta(a,b) prior can be interpreted as providing "the same amount of information as if we had obtained a successes and b failures already." The same kind of interpretation can be applied to a gamma-Poisson model for failure rates.

A popular intuitive assessment of how much information the prior contains relative to the data is obtained by inspecting the posterior mean: the posterior mean always lies between the prior mean and the mean of the data. For many Bayesian models, the posterior mean can be thought of as a weighted average of these two means. If the posterior mean is, say, 3/4 of the way from the prior mean to the mean of the data, one interprets the data as containing three times as much information as the prior contained. The example of Section 1.2 above is a case where the prior contained more information.

For several of the most popular Bayesian models that use conjugate priors, including the beta-binomial, gamma-poisson, and normal-normal, the intuitive interpretations in the two paragraphs coincide with each other and can be made rigorous.

Even so-called "non-informative" priors contain information. Three popular families of priors -- maximum entropy priors, Jeffreys priors, and reference priors -- seek to minimize the information content of the prior, for three different technical definitions of 'information'. The first maximizes the information-theoretic entropy of the prior, subject to some given constraints; the second creates a distribution the shape of which is invariance under any change of variables; the last maximizes the expected Kullback-Leibler divergence ("information gain") between prior and posterior, given some assumptions about the posterior. Jeffreys and Reference priors coincide in 1 dimension but differ in multi-dimensional problems.

Revisiting our coin that came up heads six times in a row at the beginning of the chapter, suppose we start with a Beta(10,10) prior -- a fairly strong belief the coin is approximately fair. Our posterior is a Beta(16,10) distribution. Our prior has a mean 1/2, our data have mean 1, our posterior has mean 16/26 ~ 0.615. The posterior mean moved 23% of the way -- 6/26ths -- from 1/2 toward 1, in agreement with the intuitive assessment that our prior contained 10+10=20 coinflips worth of information, to which we add 6 more with our new data.

Had we done the same experiment with a Jeffreys or reference prior, Beta(1/2,1/2), our posterior would be a Beta(6 1/2, 1/2) distribution with mean 6.5/7 ~ 0.929. The posterior mean moved 6/7ths of the way toward the mean of the data, in agreement with the interpretation that a Jeffreys prior provided "as much information as 1 coinflip, 0.5 heads and 0.5 tails."

The situation one might intuitively think of as a "no-information prior" -- Beta(0,0) -- is in fact a very strange improper prior, with all its mass concentrated at p=0 and p=1.

F.2.3 Bias

Maximum likelihood estimates (MLEs) are the "gold standard" of classical statistical estimation because of their desirable properties. Chief among these are that MLEs are asymptotically unbiased and efficient (no other asymptotically unbiased estimate has a smaller variance). Note that being *asymptotically* unbiased does *not* guarantee that an MLE based on finite sample size is unbiased. Some MLEs (like $\bar{x}/n$ for the mean of a normal distribution) are always unbiased, but others -- like the MLE for the variance of a normal distribution, $\sum(x_i - \bar{x})^2/n$, are biased. In this last case it is easy to compute a bias correction -- this is why the usual formula for sample variance is $\sum(x_i - \bar{x})^2/(n-1)$ -- but in many other cases it is not a trivial task to remove the bias. When computing maximum likelihood estimates for complicated models using small data sets, the bias problem may be severe.

MLEs are a special case of a Bayesian point estimate, with a uniform (possibly improper) prior. Bayesian point estimates are almost always biased as a result of the choice of prior: for a binomial distribution with a uniform (Beta(1,1)) or Jeffreys (Beta(1/2,1/2)) prior, observing x successes in n trials results in a point estimate of *(x+1)/(n+2)* (uniform) or *(x+1/2)/(n+1)* (Jeffreys), in contrast to the unbiased *x/n*. Generally speaking, the more information in the prior, the more strongly the Bayesian estimate is biased. In a well understood problem, this may be considered a feature, not a flaw: when we have strong prior knowledge we may want our posterior estimate to be only slightly different, and even without strong prior knowledge, we may want to prevent the estimate of a binomial probability from being unreasonably close to 0 or 1, for instance.

It is important to remember that a Bayesian update never "fails:" it always returns an answer. If you ask a question about which you have collected little or no data, the answer it gives is driven entirely by the prior. Especially in cases where it is not obvious how much information the prior contains, or an experimenter uses a standard non-informative prior without thinking about how that will affect his answer, this can lead to surprisingly bad, or at least unexpected, answers.

This is simply a limitation of having sparse data. Careful choice of prior can mitigate this issue but not avoid it entirely. Consider a rare type of accident that is only expected to occur once in 1000 site-years of exposure. No one site is going to have sufficient local experience to independently estimate its accident rate; instead, each site is going to use the nationwide average rate as a prior for a gamma-Poisson model, and update it with its local experience. How strong of a prior should each plant use?

