NASA-CR-197808

CONTINUATION OF SPACE SHUTTLE
PROBABILISTIC RISK ASSESSMENT, PHASE 3
SAIC DOCUMENT NO. SAICNY95-02-25

PROBABILISTIC RISK ASSESSMENT OF THE SPACE SHUTTLE

A STUDY OF THE POTENTIAL OF LOSING THE VEHICLE DURING NOMINAL OPERATION

VOLUME I : FINAL REPORT

PREPARED FOR

US NATIONAL AERONAUTICS AND SPACE ADMINISTRATION
HEADQUARTERS OFFICE OF SPACE FLIGHT (CODE M)
WASHINGTON, DC

BY

SCIENCE APPLICATIONS INTERNATIONAL CORPORATION
ADVANCED TECHNOLOGY DIVISION
NEW YORK, NY

28 FEBRUARY 1995

PRINCIPAL INVESTIGATOR:
JOSEPH R. FRAGOLA

CHIEF RISK ANALYST:
GASPARE MAGGIO

OTHER PRINCIPAL CONTRIBUTORS:
MICHAEL V. FRANK*
LUIS GEREZ+
RICHARD H. MCFADDEN
ERIN P. COLLINS
JORGE BALLESIO
PETER L. APPIGNANI
JAMES J. KARNS

* SAFETY FACTOR ASSOCIATES, INC.
ENCINITAS, CA
+EMPRESARIOS AGRUPADOS,
MADRID, SPAIN

Table of Contents

Volume I

1.0. Executive Summary
- 1.1. Background
- 1.2. Key Objectives
- 1.3. Key Results
  - 1.3.1. Estimated Risk of Loss of Vehicle
  - 1.3.2. Distribution of Total Estimated Risk Among Mission Phases
  - 1.3.3. Contributions to Total Estimated Risk of Shuttle Vehicle Elements
  - 1.3.4. Risk Drivers
- 1.4. Key Insights and Recommendations
  - 1.4.1. Comparison with the Results of Previous Space Shuttle Risk Assessments
  - 1.4.2. Risk Dominance of Propulsion Systems
  - 1.4.3. Risk of Mission Intervals
  - 1.4.4. Effectiveness of Space Shuttle Main Engine Protective "Redlines"
  - 1.4.5. Distribution of Risk Versus Critical Items Count
- 1.5. Ground Rules and Key Assumptions of the PRA
- 1.6. Analytical Approach
- 1.7. Sensitivity Analysis
- 1.8. Synopsis of Main Technical Report

2.0. Introduction
- 2.1. Background
  - 2.1.1. History of this PRA
  - 2.1.2. Project Organization
  - 2.1.3. Relationships with Previous Risk Assessments
- 2.2. Objectives and Principal Products
- 2.3. Ground Rules of the PRA
  - 2.3.1. Nominal Shuttle Mission
  - 2.3.2. Definition of Mission Phases
  - 2.3.3. Scope of Risk Assessment
  - 2.3.4. Assumed Vehicle Configuration
    - 2.3.4.1. Vehicle and Operations Modifications
    - 2.3.4.2. Vehicle Configuration Errors
    - 2.3.4.3. Orbiter-to-Orbiter Differences
  - 2.3.5. Mission Abort Scenarios
  - 2.3.6. Final Approach & Landing Risk
  - 2.3.7. Data Window
- 2.4. Overview of the Risk Analysis
  - 2.4.1. The Role of Scenarios (Event Sequences) in PRA
  - 2.4.2. Overview of the PRA Process

3.0. Risk Modeling
3.1. Overview of Modeling Approach
3.2. Integrated Mission Risk Model
3.2.1. Mission Master Logic Diagram
3.2.2. Space Shuttle Top-level Functional Failures
3.3. System Failure Models
3.4. Uncertainty

4.0. Data Analysis
4.1. Overview
4.2. Analytical Methods
4.2.1 Types Of Risk Data
4.2.2. Data Review and Encoding
4.2.3. Exposure Data Development
4.2.4. Failure Data Estimation

5.0. Systems Analysis
5.1. Front-Line Systems
5.1.1. Space Shuttle Main Engines and Main Propulsion System
5.1.2.1. Description
5.1.2.2. Success Criteria
5.1.2.3 Initiating Events
5.1.2.4. Accident Scenarios and Consequences
5.1.2.5. Data Analysis
5.1.2.6. SSME/MPS Risk Contribution
5.1.2. Integrated Solid Rocket Booster
5.1.2.1. Description
5.1.2.2. Success Criteria
5.1.2.3. Initiating Events
5.1.2.4. Fault Trees
5.1.2.5. Data Analysis
5.1.2.6. ISRB Initiator Risk Contribution
5.1.3. Orbiter Auxiliary Power Units
5.1.3.1. Description
5.1.3.1.1. Water Spray Boiler
5.1.3.1.2. Hydraulic System
5.1.3.2. Operation
5.1.3.2.1. Ascent
5.1.3.2.2. Orbital Operations
5.1.3.2.3. Deorbiting and Entry
5.1.3.3. Success Criteria
5.1.3.4. Initiating Events
5.1.3.5. Event Sequences
5.1.3.6. Data Analysis

5.1.3.7. APU Initiator Risk Contribution
5.1.4. Orbital Maneuvering System and Reaction Control System
5.1.4.1. Description
5.1.4.2. Phase Operations
5.1.4.3. OMS Initiators
5.1.5. Orbiter Thermal Protection System
5.1.5.1. Description
5.1.5.2. Success Criteria
5.1.5.3. TPS Study Integration
5.1.6. Risk Contribution of Orbiter Components
5.2. Supporting Systems
5.2.1. Orbiter Electric Power
5.2.1.1. Description
5.2.1.2. Success Criteria
5.2.1.3. Data Analysis
5.2.2. Orbiter Enviromental Control and Life Support (ECLSS)
5.2.2.1. Description
5.2.2.2. Risk Implications
5.2.3. General Purpose Computer and Data Management
5.2.3.1. Description
5.2.3.2. Risk Implications
5.2.4. SSME Thrust Vector Control
5.2.4.1. Description
5.2.4.2. Success Criteria
5.2.4.4. Fault Tree Model & Data Analysis
5.3. Terminal Phase Risk
5.4. Systems Analysis Summary Results

6.0. Model Evaluation & Results
6.1 Representation and Propagation of Uncertainties
6.1.1. Application of Uncertainty Bounds
6.1.2. Evaluation of Top-level Uncertainty Distributions
6.2 Base Case Risk Evaluation
6.3 Sensitivity Analysis and Redesign Effectiveness Measures

7.0. Conclusions & Recommendations
7.1. Risk Insights Based on PRA Results
7.2. Design and Operations Recommendations
7.3. Recommendations for Continuing Risk Assessment and Management Work

References

Volume II

Appendix A. Space Shuttle Integrated Loss of Vehicle Model
A.1. Integrated Space Shuttle LOV Fault Tree Model

Volume III

Appendix A. Space Shuttle Integrated Loss of Vehicle Model (cont.)
A.2. Basic Event Database
A.3. Minimal Cutsets

Volume IV

Appendix B. Systems Models and Related Data References
B.1. Space Shuttle Main Engine
B.2. Integrated Solid Rocket Booster
B.3. Orbiter Auxiliary Power Unit/Hydraulics
B.4. Electrical Power System

Volume V

Appendix C. Auxiliary Shuttle Risk Analyses
C.1. Probabilistic Risk Assessment of Space Shuttle Phase 1:
Space Shuttle Catastrophic Failure Frequency Final Report
C.2. Risk Analysis Applied to the Space Shuttle Main Engine:
Demonstration Project for the Main Combustion Chamber Risk Assessment
C.3. An Investigation of the Risk Implications of Space Shuttle
Solid Rocket Booster Chamber Pressure Excursions
C.4. Safety of the Thermal Protection System of the Space Shuttle Orbiter:
Quantitative Analysis and Organizational Factors
C.5. Space Shuttle Main Propulsion Pressurization System
Probabilistic Risk Assessment, Final Report
C.6. Space Shuttle Probabilistic Risk Assessment Proof-of-Concept Study:
Auxiliary Power Unit and Hydraulic Power Unit Analysis Report

Table of Figures

1.1. LOV Risk Uncertainty Distributions for Total Shuttle Mission
1.2. Mission Phase and Relative Element Risk Contribution
1.3. Distribution of Mean Loss-of-Vehicle Risk Among Mission Phases
1.4. Distribution of Mean Loss-of-Vehicle Risk Among Shuttle Vehicle Elements
1.5. Comparison of Current PRA Ascent Results with Previous Risk Studies
1.6. Comparison of Ascent Risk Uncertainty Distributions for Shuttle Elements
1.7. Comparison of Ascent Phase and Total Mission Risk Contributions
1.8. Relative Risk Versus Mission Intervals: Linear Risk Scale
1.9. Propagation of SSME Specific Accident Initiating Events
1.10. Comparison of Total Mission Risk Contribution to Percentage of CILs
1.11. A Generic Accident Scenario for Probabilistic Risk Assessment
1.12. Flow Chart of Shuttle PRA Data and Models

2.1. Shuttle PRA Project Organization
2.2. Relationships Between the Shutle PRA and Previous Shuttle Risk Assessments
2.3. Segments of Nominal Shuttle Mission Considered in PRA, Showing Mission Phase Boundaries
2.4. A Generic Accident Scenario for Probabilistic Risk Assessment
2.5. Task Network for a Generic Probabilistic Risk Assessment
2.6. Flow Chart of Shuttle PRA Data and Models

3.1. Accident Sequence Schematic
3.2. Shuttle Probabilistic Risk Assessment Modelling Mechanics

4.1. Overview of the Shuttle PRA Risk Data Analysis
4.2. Use of Shuttle Experience Data to Estimate Failure Frequencies

5.1. SSME Schematic
5.2. Main Propulsion System Schematic
5.3. Helium Supply System Schematic
5.4. SSME Specific Initiator Frequency Distribution
5.5. Loss of MCC Pressure Event Tree
5.6. Propagation of SSME Specific Accident Initiating Events
5.7. SSME Components Risk Contribution
5.8. HPOTP Failure Mode Risk Contribution
5.9. HPFTP Failure Mode Risk Contribution
5.10. MCC Failure Mode Risk Contribution
5.11. ISRB Diagram
5.12. ISRB Risk Contribution
5.13. APU/HYD/WSB Schematic
5.14. Auxiliary Power Unit Location

5.15. APU Risk Contribution
5.16. APU Entry/Descent Failure Mode Risk Contribution
5.17. Debris Impact Profile
5.18. Secondary Tile Loss Profile
5.19. Burnthrough Profile
5.20. Criticality Profile
5.21. TPS Min-Zone Partitioning
5.22. TPS Min-Zones Risk Contribution
5.23. Orbiter Risk Contribution

6.1. Information Flow between Analysis Computer Modules
6.2. LOV Risk Uncertainty Distributions for Total Shuttle Mission
6.3. Distribution of Mean Loss-of-Vehicle Risk Among Shuttle Vehicle Elements
6.4. Comparison of Current PRA Ascent Results with Previous Risk Studies
6.5. Comparison of Ascent Risk Uncertainty Distributions for Shuttle Elements
6.6. Mission Phase and Relative Element Risk Contribution
6.7. Relative Risk Versus Mission Intervals: Linear Risk Scale

7.1. Risk Comparison of International Launch Vehicles

Table of Tables

1.1. Summary of PRA Results: Estimated Loss-of-Vehicle Frequency
1.2. Risk Summary Statistics of Most Significant Accident Sequences
1.3. Summary of Top 20 Risk-Contributing Accident Sequences
1.4. Comparison of Current Shuttle PRA Ascent Results and Previous Studies
1.5. Estimated Loss-of-Vehicle Frequency for Base and Sensitivity Cases

2.1. Characteristics of the Nominal Shuttle Mission Considered in the PRA

4.1. Data Type General Characteristics
4.2. Example Data Types Used in Shuttle PRA

5.1. SSME Defined Redline Parameters and Limit Exceeded Definitions
5.2. Loss of MCC Pressure Event Descriptions
5.3. High Mixture Ratio in OPB Event Descriptions
5.4. Loss of Gross H2 Flow Event Descriptions
5.5. High Mixture Ratio in FPB Event Descriptions
5.6. Loss of Fuel to Both Preburners Event Descriptions
5.7. HPFTP Coolant Liner Overpressure Event Descriptions
5.8. Failure to Maintain Proper SSME Propellant Valve Position Event Descriptions
5.9. Hydraulic Lock-up Required Event Descriptions
5.10. Structural Failure of Critical SSME Components Event Descriptions
5.11. Simultaneous Dual SSME Shutdown Event Descriptions
5.12. Nominal MECO and Propellant Dump Event Descriptions
5.13. Helium System Leakage Event Descriptions
5.14. Failure to Provide POGO Accumulator Charge Event Descriptions
5.15. Leakage of SSME/MPS Propellants Event Descriptions
5.16. Summary of Hydraulic System Redundancy
5.17. APU/HYD Hub Breakup and Overspeed Event Descriptions
5.18. Large Gas or Hydrazine Leak Event Descriptions
5.19. End State Definitions
5.20. APU/HYD Hydrazine Leaks During Ascent Event Descriptions
5.21. At Least One APU/HYD Unit Fails Without a Hydrazine Leak During Ascent Event Descriptions
5.22. Fully Operational APU/HYD Units During Reentry, TAEM and Landing Event Descriptions
5.23. TPS Min-Zone Catastrophic Failure Probabilities
5.24. Systems Analysis Summary Results

6.1. Summary of PRA Results: Estimated Loss-of-Vehicle Frequency
6.2. Comparison of Current Shuttle PRA Ascent Results and Previous Studies
6.3. Summary of Top 20 Risk-Contributing Accident Sequences
6.4. Risk Summary Statistics of Most Significant Accident Sequences
6.5. Estimated Loss-of-Vehicle Frequency for Base and Sensitivity Cases

1.0. Executive Summary

1.0. EXECUTIVE SUMMARY

This document is the Executive Summary of a technical report on a probabilistic risk assessment (PRA) of the Space Shuttle vehicle performed under the sponsorship of the Office of Space Flight of the US National Aeronautics and Space Administration by Science Applications International Corporation and its subcontractors Safety Factor Associates and Empresarios Agrupados, with the participation and support of NASA Headquarters, the NASA field centers that operate the Shuttle system, and the principal Shuttle contractors. It briefly summarizes the methodology and results of the Shuttle PRA. The reader is referred to the main technical report and its appendixes for complete details.

1.1. Background

While NASA has always emphasized safety in design and operations, especially for crewed spacecraft, the *Challenger* accident brought home the need for a systematic, quantitative, and defensible way to evaluate flight risks and to identify and prioritize the factors that contribute to them so they can be targeted for improvement. In the late 1980s the successful application of probabilistic risk assessment methods to nuclear power generation, chemical processing, and other facilities and systems where technological accident risks are of concern led NASA Headquarters to consider and eventually to adopt PRA as one answer to this need. The current risk assessment of the Space Shuttle vehicle is the latest and largest of a series of NASA-sponsored probabilistic risk analyses that began with the PRA Proof of Concept studies in 1987 and the analysis of catastrophic failure frequency for the Shuttle mission that launched *Galileo* in 1989.

1.2. Key Objectives

The primary objective of this project was to support management and engineering decision-making with respect to the Shuttle program by producing...

(1) a quantitative probabilistic risk model of the Space Shuttle during flight,

(2) a quantitative assessment of in-flight safety risk,

(3) an identification and prioritization of the features of design and operations that principally contribute to in-flight safety risk, and

(4) a mechanism for risk-based evaluation of proposed modifications to the Shuttle system.

Secondary objectives were to provide a vehicle for introducing and transferring PRA technology to the NASA community, and to demonstrate the value of PRA by applying it beneficially to a real program of great international importance.

1.3. Key Results.

This section summarizes the most important results of the Shuttle PRA in a number of formats. Paragraph 1.4 discusses some of the salient implications of these results.

1.3.1. Estimated Risk of Loss of Vehicle.

Table 1.1 below and Figure 1.1 on the next page summarize the estimated risk of loss of vehicle due to all initiating events and accident sequences considered in the analysis over the entire mission from main engine ignition through wheel stop on landing. These are the principal top-level results of the PRA of the Space Shuttle. Uncertainty distributions were evaluated for the overall Shuttle risk and the element LOV probabilities.

For a variety of reasons that are discussed more fully in Appendix 1 of the main report, uncertainty is inherent in any estimate of risk. A key benefit of probabilistic risk assessment is that it defines and quantifies this uncertainty, allowing the user of the results to understand not only how risky the system is estimated to be and what factors contribute to this risk, but also how much confidence he or she should place in the estimates and where additional work is needed to make them more certain. PRA risk estimates are generally expressed as uncertainty distributions (i.e., probability density functions of accident frequency), rather than as point values. These distributions are ordinarily defined by parameters that describe the central tendency — the means, medians, or both of the probability density functions — and several percentiles that describe the extremes of the distribution. The results are stated in this way in Table 1.1, Figure 1.1, and elsewhere throughout this summary and the main technical report.

Table 1.1 gives the mean and the 5th, 50th (median), and 95th percentiles of the probability distributions of estimated loss-of-vehicle frequency for each element. Figure 1.1 on the next page presents the loss-of-vehicle frequency information in the "error bar" graphical format that illustrates the extremes of the uncertainty distributions.

Table 1.1. Summary of PRA Results: Estimated Loss-of-Vehicle Frequency

	5th Percentile	Median	Mean	95th Percentile
STS	1/230	1/145	1/131	1/76
Orbiter	1/758	1/397	1/330	1/169
SSME	1/820	1/410	1/348	1/172
ISRB	1/2591	1/1152	1/775	1/322
ET	1/86207	1/11223	1/5208	1/1460
LANDING	1/141528	1/9435	1/2433	1/629

Figure 1.1. LOV Risk Uncertainty Distributions for Total Shuttle Mission

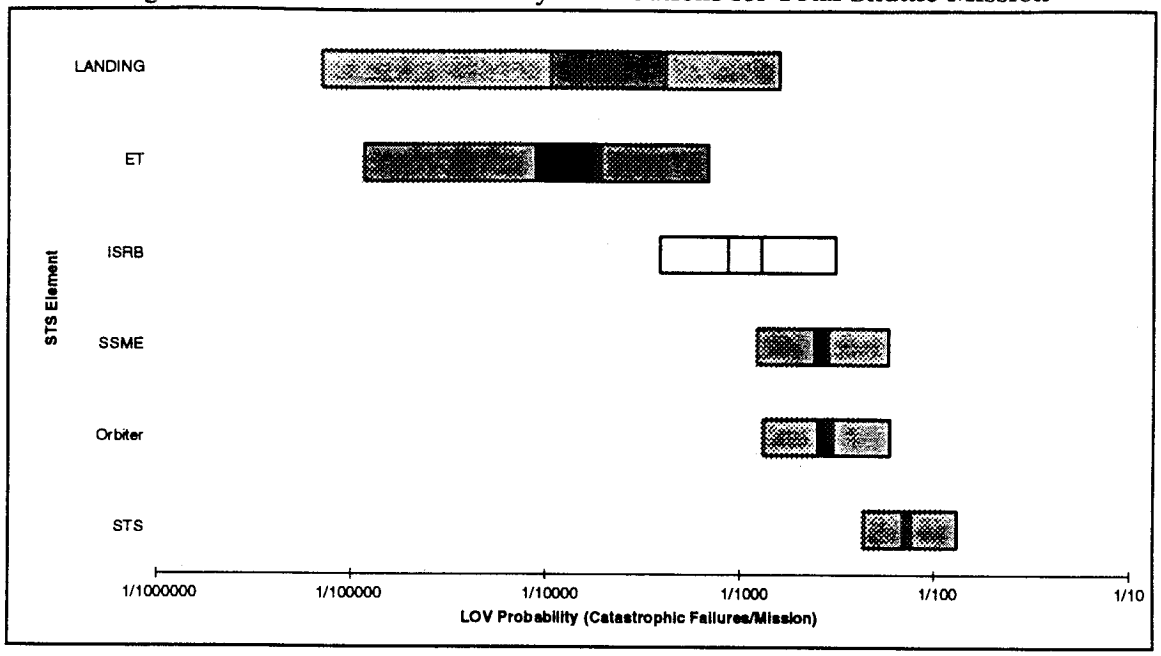

1.3.2. Distribution of Total Estimated Risk Among Mission Phases.

As discussed in detail in paragraph 2.3.1 of the main technical report, for risk assessment purposes the nominal Shuttle mission is considered to comprise three phases: *ascent*, *orbit*, and *descent*. Figure 1.2 briefly defines these phases and shows approximately how the total mean risk of loss of vehicle is distributed among them. The distributions of these contributions are presented in error-bar format in Figure 1.3 on the next page.

Figure 1.2. Mission Phase and Relative Element Risk Contribution

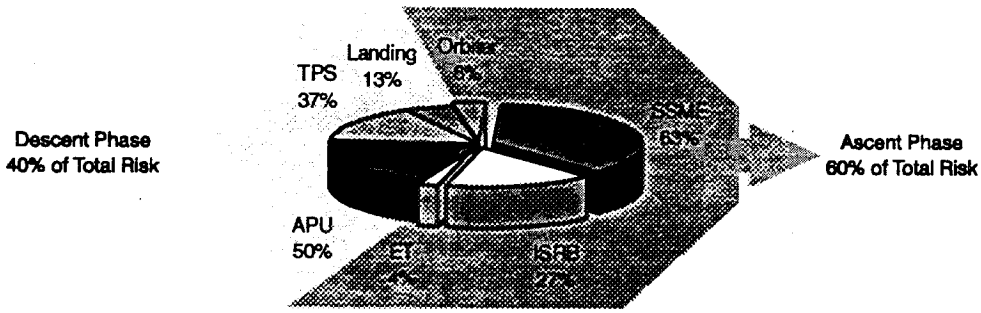

Phase Description	
Ascent	SSME Ignition Through APU Shutdown
Orbit*	APU Shutdown to De-orbit Burn
Descent	De-orbit Burn Through Wheel Stop at Landing

*Orbit phase did not contain any significant LOV risk drivers

Figure 1.3. Distribution of Mean Loss-of-Vehicle Risk Among Mission Phases.

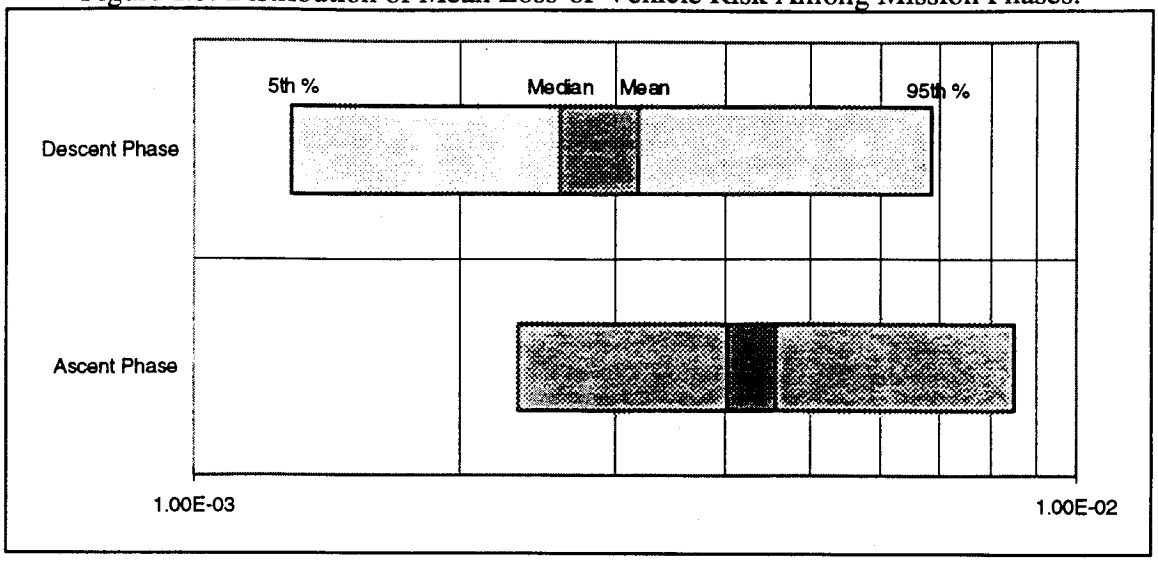

1.3.3. Contributions to Total Estimated Risk of Shuttle Vehicle Elements

Figure 1.4 depicts the approximate relative contributions of the principal elements of the Shuttle vehicle to the mean risk of loss of vehicle in pie-chart format.

Figure 1.4. Distribution of Mean Loss-of-Vehicle Risk Among Shuttle Vehicle Elements.

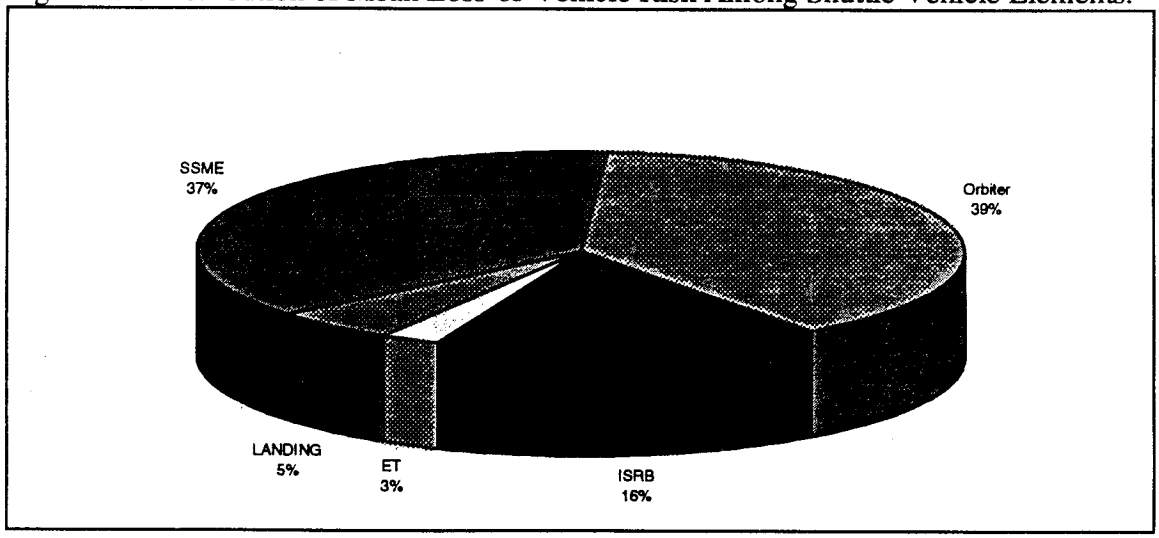

1.3.4. Risk Drivers.

The "risk drivers" of a system or operation are the factors that dominate the total risk, and consequently should be targeted for further evaluation and potentially for risk-mitigation efforts. The PRA process identifies an event or accident sequence as a risk driver when (1) its occurrence leads to loss of vehicle with little or no chance of recovery, (2) it has a high probability of occurrence, and/or (3) its likelihood or consequences are subject to so much uncertainty that it is impossible to say with confidence that it is not a risk driver.