Suppose one takes a very strong prior like Gamma(10,10000). A plant that has no accidents in 10 years will update this to a Gamma(10,10010) posterior. A plant with two accidents in 10 years -- wildly unlikely, if that site's true accident rate is close to once in 1000 years of exposure -- updates this to a Gamma(12,10010) prior, and claims that its site-specific accident rate is around once in 800 years. *Using too strong of a prior distribution means that grossly underperforming sites are not called to account for their poor performance.*

Now suppose one takes a very weak prior like Gamma(0.01,10). Now the plant with 2 accidents in 10 years has a Gamma(2.01,30) posterior, estimates a site-specific accident rate of once in 15 years, and is forced to take corrective action. But a site with no accidents in ten years has a Gamma(0.01,20) posterior, and, on the basis of only ten years of experience containing almost no real information -- we EXPECT not to see a once-in-1000-years accident in any given 10 years period -- now claims its site-specific accident rate has improved to once in 2000 years.

There is *no* prior that can completely avoid both of these two flaws. Any scheme that ensures poorly performing sites are "properly punished" will also "improperly reward" well-performing (or just lucky) sites.

The non-informative priors for the gamma-Poisson model have shape parameters near 0.5 -- the intuitive interpretation is "pretend 0.5 accidents happened at each site, in addition to however many were really observed" -- as a compromise so that sites with 1 or more accidents see some kind of significant increase in site-specific estimated rate, while sites with 0 accidents don't calculate impossibly rare site-specific rates.

F.2.4 Bayesian analysis assumes a static underlying process

Bayesian modeling is rooted in the notion that the observed data are *exchangeable*. Many classical methods are based on the similar but stronger idea that all of the observations are *independently and identically distributed*. This constitutes an assumption that the order in which the data were collected does not matter. If one flips a coin 10 times today, and flips the same coin 10 more times tomorrow, these can be pooled into a set of 20 equally-important observations.

This assumption breaks down if the underlying process has changed over the observation period. Estimating the value of real estate based on last year's (or last decade's) sales prices gives poor results if economic conditions have changed. Similarly, using the failure rate of brand-new pumps to estimate the failure rate of broken-in and well-maintained pumps, or using well-maintained pumps to estimate the failure rate of worn-out pumps, has that same issue.

When data are sparse, the temptation to pool data over an unreasonably long time period is strong. Sometimes it is justified: if one is average across thousands of pumps nationwide, perhaps it is fair to assume that pumps are constantly wearing out and being replaced, such that the overall distribution of pump ages remains static, even though any one individual pump's behavior may be different next year than last year. This is a difficult assumption to defend. Conditions nationwide may change -- in an economic downturn, facilities across the nation may defer maintenance, or a new law may be passed mandating replacement at a certain age -- or maybe a large proportion of units entered service at the same time: look at what happened to Social Security when it assumed the ratio of earners to retirees would stay approximately constant forever.

For convenience, we often use models that we know are an over-simplification of the real world. Using several years of old data to create a prior distribution for what we expect to see next year is a very common practice. It is important, when doing so, not to just directly use the distribution of past observations as one's prior, but rather to use a more diffuse prior that takes into account the possibility that conditions are the same now as they were in the past.

F.2.5 Assessing goodness of fit

Assessing whether new data are consistent with a proposed model is an important task, but it is not a task for which a single universal method exists. One (extreme) perspective is that if the prior has properly encapsulated everything we know, the posterior should always be correct: that is, one of things the posterior tells us is exactly how much we should change our belief after collecting surprising data. In principle this is true, but in practice, people commonly use less-than-perfect priors, either for mathematical convenience or due to lack of information that would have been desirable while choosing a prior.

Classical tests exist for determining whether a data set appears to have been drawn from a particular distribution, and for determining whether two data sets appear to be drawn from the same distribution or not. These tests, or their Bayesian adaptations, may be suitable for answering some questions of this type.

One particularly valuable classical test consists of fitting two models to the same data, with one model (the "reduced model") a special case of the other (the "full model"). If the larger family of distributions fits the data significantly better than the smaller, embedded, family of distributions does, this is evidence that the reduced model is inadequate for the task at hand.

This is typically done when one wishes to argue for the more complicated model. When a study reports that it has found that family income has a significant effect on academic success "after controlling for gender and race", it means that it fit a model that explains success by income, gender, and race, and shown that that model is significantly better than a model that explains success only by gender and race.

One might, for instance, assess whether a linear trend is a good fit to a scatter plot, by fitting a quadratic or cubic model to the same data set, and conclude that if the quadratic term of the larger model is statistically significant, then the simple linear model is a poor explanation of the data. Note that a non-significant result does not *prove* the simple model is correct, but it is evidence in favor of that claim.

This type of test can be adapted to almost any problem of interest. The question of data seemingly inconsistent with a prior might be approached in this way by, for instance, fitting both a simple Poisson model with Gamma prior to a set of count data, and an over dispersed Poisson model. If the later model fits much better than the former, one has a basis for arguing that there is something wrong with the first model: either you needed a more complicated model all along, or the prior and the data were not consistent, or something else.

F.2.6 Surprise

An alternative to formally testing goodness of fit (or lack thereof) is assessing whether the data are "surprising," without considering any particular alternative. This is useful, as a sanity check and to get a feel for one's data; but developing a firm rule for how surprising data must be before saying "our model is wrong" is not possible without bringing in some outside information (such as showing that another model fits the data better.)

Various people have proposed formal definitions of the notion of 'surprise'. No one definition has achieved universal acceptance. Bayarri and Berger [F-1] review the options that have been used in the past. The classical p-value has sometimes been interpreted as a measure of surprise [F-2]. Two more recent alternatives are the "s-value" [F-3] and the Kullback-Leibler divergence from the prior to the posterior, which is being vigorously promoted as a "formal Bayesian theory of surprise" [F-4].

This last proposal, grounded in the same mathematics that underlies the reference prior, may be the most likely of these to stand up to the test of time, though the emotionally charged notion of "surprise" is not likely to remain attached to it. The Kullback-Leibler divergence is more often described in drier terms like "bits of information gained" (in formal Shannon-information sense, not the informal "information contained in n observations" terms used earlier in this appendix).

Returning one last time to our coin-flip example, suppose we have a Beta(10,10) prior, and we flip a coin six times. We would be not surprised at all to see 3 heads and 3 tails, or 4-2 or 2-4; we might be mildly surprised to see 6 heads in a row. If we flipped the coin 10 times, we would not be surprised by anything between say 8-2 and 2-8; 9 heads out of 10 would be about as surprising as 6 out of 6; 10 out of 10, more surprising still -- about one bit of information more surprising, seeing something that was supposed to be a 50-50 proposition happen an additional time.

If we calculate the Kullback-Leibler divergence between prior *p(x)* and posterior *q(x)*

$$\int q(x)\ \log_2 \frac{q(x)}{p(x)} dx$$

we see that the divergence between Beta(10,10) and Beta(13,13) is 0.024 bits -- almost no surprise at all; between Beta(10,10) and Beta(14,12), 0.112 bits; between Beta(10,10) and Beta(15,11), 0.379 bits; and between Beta(10,10) and Beta(16,10),0.837 bits. The K-L divergence between Beta(10,10) and Beta(15,15) is 0.054 bits; between Beta(10,10) and Beta(18, 12), 0.654 bits; between Beta(10,10) and Beta(19,11), 1.14 bits; and between Beta(10,10) and Beta(20,10), 1.79 bits.

The question "how surprising is surprising enough to cause us to doubt that we chose an appropriate prior?" still lacks a rigorous answer.

F.3 References

F-1 M. J. Bayarri and James O. Berger, "Measures of Surprise in Bayesian Analysis," Institute of Statistics and Decision Sciences, Durham, North Carolina, 1998

F-2 A. Reinhart, "Statistics Done Wrong: the Woefully Complete Guide," No Starch Press, March 2015

F-3 J. V. Howard, "Significance Testing with No Alternative Hypothesis: A Measure of Surprise," Springer Science+Business, 2009

F-4 L. Itti and P. F. Baldi, Bayesian Surprise Attracts Human Attention," Advances in Neural Information Processing Systems, Vol. 19," Cambridge MA: MIT Press, 2006

Appendix G – Population Variability Modeling

To Be Added

Appendix H – Expert Elicitation

To Be Added

Appendix I - Failure Space Based Importance Measures

It will be convenient in the following to refer to a formula for the risk metric (e.g., top event probability) with respect to which the measures are being calculated:

$$R = \mathrm{f}(\mathrm{x}_1,\mathrm{x}_2 ,\ldots, x_i ,x_j ,\ldots, \mathrm{x}_\mathrm{n})$$

where x_k is the k^th basic event, having probability p_k , and R_0 means "R evaluated with all p's set to their nominal values." It is an aid to understanding the following formulas to bear in mind that the reduced Boolean expression for the minimal cut sets maps simply into an arithmetic expression for the rare event approximation to top event probability, and the same idea applies to any subset of the minimal cut sets. The notation "|" means "given:" A|B means "A given B."

I.1 Fussell-Vesely and Risk Reduction Worth Importance Measures

The F-V importance measure is used to determine the importance of individual minimal cut sets containing basic event x_i to the risk. F-V of event x_i is given by:

$$I_{x_i}^{FV} = \frac{Pr\left(\bigcup_j MCS_j^{x_i}\right)}{Pr\left(\bigcup_j MCS_j\right)} = \frac{Pr\left(\bigcup_j MCS_j\right)}{R_0}$$

where: I^{FV} is the F-V importance for event x_i,

$Pr(\bigcup_j MCS_j^{x_i})$ is probability of the union of the minimal cut sets containing event x_i ;

$Pr\left(\bigcup_j MCS_j\right) = R_0$ (the probability of the union of ALL of the minimal cut sets) is the baseline risk.

The simple interpretation of the FV is that it is the fraction of total risk involving x_i. Corollary interpretations are (1) that the FV is the conditional probability that at least one minimal cut set containing event x_i will occur, given that the system has failed, or (2) the fraction by which risk would decrease if $Pr(x_i)$ were reduced to zero. The latter interpretation points to another way of calculating FV:

$$I_{x_i}^{FV} = \frac{R_0 - R|Pr(x_i) = 0}{R_0}$$

where $R|Pr(x_i) = 0$ is the value of the risk metric when the probability of event x_i is set to zero. In this calculation, in the numerator, we are subtracting off the contribution from minimal cut sets that do NOT contain x_i, leaving the minimal cut sets that DO contain x_i.