Table 1.2. summarizes the risk statistics for the most important Shuttle flight risk drivers identified by the base-case risk assessment. (Please refer to paragraph 1.6 for an explanation of the terms "accident sequences" and "initiating events.")

Table 1.2. Risk Summary Statistics of Most Significant Accident Sequences (Sequences shown in Table 1.3)

	Top 10 Accident Seq.	Top 20 Accident Seq.		Top 10 Accident Seq.	Top 20 Accident Seq.
Percent of Total Risk	38.88%	61.17%			
Orbiter	47.49%	41.98%	Auxiliary Power Units	39.18%	28.99%
			Thermal Protection System	8.31%	12.99%
SSME	45.48%	45.51%	Turbomachinery	37.01%	29.95%
			Combustion Devices	8.47%	15.56%
ISRB	7.03%	12.51%	Redesigned Solid Rocket Motor	7.03%	8.73%
			Solid Rocket Booster	-	3.78%

Table 1.3. Summary of Top 20 Risk-Contributing Accident Sequences.

Rank	Accident Description	Mean LOV Prob	Mean Percent Contrib. to Total Risk
1	SSME HPOTP Bearing Failure Due To Spalling, Pitting, Wear Or Corrosion	4.52E-04	5.89%
2	Two Leakage Induced Orbiter APU Failures During Re-entry/Descent and Failure To Land Using One APU	4.28E-04	5.57%
3	Two Orbiter APUs Fail To Start Or Run During Re-entry/Descent Due to Common Cause Failure and Failure To Land Using One APU	3.99E-04	5.20%
4	All Three Orbiter APUs Fail To Start Or Run During Re-entry/Descent Due to Common Cause Failure	3.43E-04	4.47%
5	SSME MCC Manifold Weld Failure	2.53E-04	3.29%
6	SSME HPFTP Turbine Blade Failure	2.51E-04	3.27%
7	Catastrophic Failure Of Right Side TPS, Fwd Mid Edge (624 Tiles)	2.48E-04	3.23%
8	Common Cause Failure of ISRB Igniter Joint S&A Primary and Secondary Gasket Seals	2.10E-04	2.73%
9	SSME HPOTP Failure Due To Cavitation Damage	2.01E-04	2.62%
10	SSME HPFTP Impeller/Diffuser Failure	2.01E-04	2.62%
11	Propellant Fails To Ignite In One Of The ISRBs	2.00E-04	2.60%
12	All Three Orbiter APUs Fail To Run During Ascent Due to Common Cause Failure	1.92E-04	2.50%
13	Catastrophic Failure Of Left Side Near Main Landing Gear TPS (780 Tiles)	1.87E-04	2.43%
14	Two or more ISRB Holddown Studs Hang-up	1.78E-04	2.31%
15	Failure In SSME MCC EDNi Liner Closeout Structure	1.76E-04	2.29%
16	Catastrophic Failure Of Forward Right Side Near Main Landing Gear TPS (676 Tiles)	1.75E-04	2.28%
17	SSME MI Lox Post Structural Failure	1.51E-04	1.97%
18	Structural Failure Of SSME LPOTP	1.51E-04	1.97%
19	SSME HPOTP Turbine Blade Failure	1.51E-04	1.97%
20	SSME FPB Faceplate Failure Due To Erosion	1.51E-04	1.97%

1.4. Key Insights and Recommendations.

1.4.1. Comparison with the Results of Previous Space Shuttle Risk Assessments.

Although the probabilistic risk assessment that is the subject of this report is the first full-mission risk assessment to be performed on the Shuttle vehicle to date, NASA and SAIC have conducted a number of previous risk analyses on various aspects of the vehicle and mission. Table 1.4 summarizes the results of two of these assessments conducted for the ascent phase of a Shuttle mission and compares them with the relevant base-case results of the current PRA.

Table 1.4. Comparison of Current Shuttle PRA Ascent Results and Previous Studies

	5th Percentile	Median	Mean	95th Percentile
STS PRA	1/429	1/248	1/219	1/118
PRA Phase 1 Study	1/223	1/90	1/73	1/31
Galileo Study	1/350	1/78	1/55	1/18

Figure 1.5 shows that the uncertainty distribution developed by this PRA for the ascent phase agrees with but is not entirely bounded by those estimated for analogous conditions in the earlier studies. However, the new results have considerably narrower bounds of uncertainty than the old ones. There are two main reasons for this situation. First, much of the data underlying the current PRA is based on statistical analysis of Shuttle flight and test experience; the additional failure free experience accumulated since the earlier studies necessarily narrows the uncertainty bounds of the risk estimates. Second, the current PRA has analyzed the risk-driving systems in much greater detail than the earlier analyses. In many cases, but not all, a deeper analysis reduces the uncertainty in the results. The top-level analysis of earlier studies also tended to produce conservative results which explains why the previous results are skewed toward high failure frequencies.

Figure 1.5. Comparison of Current PRA Ascent Results with Previous Risk Studies

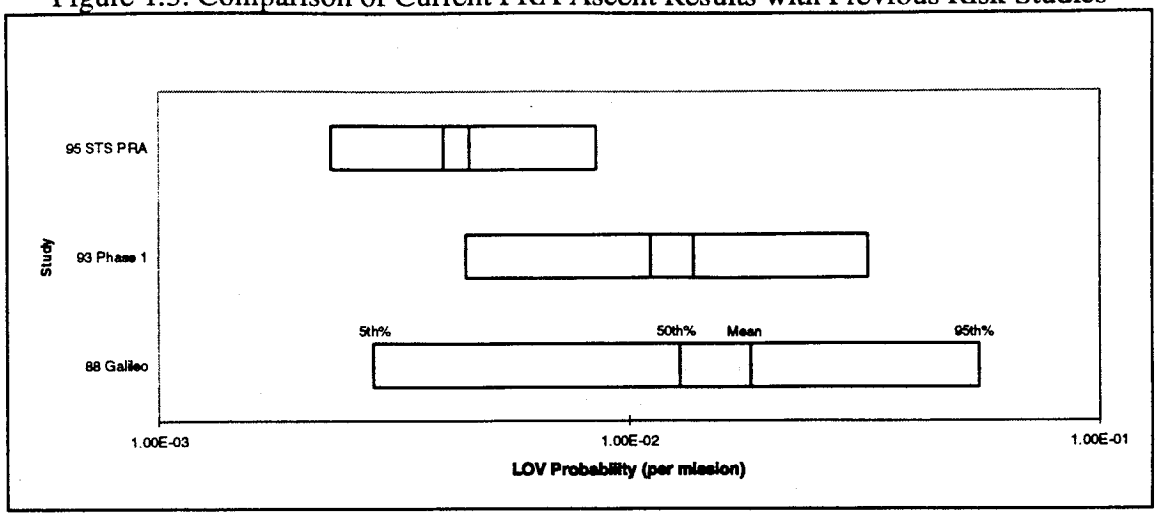

Figure 1.6. Comparison of Ascent Risk Uncertainty Distributions for Shuttle Elements

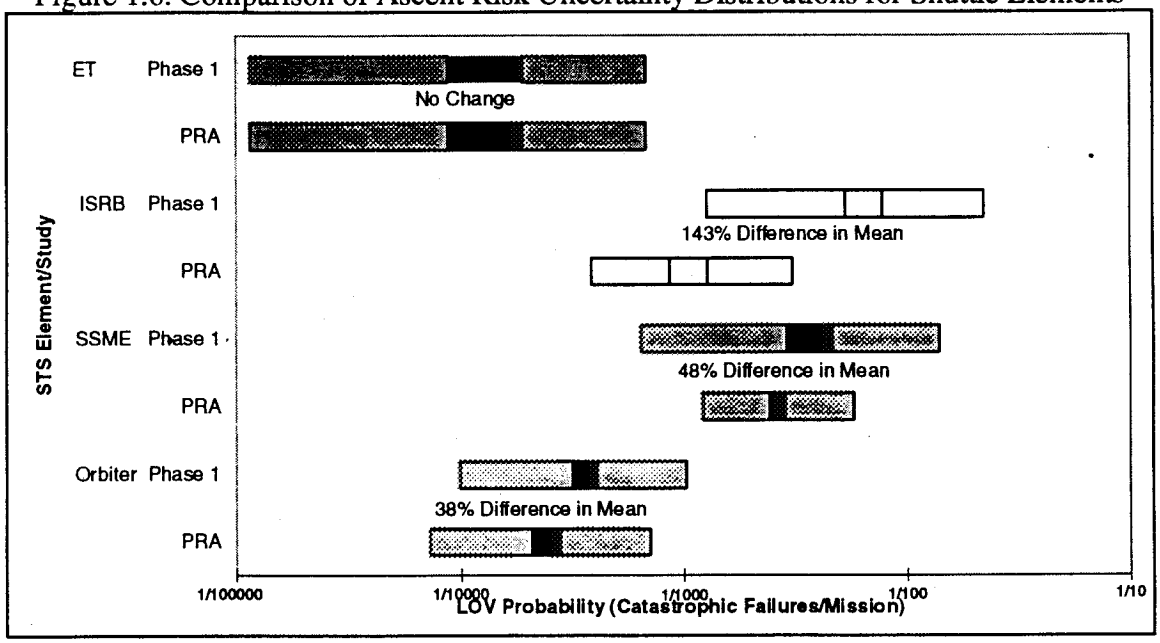

1.4.2. Risk Dominance of Propulsion Systems.

The foregoing information indicates that the main ascent propulsion systems of the shuttle vehicle contribute the majority of flight risk although they operate only during the ascent phase for a little over eight minutes. The combined mean contribution of the solid rocket motors and the orbiter main propulsion systems to loss-of-vehicle risk is nearly equivalent to the total contribution of all other systems considered in the PRA. This is to be expected because the ascent propulsion systems are critical, highly-stressed, high-energy systems with little redundancy. As a result, many of their failure modes are unrecoverably catastrophic to the vehicle through such mechanisms as burn-throughs of hot-gas pressure boundaries in both the SSMEs and the RSRMs, and violent turbomachinery disassembly in the SSMEs. The orbiter auxiliary power units and the thermal protection tiles are also susceptible to catastrophic failures, but failures of these systems that can precipitate loss of vehicle are much less likely than those of the propulsion systems. The orbiter also contributes a significant degree of the risk, mostly during re-entry and descent related maneuvers. The remaining systems contribute relatively little to total flight risk for one or both of the following reasons: (1) initiating events that can lead to loss of vehicle are extremely rare in these systems (e.g., structural failures); or (2) failures of these systems are relatively inconsequential to flight safety under the ground rules of the PRA because redundancy, functional diversity, and long time-to-effect permit recovery or a safe mission abort rather than loss of vehicle (e.g., electric power and environmental control and life support system failures).

Figure 1.7. Comparison of Ascent Phase and Total Mission Risk Contributions

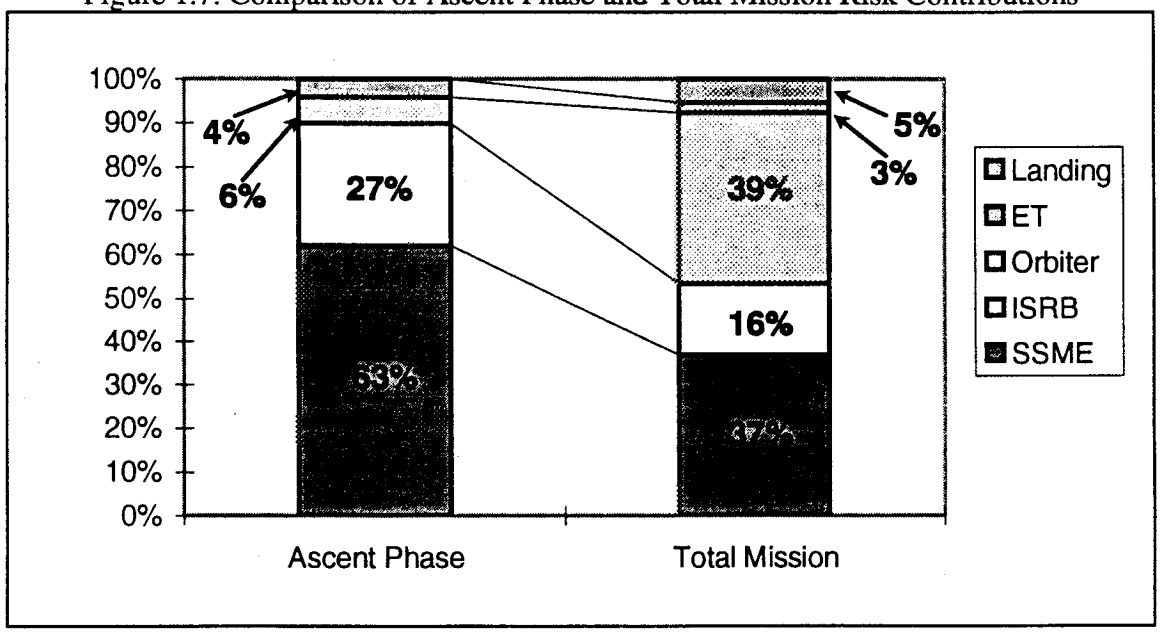

1.4.3. Risk of Mission Intervals.

As one would expect from the above discussion, the PRA shows that the ascent phase of the mission has the highest risk concentration or failure rate of the three phases. The plots of relative loss-of-vehicle risk versus mission time and events in Figure 1.8 illustrate this observation most effectively. The relative risk is plotted linearly in order to highlight the importance of the ascent phase when considering the risk per unit time. The orbiter has a relatively high overall risk because it must operate many times longer than the propulsion systems. Basically, the longer operation time gives the orbiter more time to fail. The risk or LOV probability for each interval is evaluated by multiplying the LOV frequency by the time of duration.

Figure 1.8. Relative Risk Versus Mission Intervals: Linear Risk Scale.

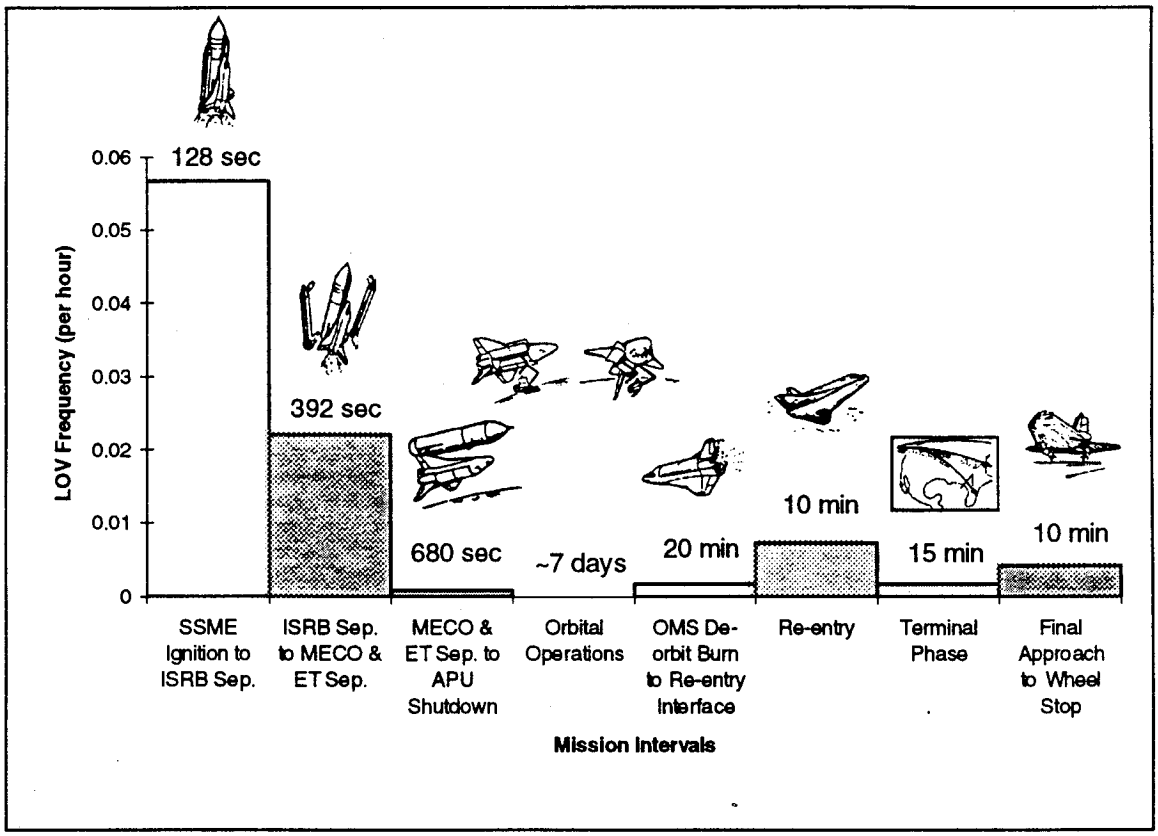

1.4.4. Effectiveness of Space Shuttle Main Engine Protective "Redlines."

The SSME engine controllers monitor a number of engine parameters such as temperatures, pressures, and vibration readings during operation, and shut down the affected engine when they exceed safe limits (i.e., "redlines"). The PRA confirms the effectiveness of the engine health monitoring subsystem. The great majority of the accident sequences initiated by

problems within the main engine that would otherwise progress to disruptive engine failure and loss of vehicle are successfully interrupted by a redline shutdown. This is illustrated in Figure 1.9 which shows the initial initiating accident event distribution for the SSME cluster and the final risk contribution which is dominated by critical structural failures with no opportunity for mitigation. That is only 7.8% of the initiators are due to critical structural failures but these contribute to 99.96% of the residual risk. This issue will be discussed to a greater extent in section 3.3.1.2.

Figure 1.9. Propagation of SSME Specific Accident Initiating Events.

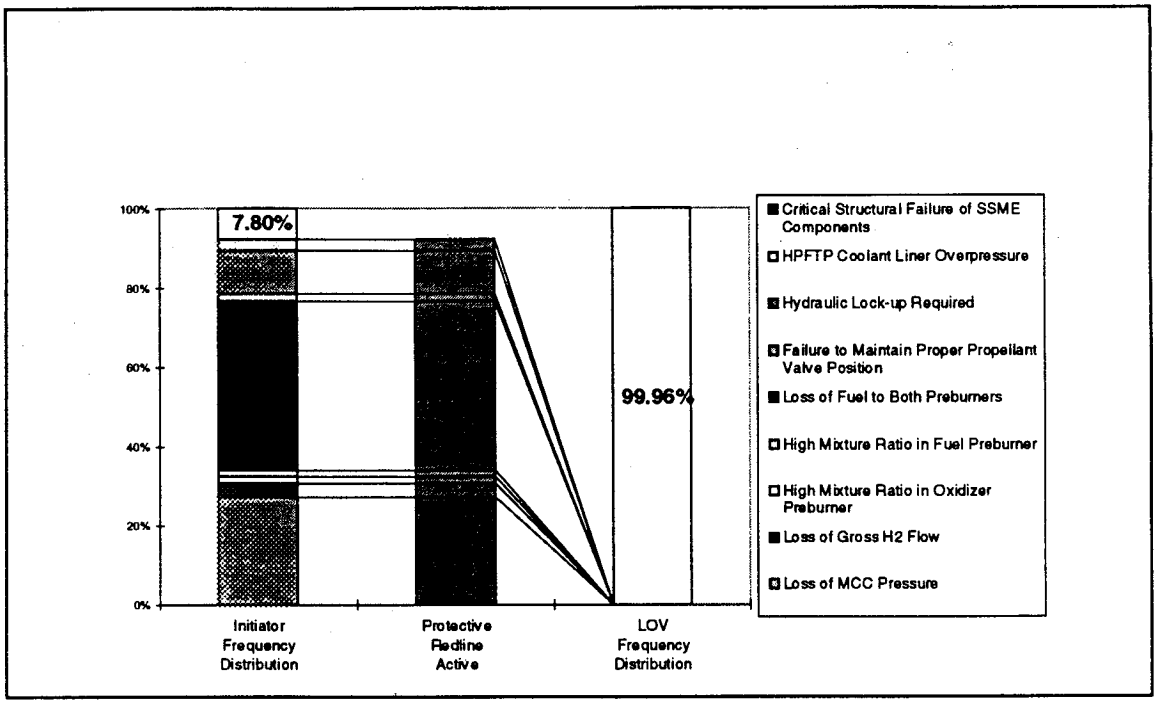

1.4.5. Distribution of Risk Versus Critical Items Count.

The number of items associated with a Shuttle element or system on the Critical Items List (i.e., the "CIL count") is often taken to be a top-level indicator of the contribution of that part of the vehicle to total risk. It turns out from the current study (and also from the earlier ones that broke down total risk into its contributing factors) that the relative CIL count is only grossly correlated with the actual distribution of risk. Figure 1.10 illustrates this observation clearly by comparing the proportion of CIL items per vehicle element with the proportion of estimated risk from the PRA. The orbiter dominates the CIL simply because as the most complex element, it accounts for the majority of Criticality 1 failure modes under the conservative rules of the NASA Failure Modes, Effects, and Criticality Analysis (FMECA) procedure. However, the PRA results show that the contribution of non-SSME orbiter failure modes to flight risk is only about four percent of the contribution that would be expected from the CIL count. It is apparent that the CIL count is a potentially misleading measure of risk.

Figure 1.10. Comparison of Total Mission Risk Contribution to Percentage of CILs

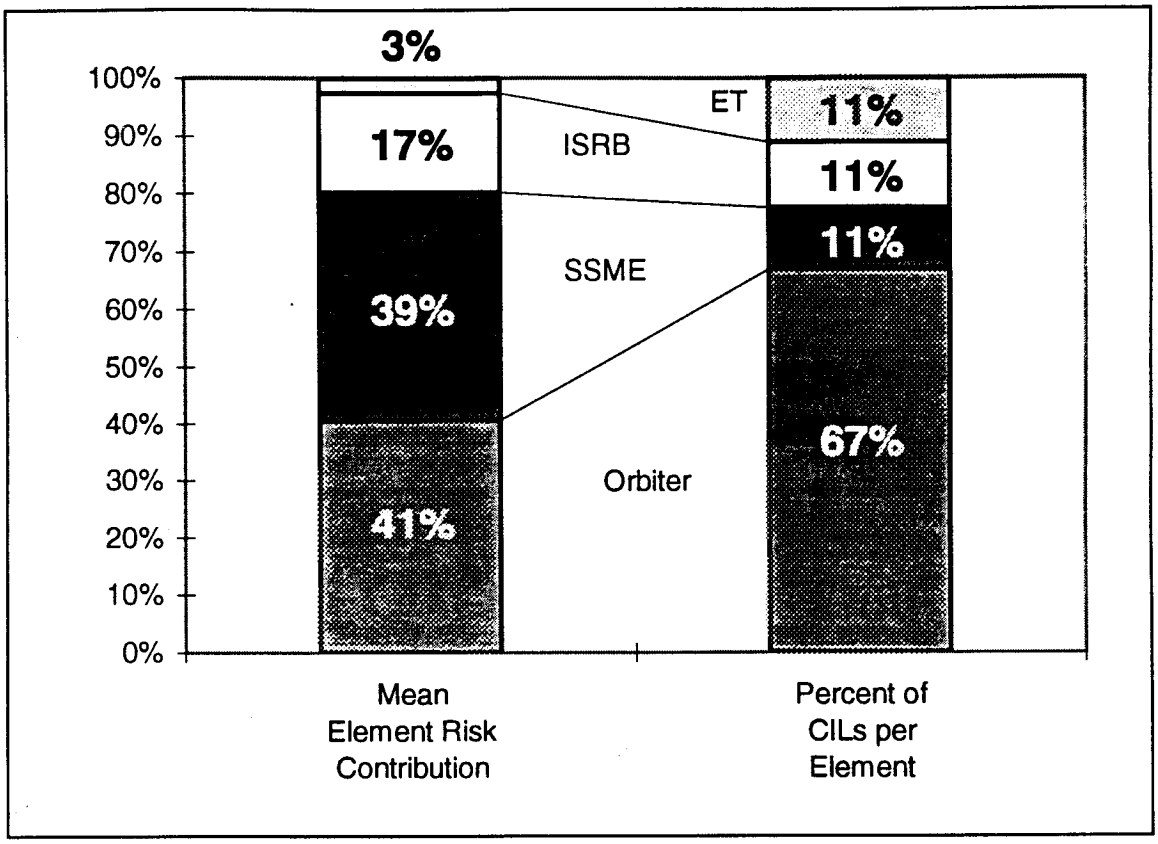

1.5. Ground Rules and Key Assumptions of the PRA.

In the interest of maximizing cost-effectiveness by focusing on the major risk drivers while de-emphasizing secondary risk issues, the risk assessment team established the ground rules and key simplifying assumptions listed below with the concurrence of NASA.

(1) The scope of the risk assessment included only accident scenarios initiated within the Shuttle vehicle and leading to loss of the vehicle during the mission phases from main engine ignition at launch through wheel stop on landing.

(2) Accident scenarios leading to mission aborts for which an abort procedure has been established were not modeled beyond the condition that triggers the abort. (From the standpoint of mission-level risk, this is equivalent to assuming that all proceduralized aborts are successful.)

(3) Only those accident scenarios that contribute substantially to the total estimated risk of loss of vehicle (as determined by a preliminary screening risk assessment using conservative

assumptions) were analyzed in detail. Other, smaller risk contributors were represented by conservative estimates.

(4) Shuttle vehicles were assumed to be configured and operated as they were at the beginning of the analysis in early January 1994; neither actual nor planned changes in design or operating practices subsequent to that cutoff date were considered.

(5) The Shuttle vehicle was assumed to be in the "as designed" configuration at main engine ignition; i.e., pre-launch configuration errors were not considered.

(6) The risk impacts, if any, of differences among the orbiters in the fleet (e.g., the presence or absence of modifications for long-duration missions) were not considered.

(7) Relevant information from existing, technically sound risk assessments was utilized wherever possible.

(8) The mean risk of loss of vehicle during the landing phase of the mission was assumed to be equal to 3.0% of the base-case ascent-phase risk estimated by the 1993 "Space Shuttle Catastrophic Failure Frequency" study (see Appendix C.1).