The closely-related Risk Reduction Worth (RRW) is a measure of the change in risk when a basic event probability (e.g., unavailability of a hardware device) is set to zero. It measures the amount by which risk would decrease if the event would never occur. The RRW measure is calculated as the ratio[9] of the baseline expected risk to the conditional expected risk when the probability of event x_i is set to zero (assuming that the hardware device is "perfect"):

$$I_{x_i}^{RRW} = \frac{R_0}{R|Pr(x_i) = 0}$$

where $I_{x_i}^{RRW}$ is the risk reduction worth for event x_i .

It should be clear that FV and RRW will produce essentially the same ranking: event lists ordered by decreasing FV and decreasing RRW are the same. In fact, it is straightforward to show that

$$I_{x_i}^{FV} = 1 - \frac{1}{I_{x_i}^{RRW}}.$$

I.2 Birnbaum (B) and Risk Achievement Worth (RAW)

The B is the rate of change of the expected risk as a result of the change in the probability of an individual event. Mathematically, the B importance of event x_i is

$$I_{x_i}^{B} = \frac{\partial R}{\partial x_i}.$$

In many cases, B can be calculated as:

$$I_{x_i}^{B} = (R|Pr(x_i) = 1) - (R|Pr(x_i) = 0)$$

where $R|\Pr(x_i) = 1$ (0) is the risk metric calculated with Pr(x_i) set to 1 (0).

In general, the B of a basic event x_i does not depend on the probability of x_i; it depends on the probabilities of the other basic events in the cut sets in which x_i appears.

Risk Achievement Worth (RAW) is a measure of the change in risk when the probability of a basic event (e.g. unavailability of a component) is set to unity. Analogously to RRW, the calculation is typically done as a ratio:

$$I_{x_i}^{RAW} = \frac{R|Pr(x_i) = 1}{R_0}.$$

Again analogously to RRW, some PRA codes calculate an interval measure corresponding to RAW, the "Risk Increase Interval," which is the *difference* between the conditional expected risk when event x_i is set to unity, and the baseline risk.

[9] Instead of ratio, some PRA codes calculate "Risk Decrease Interval," which is the *difference* between baseline risk and the conditional risk when event x_i is set to zero.

Both RAW and RRW correspond to drastic sensitivity studies, displaying how much difference it makes when a basic event probability is maximized (RAW) or minimized (RRW). This kind of information points to properties of the model, and perhaps the system; for example, a high RAW can result from a component for which there is relatively little backup, such as a single item that is required to succeed regardless of whether anything else succeeds or fails. But as discussed in [I-1], measures such as RAW are difficult to use in quantitative reasoning processes.

I.3 Computing B, FV, RAW, RRW

It is straightforward to compute FV and RRW within the rare event approximation, given a Boolean expression for the top event in properly reduced form. If the basic event names are replaced by their probability values, AND by multiplication symbols, and OR by addition symbols, one has an expression for top event probability (again, within the rare event approximation). If a more precise answer is required, better approximations can be applied (such as the min cut upper bound).

Strictly speaking, evaluating RAW calls for actually restructuring the expression. Computing the RAW of a basic event calls for setting that event to "TRUE" (typically, the corresponding component to "failed" or perhaps "unavailable") and re-reducing the top event expression. Consider computing the numerator of the RAW of event A in an expression including

$$A*B*C + X*B*C + \dots .$$

If we simply set A to a value of 1, we will still include the contribution of X*B*C, which, strictly speaking, we should not. Setting A to "TRUE" and re-reducing leaves us with

B*C + … , the "X*B*C" having been absorbed.

However, computing B(A) gives us

R(A=1)-R(A=0) =[B*C + X*B*C + …] – [X*B*C + …] = B*C (plus perhaps other terms) .

I.4 Differential Importance Measure for Basic Events and Parameters

The importance measures discussed previously are defined to deal with basic event probabilities *one event at a time*, and, as formulated, they do not reflect the influence of the underlying parameters in the models of event probability: they do not measure the importance of changes that affect component properties or failure modes. They also lack an additive property that some workers consider desirable. For these reasons, the "differential importance measure, DIM, was introduced.

I.4.1 Definition of DIM

Let R be the risk metric of interest expressed as a function of basic events or fundamental parameters of the PRA model as shown below:

R= f(x_1,x_2 ,..., x_i ,x_j ,..., x_n) where x_i is the generic parameter such as basic event probability of a component x_i or the failure rate of a component x_i .

The differential importance measure of x_i is defined as

$$I_{x_i}^{DIM} = \frac{dR_{x_i}}{dR} = \frac{\frac{\partial R}{\partial x_i} dx_i}{\sum_j \frac{\partial R}{\partial x_j} dx_j}$$

DIM reflects the fraction of the total change in R due to a change in parameter x_i. It can be shown that DIM is additive, that is,

$$I_{x_i \cup x_j \ldots \cup x_k}^{DIM} = I_{x_i}^{DIM} + I_{x_j}^{DIM} + \cdots + I_{x_k}^{DIM} \ .$$

I.4.2 Calculations of DIM

With respect to calculation of DIM for a parameter of the PRA model, there are two computational inconveniences:

1. The DIM can be calculated only if the expression for the risk is in parametric form, which is not a standard output form generated by the PRA codes.

2. There is no available computer program for use.

However, one can compute DIM for basic events using the F-V and RAW importance measures. The latter measures are often generated by standard PRA codes by applying formulas developed in the previous subsection.

As noted, calculation of DIM deals with change in R (its differential). Since the change depends on how the values assigned to a parameters are varied, DIM can be calculated in different ways. Two possibilities are:

1. Assume a uniform change for all parameters (i.e., $\delta x_i = \delta x_j = \delta x_k \ldots$). Under this operation, parameters are ranked according to the effect they produce on R when they undergo small changes that are the same for all. This has meaning when parameters of the model have the same dimensions (e.g., the risk metric is expressed in terms of basic event probabilities only). DIM for parameter x_i is calculated as follows:

$$I_{x_i}^{DIM} = \frac{\frac{\partial R}{\partial x_i} dx_i}{\sum_j \frac{\partial R}{\partial x_j} dx_j} = \frac{\frac{\partial R}{\partial x_i}}{\sum_j \frac{\partial R}{\partial x_j}}$$

2. Assume a uniform *percentage* change for all parameters ($\frac{\delta x_i}{x_i} = \frac{\delta x_j}{x_j} = \frac{\delta x_k}{x_k} \ldots$). Under this operation, PRA parameters are ranked according to the effect they produce on R when they are changed by the same fraction from their nominal values. This ranking scheme, which is applicable to all analysis conditions, can be calculated from:

$$I_{x_i}^{DIM} = \frac{\frac{\partial R}{\partial x_i} dx_i}{\sum_j \frac{\partial R}{\partial x_j} dx_j} = \frac{\frac{\partial R}{\partial x_i} \frac{dx_i}{x_i} x_i}{\sum_j \frac{\partial R}{\partial x_j} \frac{dx_j}{x_j} x_j} = \frac{\frac{\partial R}{\partial x_i} x_i}{\sum_j \frac{\partial R}{\partial x_j} x_j}$$

The relation between DIM and F-V, RAW, and BM are shown in the Table below. These relationships hold only when the risk metric is (1) linear, and (2) expressed in terms of basic events only.

Table I-1. Relation between DIM and F-V, RAW, and Birnbaum Importance Measures.

	Relation between DIM and …		
	F-V	**RAW**	**B**
Constant Increment: $I_{x_i}^{DIM} =$	$\frac{\frac{I_{x_i}^{F-V}}{Pr(x_i)}}{\sum_k \frac{I_{x_k}^{F-V}}{Pr(x_k)}}$	$\frac{\frac{I_{x_i}^{RAW} - 1}{1 - Pr(x_i)}}{\sum_k \frac{I_{x_k}^{RAW} - 1}{1 - Pr(x_k)}}$	$\frac{I_{x_i}^{B}}{\sum_k I_{x_k}^{B}}$
Constant Percentage Increment: $I_{x_i}^{DIM} =$	$\frac{I_{x_i}^{F-V}}{\sum_k I_{x_k}^{F-V}}$	$\frac{\frac{I_{x_i}^{RAW} - 1}{\frac{1}{Pr(x_i)} - 1}}{\sum_k \frac{I_{x_k}^{RAW} - 1}{\frac{1}{Pr(x_k)} - 1}}$	$\frac{I_{x_i}^{B} Pr(x_i)}{\sum_k I_{x_k}^{B} Pr(x_k)}$

I.5 References

I-1 M. C. Cheok , G. W Parry, R.R. Sherry, "Use of Importance Measures in Risk Informed Regulatory Applications," Reliab Enggn Syst Safety, 1998

Appendix J - Prevention Worth

The measures of "basic event" importance summarized above are not really about just the basic events; each basic event measure reflects properties of the union of cut sets that contain the designated basic event. So the "importance" of a basic event is really an attribute of the cut sets (the failure scenarios) in which that event appears. All of the measures introduced above are couched in failure space: they reflect contributions to risk, or the sensitivity of risk to changes in basic event characteristics. Analogous measures can be defined in success space: we can examine the properties of the union of *path* sets containing a designated component. One such measure is "Prevention Worth," defined as

$$PW^i = P\left(\bigcup_j MPS_j^i\right),$$

in which PW stands for "Prevention Worth," i indexes basic events, and MPS_j^i are the minimal path sets containing basic event i. This is a bit like the numerator of the F-V measure, substituting path sets for cut sets; but since success path probabilities are generally of order unity, the rare event approximation cannot be used to calculate the probability associated with a union of success paths. However, in many cases, we can approximate the right-hand side by computing the probability of *failure* of that union of path sets, and subtracting it from unity:

$$PW^i \cong 1 - P\overline{\left(\bigcup_j MPS_j^i\right)}$$

Finally, for ease of interpretation, it is useful to introduce the "NINES" index, which, for a given PW, is calculated as

$$NINES(PW) = -\log(1 - PW).$$

This says how many "nines" of reliability are afforded by the union of path sets considered: for example, a reliability of 0.999 is said to provide three nines of reliability.