1.6. Analytical Approach.

The fundamental approach used in this Shuttle risk assessment is that of scenario-based probabilistic risk assessment. The following paragraphs outline the PRA process in very general terms. Please refer to sections 3.0 through 5.0 of the main technical report for details on the risk modeling of specific Shuttle systems and the Shuttle vehicle as a whole.

Probabilistic risk assessment is a multi-disciplinary complex of techniques that integrates probabilistic reliability-availability engineering and analysis with mathematical statistics, decision theory, systems engineering, conventional engineering analysis, and even cognitive psychology. It provides a systematic methodology for quantitatively evaluating the performance of a complex technological system, usually under abnormal conditions, in order to:

- evaluate the risks to people, property, and/or the success or productivity of a facility, mission, or project resulting from potential mishaps such as equipment failures, human errors, programmatic delays, and external events;

- estimate the consequences of anticipated mishaps; and

- identify and prioritize the design, construction, operation, maintenance, and management factors which contribute to risk.

Figure 1.11. A Generic Accident Scenario for Probabilistic Risk Assessment

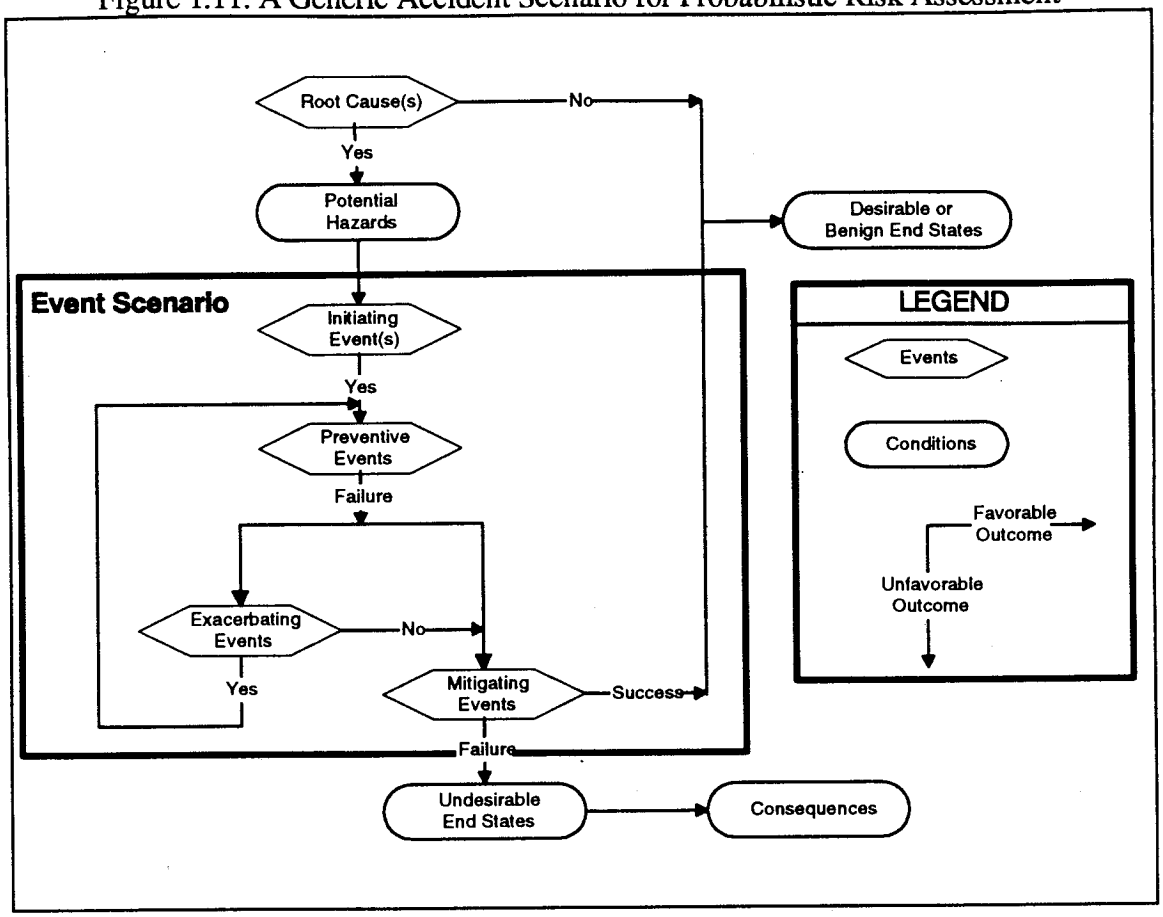

The concept of scenarios is basic to any understanding of the PRA process. As the name implies, a scenario is simply the chronological "story" of a sequence of events that is triggered by some incident and proceeds through intervening events to an end state. (In fact, a scenario is made up of "event sequences," if the end state of a sequence is an accident, then it is referred to as an "accident sequence.") In all but the simplest systems, there are several alternative sequences of events that can follow an initiator, depending on the outcomes of the intervening events. Figure 1.11 depicts an accident scenario in the most generic form, including some of the terminology used to describe the elements of scenarios. The key terms are (1) initiating events (or trigger events), which — in conjunction with pre-existing potential hazards — begin the scenario; (2) pivotal events, which have the potential to change the course of the scenario, and can have preventive, exacerbating, or mitigating effects; and (3) end states, which can have desirable, benign, or unfavorable consequences.

PRA is simply a systematic technique to for evaluating the probabilities and consequences of the various scenarios that can occur in a process or system. In general terms, this evaluation is accomplished by creating a hierarchy of risk models — master logic diagrams, functional event sequence diagrams, event trees, and fault trees, supported by phenomenological and

statistical analyses — which integrate information on the configuration and the normal and abnormal operation of the system with data on the likelihood of initiating events and the success or failure of pivotal events. Figure 1.12 on the next page is a top-level flow chart of the Shuttle PRA showing how information sources and models interact to produce the results.

1.7. Sensitivity Analysis

This analysis is unique not only because it represents the first complete Shuttle risk model but it is also unique in the way in which some systems were analyzed. For the RSRMs, leak check data was incorporated into the model in an attempt to compensate for a limited amount of operational data and to distinguish the Shuttle RSRMs from the historical class of solid rockets which have no leak check. Common cause failures of the APUs were investigated and were found to be significant contributors to APU risk given their proximity and shared support systems. A sensitivity analysis was conducted to determine the extent to which these issues influenced the final result. The mean estimates for the various sensitivity cases are shown in Table 1.5; note that all of the mean estimates fall within the established base case uncertainty distribution (Table 1.1).

Table 1.5. Estimated Loss-of-Vehicle Frequency for Base and Sensitivity Cases.

Sensitivity Case Description		STS	Orbiter	SSME	ISRB	ET	Landing
Base Case	Mean Mission	1	1	1	1	1	1
As Described in Report	LOV Probability	131	330	348	775	5208	2433
	Risk Percentage		39%	37%	17%	2%	5%
Sensitivity Case 1	Mean Mission	1	1	1	1	1	1
Additional Credit for Successful	LOV Probability	106	330	348	337	5208	2433
ISRB Leak Checks Not Considered	Risk Percentage		32%	30%	31%	2%	4%
Sensitivity Case 2	Mean Mission	1	1	1	1	1	1
Common Cause Failures not	LOV Probability	150	508	348	775	5208	2433
Considered for APUs	Risk Percentage		29%	43%	19%	3%	6%
Sensitivity Case 3	Mean Mission	1	1	1	1	1	1
Both ISRB Leak Checks and APU	LOV Probability	119	508	348	337	5208	2433
Common Cause Failures Neglected	Risk Percentage		23%	34%	35%	2%	5%

1.8. Synopsis of Main Technical Report.

The technical report from which this summary is derived addresses the following topics:

- Introduction to the Shuttle PRA: history, objectives, ground rules and assumptions, overview of analysis
- Risk modeling of the Shuttle vehicle and mission, and of its major systems
- Evaluation and results of the risk model
- Risk data analysis
- Insights and recommendations
- Appendixes: details of risk models, basic events database, references.

Figure 1.12. Flow Chart of Shuttle PRA Data and Models.

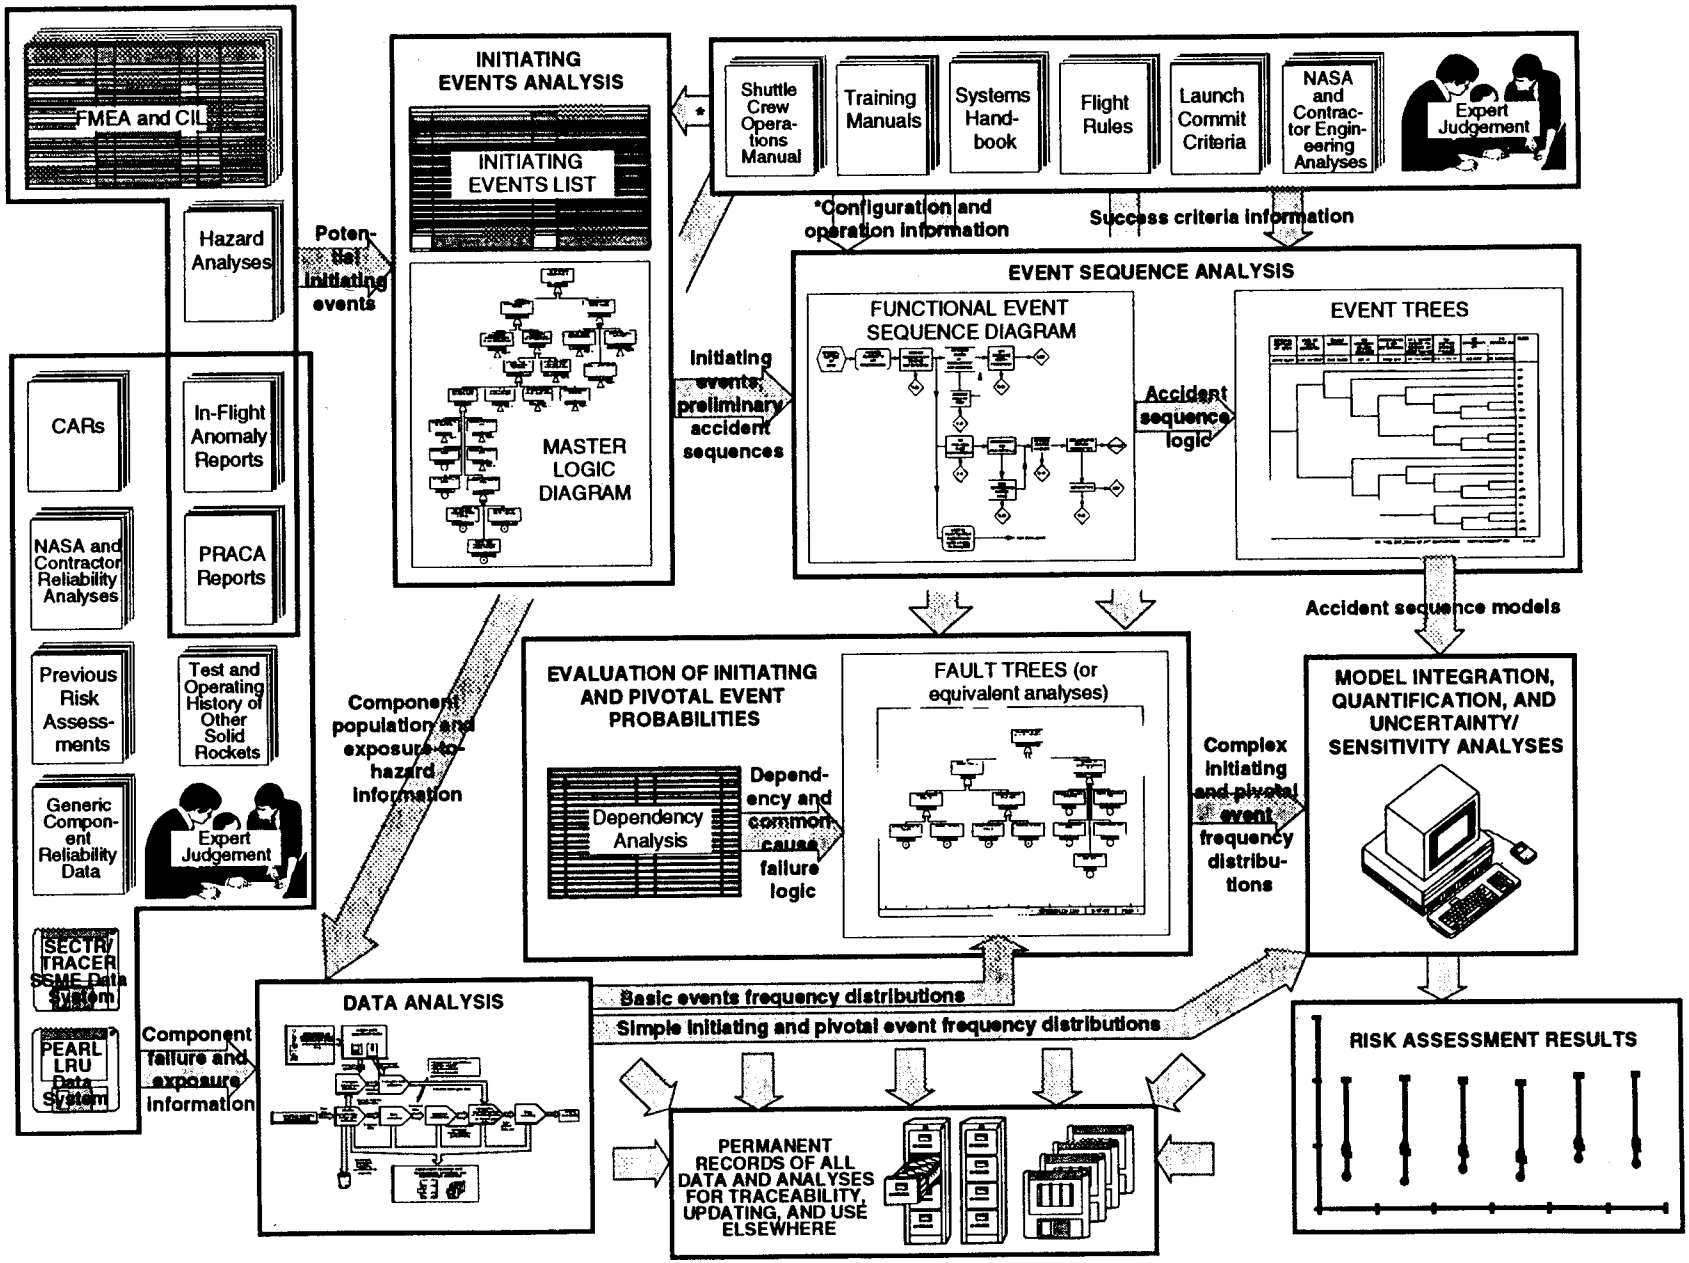

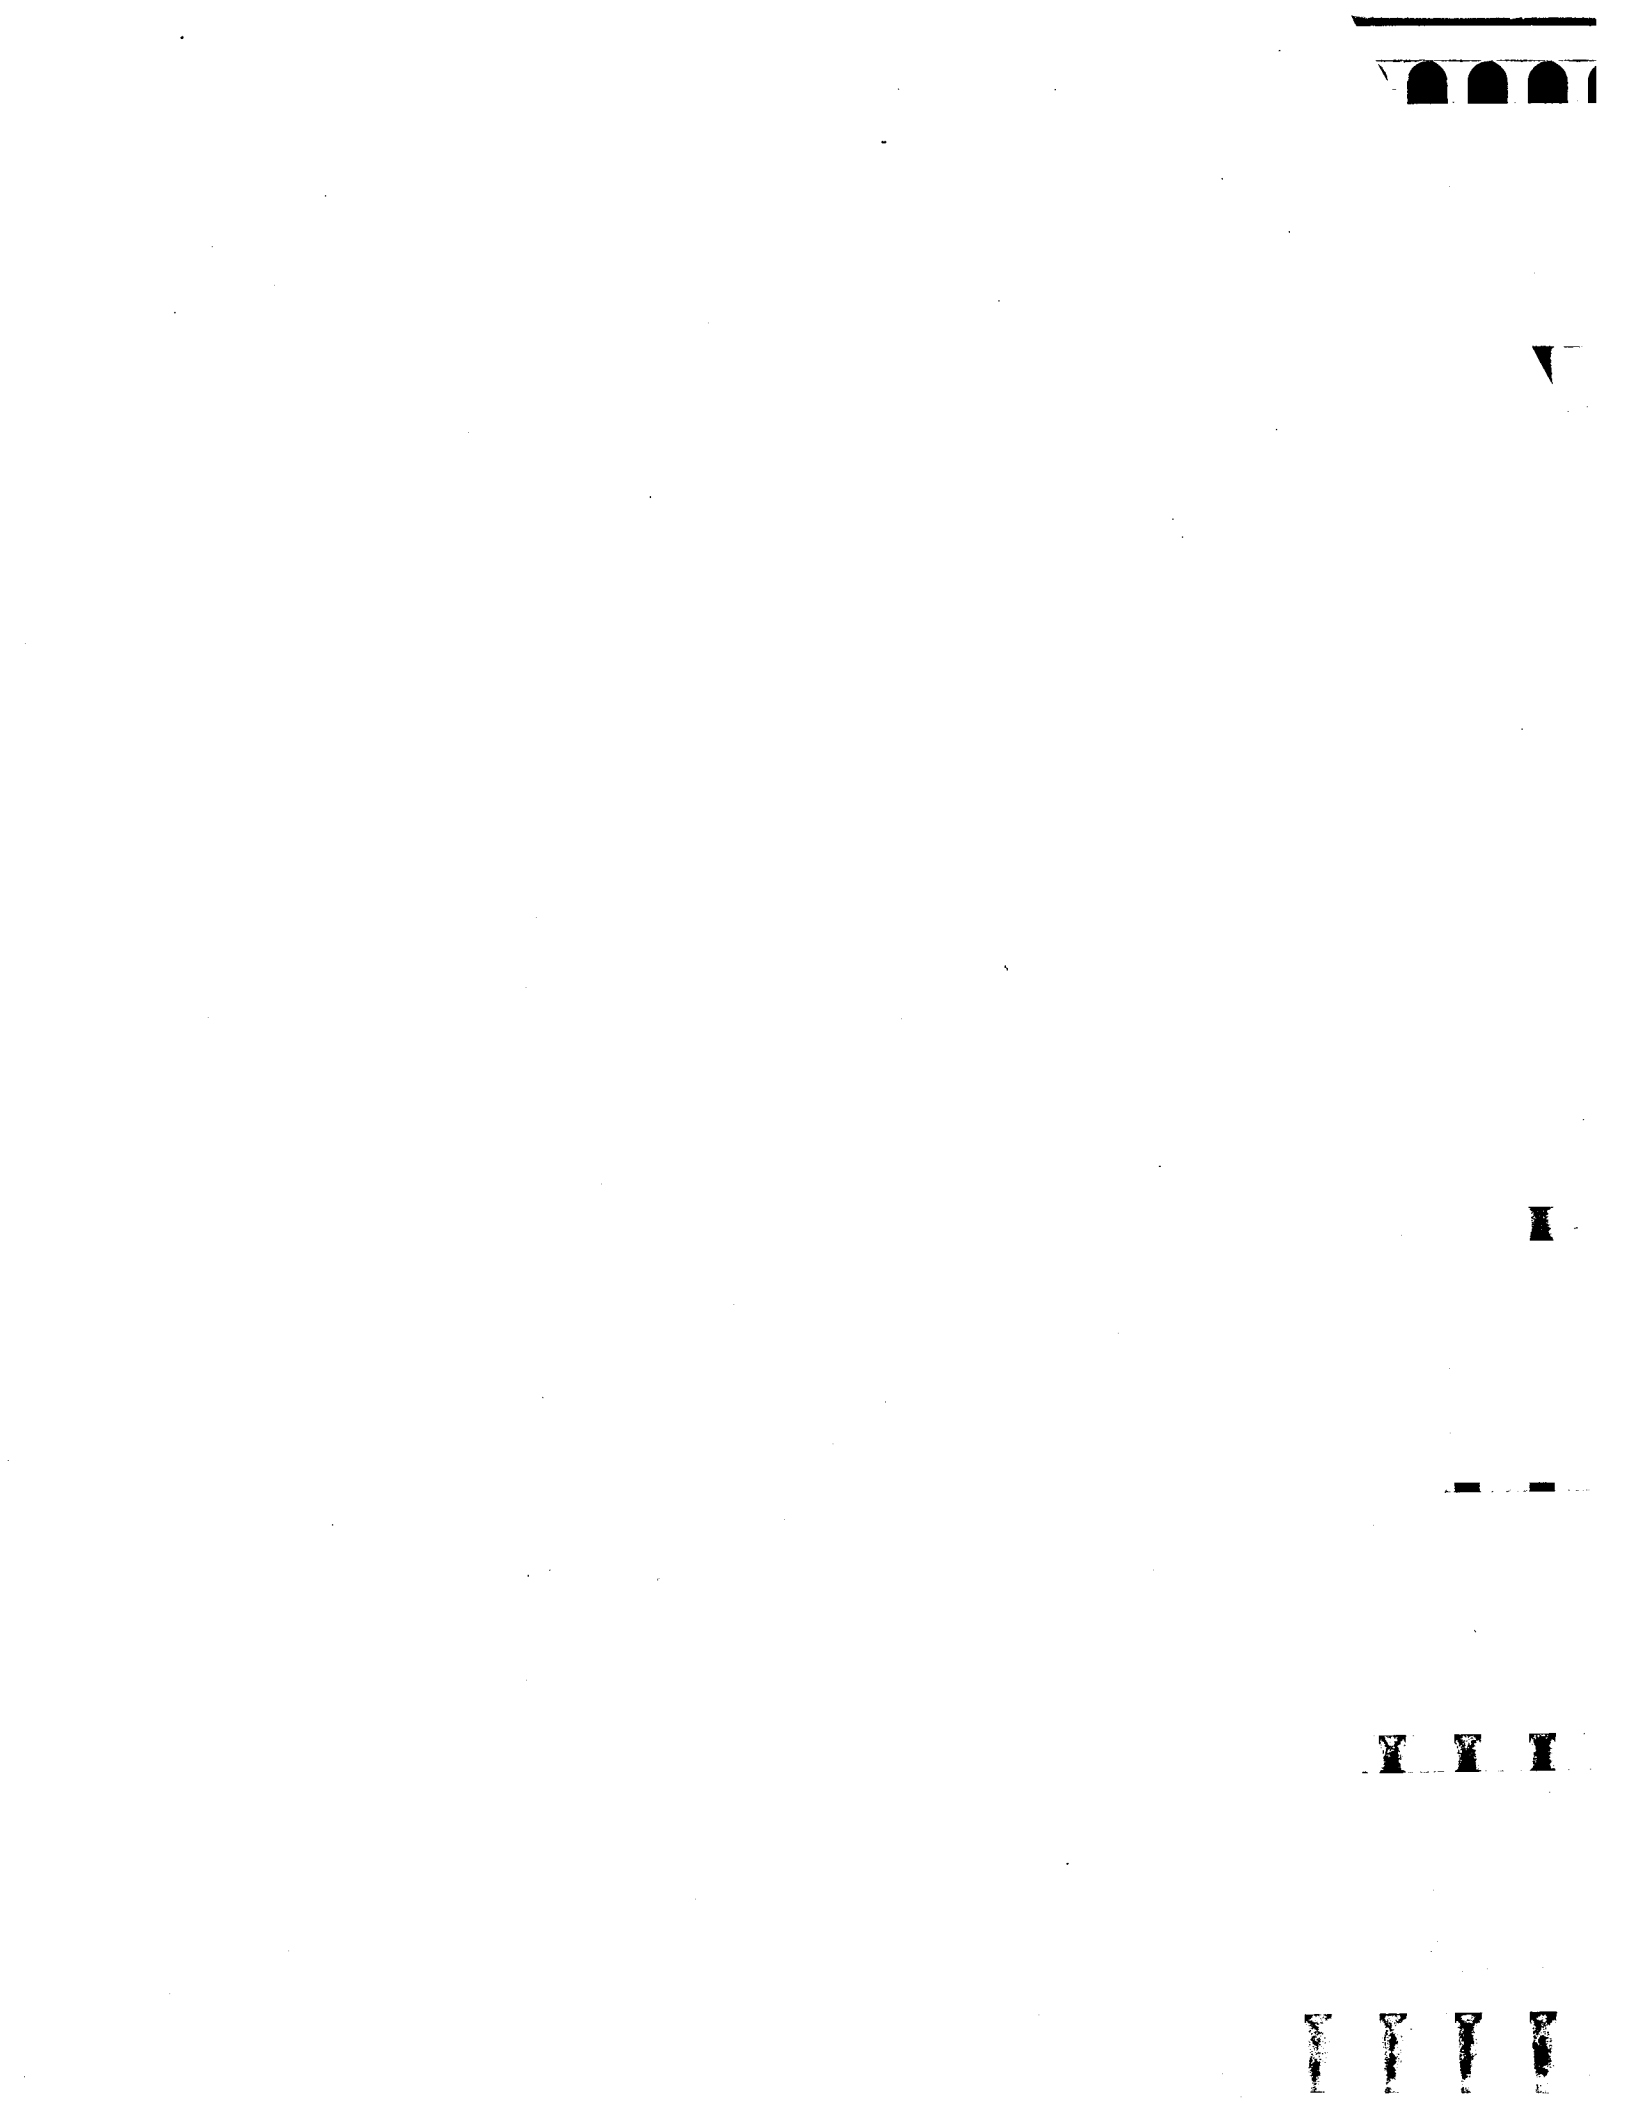

2.0. INTRODUCTION.

2.1. Background.

While NASA has always emphasized safety in design and operations, especially for crewed spacecraft, the Challenger accident brought home the need for a systematic, quantitative, and defensible way to evaluate flight risks and to identify and prioritize the factors that contribute to them so they can be targeted for improvement. In the late 1980s successful experience in a variety of environments where technological risks are of concern led NASA Headquarters to adopt probabilistic risk assessment methods as one answer to this need. The current risk assessment of the Space Shuttle vehicle is the latest and largest of a series of NASA-sponsored probabilistic risk analyses that began with the PRA Proof of Concept studies in 1987 and the analysis of catastrophic failure frequency for the Shuttle mission that launched Galileo in 1989.

2.1.1. History of this PRA.

The current risk assessment is part of an integrated Space Shuttle flight PRA program that started in mid-1993 under the sponsorship of the NASA Headquarters Safety and Mission Assurance Office and has since been adopted by the Office of Space Flight. It is structured as a three-phase program comprising the following projects:

Phase 1: an update of the 1989 Independent Assessment of Shuttle Accident Scenario Probabilities for the *Galileo* Mission (hereafter called the "*Galileo* RTG risk assessment" in the interest of brevity);

Phase 2: a risk assessment of the Space Shuttle Main Engine (SSME); and

Phase 3: an integrated launch-to-landing flight risk assessment.

The current analysis constitutes Phase 3. Phase 1 was completed in 1993. Under Phase 2, the NASA-SAIC team completed a human reliability analysis of the SSME Controller software build and load process and a risk assessment of the main combustion chamber of the SSME. The earlier phases are described in detail in the reports cited in Appendix C.

2.1.2. Project Organization.

The Space Shuttle PRA has been a true multi-disciplinary, multi-organization undertaking. The core risk analysis team was composed of representatives of the Shuttle Program Office at NASA Johnson Space Center (JSC), SAIC Advanced Technology Division, and SAIC's subcontractors Safety Factor Associates and Empresarios Agrupados. The PRA has also taken advantage of the talents and experience of design engineering, reliability engineering, risk assessment, and operations experts from JSC, Marshall Space Flight Center (MSFC), NASA Headquarters, and NASA's contractors Rocketdyne Division of Rockwell International, Thiokol Corporation, US Boosters, Inc., and Loral Space Information Systems. Figure 2.1 is an organization chart showing the principal organizations and key personnel involved in the project.

Figure 2.1. Shuttle PRA Project Organization.

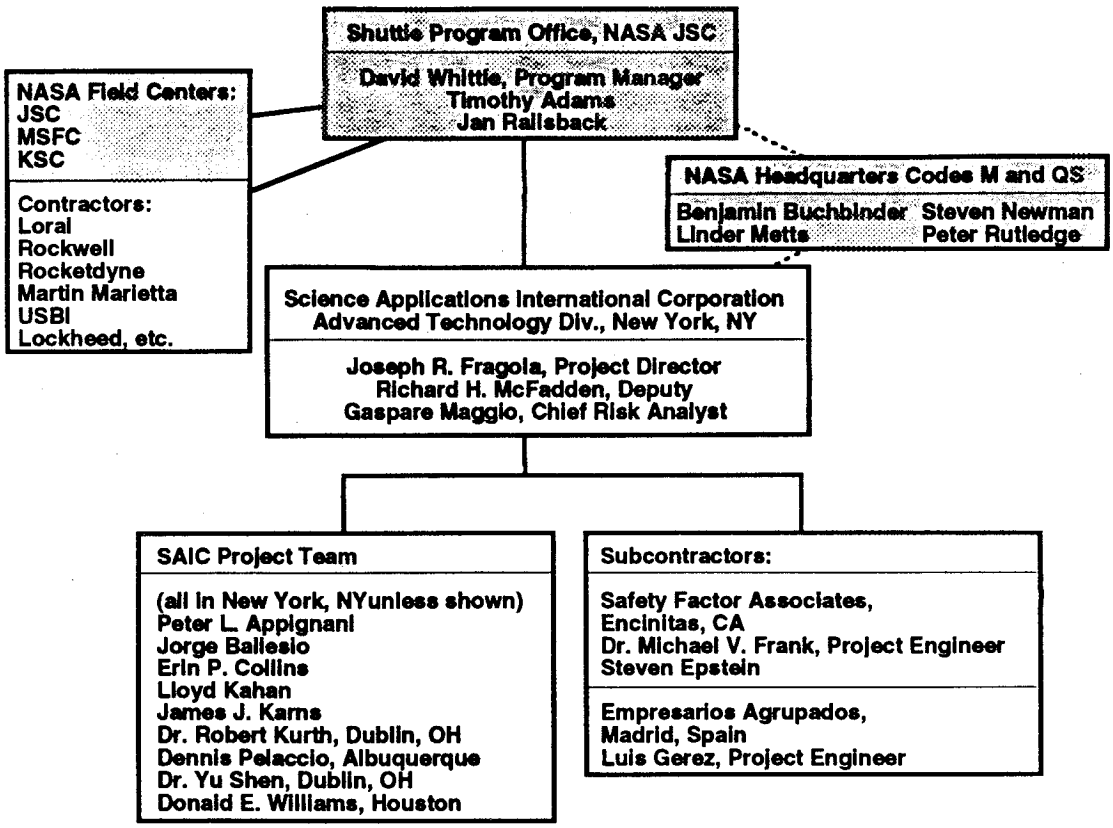

2.1.3. Relationships with Previous Risk Assessments.

In the interest of cost-effectiveness, the current PRA builds on the work done in a variety of earlier risk assessments. Figure 2.2 summarizes the connections among these analyses.

2.2. Objectives and Principal Products.

The primary objective of this project was to support management and engineering decision-making with respect to the Shuttle program by producing...

(1) a quantitative probabilistic risk model of the Space Shuttle during flight,

(2) a quantitative assessment of in-flight safety risk,

(3) an identification and prioritization of the features of design and operations that principally contribute to in-flight safety risk, and

(4) a mechanism for risk-based evaluation of proposed modifications to the Shuttle system.

The secondary, longer-term objectives were to provide a vehicle for introducing and transferring PRA technology to the NASA community, and to demonstrate the value of PRA by applying it beneficially to a real program of great international importance.

Figure 2.2. Relationships Between the Shuttle PRA and Previous Shuttle Risk Assessments

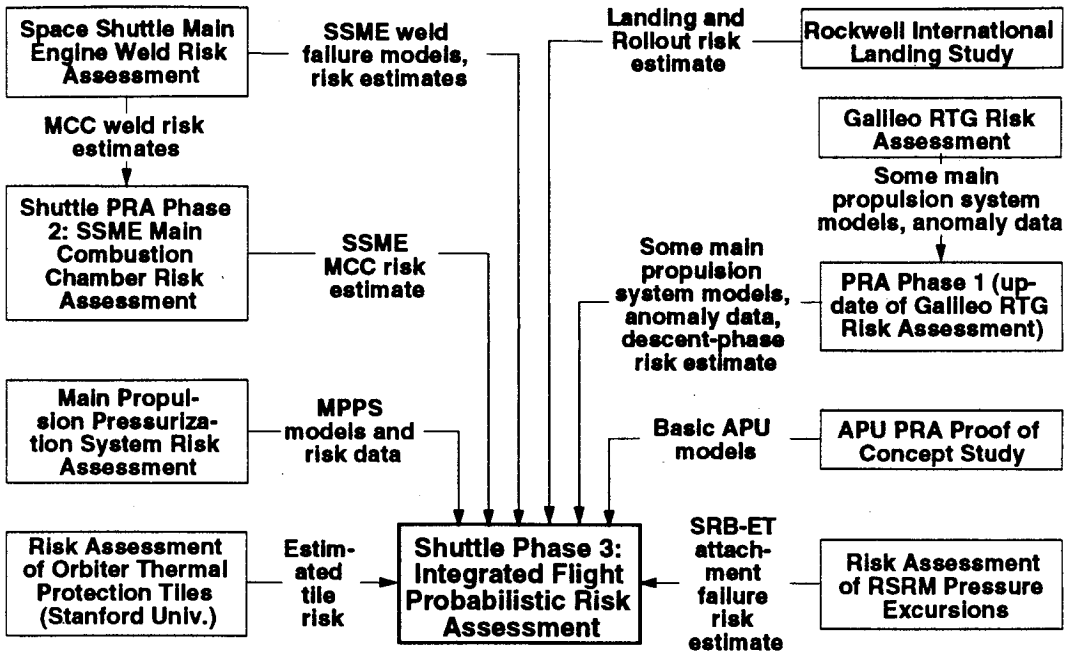

2.3. Ground Rules of the PRA.

In the interests of limiting the scope of the analysis to match the resources available and of maximizing cost-effectiveness by focusing on the major risk drivers while de-emphasizing secondary risk issues, the risk assessment team established the ground rules and simplifying assumptions stated below.

This PRA does not — and does not purport to — capture all of the risk of a Shuttle mission under the ground rules described below. For instance, pre-flight configuration errors[1] and mishaps during proceduralized aborts, were not within the current scope of the PRA. Nevertheless the PRA team is confident that this analysis both captures the majority of the total risk, and ranks the principal contributors to risk correctly within the bounds of uncertainty established by the available data and phenomenological analyses.

2.3.1. Nominal Shuttle Mission.

The principal events making up the standard or "nominal" Shuttle mission on which the risk assessment is based are shown schematically in Figure 2.3 on the next page. Table 2.1 on page 2.5 summarizes the characteristics of the nominal mission.

---

[1]Passive configuration anomalies such as undetected cracks in systems or undetected breaches in pressure boundaries are considered

Figure 2.3. Segments of Nominal Shuttle Mission Considered in PRA, Showing Mission Phase Boundaries.

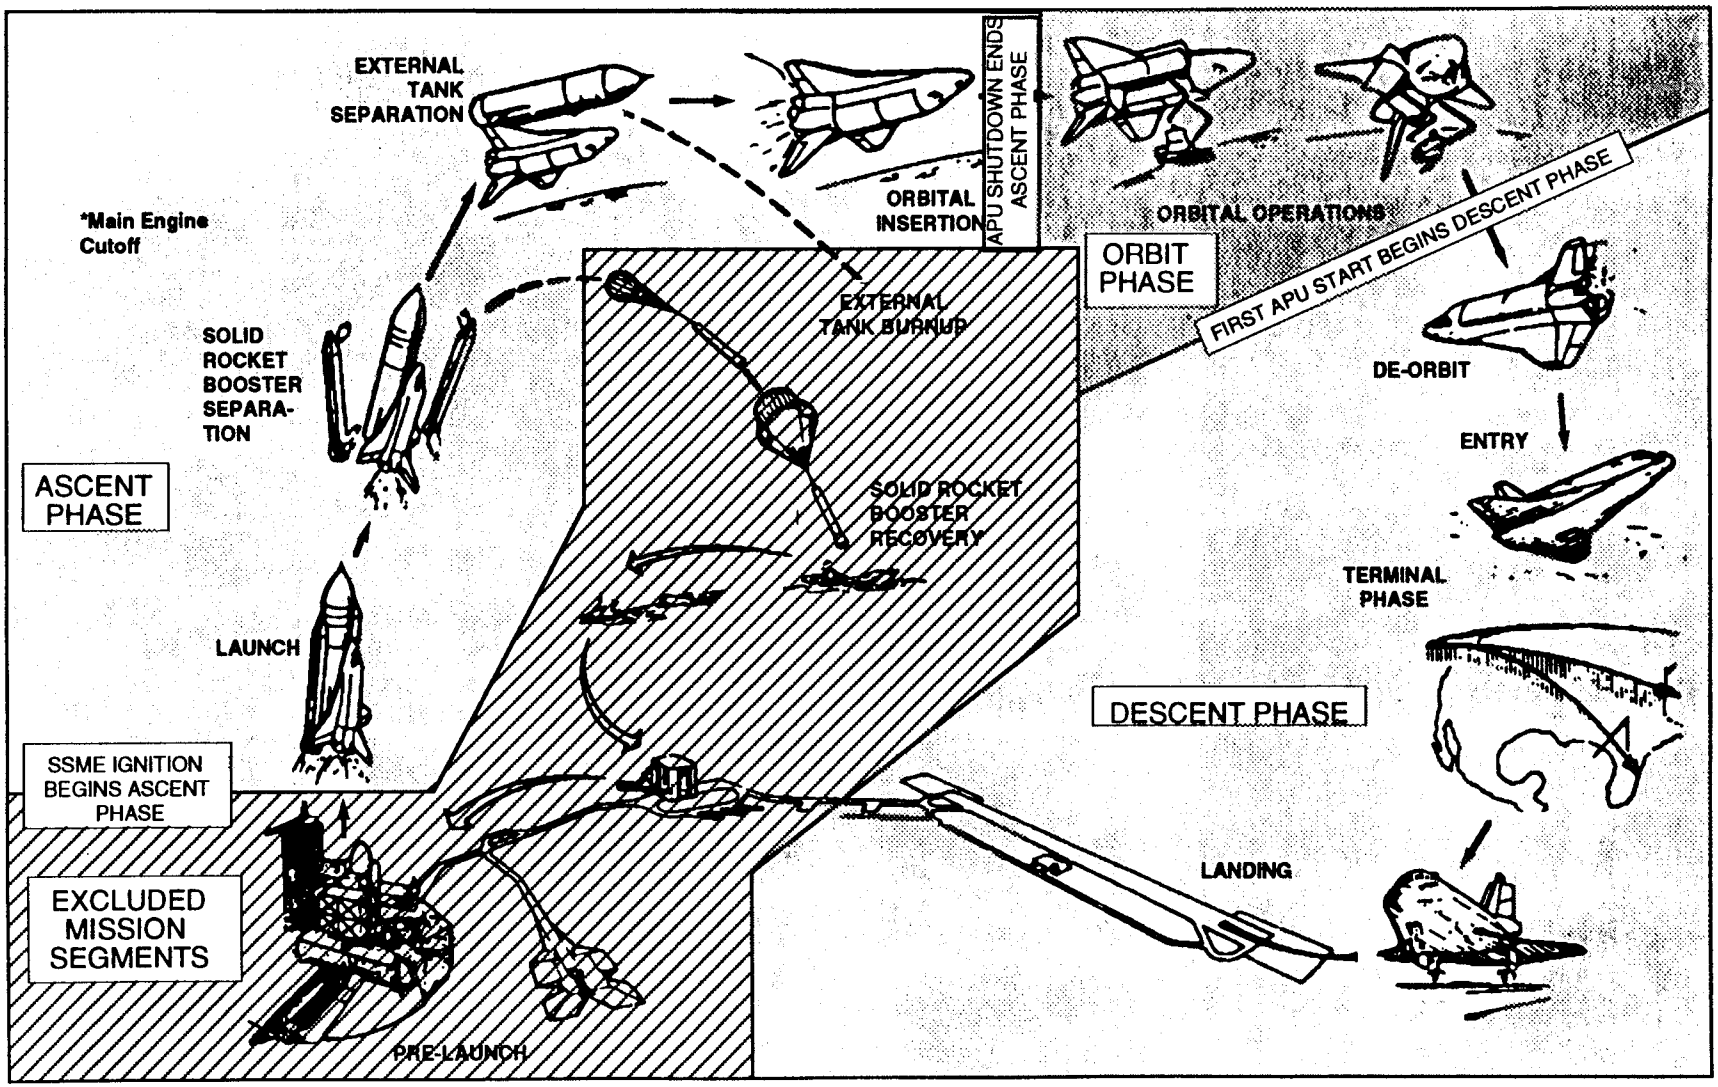

Table 2.1. Characteristics of the Nominal Shuttle Mission Considered in the PRA.	
Characteristic	Specifications
Launch location	Kennedy Space Center
Vehicle configuration at launch	Orbiter, external tank, two solid rocket boosters (SRBs) with Redesigned Solid Rocket Motors (RSRMs)
Vehicle mass at launch	Approximately 4.5E6 lbm; orbiter approximately 2.3E5 lbm including payload
Maximum total thrust during launch	7.78E6 lbf
Maximum main engine power during launch (percent of nominal)	104%
Time of maximum dynamic pressure ("Max Q")	t=30 to 60 sec (time after liftoff)
Time of SRB separation	t=120 sec
Main engine cutoff and external tank separation:	
Time	t=510 sec
Altitude	59 nm
Velocity	25,600 ft/sec
Orbital insertion	Direct, one Orbital Maneuvering System (OMS) burn
Launch azimuth	90°
Orbit:	
Altitude	100-312 nm
Inclination	28.5°
Mission duration	4-16 days
Entry interface:	
Altitude	400,000 ft
Distance from landing site	4200 nm
Velocity	25,000 ft/sec
interface:	
Altitude	83,000 ft
Distance from landing site	52 nm
Velocity	2500 ft/sec
Landing:	
Location	Kennedy Space Center
Final sink rate	3 ft/sec
Speed at touchdown	195-205 KEAS depending on weight
Touchdown	2500 ft past threshold
Source of data: Shuttle Crew Operations Manual, SCOM 1.0, November 1991, Section 1.1.	

Table 2.2. Boundary Events of Mission Phases.		
Phase	Beginning Event	Ending Event
Ascent	SSME start (ignition verified, closed-loop control activated)	Auxiliary power units (APU) are shutdown following orbit insertion
Orbit	Post-APU shutdown	First auxiliary power unit (APU) started for descent
Descent	Post-first APU start	Wheel stop after landing

2.3.2. Definition of Mission Phases.

To avoid having to carry models of systems that are inactive or no longer present through the entire risk model, the nominal mission is divided into three phases referred to as *ascent*, *orbit*, and *descent*, and the three phases are represented by distinct but interrelated sub-models within the integrated risk model. The phases are defined by the boundary events listed in Table 2.2.

2.3.3. Scope of the Risk Assessment.

The scope of the risk assessment includes only accident scenarios initiated within the Shuttle vehicle and leading to loss of the vehicle during the mission phases from main engine ignition at launch through wheel stop on landing.

Initiating events occurring before SSME ignition (e.g., mishaps in fueling) or after wheel stop on landing are excluded from consideration, as are initiators occurring outside the Shuttle vehicle (e.g., ground support equipment failures or accidents, and external events such as severe weather, launch pad fire, or collision of the Shuttle with another aircraft). This ground rule is intended to limit the analysis to a tractable scope while ensuring that the dominant risk-driving factors that are within the control of the Shuttle program are included.

2.3.4. Assumed Vehicle Configuration.

2.3.4.1. Vehicle and Operations Modifications.

Shuttle vehicles are assumed to be configured and operated as they were at the beginning of the analysis in early January 1994; neither actual nor planned changes in design or operating practices subsequent to that cutoff date are considered. This ground rule is intended to establish a baseline system configuration in order to avoid misleading risk comparisons among unlike systems. (The risk models are designed to be flexible enough to accommodate such modifications later.)

2.3.4.2. Vehicle Configuration Errors.

Shuttle vehicles are assumed to be in the "as designed" configuration at main engine ignition. That is, pre-launch configuration errors such as setting control switches erroneously or installing the wrong hardware components or an incorrect version of the engine control software are not considered in this analysis.

2.3.4.3. Orbiter-to-Orbiter Differences.

The risk implications, if any, of differences among the Orbiters in the fleet (e.g., the presence or absence of modifications for long-duration-missions) are not considered in the current PRA.

2.3.5. Mission Abort Scenarios.

All event scenarios that lead to a mission abort for which there is an established procedure in the Flight Data File are terminated at the end state that triggers the abort. That is, the current PRA considers the conditional likelihood (i.e. the probability of requiring an abort multiplied by the probability that it would be unsuccessful) that the vehicle will fail to survive proceduralized aborts as second order and conditional risk and therefore this abort risk has been considered out of the present scope. However the current model would allow for the abort risk to be readily incorporated at some future time.

2.3.6. Final Approach & Landing Risk.

Preliminary screening risk assessments indicate that the loss-of-vehicle risk during the landing phase of the mission is a relatively small — although still potentially significant — contributor to total flight risk. However, assessing this risk through detailed modeling would be very resource intensive because a significant human factors analysis effort and human reliability analysis would have been needed to evaluate the risk effects of critical flight crew actions. To conserve resources for the modeling of event sequences of far greater risk importance in the necessary detail, the landing phase is represented by a risk estimate of 3.0% of the base-case ascent-phase risk developed by Phase 1 of the PRA (the update of the *Galileo* RTG risk assessment), rather than by a detailed risk model. This estimate is the consensus of a number of experts from the NASA Headquarters Office of Space Flight, Johnson Space Center, SAIC, and Safety Factor Associates and is consistent with the judgements made in an allied study performed by Rockwell.

2.3.7. Data Window.

Relevant Shuttle flight and test experience data from the beginning of the Shuttle program though Mission 67 in July 1994 are used in the risk data analysis in order to estimate the frequencies of initiating and pivotal events, although in some cases early data is discounted[2] to allow for reliability growth. In addition to Shuttle data, the analysis for the Redesigned Solid Rocket Motor incorporates some relevant experience data from other large solid-fuel rocket programs. Refer to Section 5.0 for details.

2.4. Overview of the Risk Analysis.

This section provides an overview of the technical approach used in the Shuttle PRA. Sections 3.0 and 4.0 provide additional details on risk modeling methods and data analysis respectively.

2.4.1. The Role of Scenarios (Event Sequences) in PRA.

The fundamental approach used in this Shuttle risk assessment is that of scenario-based probabilistic risk assessment. The concept of scenarios is basic to any understanding of the PRA process. As the name implies, a scenario is simply the chronological "story" of a sequence of events that is triggered by some incident and proceeds through intervening events to an end state. (In fact, a scenario is often called an "event sequence," and if it deals with an accident, an "accident sequence." However in this risk assessment an event sequence is a particular path through the event tree and therefore each scenario has associated with it a set of event sequences all stemming from the same initiator.)

---

[2]Here the word discounted is used in the economics sense meaning "reduced in impact" not in the common sense meaning of "disregarded".

Figure 2.4. A Generic Accident Scenario (or Sequence) for Probabilistic Risk Assessment.

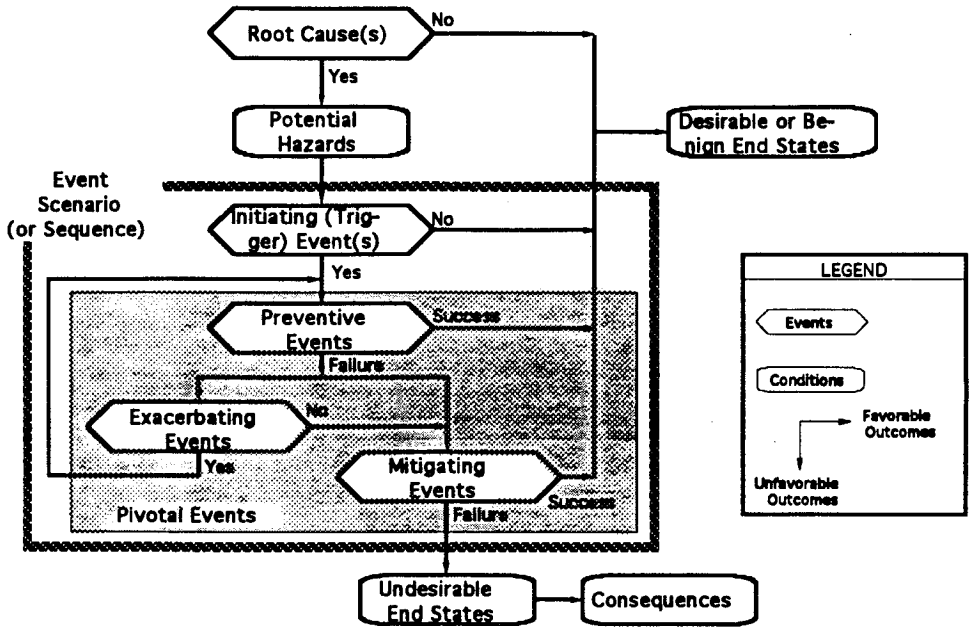

Figure 2.4 depicts an accident scenario in the most generic form, including some of the terminology used to describe the elements of scenarios. The key terms are (1) *initiating events* (or trigger events), which — in conjunction with pre-existing potential hazards — begin the scenario; (2) *pivotal events*, which have the potential to change the course of the scenario, and can have preventive, exacerbating, or mitigating effects; and (3) *end states*, which can have desirable, benign, or unfavorable consequences. In all but the simplest systems, there are several alternative sequences of events that can follow an initiator, depending on the outcomes of the intervening events; each such path is considered a part of the associated with the particular initiating event or binned (i.e. grouped) initiating event set. PRA is simply a systematic technique to evaluate the probabilities and consequences of the various scenarios that can occur in a process or system as well as their associated uncertainties.

2.4.2. Overview of the PRA Process.

How does one perform a PRA? A comprehensive probabilistic risk assessment of a complex technological system such as the Space Shuttle vehicle comprises the generic tasks described below and laid out graphically in Figure 2.5. Of course, the objectives of the risk assessment, the characteristics of the system under consideration, and the resources available determine the specific strategy of the analysis and intensity with which each of these tasks is attacked. Some tasks may even be eliminated entirely if not needed. In the current Shuttle PRA, for example, it was not necessary to perform a human reliability analysis because preliminary screening of the event scenarios showed that human-mediated events would be risk-significant only during the descent phase, which (by ground rule) was represented by an expert-judgement-based risk estimate rather than modeled in

Figure 2.5. Task Network for a Generic Probabilistic Risk Assessment.

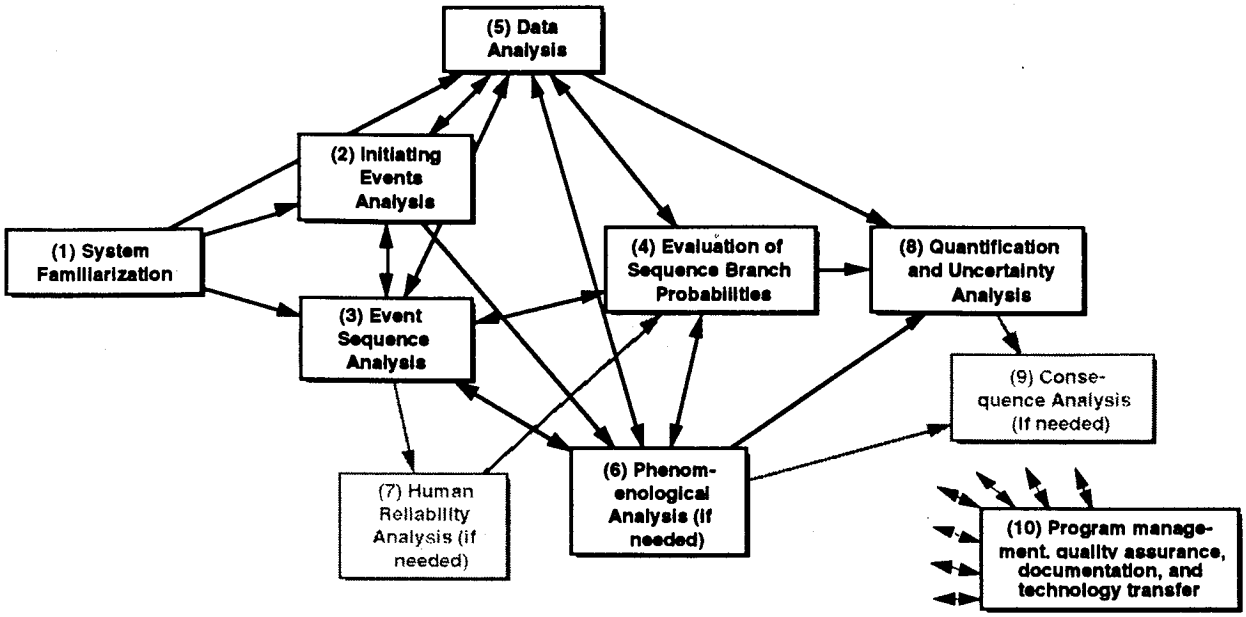

detail. Nor was a consequence analysis necessary (in the conventional sense of a quantitative evaluation of the consequences of accident sequence end states) because only one consequence — loss of vehicle — was of interest. On the other hand, the system familiarization, initiating events analysis, event sequence analysis, and data analysis tasks required intensive work. All of the tasks which make up a comprehensive PRA are described below in the interest of completeness, but the tasks or subtasks which did not need to be performed in the current Shuttle PRA are denoted by gray type in the task descriptions and in Figure 2.5.