Table K-1 shows the results for the case of the simple problem used in Appendix L to introduce prevention analysis.

Table J-1. Comparison of PW with RAW and F-V

Importance Measure	**Element (From Figure L-1)**				
	N2	A	B	C	D

F-V	1.0	~1	0.001	0.001	0.001
RAW	10^4	10^3	1.1	1.1	1.1
NINES	4	2.9996	1.9586	1.9586	1.9586

For a given element, this measure reflects the safety significance of the success paths containing that element. Put another way: each element potentiates the success paths that contain it, and its Prevention Worth is measured by the worth of the totality of those success paths. N2 has the highest Prevention Worth: for the numbers assumed, the path that contains N2 is "worth" more than all of the other success paths put together. This is true because all of those other success paths contain A, whose failure probability (success probability) is greater (less) than that of N2.

Other insights are available from the table. N2 has a F-V of 1, independently of its failure probability, because it appears in every minimal cut set. But all of the other measures tabulated depend, to some extent, on the nominal failure probabilities assigned to the associated basic events.

Prevention Worth was formulated originally for the purpose of illustrating the benefits of thinking both in success space and in failure space, rather than focusing exclusively on failure space. However, although the measure arguably provides an interesting perspective on the role played by events in the model, as far as the authors are aware, no commercial PRA software computes Prevention Worth. Unfortunately, for realistic problems, the calculations are rather difficult; one needs first to parse out the success paths containing the element(s) of interest, and then (in order to use the above approximation to compute PW) evaluate the complement of that expression in order to approximate PW.

Appendix K - Top Event Prevention Analysis

Consider the problem of determining the allocation of resources to activities aimed at maintaining and verifying the performance and reliability of safety equipment (“special treatment,” as it is called in the nuclear industry). This is important both to facility operators and to their regulators. To see why basic event importance measures are not necessarily a reliable guide to solving this problem, consider the example presented in the figure below. The system shown is supposed to supply compressed nitrogen (or air) to another system downstream. In order to succeed, we need either to supply air from one of the compressors via the receivers and the air dryers shown in Figure L-1, or to supply compressed nitrogen from the tanks shown in the upper portion of the figure. A simplified fault tree is shown on the left, showing that the top event is an AND of the failure of these two options (“Air” and “N2”). For simplicity, the compressed-nitrogen option is modeled as a single event “N2.” All components in that leg are logically in series, so no information is lost by this, unless there is some linkage between components in that segment and components in functionally redundant segments. In a real system, this is a real possibility, but the present illustration does not require us to address it. Similarly, the Air Dryer segment is modeled as a unit, and each compressor-receiver pair is modeled as a unit. Again, shared dependency of the compressors (e.g., of power supply) is a real possibility, but the present illustration does not require us to address it.

There are two minimal cut sets of the fault tree shown: N2 * A and N2 * B * C * D. Notional basic event probabilities are assigned on the fault tree figure itself, and based on these, the F-V and RAW are tabulated below the system diagram. One sees that N2 and the Air Dryer have large values for both RAW and F-V, while the compressors do not. This is a result of the compressors being mutually redundant: if B fails, you still probably have C and D; if C fails, you still probably have B and D; and if D fails, you still probably have B and C. This is an example of the “portfolio” effect mentioned above. It would be inappropriate (but not unprecedented) to conclude from these F-V and RAW values that the compressors are not “important.” This example is simple enough to see through without much machinery, but not all applications have that property.

Instead of trying to determine “special treatment” from importance measures, consider a different approach, called “Top Event Prevention.” Within extant versions of this approach, one first formulates a prevention criterion to be satisfied by the *complement* of equipment to be considered “special.” A simple example is to require single-failure tolerance in the complement of credited equipment: require the function(s) to succeed despite any single failure. Next, one applies an algorithm to identify subsets of the equipment potentially available, each subset having the property of satisfying the prevention criterion. In the lower left portion of the figure, we see the mechanics and the results of applying the single-failure criterion to the problem given. Start with the minimal cut sets given in the lower left of the figure (under “Top Event”). Evidently, any subset satisfying the single-failure criterion must contain both N2 and A; if a subset contains only A, and not N2, then the single failure of A fails the function, and vice versa. The second cut set requires us to work out some combinations: any single-failure-tolerant subset of the elements in a cut set must contain at least two of the elements in each cut set, and the logic expression for the six possibilities for the second cut set is shown. Since we need to “prevent” all of the cut sets, in order to obtain the prevention sets for the system, we “AND” together the prevention sets for each minimal cut set, and reduce the resulting expression. The resulting “minimal prevention sets” (the sets of events that collectively satisfy the prevention criterion) are shown in the lower right. It is straightforward to verify by inspection that each prevention set satisfies the prevention criterion.

A noteworthy feature of these prevention sets is that they all contain at least one compressor, a result that the importance-measure-based heuristic does not achieve. In general, prevention analysis always yields solutions that comprise unions of complete success paths, a result that is not to be expected from importance-measure-based reasoning.