Similarly, the requirements of a specific risk assessment — and practical considerations such as staggered availability of required data and analytical resources — ordinarily dictate that some tasks be broken up into subtasks that are worked separately. In the current Shuttle PRA, the initiating events, event sequence, and risk data analysis tasks for each of the major risk-driving systems were performed as distinct subtasks, and the resulting models then integrated into a full-mission risk model.

Note that many of the PRA tasks are mutually interactive, as denoted by double-headed arrows in Figure 2.5. For instance, the event sequence modeling task defines the desired outputs of the data analysis task, but the availability of acceptable data at acceptable cost ultimately determines what can practically be modeled. Also, tasks are often performed iteratively. For example, typically there are several iterations of event sequence modeling, beginning with preliminary, top-level accident sequence models (based on worst-case assumptions and conservative event frequency estimates) that determine which sequences probably will turn out to be significant contributors to the total risk and thus need to be modeled in greater detail, and then proceeding through several stages of refinement.

(1) Systems Familiarization. The PRA team becomes thoroughly familiar with the design, operation, and environment of the systems under consideration.

2) Initiating Events Analysis. Using hazard analyses and other relevant analyses such as FMEAs (preferably already existing as in the case of the Shuttle), the analysts define all of the credible mishap-initiating events, including (where applicable) spontaneous equipment failures, external events, and human errors by operators or maintenance technicians, and integrate the results into an initiating events list and a master logic diagram. The master logic diagram appears as the top-level of the integrated fault tree model (Appendix A) and in fact may be separated and considered as a type of reduced fault tree that incorporates all of the initiating and exacerbating events that can lead to the accident end state of interest while neglecting preventive and mitigating events.

(3) Event Sequence Analysis. The sequences of system responses to the various initiating events (i.e, the event sequences), including both automatic and (where applicable) human-mediated functions, are defined. Event sequence models for the sequences of events which lead to the accidents of primary concern, usually in the form of functional event sequence diagrams (elaborations on the generic event sequence diagram in Figure 2.4) are created, and then event trees (specialized decision trees that are logically equivalent to event sequence diagrams but easier for computer-based tools to analyze) are built using personal-computer-based analytical tools such as ETA™ and NASA's EC-TREE.

(4) Evaluation of Initiating and Pivotal Event Probabilities. Using fault trees or other applicable reliability analysis techniques in conjunction with the results of the risk data analysis of task (5), the analysts assign probabilities of occurrence to the initiating events identified in task (2), and success or failure probabilities to the pivotal events of the event sequences developed in task (3) (i.e., to the branch points in the event trees). Dependent (common-cause and cascading) failure effects (i.e., situations where a single condition could disable several nominally independent, redundant subsystems or one failure leads to others) are evaluated and incorporated into the fault tree models. This task utilizes computer analytical tools such as CAFTA™ and RBDA™.

(5) Data Analysis. Using computer data analysis and aggregation tools such as CARP™, the analytical team assembles a risk data base in order to assign numerical probabilities to the various initiating events, system responses, and event sequence end-state consequences. For a risk assessment on a hardware system such as the Shuttle, this data base contains the time failure rates and/or failure-on-demand probabilities of the risk-critical components of the system, with uncertainty bounds for at least those components whose failure rates/probabilities are significant to the outcome. The risk data base is preferably developed from the operating experience of the system being assessed, but if system-specific experience data is insufficient, the analysts use "surrogate" data derived from the experience of similar components in analogous applications elsewhere, analytical or test results, applicable generic data sets, or expert judgement. It is essential to coordinate this task closely with the initiating events, event sequence, and fault tree analyses in tasks (2)-(4) respectively to ensure consistency.

(6) Phenomenological Analyses. Any phenomenological analyses needed to support the probabilistic analysis are performed, usually by design and systems engineers supporting the PRA team. An engineering analysis to determine the minimum acceptable performance ("success criteria") of a risk-critical system or component is an example.

(7) Human Reliability Analysis. The analysts estimate the probabilities of failure or success of human-mediated events within the event sequences, using human reliability analysis techniques and computer tools such as ORCA™, and incorporate them into the event tree models of the accident sequences which have human activities as major contributing factors. Depending on the specifics of the analysis, this task may involve estimating the probabilities of human-mediated initiators and the likelihood that human intervention will successfully interrupt event sequences. (As noted above, this task was not within the scope of the current the Shuttle PRA.)

(8) Quantification and Uncertainty Analysis. The models are evaluated quantitatively in order to derive the probabilities of the event sequence end states of interest, and to prioritize the various initiating events and system responses in terms of their contribution to risk. The statistical and technical uncertainties associated with the input data and modeling assumptions are propagated through the models to evaluate the uncertainties of the end states. This task utilizes such computer tools as RMQS™, UNCERT™, and the uncertainty distribution propagation function of EC-TREE.

(9) Consequence Analysis. If necessary, quantitatively evaluate the consequences associated with the risk assessment end states. In a probabilistic safety assessment, these would normally be the consequences of significant accident conditions, such as the potential public health hazards from off-site release of toxic materials or the costs of facility damage and lost production. (In the Shuttle PRA only one consequence — loss of the Shuttle vehicle — was of concern, so this analysis was not required.)

(10) Program Management, Documentation, Quality Assurance, and Technology Transfer. In this task the project is managed technically and administratively; technical reports and briefing materials are prepared; engineering quality assurance and configuration management are maintained; and — perhaps most important — the data, models, results, and supporting technical information developed during the PRA are organized into traceable archival records. In combination, the final technical report and the archival records allow current users to understand the risk assessment and how to use it appropriately, and future users to update the data base and modify the risk models to accommodate accumulating operating experience and changes in equipment and operations.

Figure 2.6 is a top-level flow chart showing how the various risk models, information sources, and data streams involved in the Shuttle PRA fit together. (The less important data sources and the detailed structure of the models have been omitted in the interest of readability.)

Figure 2.6. Flow Chart of Shuttle PRA Data and Models.

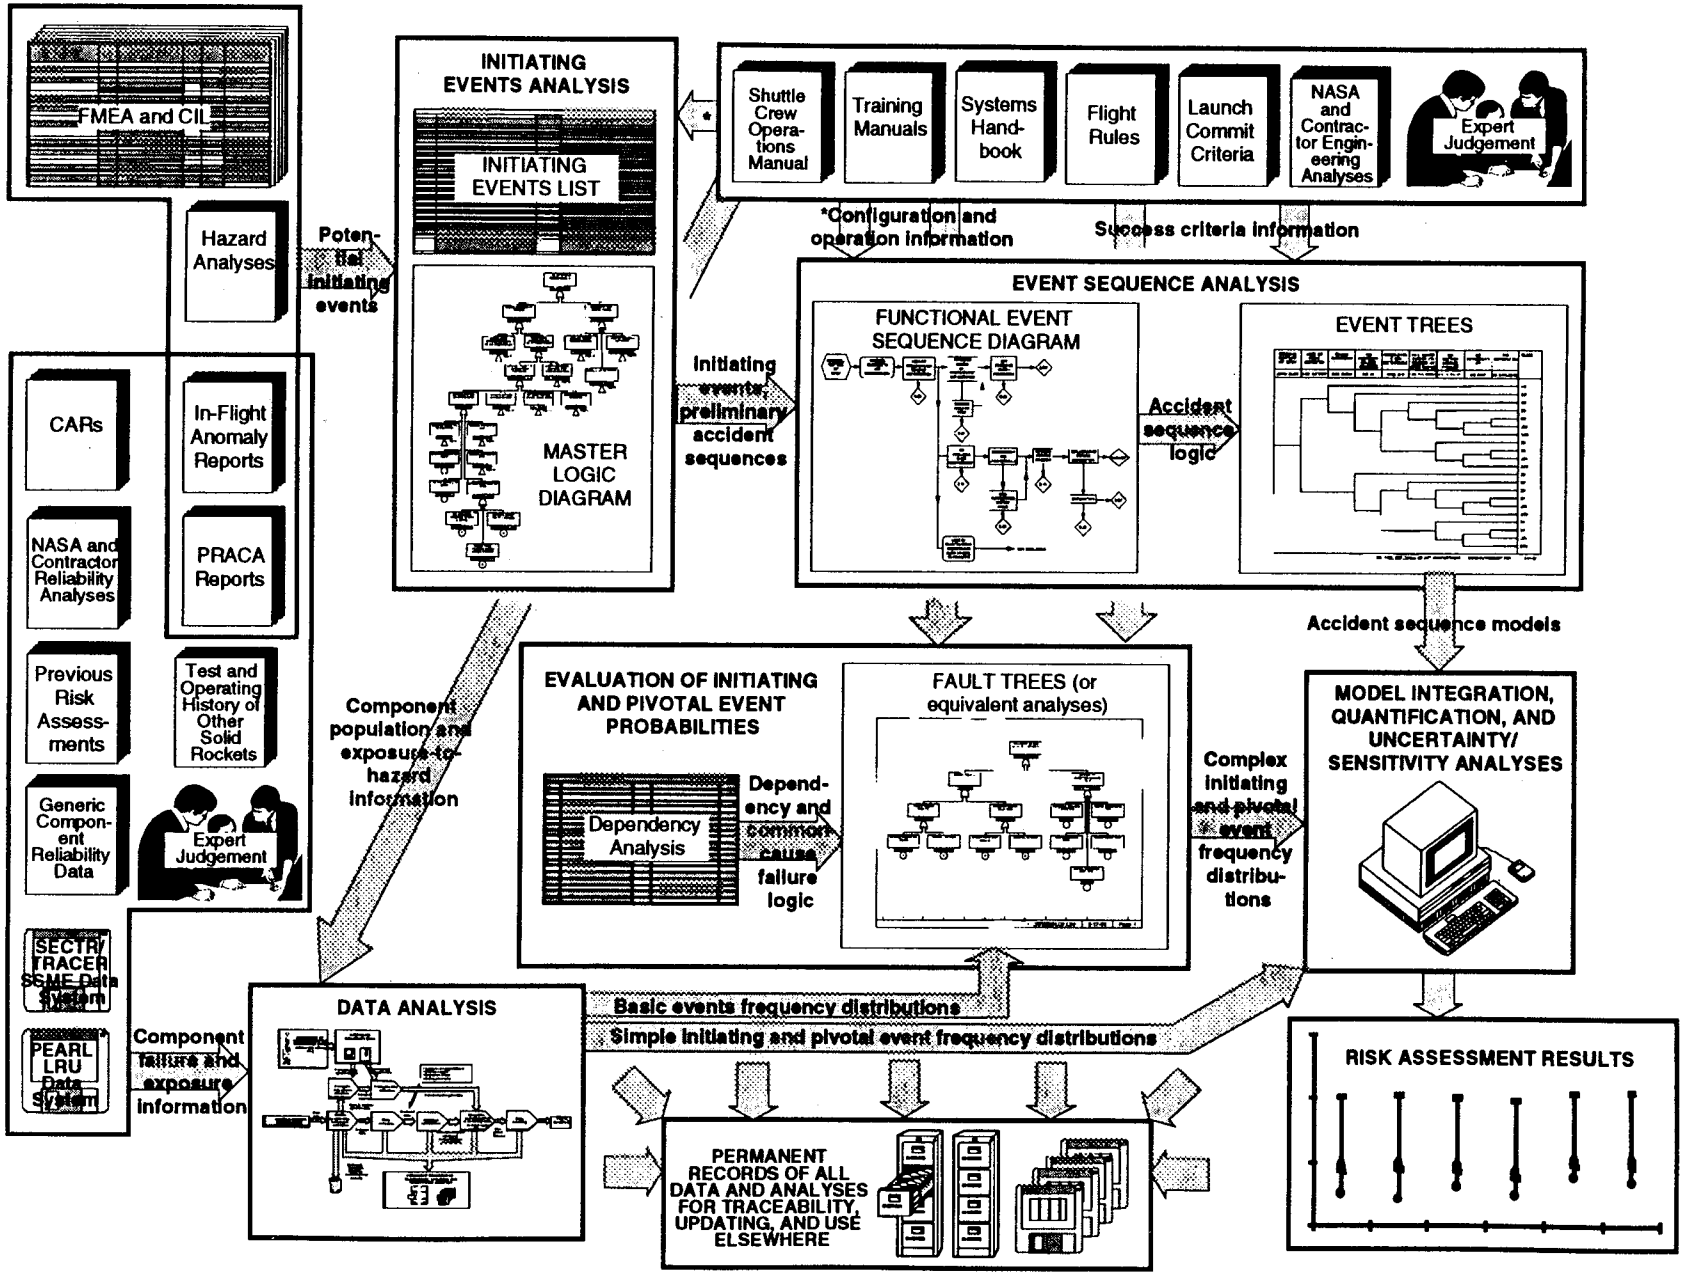

3.0 RISK MODELING

3.1. Overview of Modeling Approach

The basis of PRA is the development of scenarios. Scenarios may be thought of as strings of events which lead to consequences which are undesired. Each scenario begins with a set of "trigger events", sometimes called an initiating event category, and ends with an end state, sometimes called a consequence. A trigger event is any abnormality, malfunction, or failure (whether it be human, hardware, software, process, etc.) that causes a deviation from desired operation. In this assessment, for example, one end state of interest is LOV. End states are defined by the decision-maker. (What is the quantity of interest to the decision-maker?) In between trigger events and end states are "pivotal" events which determine whether and how a trigger event propagates to an end state. Each scenario is defined by one trigger event (or alternately a class or bin of trigger events), one or more pivotal events, and one or more end states. Pivotal events may be protective, mitigative, aggravative, or benign. Scenarios, therefore, may be conceptually represented as follows:

Figure 3.1. Accident Sequence Schematic

Scenarios may be developed and documented by a variety of diagrammatic forms. One of the features of a probabilistic risk assessment process is that the exact diagrammatic form is not unique. Different analysts may select different forms to help themselves both to better develop and display the model. Part of the "art" or creativity in performing a probabilistic risk assessment is the selection of the diagrammatic forms that best aid in both the model development and model presentation functions. The specific set of diagrams chosen depends on the objectives and scope of the analysis as well as the audience for the results. Experience with many risk assessments helps the analyst make a good choice of diagrams.

In safety and reliability PRAs, the most popular presentation forms are event trees, fault trees, and event sequence diagrams. Human actions and software errors as well as hardware malfunctions, and physical and chemical process/phenomenological events should be included in the scenarios. Dependent events and common cause failures are often important to overall risk and are usually modeled in event trees and fault trees.

It is typical to depict an overview of the system response to an initiating event in a diagram called a functional event sequence diagram (FESD). An FESD represents scenarios in terms of initiating events, pivotal events, and damage states. Construction of an FESD makes use of an inductive

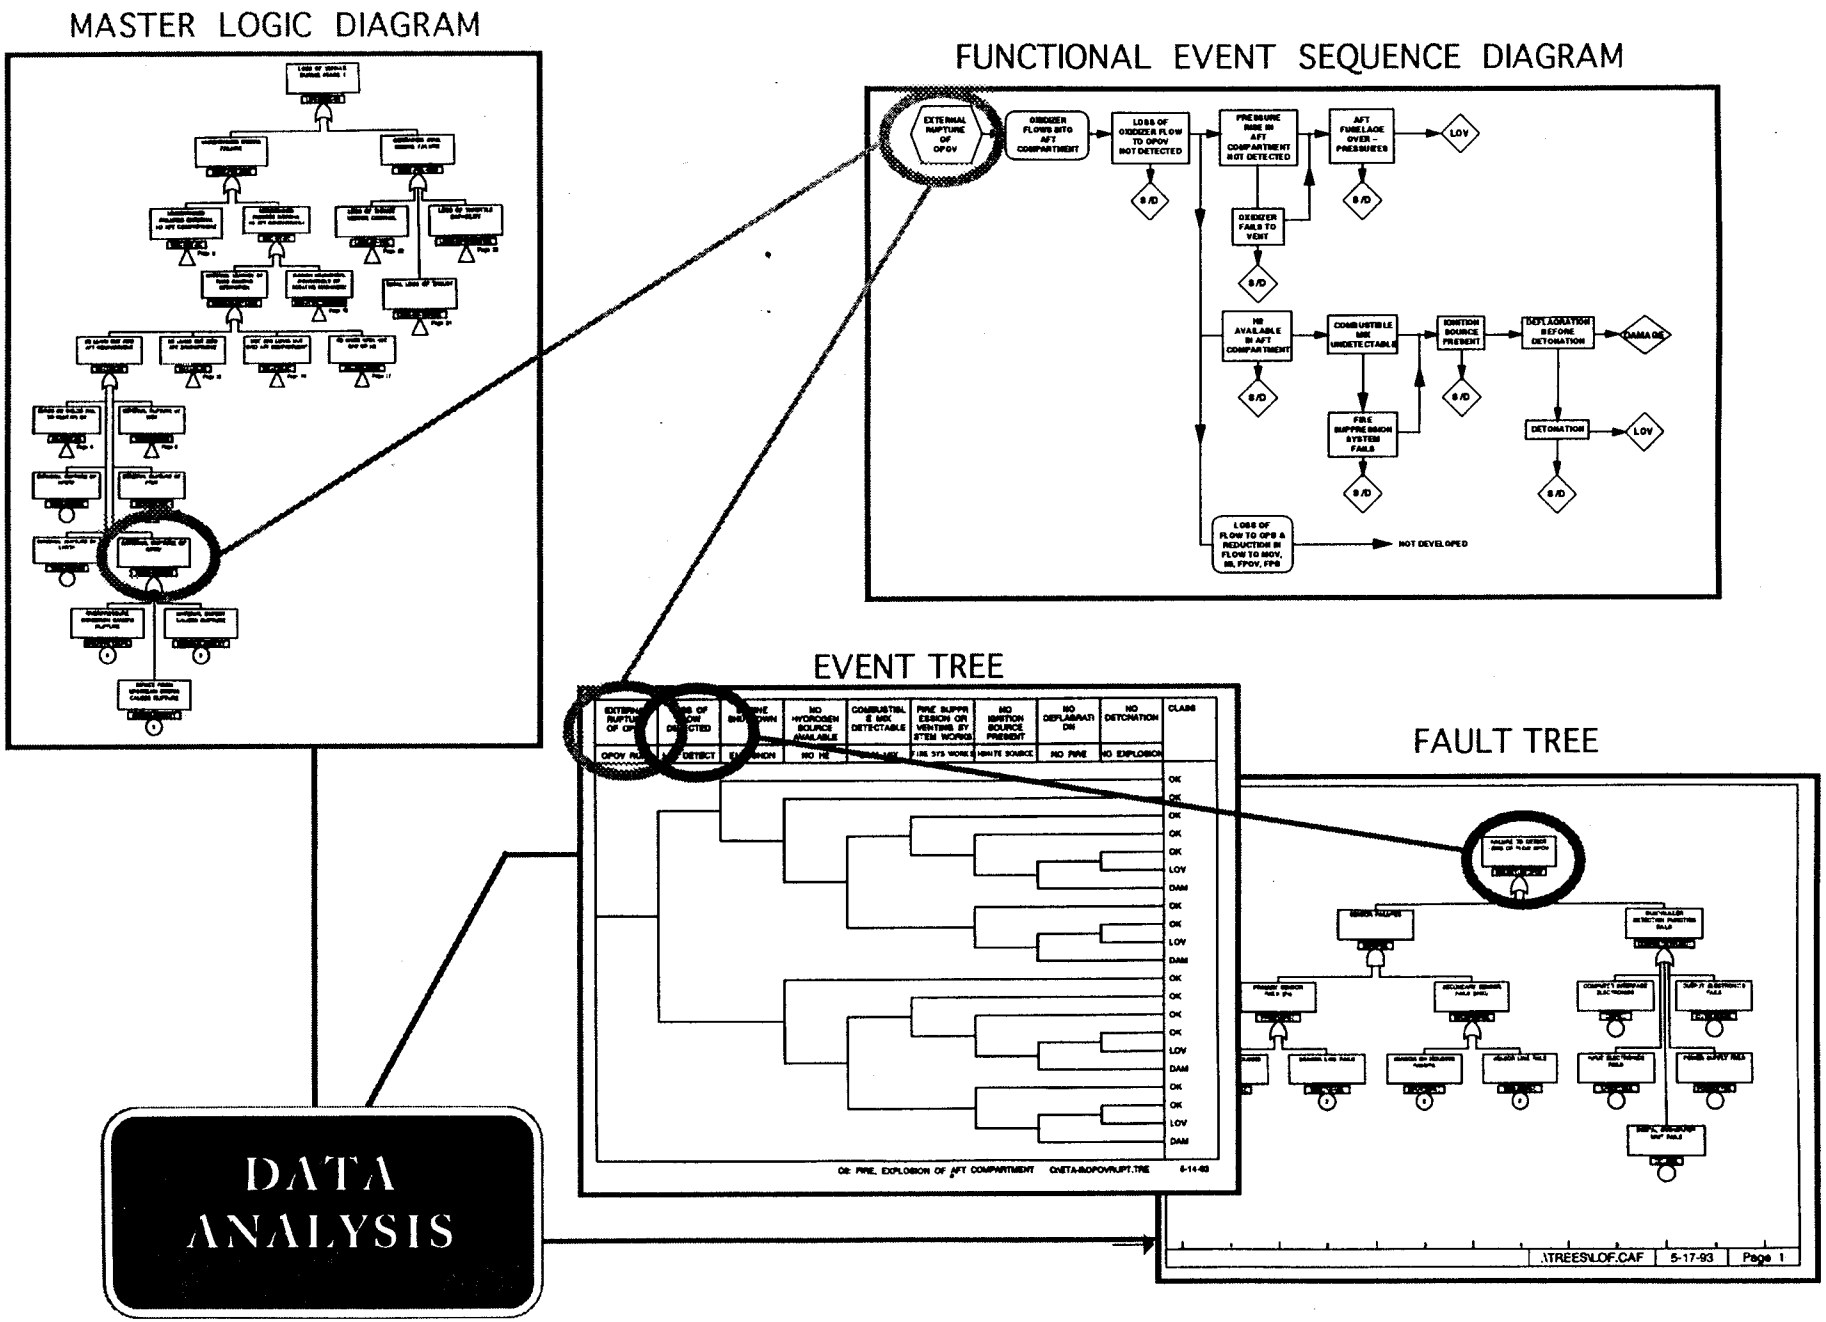

Figure 3.2. Shuttle Probabilistic Risk Assessment Modelling Mechanics

reasoning process. That is, after a trigger event is identified, the rest of the events are developed by asking and answering the question "What can happen next?". As shown in Figure 3.2, an FESD is a series of boxes with attached lines. The boxes are events that are constructed so that they can be considered to have binary outcomes (success/yes or failure/no). An FESD developed in this way has been found to been effective tool for capturing the knowledge of system experts. The scenarios of an event sequence diagram are usually converted to an event tree. An event tree is a decision tree that is also limited to binary outcomes for each event. It usually contains the same information as an FESD but is more amenable to computerized development of the needed algebraic equations. Each decision node in an event tree requires the establishment of an associated probability of occurrence. The boolean models used to develop such probabilities (if development is required) are often fault trees. Event Trees and Fault Trees are complementary techniques that are often used together. Together they map the system response from initiating event through damage states. Together they delineate the necessary and sufficient conditions for the occurrence of each damage state. They also form the basis of the algebraic equations that are ultimately used to obtain the frequency with uncertainties of the damage states. Construction of a fault tree is a deductive reasoning process which proceeds systematically to answer the question "How can the top event have occurred?". Thus, fault trees are often useful in developing the hierarchy of events. This development is often used to provide more resolution (or detail) to events of event trees to facilitate quantification. Because of the complementary nature of using both inductive and deductive reasoning processes, combining event trees and fault trees often produces a more complete, concise, and clearer development and documentation of scenarios than using either one exclusively.

3.2. Integrated Mission Risk Model

Event trees and fault tree models are developed at the sub-system level and then grouped into functional failure categories and integrated to obtain the overall Shuttle risk. Only in this way can meaningful comparisons be made between the risk of various systems.

3.2.1. Mission Master Logic Diagram

A master logic diagram is a convenient method for developing a set of initiating event categories that can be shown to be reasonably complete. It is a hierarchical depiction of ways in which system perturbations can occur. Completeness in attempting to predict all such perturbations in every detail is quite impractical. However, by a functional categorization of perturbations to the system that eventually leads down to a component characterization for each function, a team of analysts can usually capture all but the most indiscernible events. An MLD starts with a top event that is a damage state of interest (e.g., catastrophic failure of the engine). Events that are *necessary but not sufficient* to cause the top event are enumerated in ever more detail as lower levels of the hierarchy are built. Typically, the top levels are functional failures (e.g., failure of propulsion, failure to control etc.). The lower levels are subsystem and component failures that contribute to the functional failures. The diagram continues toward more detailed events as long as each

identified initiating event category has a different system response. Ultimately, a level of detail will be reached such that enumerated events have the same system response. Development of the diagram stops at the interface between these levels of detail.

3.2.2. Space Shuttle Top-level Functional Failures

The methodology applied in performing the PRA allows for the assignment of risk to individual system elements, components and individual failure modes. A master logic diagram (MLD) was developed to identify possible anomalous conditions which could lead to a loss of vehicle (LOV). The master logic diagram was developed from a top-level functional nature to specific functional failures which prompted different element responses although many failure modes could induce the initial anomaly. Five top-level Shuttle functional failures which have the potential to lead to LOV were identified; these are:

- *Failure to Provide Proper Propulsion*

A minimum amount of propulsion is necessary to maintain a favorable vehicle trajectory and assure the achievability of abort contingencies. The minimum amount of thrust needed from each propulsion system on the Shuttle will be discussed as "success criteria" in the appropriate sections.

- *Failure to Maintain Proper Vehicle Configuration*

Through out the mission there are configurational tolerances which must be maintained to assure a safe return of the vehicle. For example, during ascent rocket nozzles positions must be gimbaled to direct thrust such that an acceptable vehicle attitude is maintained. Attitude control is maintained by proper control surface configuration during re-entry. Untimely separation of elements or other gross configurational failures (e.g. activation of SRB recovery system before separation) are also considered under this heading.

- *Failure to Contain Energetic Gas or Debris*

In a high energy system such as the Shuttle there is always a possibility that energetic gases or high kinetic energy debris is not contained within its designated boundaries. This functional failure is of particular concern to systems such as propulsion and power generation. In most instances contributing failure modes are those designated as criticality 1 however some otherwise benign failure modes may also lead to such energetic discharges due to the failure of protective systems.

- *Failure to Maintain Orbiter Environment*

Certain pressures and temperatures must be maintained within the Orbiter to assure that both hardware and crew are able to perform there allotted functions. This is especially critical during re-entry where ambient temperatures may reach 2200°F.

- *Failure to Negotiate Adverse External Events*

Although careful planning is performed to assure that the Shuttle does not encounter any adverse environmental elements (e.g. lightening strikes, wind shear, exceptionally high energy space debris, etc.) the possibility of catastrophic consequences may still be possible though unlikely. This particular failure category was out of scope but is included for purposes of completeness.

One or some of these top-level functional failures may be associated with each of the Shuttle elements. Each of these element level functions must be supported by dedicated sub-systems. It is at this subsystem level that failure initiating events are defined. Failure initiating events are subsystem failures which elicit a similar corrective response from the system if one is possible at all. In some cases the system does not have an active response mode and in such cases pre-flight tests, inspections, and maintenance is depended upon to obviate such occurrences. In other words the initiating event or anomaly leads directly to LOV. Functional failures involving the containment of energetic gas or debris have little in the way of protective systems once the passive design features have been compromised. The adequacy of passive design features is critical to the safety of the Shuttle since little can be done after the Shuttle lifts off the pad.

3.3. System Failure Models

After the identification of the initiating events, their credibility or frequency must be determined. This measure of probability of occurrence is utilized as a screening mechanism to focus the analysis to risk driving events, therefore if the occurrence of an initiator is assessed to be highly unlikely it will not be explicitly modeled in the study. Initiators which survive this initial risk screening exercise are modeled using functional event sequence diagrams (FESD). FESDs were used in two capacities depending on the consequential nature and mitigative mechanism of the initiating event. When the initiating event was considered to be a catastrophic occurrence, that is the initiating event leads directly to LOV, an FESD was developed to illustrate the failures of passive protective design features which lead to the initiating event (i.e. seal failures leading to a hot gas leak). In other cases the initiating event evokes an active mitigative response from the system. For this case an FESD was developed to demonstrate how a faulty system response to the initiating event could lead to LOV.

In developing an FESD, the analyst captures dynamic design characteristics of the system under study. For a complex system such as the Space Shuttle, the process can become quite involved. Many sources of information were examined during the course of this study to assure an accurate representation of the system mechanics; some of the sources studied include:

- Space Shuttle Element Training Manuals
- Space Shuttle Element Operational Descriptions
- Space Shuttle Crew Operations Manual
- Space Shuttle Operational Flight Rules

- Space Shuttle Hazard Analyses
- Space Shuttle FMEA/CILS
- Space Shuttle PRACA Reports

The FESDs served as tools of communication between risk analysts and system engineers. The FESDs were presented to system experts who reviewed them for accuracy and suggested potential sources of data for the quantification of pivotal events. These multi-faceted exchanges resulted in a number of accomplishments:

1 NASA and Shuttle contractors were familiarized with some aspects of PRA
2 The risk analysts were assured that the FESDs represented an accurate depiction of system design characteristics
3 Sources of data and the methods necessary for analysis began to be identified

At this point in the study data analysis and system modeling became concurrent activities. Having established and documented, via FESDs, the fundamental dynamics of the Shuttle the FESD models were restructured to a level which corresponded to the level of available data. In the meantime, available data was analyzed according to the constraints dictated by the scope of the model. In order to introduce the data into the models, the FESDs were converted into quantifiable entities. For the purposes of quantification, the format chosen is unimportant as long as the same set of boolean equations are represented. However experience has shown that active system responses are better represented in an event tree format. The reason for this is that the events are delineated more or less in their natural order of occurrence. Moreover, success paths are also represented which is important since some mitigative actions are triggered by the successful operation of protective systems. Nowhere is this better demonstrated than in the SSME where a closed loop control process is driven by active monitoring by the computer controller. Any detected deviation in engine performance beyond a predescribed limit triggers a sequence of events with the occurrence of each subsequent event depending on the success or failure of the previous one.

In certain cases, the failure of a protective or mitigative system necessitates the development of a supporting fault tree. The fault tree breaks the system failure down to groups of component failures which would cause the top event in the event tree. These groups of component failures are known as cutsets in the PRA with the constituent component failures labeled as basic events. The basic events constitute the interface level at which much of the failure probability estimates are introduced into the model.

3.4. Uncertainty

In PRA, the fundamental viewpoint is probabilistic. The complexity of the potential scenarios (as indicated in the previous sections) demand that the uncertainties in knowledge of these processes be accounted for. Uncertainty may originate from the inherent variation of a physical process

over many similar trials or from the limited amount of explicit experience in the particular phenomena of interest. Two example process variables that exhibit variability are wind direction and propellant burn temperature. To explain further, consider the mind experiment in which many launch vehicles (e.g. shuttle engines) were to hypothetically undergo repeated abnormal ignition or explosion accident scenarios, the temperature of the fireball would be expected to vary by virtue of the stochastic nature of these burn processes and the direction of toxic gases would vary by virtue of the stochastic nature of wind direction.

Uncertainty also refers to our state-of-knowledge about a parameter or variable. Some parameters could be accurately represented by probability distributions if sufficient research or experimentation could be performed. For example, the failure rate of a particular component on a launch vehicle could be accurately known if we could perform a sufficient number of trials that demand the operation of the component. Since this experimental evidence is unavailable, the uncertainty in the failure rate must be represented by increasing the variance of the representative probability distribution. Uncertainty is also produced by virtue of limitations in the ability to measure the failure rate. Thus, uncertainties arise from such things as inaccuracies in modeling and data, applicability of data to the situation of interest, and incomplete knowledge of the physical processes. It has been found from previous studies that the uncertainties associated with the use of available experimental data, the calculational models assumed, simplifying assumptions, and the values of variables used as input to the calculations are important sources of uncertainty within the risk assessment. Uncertainty is a probabilistic concept that is an inverse function of the "amount of knowledge" available to the analyst.

While construction of a risk model is a top-down process, quantification of a risk model takes place beginning at the lowest level of detail of the model. Thus, uncertainties are developed for model parameters at that level. Any particular scenario may or may not occur during any operating time interval, modeling of physical and chemical processes may be approximate, and the values of the parameters of the models may not be precisely known. A PRA framework allows treatment of the uncertainties. Characterization of both variabilities and knowledge uncertainties via probability distributions forces the decision-maker to come to grips with quantitative statements about the bounds and limits of knowledge that contribute to the assessment. Quantification of the uncertainties in the context of a scenario based risk model provides the means to identify the aspects of the problems that are most important to risk. The PRA, therefore, can help identify the emphasis for future design, testing, and research.

4.0. PRA DATA ANALYSIS

This section of the technical report describes the acquisition, evaluation and analysis of Shuttle-specific and applicable surrogate performance and failure data for use in quantifying the Space Shuttle PRA risk models.

4.1. Overview

The Space Shuttle PRA, as described in the previous report sections, has been undertaken to evaluate and quantify the risk of catastrophic Shuttle failure. While estimates of Shuttle failure frequency have been made during earlier phases of this study, deeper understanding of the dominant contributors to this frequency require risk modeling. Quantification of risk in terms of the severity of the consequences and the likelihood of occurrence provides flight operations and management with an important decision-making tool. By using the results of a quantitative risk analysis, such questions as "Which of several candidate systems pose the most risk?", "Are risk reduction modifications necessary?" and "What modifications would be most effective in reducing risk?"[1] may be addressed. Answering such questions requires that systems, subsystems and components of the Shuttle are modeled and their modes of failure investigated. To evaluate the likelihood of occurrence of the incidents postulated in the risk model, the analyst must know how frequently the contributory failure incidents are likely to occur. Consequently, failure rate data for the equipment involved in the incidents is essential to the risk analysis.

PRA data analysis is the process of developing an organized set of estimated failure frequency and (where applicable) maintenance unavailability[2] distributions for systems, subsystems and/or components whose failure or unavailability is included in a potential failure path in the probabilistic risk models. The resulting data set is referred to as a "basic event" data base because data is gathered for those events at the lowest or most basic level of the risk models.

The Shuttle PRA required the following types of basic event data:

(1) Descriptions of the dominant failure modes of the systems, subsystems and components that constitute the basic events;

---

[1] from Guidelines for Process Equipment Reliability Data with Data Tables, Center for Chemical Process Safety of the American Institute of Chemical Engineers, New York, 1989.

[2] Unlike, say, a nuclear power plant, the Shuttle has no standby safety systems that can be out of service for preventive maintenance when needed to interrupt an accident sequence. Thus it was not necessary to develop maintenance unavailability data for the current PRA. However, this will definitely not be true for the PRA of a long-term facility such as the Space Station and may also no longer be true if ground support systems are added to the scope of the Shuttle risk assessment.

(2) Mean values of failure rates for the dominant time-related failure modes and/or failure-on-demand probabilities for dominant demand-related failure modes (e.g., of one-shot, intermittently-operated, or standby equipment);

(3) Uncertainty bounds (typically the 5th and 95th percentiles) of lognormal failure rate and failure-on-demand probability distributions for risk-significant items.

The sections which follow discuss the development of this data in greater detail.

4.2. Analytical Methods

Since the objective of the PRA data task is to obtain information on Shuttle equipment failure frequencies, the most logical source of data is the record of Shuttle experience. While this is certainly the preferred data source, it must be recognized that the PRA requires data even for those pieces of equipment which have experienced no to few failures over many Shuttle missions. Further, it is generally advisable to compare the specific experience of the vehicle under study with generic experience for similar equipment to recognize when and where there are deviations and strive to understand why. For these reasons, as Figure 4.1 shows, the Shuttle PRA data analysis collected information from a variety of data sources. The nature of the raw information and its treatment to yield the required basic event data differed from one source to another; this is depicted by the different pathways shown in Figure 4.1 from raw source output to risk model input.

One of the first steps of any PRA data study is to establish the official study data "window", or the timeframe of information to be considered relevant to the study. For the Shuttle PRA data set, the data window ranged from the first Shuttle flight to the STS-62 Mission dated March 4, 1994 . This corresponds to the time at which the initial failure related Shuttle data was received at SAIC. Information outside this time boundary was not available to or was discarded from the data analysis for the sake of consistency across components, subsystems and systems. The subsections below describe the types of data and the processes used to convert it to the required statistics and explanatory details.

Figure 4.1. Overview of the Shuttle PRA Risk Data Analysis.

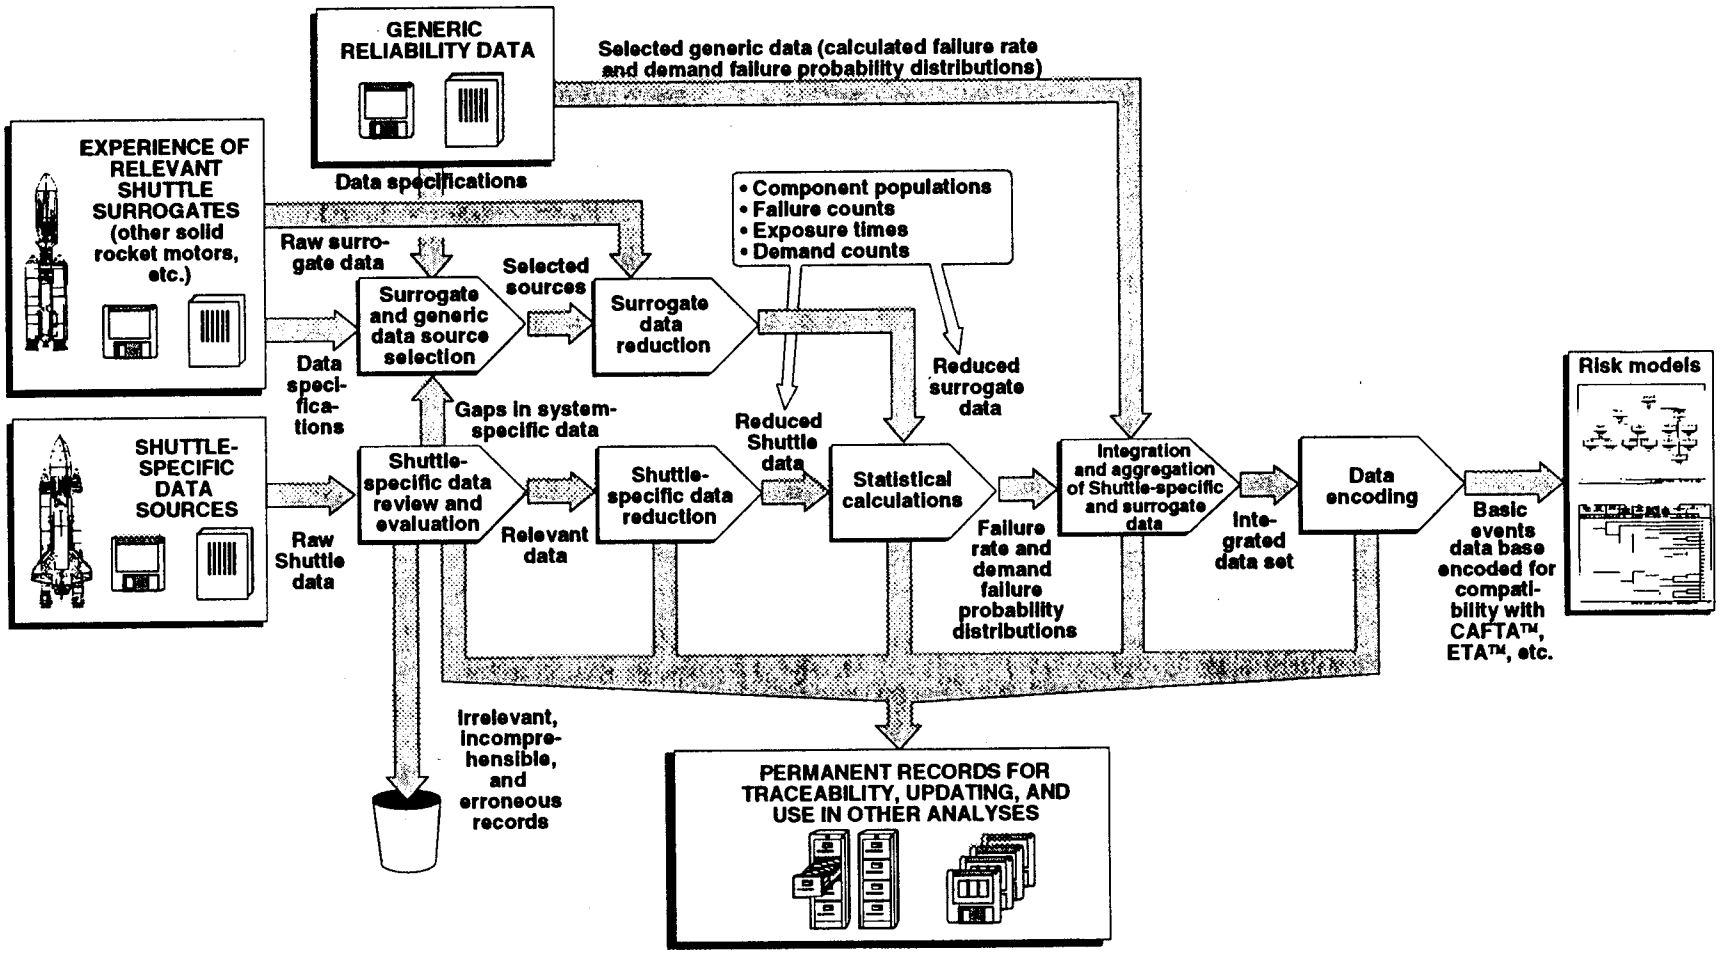

4.2.1 Types Of Risk Data

The data utilized for the quantification of the PRA models was not limited to information obtained from the Shuttle program. Although the Shuttle specific data is the most relevant to this effort, the quantity of data available was at times not statistically significant, therefore many potential sources of data were studied in an effort to be as comprehensive as possible. The information available to quantify the PRA model may be classified according to the following taxonomy:

1. *Shuttle Flight-Specific*: Data collected during actual Shuttle flight or related operations.

2. *Shuttle Test-Specific*: Data collected from tests using actual Shuttle components.

3. *Launch Vehicle-Specific*: Data collected during the flight operations of launch vehicles with components similar to those found on the Shuttle.

4. *Shuttle Surrogate (a.k.a. Generic)*: Data collected from the experience of components with design characteristics similar to those found on the Shuttle, but not necessarily in a launch vehicle environment. This data is obtained from various industries documented in reference books; wide range and amount of experience but requires careful application to Shuttle due to significant differences in design, application and environment.

Each of the above categories of data has advantages and disadvantages in terms of its use for the Shuttle PRA. The relevance of a data set to the PRA model can be characterized by these main factors:

> Tolerance or Applicability
> Sample Size
> Degree of Failure.

These factors are discussed further below and are shown in Table 4.1 in relation to the Data Type categories described above.

The applicability of the data refers to the degree of correspondence between the design characteristics and operating environment of the component from which the data resulted to that of the component in question. Components used on an actual Shuttle flight would produce data with an applicability factor of 100% or a rating of excellent. Data resulting from tests, even of Shuttle grade components, would rate a lower factor of applicability due to different operational environments. The applicability would be still lower for non-Shuttle grade equipment test information due to the difference in certification requirements. Launch vehicle equipment is only a fair match to Shuttle equipment, but is still preferable to non-launch vehicle component data, whose correspondence to the design and environment of Shuttle equipment can be quite poor. The fact that Shuttle equipment is reliable to begin with means that the pool of directly applicable

Table 4.1. Data Type General Characteristics

Data Type	Applicability	Sample Size	Degree of Failure
Shuttle Flight Specific	Excellent	Small - Moderate	Low
Shuttle Test Specific	Good - Fair	Moderate	Medium - High
Launch Vehicle Specific	Fair - Less than Fair	Large	High
Space Shuttle Surrogate	Fair - Poor	Very Large	High

component failure information is rather small. The data analyst therefore may have no alternative in some cases than to access albeit less applicable information to at least scope out the range within which Shuttle data may fall.

The sample size of the data set indicates the population of components from which the data was taken and correspondingly, the failure history available. If Shuttle flight specific equipment is used as the data population, the sample size will be limited. By including Shuttle test information, the sample size is increased, but the applicability is decreased. Trading sample size (to improve statistical validity) for reduced applicability is a continual issue in data base development, as evidenced by the comparison of the Applicability and the Sample Size columns in Table 4.1 for the four Data Types.

Degree of failure is a measure of the severity of a failure represented in the experience data set. For example, a crack in a pipe would be considered to exhibit a low degree of failure where a pipe rupture would indicate a high degree of failure. Table 4.1 shows that Shuttle Flight Specific data set contains information with a low degree of failure, meaning that the data included therein are unlikely to directly correspond to the failure severity or mode represented in the risk models, because basic events most often reflect complete or catastrophic equipment failure, meaning total loss of that equipment's function. Conversely, the Space Shuttle surrogate history, if only due to its larger sample size, contains data representing a much higher degree of failure. The low degree of failure events can still be utilized in the data base, however, by considering their propensity to propagate toward the failure occurrence of concern in the risk models. The crack in the pipe could be seen as an initial step in the progression toward a pipe rupture. Should the crack continue unmitigated under stress imposed by system operation, it would ultimately reach a point of unstable

growth leading to pipe rupture. Therefore, the existence of the crack may be considered as a precursor to actual failure and, with the proper consideration by the data analyst, as information for inclusion in the failure rate estimate. The inclusion of precursor events increases the statistical sample size, but the effects on uncertainty may vary, either increasing the uncertainty due to the addition of lower applicability data or decreasing it if the phenomenological failure mechanism is understood fairly well.

Generic data has the benefit of providing significant equipment experience data from a range of industries, but the design, application, and environment of the equipment had to be seriously considered prior to the data's use toward the evaluation of Space Shuttle flight risk. Further, since the uncertainty estimate for the mean failure rate and demand failure probabilities provided in generic data sources reflects the confidence in the data itself not its applicability to the Shuttle environment, uncertainty bounds assigned to this data for the Shuttle PRA had to be revised to reflect the tolerance uncertainty. Given these caveats, generic data were used for high reliability components which had a minimal impact upon the risk of the Shuttle.

Ideally, the data analyst would prefer to construct failure rate estimates using data bases with an excellent applicability, a very large sample size, and representing a high degree of failure. Fortunately for the users of the technology and unfortunately for the data analysts, however, operating experience provides few such instances. Further, as Table 4.1 shows, it is unlikely that one data source would contain all these attributes. Therefore, to construct the Shuttle PRA data base, it was necessary to combine data from the four primary data types, compensating for their varying levels of applicability, sample size, and failure degree to obtain the best estimate possible for each basic event. Where these sources were not available or applicable, it was necessary to use generic data or expert opinion.

Expert opinion was continually supplied by the contractors which support the Shuttle program. Their input was not only solicited for reviewing the models but also to perform a sanity check on the failure estimates determined from the available data. In some instances, the data available was not statistically significant and the component was unique to the Shuttle. In these cases, the component experts were asked for their heuristic input such that the "best" estimate possible could be modified to more accurately depict the experience divulged by the system experts.

4.2.2. Data Review and Encoding

SAIC's previous experience in Shuttle data collection indicated a starting place for Shuttle Flight Specific information to request from NASA, namely the Problem And Corrective Action (or PRACA) reports. Still, these requests led to the discovery of other data sets which might complement or supplement the PRACA records. These sources are listed by data type category in Table 4.2.

Table 4.2. Example Data Types Used in Shuttle PRA

Data Type	Data Description
Shuttle Flight Specific	Field PRACA In-Flight Anomaly Data
Shuttle Test Specific	Test PRACA SSME Premature Cutoff Database Thiokol Leak Check Results USBI SRB Component Reliability Data
Launch Vehicle Specific	Solid Rocket Motor Failures: Castor IV Minuteman III, Stage 1 Poseidon C3 Titan SRM
Space Shuttle Surrogate	NPRD-3 NPRD-91

Field PRACA and Anomaly reports were obtained from NASA in electronic format to facilitate review and storage, but significant and intensive review was still required due to the narrative format of the reports and the fact that problem identification and resolution could occur significantly later in time from the initial post-mission anomaly reporting. The first level of review involved the removal from further consideration of records which did not conform to the success criteria established for the models. Thousands of failure reports were reviewed either manually or through a computer-aided method developed by SAIC to streamline the screening process. This approach involved an initial automated review using the standard record attribute codes. In other words, the automated process would scan the computerized record set to identify codes predetermined by the data analysts to be important to the PRA data needs, for example, those record entries which related to HPOTP failures of Criticality 1. The records which survived this initial screening process were then individually inspected by a risk data analyst, who assessed their relevance to the PRA. The automated screening process did not prove to be very effective with PRACA. This was because record codes which could have more effectively screened the records were missing, incorrect, or inconsistent in many cases. For example, the failure criticality code, which is a measure of worst case consequence, was inconsistently reported when available at all. Further, the criticality level associated with several failure modes was changed, particularly following the Challenger accident, making the significance of the failure mode code inconsistent across the timeframe of the PRACA data set. As a result, the data analysts were required to conduct a thorough review of the PRACA records in their entirety, rather than just the attribute codes, to make an assessment as to the possible mission consequences of each anomaly. Since the PRACA reports were created as an anomaly reporting system and were not developed with risk quantification in mind, every anomaly was recorded, no matter what its impact on either component performance or mission consequences might be. This meant that the amount of data to be analyzed varied significantly from one subsystem to the next and that PRACA data was ultimately used in a slightly different manner in each subsystem. These PRACA data application issues will be discussed in each of the subsystem sections of this report.

Fortunately, it was determined by the data analysts through the course of the study that the Shuttle contractors maintained their own equipment failure data bases whose reporting was more consistent with the objectives of the PRA. These data sets proved to be a valuable resource, not only due to their consistency and completeness, but because individual contractor representatives with detailed understanding of the data sets were available to explain each data base's specific nomenclature and provide additional information as required.

The records determined to be relevant to the PRA scope were then analyzed further to postulate the impact of each incident described had it occurred during an actual Shuttle flight. Most records did not explicitly describe operational failures, but rather anomalies or test-related failures which could imply possible flight failures. While it was understood that these anomalies and test failures could not be counted as full operational failures, they did provide some information on equipment performance and therefore could not be discarded from the PRA data base. Instead, it was decided to consider them as what amounted to fractions of equivalent flight failures by applying a "potentiality factor". Where anomalies were concerned, the determination of the potentiality factor was based on the particular type of failure mechanism which could potentially drive the anomaly to catastrophic failure. In the case of test-related failures, the potentiality factor was indicative of the estimated correlation between the test and flight environments. By multiplying the anomaly or test failure by the potentiality factor, the PRA data base was able to give credit for non-flight experience and to enrich the data by including relevant fractions of failures. Further details on the use of potentiality factors are provided in the Systems Analysis section of this report.

As a result of the data review task, failure instances and modes relevant to basic events modeled for the Space Shuttle subsystems were extracted from the raw data bases. Their use as numerators for time failure rates and demand failure probabilities is described further in section 4.2.4.

4.2.3. Exposure Data Development

Knowing the number of failures experienced by the Shuttle equipment was not sufficient for the development of data to support the Shuttle PRA basic events. It was also necessary to determine the amount of operational time or the number of demands the equipment had been subjected or 'exposed' to, known as the exposure time or hazard exposure. This information was obtained, calculated and used differently depending upon the Shuttle subsystem.

For the main engines (SSMEs), the ACTS data base obtained from Marshall Space Flight Center was used to obtain the total flight and test stand exposure time. This data set listed all flights and tests conducted up to March 1994 and was conveniently formatted in a spreadsheet, including the number of seconds of operating time per flight or firing time per test. Thus, the process of summing the exposure time by flight or test or total was simplified.

In the case of the Auxiliary Power Units (APUs), a nominal operating time per mission was estimated and was multiplied by the number of missions within the study data window.

The ISRB exposure was demand based rather than time based for two major reasons. Firstly the relatively short operation time of the ISRB makes the use of demand related data applicable even in instances where the component is subjected to the failure inducing environment for the duration of the ISRB operation time. Secondly a large number of risk contributing components must perform their function for an instant (i.e. pyrotechnics) and therefore are their failures are truly demand related. This attribute was a double edged sword, on the positive side the number of failures and the number of demands were found in the same source of data; on the negative side, since the failure mechanism does not act over a period of time there is a lack of precursory data which may be used to gain further insight and therefore better estimates of the potential failures.