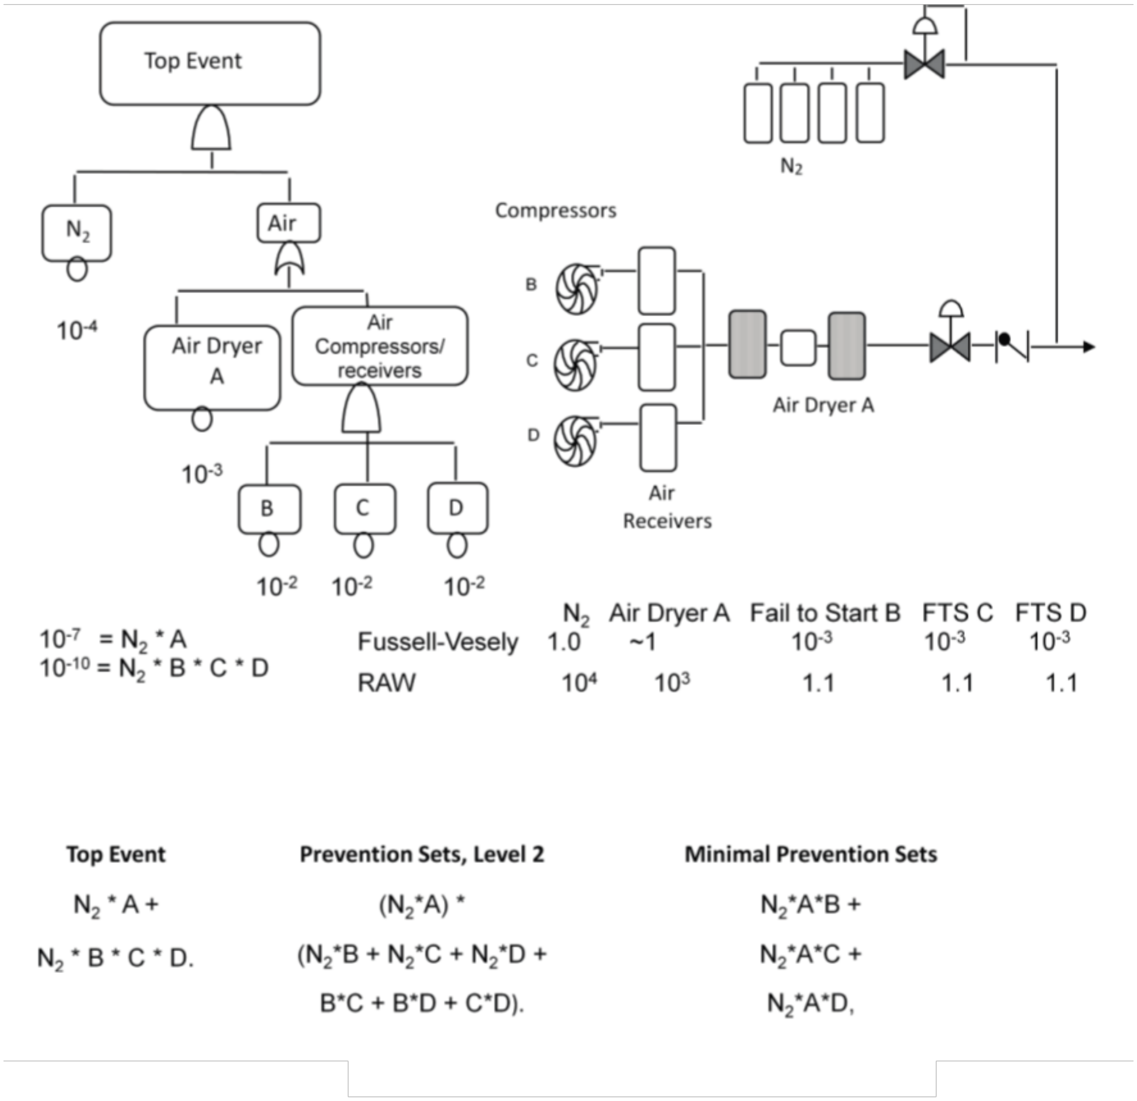

$10^{-7} = N_2 * A$
$10^{-10} = N_2 * B * C * D$

	N_2	Air Dryer A	Fail to Start B	FTS C	FTS D
Fussell-Vesely	1.0	~1	10^{-3}	10^{-3}	10^{-3}
RAW	10^4	10^3	1.1	1.1	1.1

Top Event	Prevention Sets, Level 2	Minimal Prevention Sets
$N_2 * A +$	$(N_2*A) *$	$N_2*A*B +$
$N_2 * B * C * D.$	$(N_2*B + N_2*C + N_2*D +$	$N_2*A*C +$
	$B*C + B*D + C*D).$	$N_2*A*D,$

Figure K- 1. Top Event Prevention (Simple Example) (After (K-1, K-2])

Each prevention set satisfies the prevention criterion, without credit for any other elements. As shown in the figure, the method has given us three options. In a real application, one could simply choose any one of the options, and assure that sufficient resources are allocated to every component in that set to achieve a good quantitative outcome (for example, test the active components at some regular interval).

The importance-measure-based heuristic does not, in general, point to unions of complete path sets. This is not to say that importance measures are “wrong;” they provide information about what the model is saying. But they do not answer questions that need to be addressed at the portfolio level.