Another important aspect of exposure information necessary for the estimation of failure probability or frequency per flight is the operation time or number of demands imposed on the components under study. Nominal operational exposure estimates were obtained from flight operations and training manuals.

The exposure data was then used along with the number of failures to form time-related failure rate and demand failure probability estimates, as described in the following section.

4.2.4. Failure Data Estimation

Most of the equipment failure rates and demand failure probabilities used for the Shuttle PRA risk model quantification were developed from Shuttle operations and test experience. As was shown earlier in Figure 4.1, where such data was not sufficient or available at all, data from surrogate or generic sources were used.

The following equation summarizes the fundamental method for the calculation of the probability per mission or frequency of failure from Shuttle operation and test experience for the example of time-related failure rates (the equation for failure-on-demand probability is analogous.)

$$\textit{Failure Frequency} = \frac{\textit{Number of Failures}}{\textit{Hazard Exposure}} \times \textit{Mission Specific Operation Time}$$

where: *Failure frequency* is the estimated number of failures expected in one mission;
Number of failures is the total count of flight equivalent failures of a given mode within the population of items of concern during the exposure to hazard;
Exposure to hazard is the total unit-time or unit-cycles to which the population is subjected;
Mission specific operation time is the duration of time for which the component is exposed to the particular hazard during a nominal mission (for demand related failures this would be the number of demands made upon component during a nominal mission)

For obvious reasons, the information needed to develop failure counts is often called "numerator data," while the information that provides exposure estimate is called "denominator data." Normally

it is necessary to acquire, reduce, and integrate a number of distinct data sets in order to obtain one or all types of information, and the Shuttle PRA was no exception. Figure 4.2 illustrates the most important Shuttle PRA reliability data sources and how they contribute to the evaluation of the failure frequency equation above.

This equation yields a maximum-likelihood estimator of the failure rate. Assuming that failure rates and demand failure probabilities are lognormally distributed (as we did in this PRA; see below for further discussion), this estimate was taken as the mean of the distribution. Standard statistical techniques were then used to develop the remaining distribution parameters for items that appeared to have sufficient risk importance to require the use of a distribution rather than a point estimate. The development of these distributions through application of uncertainty bounds and the criteria for risk significance will be discussed in section 6.1.1.

Figure 4.2. Use of Shuttle Experience Data to Estimate Failure Frequencies

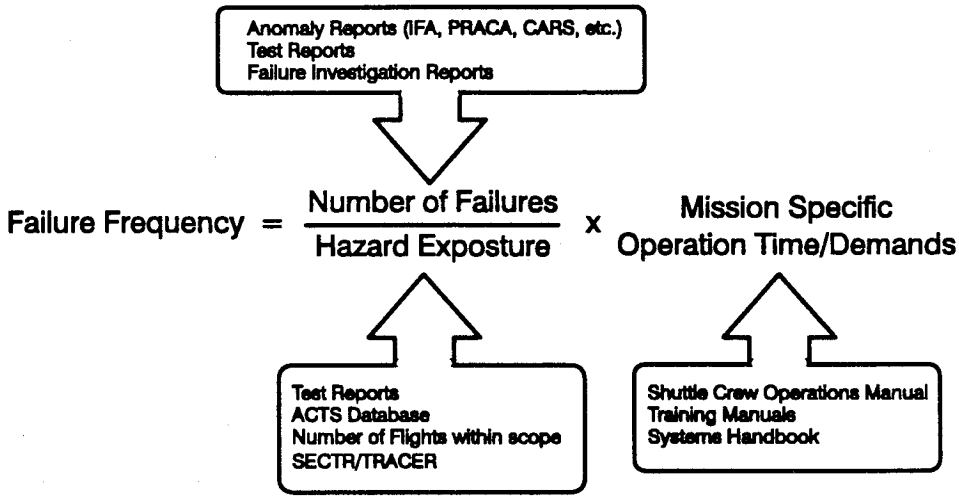

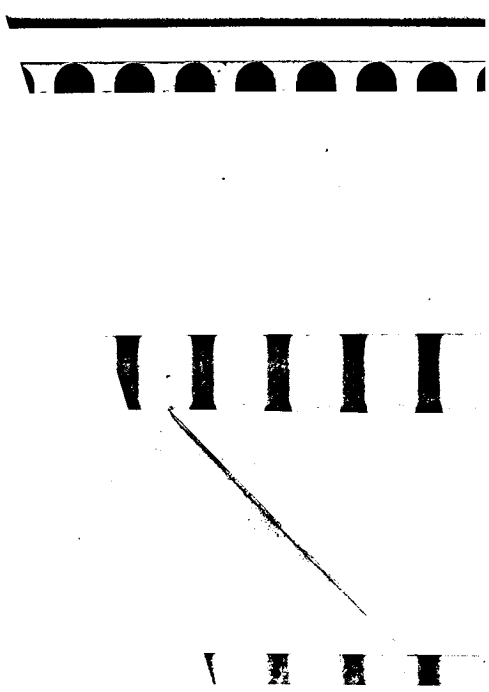

5.0 Systems Analysis

5.1. Front-Line Systems

5.1.1. Space Shuttle Main Engines and Main Propulsion System

The SSME and MPS comprise the primary propulsion system for the Space Shuttle. Although the MPS is actually part of the Orbiter it was analyzed along with the SSME because of the functional connection between the two systems. Any risk due to a malfunction of the MPS, however, was allocated to the Orbiter. Three SSME's in the aft compartment of the Orbiter produce approximately 490,000 lbs (104% throttle) each using propellant, liquid hydrogen and liquid oxygen, supplied by the plumbing connecting the SSME to the external tank. The MPS is composed of the Propellant Management System (PMS) and the Helium Supply System (HSS).

For the purposes of the PRA the combined systems operate from SSME ignition to Main Engine Cut-off (MECO) for a nominal duration of 520 seconds[1]. After MECO only the HSS and parts of the PMS and SSME must operate to perform a propellant dump sequence to evacuate the plumbing of residual H2 and O2. In addition a vacuum inerting process is performed to insure the systems are free of all traces of the propellants to avoid deterioration which may occur due to extended exposure to the propellant chemicals.

5.1.2.1. Description

The SSME is a liquid hydrogen/liquid oxygen engine that employs a two-stage combustion cycle. In the first stage, a fuel rich mixture is partially burned in two preburners. The resulting fuel-rich hot gas streams are first used to drive high pressure turbopumps. The fuel-rich streams are then injected into a main combustion chamber along with coolant fuel and the required oxidizer for burning at a controlled mixture ratio of 6.0. A simplified schematic of engine operation in figure 5.1. Description of individual components and functions will be supplied as their contribution to failure initiators is noted.

The three SSME receive their supply of propellants from the two 12 inch manifolds (one for H2 and one for O2) connected to the low pressure turbopumps. A computer Controller actively monitors various system parameters to assure that the engine is performing within specifications and to control the position of the propellant valves to maintain the proper conditions (this process is termed closed loop control). There are seven specific parameters for which redline values have been defined, these are shown in Table 5.1. If any of these redlines are exceeded the Controller calls for an emergency

---

[1] The actual operation time may vary due to changes in mission characteristics, the variation has a negligible affect on the risk posed by the system and therefore the nominal point estimate is used throughout the study.

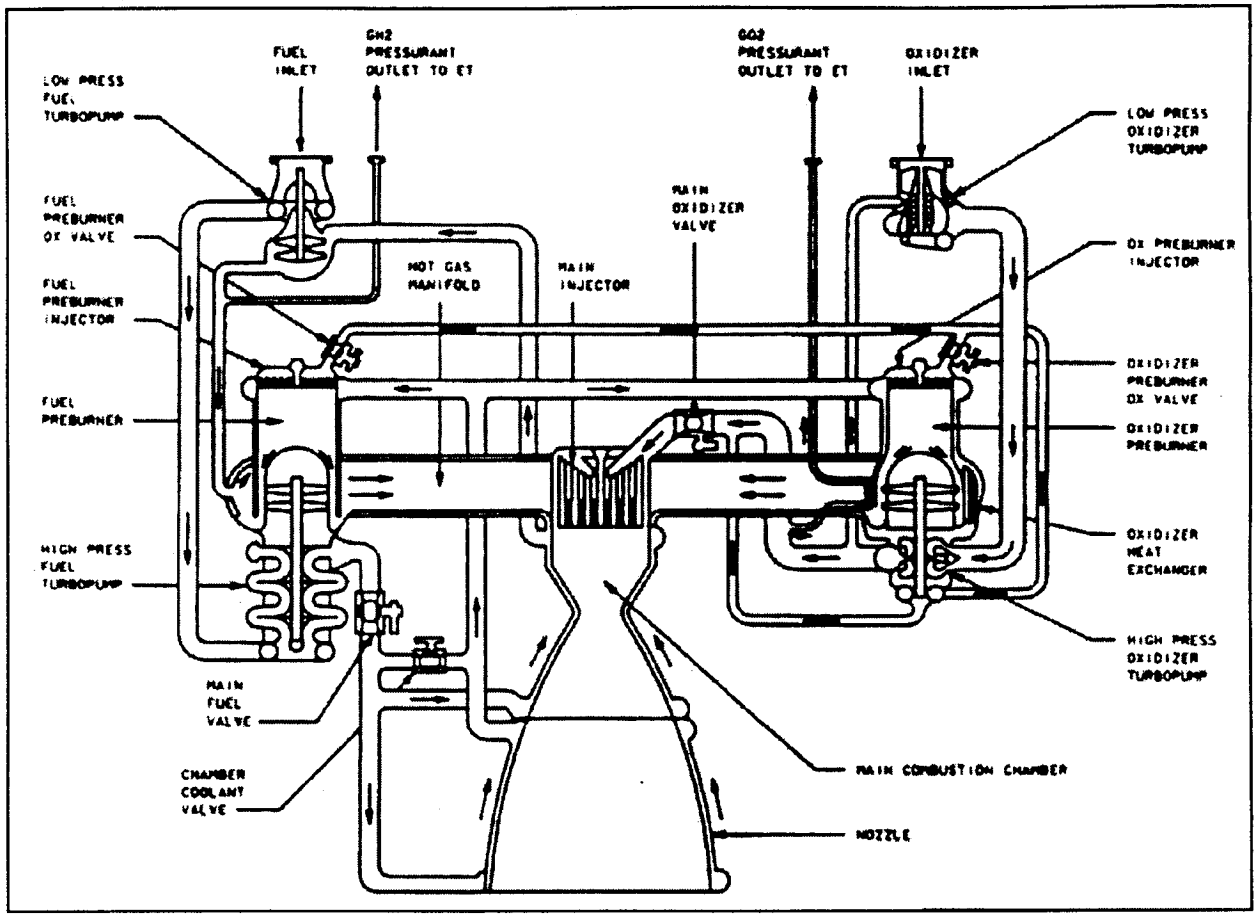

Figure 5.1. SSME Schematic

shutdown at which point the propellant valves are sequentially closed. Once the valves are closed the engine is purged with helium from the Helium Supply System. Hydraulic pressure for the five propellant valves in each SSME is supplied by an APU driven hydraulic system. A separate hydraulic system exists for each SSME, a drop in hydraulic pressure will cause the Controller to lock-up the engine; all valves are set to their last commanded position and closed loop control functions are suspended. However, redlines are still monitored and if any are exceeded the engine will be shutdown pneumatically using the helium supply as the driving gas. If no redlines are exceeded the engines continue to fire until MECO is commanded at which time the same process described for emergency shutdown is performed to shutdown all three engines.

Table 5.1. SSME Defined Redline Parameters and Limit Exceeded Definitions

PARAMETERS	CONTROLLER SHUTDOWN PROCESSING[2]	REASONABLENESS TESTING[3]	SHUTDOWN LIMITS[4]
HPFTP Coolant Liner Pressure Sensors A & B (PSIA)	1) A & B Fail Limit 2) A Fails Limit B Fails Reasonableness 3) A Fails Reasonableness B Fails Limit	>4500 PSIA <1800 PSIA	>Controller calculated limit[5] No Lower Limit
HPFTP Turbine Discharge Temp. Sensors A & B (°R)	1) A & B Fail Limit 2) A Fails Limit B Fails Reasonableness 3) A Fails Reasonableness B Fails Limit	>2650°R 810°R	CH A >1850°R CH B >1960°R No Lower Limit
HPOTP Turbine Discharge Temp. Sensors A & B (°R)	1) A & B Fail Limit 2) A Fails Limit B Fails Reasonableness 3) A Fails Reasonableness B Fails Limit	>2650°R <150°R	>1760°R <720°R
HPOTP Secondary Seal Pressure Sensors A & B (PSIA)	1) A & B Fail Limit 2) A Fails Limit B Fails Reasonableness 3) A Fails Reasonableness B Fails Limit	>300 PSIA <4 PSIA	>100 PSIA No Lower Limit
HPOTP Intermediate Seal Pressure Sensors A & B (PSIA)	1) A & B Fail Limit 2) A Fails Limit B Fails Reasonableness 3) A Fails Reasonableness B Fails Limit	>650 PSIA <0 PSIA	No Upper Limit <170 PSIA
MCC Chamber Pressure Sensor Average Sensors A & B (PSIA)	1) A & B Fail Limit 2) A Fails Limit B Fails Reasonableness 3) A Fails Reasonableness B Fails Limit	|Brg1-Brg2|≥125PSI PC Channel Avg ≥3500 PSI ≤1000 PSI	No Upper Limit <170 PSIA

[2]If the two sensors are qualified, both vote for shutdown for action to proceed. If only one sensor is qualified, the Controller must rely on only that vote.

[3]Reasonableness testing is required to qualify the sensors for application to shutdown logic. The values chosen screen out sensors with identified sensor problems.

[4]Redline parameters are monitored to assure that the engine is performing within safe operating conditions. Limits are set to guard against uncontained SSME damage. The limits are based upon test stand data, flight experience, and engineering analysis. The engine redline design criteria was defined by MSFC and approved by Level II.

[5]The redline limit is calculated in real time by controller software and is a function of engine power level. The limit at 104% power level is approximately 3675 PSIA.

If an engine is shutdown prematurely, the redlines in the other two engines will be inhibited thereby precluding another premature shutdown. The safety implications of inhibiting the redlines is not considered in this study since a single engine shutdown is defined as an abort scenario for the purposes of this analysis. There are instances where two engines may shutdown virtually simultaneously which would lead to certain catastrophic consequences if it were to happen in the early stages of ascent; these scenarios will be discussed later in this section.

The Propellant Management System consists of two 17-inch-diameter propellant feedline manifolds located in the Orbiter aft compartment. Each manifold interfaces with the ET, one with the liquid hydrogen supply and the other with the liquid oxygen supply. Inside the aft compartment the manifolds diverge into three 12-inch feedlines, one for each SSME. Both manifolds interface with an 8-inch fill/drain line containing an inboard and outboard fill/drain valve in series. The manifolds and fill/drain lines contain a number of valves which are cycled during prelaunch and after MECO to dump the residual propellants. The most important valves to this study are the feedline disconnect valves which close automatically prior to external tank separation, the prevalves in each 12-inch feedline, the fill/drain valves, the back-up liquid hydrogen dump valves and the relief valves connected to 1-inch lines emanating from the 8-inch lines.

Once MECO has been confirmed at approximately 8 minutes 30 seconds MET, the GPCs execute the external tank separation sequence. The sequence takes approximately 18 seconds to complete and includes opening the feedline relieve isolation valves, arming the external tank separation pyro initiator controllers, closing the liquid hydrogen and liquid oxygen feedline 17-inch disconnect valves, turning the external tank signal conditioners' power off (deadfacing), firing the umbilical unlatch pyrotechnics, retracting the umbilical plates hydraulically, and gimbaling the SSMEs to the MPS dump sequence position.

Ten seconds after main engine cutoff, the RTLS liquid hydrogen dump valves are opened for 80 seconds to ensure that the liquid hydrogen manifold pressure does not result in operation of the liquid hydrogen feedline relief valve.

After MECO confirmed plus 20 seconds, the GPCs interconnect the pneumatic helium and engine helium supply system by opening the three out/open interconnect valves if the MPS He INTERCONNECT LEFT, CTR, RIGHT switches on panel R2 are in the GPC position. This connects all 10 helium supply tanks to a common manifold, and it ensures that sufficient helium is available to perform the liquid oxygen and liquid hydrogen propellant dumps.

After external tank separation, approximately 1,700 pounds of propellant are still trapped in the SSMEs, and an additional 3,700 pounds of propellant remain trapped in the orbiter's MPS feedlines. This 5,400 pounds of propellant represents an overall center-of-gravity shift for the orbiter of approximately 7 inches. Non-nominal center-of-gravity locations can create major guidance problems during entry. The residual liquid oxygen, by far the heavier of the two propellants, poses the greatest impact on center-of-gravity travel.

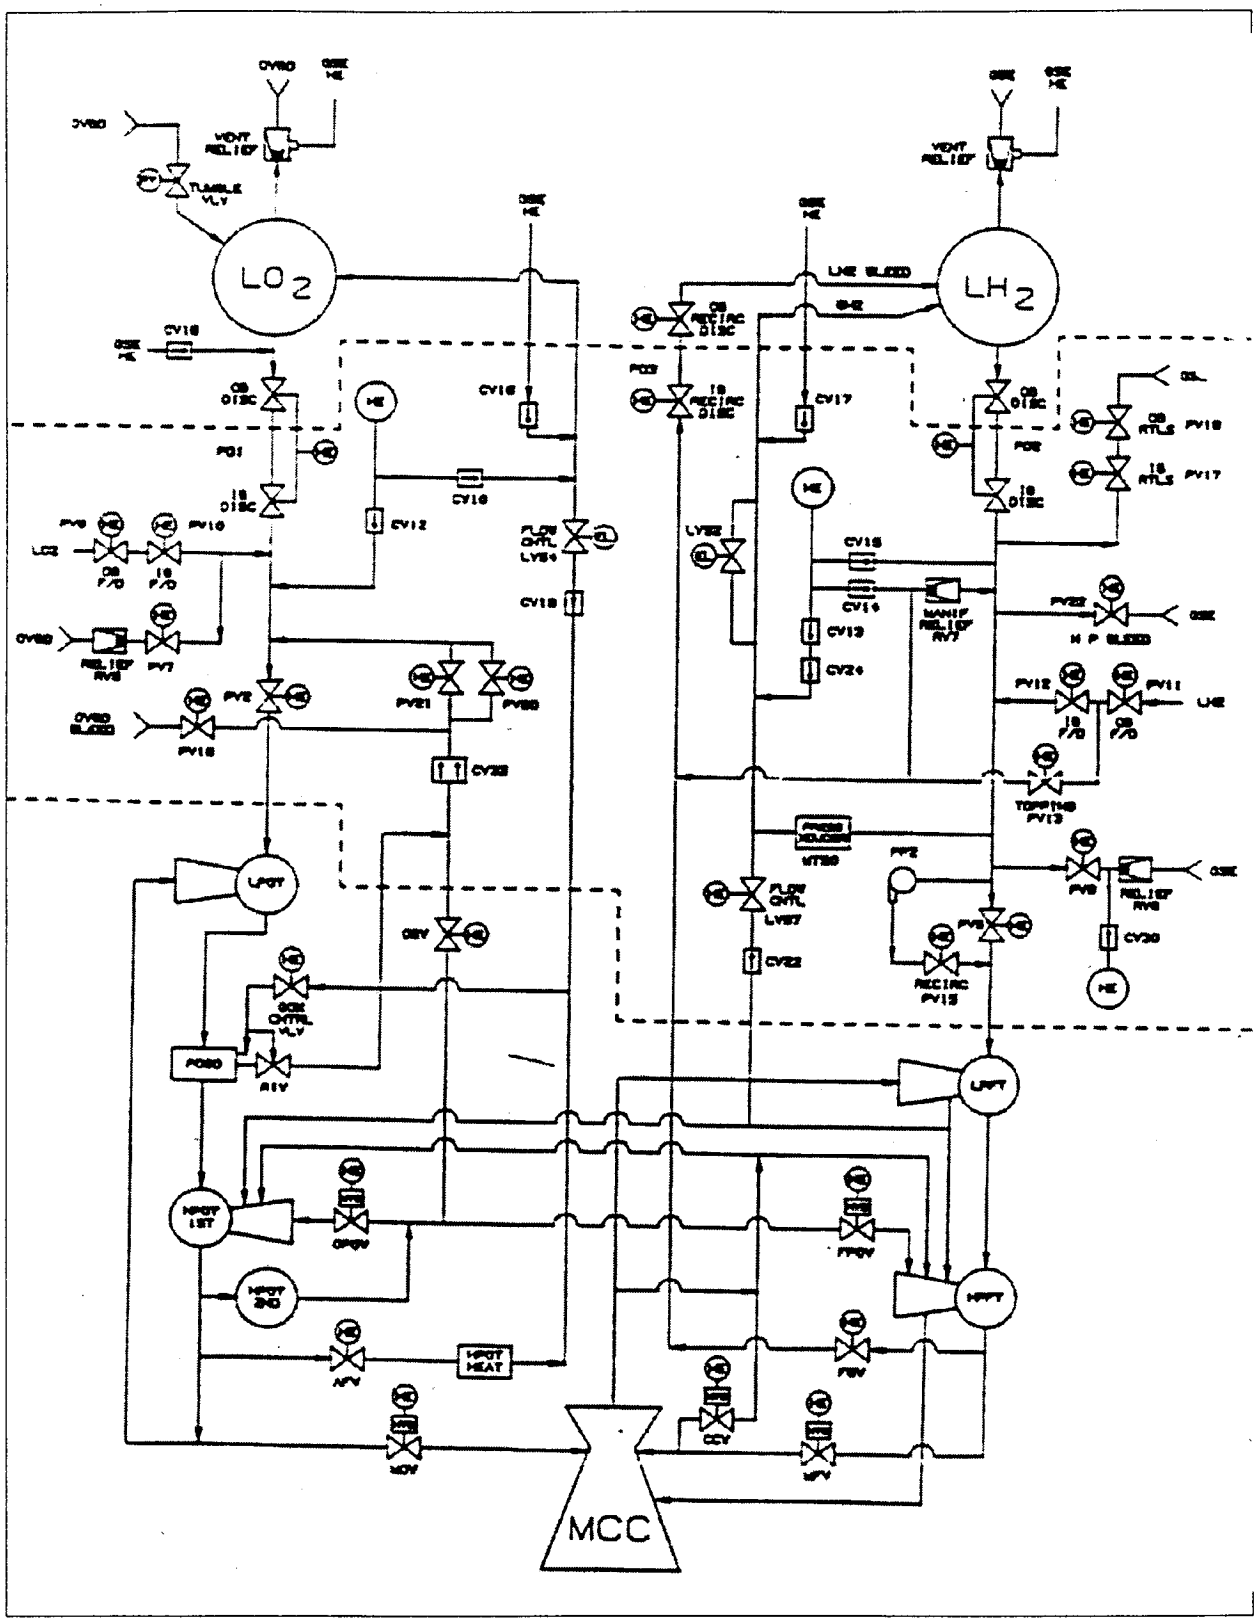

Figure 5.2. Main Propulsion System Schematic

A hazard from the trapped liquid hydrogen occurs during entry, when any liquid or gaseous hydrogen remaining in the propellant lines may combine with atmospheric oxygen to form a potentially explosive mixture. In addition, if the trapped propellants are not dumped overboard, they will sporadically outgas through the orbiter liquid oxygen and liquid hydrogen feedline relief valves, causing slight vehicle accelerations.

The MPS propellant dumps (LO2 and LH2) occur simultaneously. The method of initiating the dump depends on the type of mission. Both dumps are completely automatic once initiated. The helium subsystem is used during the MPS dump to help expel the propellants from the manifolds. To support this, the GPCs command the left, center, and right helium interconnects to out/open at MECO plus 20 seconds. This occurs provided the helium interconnects are in the GPC position.

For standard insertion flights, the MPS dump starts at OMS-1 TIG which usually occurs at MECO plus 2 minutes, provided the MPS PRPLT DUMP SEQUENCE switch on panel R2 is in the GPC position. This dump takes 2 minutes and 1 second to complete.

For direct insertion flights, the MPS dump is started manually, by taking the MPS PRPLT DUMP SEQUENCE switch to START. This is performed manually at MECO plus 2 minutes. The earliest that the manual MPS dump can be performed is MECO plus 20 seconds. The only reason that the crew may need to start the dump prior to MECO plus 2 minutes is if the manifold pressure rises unexpectedly. The manual dump takes 2 minutes and 21 seconds to complete. The STOP position of MPS PRPLT DUMP SEQUENCE is functional but is never used for either dump case.

For the LO2 dump, the computers command the two liquid oxygen manifold repressurization valves to open (the MAIN PROPULSION SYSTEM MANF PRESS LO2 switch on panel R4 must be in the GPC position), command each engine controller to open its SSME main oxidizer valve (MOV), and command the three liquid oxygen prevalves to open (the LO2 PREVALVE LEFT, CTR, RIGHT switches on panel R4 must be in the GPC position). The liquid oxygen trapped in the feedline manifolds is expelled under pressure from the helium subsystem through the nozzles of the SSMEs. This is propulsive and typically provides about 9-11 feet-per-second of delta V.

The pressurized liquid oxygen dump continues for 90 seconds. At the end of this period, the GPCs automatically terminate the dump by closing the two liquid oxygen manifold repressurization valves, wait 30 seconds, and then command the engine controller to close their SSME main oxidizer valve. The three liquid oxygen prevalves remain open during the orbit phase of the flight.

Concurrent with the liquid oxygen dump, the GPCs automatically initiate the MPS liquid hydrogen dump. The computers command the two liquid hydrogen manifold repressurization valves to open (the MAIN PROPULSION SYSTEM MANF PRESS LH2 switch on panel R4 must be in the GPC position) and command the two liquid hydrogen fill and drain valves (inboard and outboard) to open.

The liquid hydrogen trapped in the orbiter feedline manifold is expelled overboard under pressure from the helium subsystem through the liquid hydrogen fill and drain valves for 6 seconds. The inboard fill and drain valve is closed, the three liquid hydrogen prevalves are opened, and liquid

hydrogen flows through the topping valve, between the inboard and outboard fill and drain valves, and overboard through the outboard fill and drain for approximately 88 seconds. The GPCs automatically terminate the dump by closing the two liquid hydrogen manifold pressurization valves and 21 seconds later closing the liquid hydrogen topping and outboard fill and drain valves.

At the end of the liquid oxygen and liquid hydrogen dumps, the GPCs close the helium out/open interconnect valves provided the HE INTERCONNECT LEFT, CTR, RIGHT switches on panel R2 are in the GPC position. After the MPS dump is complete, the SSMEs are gimballed to their entry stow position with the engine nozzles moved inward (toward one another) to reduce aerodynamic heating. Although the gimbals move to an MPS dump positions during the external tank separation, the I-loads are currently the same as the entry stow position. At this time, the BODY FLAP lights on panel F2 and F4 turn off. This is the crew's indication that the MPS dump is complete.