It is straightforward to extend the calculations illustrated above to address prevention criteria that call for quantitative reliability estimates, rather than essentially barrier-counting, although that form of the algorithm is not a true global reliability optimizer. However, it illustrates the more general process of choosing not only what items of equipment (operator actions, instrumentation, …) need to be credited, but also what assumptions, initial conditions, and so on need to be assured (and perhaps monitored during the operational phase) in order to provide reasonable assurance of the claims presented in the claims tree of Figure 3-1. This iterative process of self-consistently determining this portfolio of items is illustrated in Figure L-2 below.

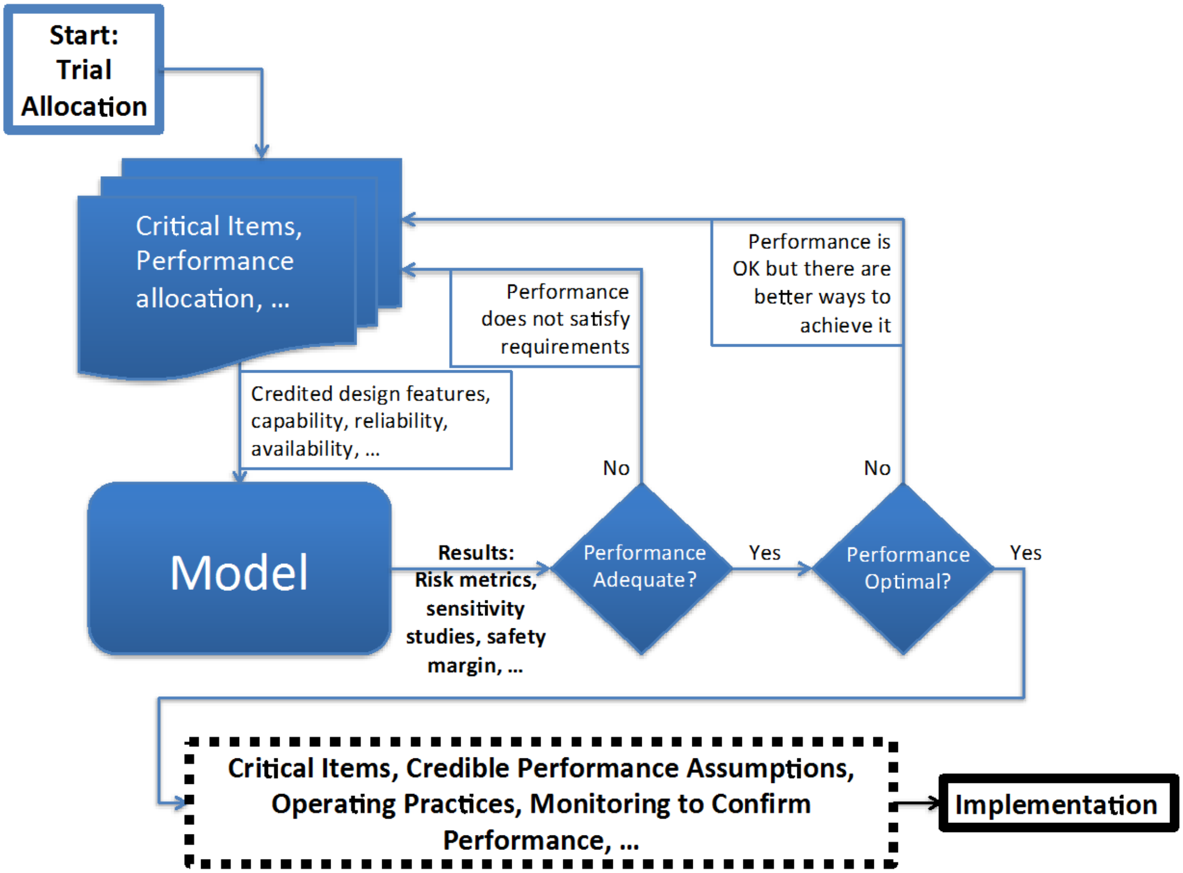
Start: Trial Allocation
Critical Items, Performance allocation, …
Performance is OK but there are better ways to achieve it
Performance does not satisfy requirements
Credited design features, capability, reliability, availability, …
No
No
Model
Results: Risk metrics, sensitivity studies, safety margin, …
Performance Adequate?
Yes
Performance Optimal?
Yes
Critical Items, Credible Performance Assumptions, Operating Practices, Monitoring to Confirm Performance, …
Implementation

Figure K- 2. Process for Confirming Overall Performance Based on Items Credited in the Assurance Case

K.1 References

K-1 W. Brinsfield, J. Voskuli, "Focusing the Scope of Fire PRA Human Reliability Analysis Using Top Event Prevention (TEP)," PSA 2015, 2015

K-2 D. P. Blanchard and R. W. Youngblood, "Risk-Informed Safety Margin Characterization Case Study: Use of Prevention Analysis in the Selection of Electrical Equipment to Be Subjected to Environmental Qualification," PSAM 12, 2014

DRAFT

Appendix L – Human Reliability

To Be Added