Approximately 19 minutes into the mission and after the MPS dump, the flight crew initiates the procedure for vacuum inerting the orbiter's liquid oxygen and liquid hydrogen lines. Vacuum inerting allows any traces of liquid oxygen or liquid hydrogen trapped in the propellant lines after the propellant dumps to be vented into space.

The liquid oxygen vacuum inerting is accomplished by opening the liquid oxygen inboard and the outboard fill and drain valves. The are opened by placing the MAIN PROPULSION SYSTEM PROPELLANT FILL/DRAIN LO2 OUTBD, INBD switches on panel R4 to the OPEN position.

For liquid hydrogen vacuum inerting, the liquid hydrogen inboard and outboard fill and drain valves are opened by placing the MAIN PROPULSION SYSTEM PROPELLANT FILL/DRAIN LH2 OUTBD, INBD switches on panel R4 to OPEN. The external tank gaseous hydrogen pressurization manifold is also vacuum inerted by opening the hydrogen pressurization line vent valve by placing the MAIN PROPULSION SYSTEM H2 PRESS LINE VENT switch on panel R4 to OPEN.

After a one minute inert period, the switch is taken back to the GND position, which closes the valve. The hydrogen pressurization vent line valve is electrically activated; however, it is normally closed (spring loaded to the close position), and removing power from the valve solenoid closes the valve.

The liquid oxygen and hydrogen lines are inerted simultaneously. Approximately 18 minutes is allowed for vacuum inerting. At the end of the 18 minutes (OSM-2 TIG minus five minutes), the pilot closes the LO2 and LH2 outboard fill drain valves by placing the MAIN PROPULSION SYSTEM PROPELLANT LH2 AND LO2 FILL/DRAIN OUTBD switches on panel R4 to CLOSE. The procedure has the pilot wait ten seconds to insure the valves full close before taking the switches to GND. Taking the switches to GND removes power from the close solenoids. Although the power is removed from the solenoids, they remain in their last position (closed) since the fill drain valves are bi-stable valves. Also at this point the pilot removes power from the open solenoids of the LO2 and LH2 inboard fill drain valves. This done by placing the MAIN PROPULSION SYSTEM PROPELLANT LH2 AND LO2 FILL/DRAIN INBD switches on panel R4 to GND. Remember that these valves were opened during the vacuum inert initiate procedure. Placing the switch in GND only removes power from the solenoid; the valves remain open. These valves are left open to prevent a

pressure buildup between the inboard and outboard valves. Finally, the PNEUMATICS He ISOL is taken to the GPC position since there is no longer a need to operate the pneumatic valves. This action removes power from the valve, causing it to close.

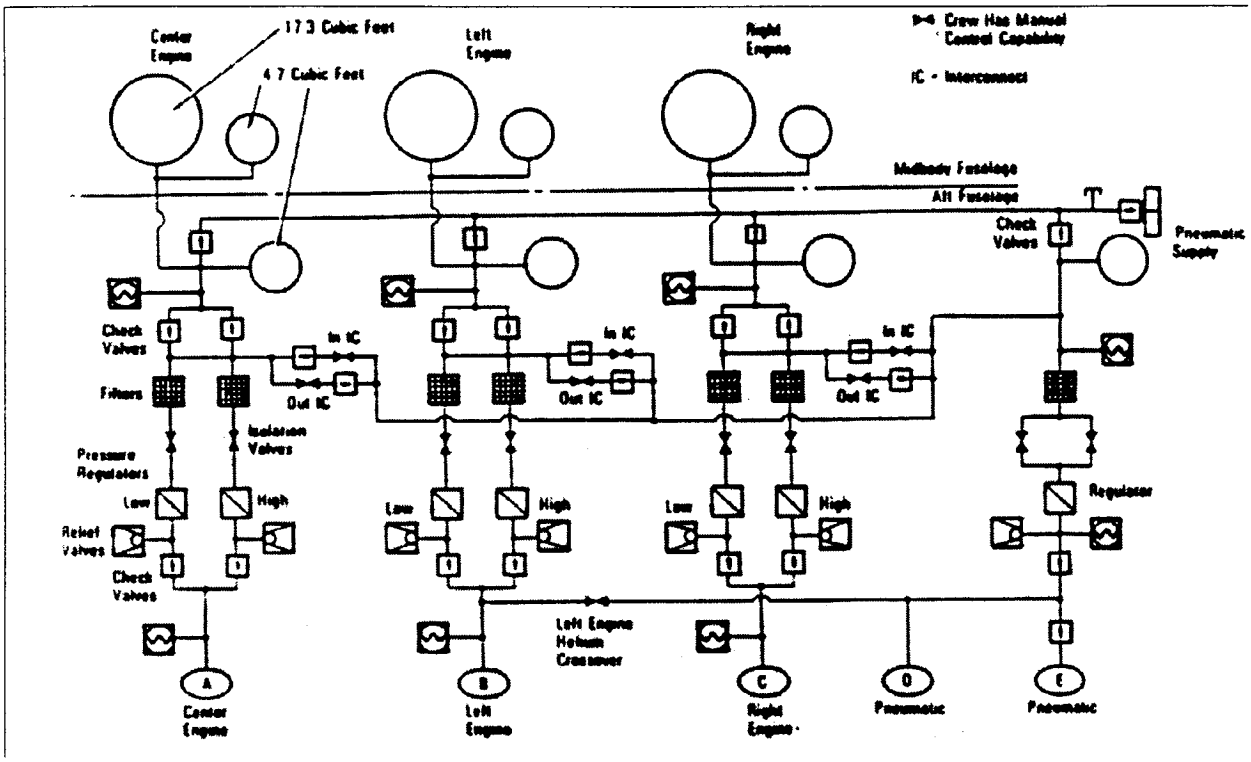

Figure 5.3. Helium Supply System Schematic

The helium system consists of seven 4.7-cubic-foot helium supply tanks and three 17.3-cubic-foot helium supply tanks, and associated regulators, check valves, distribution lines, and control valves. The helium system is used for in-flight purges within the engines, and it provides pressure for actuating engine valves during emergency pneumatic shutdowns. It also provides pressure to actuate the pneumatically operated valves within the propellant management system. During entry the remaining helium is used for entry purge and re-pressurization.

Each of the larger tanks is plumbed to two of the smaller supply tanks, forming three clusters of three tanks. Each set of tanks normally provides helium to only one engine, however, cross-ties exist such that helium from one system may be routed to support another engine. This may be necessary if a leak is detected and isolated in one of the systems. Such cross-overs are instituted by the crew which controls the positions of the cross-ties from the cockpit.

5.1.2.2. Success Criteria

There are two top-level STS functions which failures of the SSME/MPS can challenge; these are failure to provide proper propulsion and failure to contain energetic gas or debris. The failure to provide proper propulsion function is not the loss of propulsion from one engine, although such an incident may have catastrophic consequences during the abort process the direct effects are considered benign. A dual SSME shutdown is considered to have a direct catastrophic effect if it occurs before the "droop 109" call[6], which was determined to occur approximately 5.8 minutes after lift-off. Any event which leads to a three engine shutdown is considered to be a LOV regardless of the time at which it occurs.

Any discharge of energetic gas or debris is considered to lead directly to LOV. Such a discharge may be the result because of random critical structural failures or structural failures caused by operation beyond the redline limitations discussed earlier. No allowance is given for safe extended operation beyond the redline limits therefore any failure to shutdown an engine is equivalent to the occurrence of a critical structural failure. Therefore any event which may lead to a redline shutdown is considered as a potential catastrophic accident initiating event.

Failure to shutdown an engine for emergency purposes or at MECO, specifically closure of the OPOV and purging the OPB, is considered critical to the safety of the mission. The OPOV must be closed to avoid a lox-rich cutoff of the engine which can cause considerable damage to all combustion devices; during a pneumatic shutdown the OPOV must close before the helium driving the actuation piston can be routed to the other valves. Although the ratio of liquid propellants is such that the oxygen will be depleted before the hydrogen in the event of an engine which does not shutdown at MECO, the incident is considered LOV due to the uncertainties involved.

Soon after MECO, the plumbing of the Propellant Management System must be evacuated of residual propellants to avoid a build up of pressure as the liquids evaporate. Failure to dump these propellants overboard is considered a LOV event. The leakage of both oxygen and hydrogen is assumed to be a direct catastrophic event.

Helium is necessary to purge the HPOTP intermediate seal which separates the fuel-rich mixture in the preburner from the oxygen in the pump. The helium is also used for shutdown purges and in some cases acts as the driving gas to shutdown the engine. Leakages of helium are potentially critical due to the failure to provide these necessary functions.

---

[6]The "droop 109" call by the mission controller signifies that the Shuttle has attained an energetic state which should make an abort with only one engine operational possible.

5.1.2.3 Initiating Events

The following anomalous conditions were identified as initiating events which could cause a redline condition:

- Loss of MCC Pressure
- Loss of Gross H2 Flow
- Loss of Fuel to Both Preburners
- High Mixture Ratio in Fuel Preburner
- High Mixture Ratio in Oxidizer Preburner
- HPFTP Coolant Liner Overpressure
- Failure to Maintain Proper Propellant Valve Positions
- Hydraulic Lock-up Required

Failure modes which can cause an energetic discharge leading directly to LOV during mainstage which are all under the category of Criticality 1 failures[7] are grouped under the initiator:

- Critical SSME Structural Failures

There were two initiators identified for conditions related to abnormal SSME shutdown scenarios:

- Simultaneous Dual SSME Shutdown
- Failure to Perform Nominal MECO & Dump

The propellant management system may initiate a failure in one of two manners:

- Combustible Leakage of MPS Propellants
- Failure to Maintain Positions of MPS Propellant Supply Valves

The helium supply system related accident sequences may be initiated by:

- Leakage of the Helium Supply System
- Failure to Provide Charging for POGO Accumulator

The initiating events above which may cause a redline are phenomenological in nature and may be caused by various component failures. The occurrence of a redline served as an indication of one of these initiating events and was used to evaluate the initiator frequencies. Ref. 32 provided the entire history of redline occurrences for both flight and test. SSME related initiators which resulted in abnormal operation and likely shutdown were studied using only relevant engine shutdown incidences both in flight and during testing. Relevant failures were those which were considered to pose a non-negligible possibility of catastrophic failure to the vehicle during the time considered within the scope of the analysis (scope of analysis is from SSME confirmed ignition, T-3 seconds, to wheel-stop after landing. aborts are out of scope). Redline sensor failures causing an engine shutdown were neglected

---

[7]There are three primary grades of failure defined in the FMEA/CILS: Criticality 1, 2 & 3 which are failures which may cause Loss of Vehicle/Crew, Loss of Mission or Degraded Operation respectively

because they pose no credible risk to the vehicle since the engine is not actually operating beyond the redline limits. Other reasons for discounting redline cutoff incidents included shutdowns due to test facility malfunctions, operation at 111% or manual shutdowns due to fire (fires were considered in the propellant leakage initiator, the frequency is determined using an alternative method). Shutdowns due to FASCOS were also discounted since FASCOS is not active during flight. In addition to this straight forward discounting the database also provided engine configuration which was also used as a semi-screen; more modern configurations were weighted more than older configurations to account for reliability growth[8]. The records which were used in evaluating the initiator frequencies and their respective configuration applicability factor are shown in Appendix B.1.

SSME structural failures which lead directly to an energetic discharge and subsequent LOV were analyzed using actual field anomalies. Fortunately, no catastrophic failures of the SSME have occurred during flight, although inspections both prior and after missions have found components which show deterioration. If these were allowed to continue they would have eventually lead to a catastrophic structural failure of the SSME. Test failures were utilized to examine the agreement between estimates and observed test catastrophic occurrences. The reason for not using test failures explicitly is a matter of data applicability. The mechanism which drives component structural failure is material deterioration or the rate at which a material defect propagates. The rate at which a defect propagates is a function of many variables including operating environment, time between inspections, duration of exposure to stresses, transient application of stresses, age of components, etc. Both the number of variables and the differences between their values in flight vs. testing did not make the test data an extremely reliable source for actual flight failure rate for structural components. In addition, the sample size of test failures was relatively small.

Another reason for using field anomalies rather than actual test failures was that although there are relatively few catastrophic structural failures during testing there are no mitigative events between the initiator and catastrophic failure during flight. The inspection process may actually be considered as the mitigation event in this case and as such counting up the number of times that this mitigative event has been called upon should provide some indication of potential failure. The "potentiality" of the actual failure occurring was taken as 1%- 2% of the number of inspection "squawks". This estimate was chosen from considering the probability of detecting a defect during an inspection and the likelihood of an undetected defect propagating to failure before the next inspection. The fact that the same potentiality factor is used for multiple components is consistent with actual practice since Shuttle engineers have heuristically determined what the optimum inspection interval is for each component. In this case, PRACA records which were tagged as field occurrences were reviewed and those anomalies identified as potential structural failures were used to determine the estimated frequency of an in-flight critical structural failure for various SSME components.

---

[8]Studies conducted at both Rocketdyne (Ref 8) and at MSFC (Ref 9) have demonstrated a modest but consistent growth in the reliability of the SSME with each subsequent configurational change.

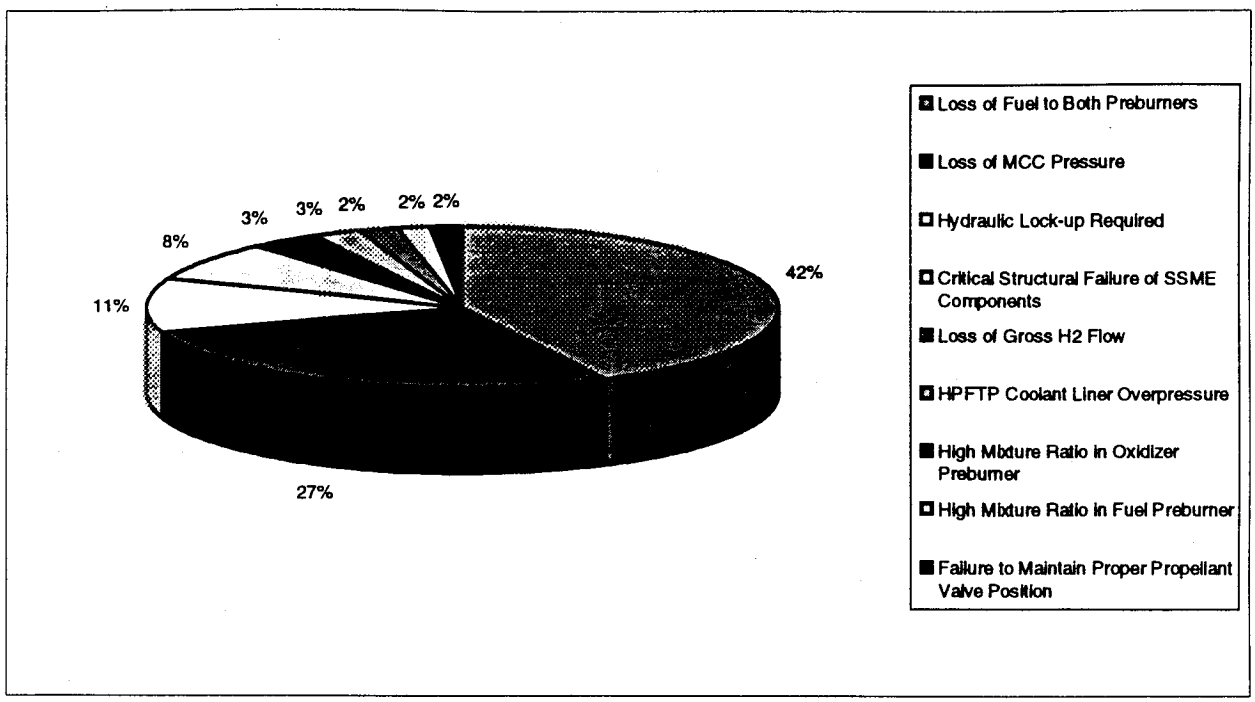

Figure 5.4. SSME Specific Initiator Frequency Distribution

The mean distribution of initiator frequencies is shown in Figure 5.4. Note that the top three initiating event contributions to frequency of occurrence are, loss of fuel to both preburners (42%), loss of MCC pressure (27%) and hydraulic lock-up required (11%). The loss of fuel to the preburners had two major causes: an erroneous flowmeter calibration coefficient which caused the Controller to starve the engine of fuel by slowing down the HPFTP and excessive leakage from the nozzle coolant tubes. Loss of MCC pressure was also caused by control loop problems and main injector erosion which is less of a problem in the current configuration. It was found that hydraulic lock-up was required largely because of failures of the hydraulic supply system rather than any one SSME failure; the catastrophic failure however occurs in the SSME.

Normally once an emergency shutdown has been commanded, the redline limits on the other two engines are inhibited to avoid a second shutdown. This is done to assure that an abort contingency window is available through out the ascent trajectory. Also the remaining SSMEs are throttled up to 109% from 104% to partially account for the loss of one engine; operating at 109% increases the likelihood of a redline exceedance since the redline margins are degraded. However, in actuality the redline inhibit command is issued once the MCC pressure in the engine which is shutting down reaches 30% of nominal, which takes approximately 3 seconds. Although single engine shutdowns are not considered catastrophic events, their frequency was determined to be high enough to consider a second engine initiating a shutdown within the 3 second window. The likelihood of a second shutdown is even higher given a common cause failure of two APUs in which case two engines go into hydraulic lock-up. When an engine is in hydraulic lock-up, redline limits are still active and the possibility of a valve drifting and causing a redline exceedance, according to system experts, is as high as 20%.

There is a number of ways in which a failure may cause an abnormal shutdown but only two critical functions need to be performed to avoid catastrophe. Firstly, the OPOV must be closed; failing to do so will cause a LOX rich shutdown which was considered to lead to burnthrough. Secondly, the oxidizer preburner must be purged with helium to avoid subsequent mixing of residual oxygen and hydrogen which is considered to lead to a catastrophic explosion.

The frequency of the Propellant Management and Helium Supply System initiators were obtained from previous studies. The probability of leakage as well as the effectiveness of mitigation processes was acquired from the *Space Shuttle Main Propulsion Pressurization System Probabilistic Risk Assessment* completed by Lockheed in 1988 (Ref. 28).

5.1.2.4. Accident Scenarios and Consequences

In the case of the SSME the control system, comprised of the sensors, harnesses, controller, actuators and propellant valves, plays a pivotal role in administering the system response to some of the initiating events. The first nine initiators from Table I all evoke a response from the SSME involving valve configuration adjustments and/or performance of an emergency shutdown. These pivotal events are separated into two groups: protective and mitigative events. Protective events are those involving active monitoring and control. Mitigative events are called upon when the protective systems were unsuccessful in negotiating the abnormalities introduced by the occurrence of the initiating event. Mitigative systems are in place to divert the accident sequence to a more desirable or non-catastrophic consequence. For the SSME this non-catastrophic end state is a single engine shutdown as opposed to an uncontained energetic discharge. An uncontained energetic discharge is assumed to occur if:

(1) *A redline condition exists and is not recognized.* It should be understood that all initiators for which a redline system exists are by definition redline condition incidents. Therefore the question in the event trees is not whether a redline will exist but rather does the system recognize and act upon the existing condition.

(2) *The engine is not able to perform the functions necessary to safely shutdown the engines during an emergency or under nominal conditions.* The main contributors to this catastrophic accident scenario are failure of the OPOV to close and failure to purge the oxidizer preburner.

(3) *A propellant evacuation is unsuccessful due to both the nominal dump and vacuum inerting failing to rid the MPS/SSME of residual propellants.*

(4) *A sudden structural failure occurs in a high pressure vessel or high angular momentum component.* These would be termed criticality 1 events meaning that no protective or mitigative systems exist during engine operation to preclude the accident sequence from going directly to LOV. All such component failures are grouped under the initiator heading of "Structural Failure of SSME Components Leading to LOV".

(5) *Uncontained release of propellant.* A conservative stance is taken on this issue, the assumption is made that all significant leakage of propellant lead to LOV; in other words an ignition source is always present. This is perhaps the initiator in the SSME/MPS with the largest uncertainty due to the fact that test conditions are very much different from the conditions in the aft-compartment. Also maintenance records offer little insight into in-flight frequency since there is no established correlation

between different leak rates and potential for catastrophic failure.

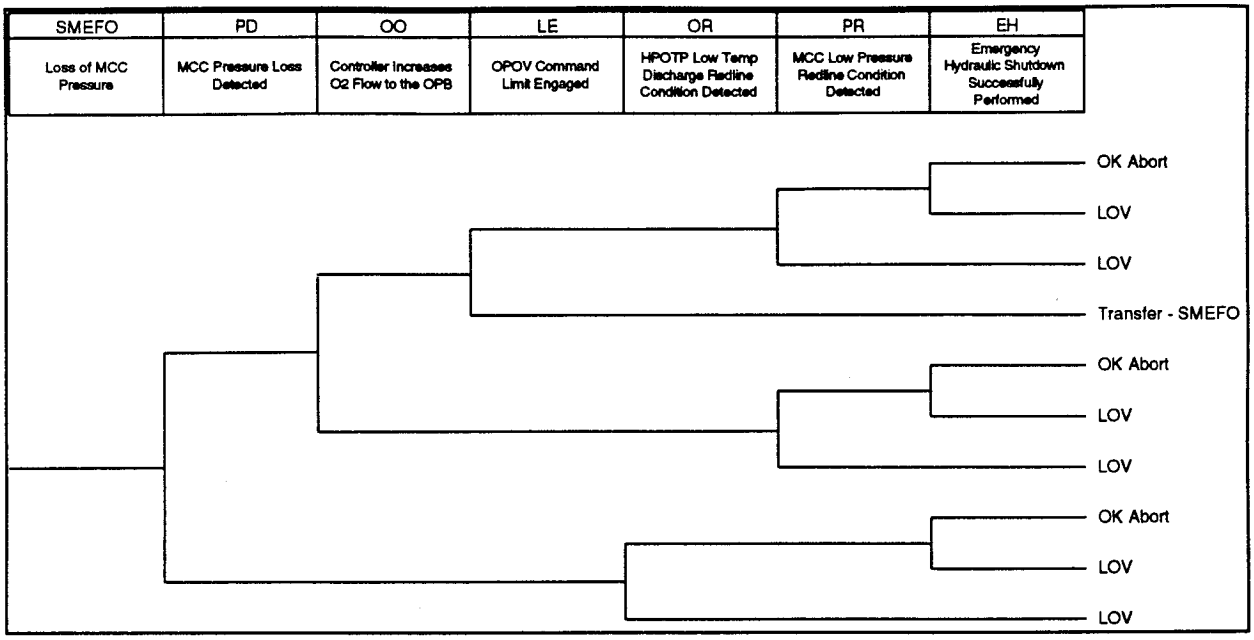

Figure 5.5. Loss of MCC Pressure Event Tree

The event sequences begin with the initiator, which is a grouping of component failures or other transient conditions which evoke a similar system response, the various protective and mitigative systems or actions are then tested. Each protective and mitigative action either succeeds in required task or fails. Success paths of each action are designated by the upward branch in the event tree while failure or nonoccurrence is shown by a downward branch. A failure branch may be supported by a fault tree if the protective or mitigative system requires a number of components to be successful. For example, in examining the loss of MCC pressure event tree (Figure 5.5), the initiating event is a decreased O2 flow beyond the compensating capability of the engine. The first system which must operate in such an event is the MCC pressure sensing system (MCC pressure is used to calculate O2 flow rate). This system requires that two redundant pressure transducers and their respective Controller channels be qualified and operating nominally. If the MCC pressure drop is successfully sensed and processed the Controller will then attempt to increase the O2 flow in the system by opening the OPOV. The OPOV will continue to open until such a time at which the position reaches the OPOV Command Limit programmed in the Controller software. If the OPOV Command Limit is not engaged, the situation is equivalent to having a high mixture ratio in the oxidizer preburner and the sequence is transferred to the High Mixture Ratio in Oxidizer Preburner event tree. Engaging the OPOV Command Limit or failing to open the OPOV in the first place will cause a drop in MCC pressure from lack of oxygen. The drop will be detected by the MCC pressure sensing system (it is assumed to be successful at this point since it was used to detect the initial pressure drop) which has a redline limit of 100psi below nominal. If the initial pressure drop was not detected there is another redline on the HPOTP turbine discharge temperature to avoid a temperature

drop which can lead to ice formation on the heat exchanger coils. Once a redline is detected the only remaining question is the successful shutdown of the engine.

The event trees for the SSME and MPS accident scenarios are shown in Appendix B.1. The following tables give the associated codes and explanations for the events constituting the event trees.

Table 5.2. Loss of MCC Pressure Event Descriptions

SMEFO	Decrease in O2 flow caused by control loop failure or main injector erosion results in low MCC pressure.
PD	Two pairs of redundant pressure sensors monitor the MCC pressure, detection of the pressure drop is transmitted to the Controller.
OO	Controller attempts to compensate for loss of MCC pressure by opening the OPOV thereby increasing the O2 flow to the OPB which increases the power to the HPOTP such that it pumps more O2 into the MCC.
LE	The OPOV has a command limit programmed into the Controller software which is the largest OPOV position allowable without triggering the HPOTP high discharge temperature redline due to a high mixture ratio in the OPB.
OR	If the pressure drop in the MCC is not detected, the decrease in O2 in the OPB should trigger the HPOTP low discharge temperature redline. Failure to do so will cause ice formation which will rupture the heat exchanger causing a catastrophic mixture of H2 and O2.
PR	If the pressure in the MCC drops below 100psi of nominal a redline is triggered and the Controller issues an emergency shutdown command. This event is considered 100% successful since the pressure detection system operated successfully in detecting the initial pressure drop.
EH	Once an emergency shutdown is commanded the processes for satisfying the command must be performed successfully to avoid catastrophic failure. The most critical functions are closing the OPOV and purging the OPB.

Table 5.3. High Mixture Ratio in OPB Event Descriptions

SMEMO	A high mixture ratio in the OPB may be caused by a control loop failure which erroneously increases O2 to the OPB or an OPOV malfunction with the same effect. Exceeding the OPOV command limit in the SMEFO event tree will also cause a high mixture ratio in the OPB.
OR	A high mixture ratio should trigger the HPOTP high discharge temperature redline. Failure to detect the redline condition will cause catastrophic failure of the OPB due to high thermal stresses.
EH	Once an emergency shutdown is commanded the processes for satisfying the command must be performed successfully to avoid catastrophic failure. The most critical functions are closing the OPOV and purging the OPB.

Table 5.4. Loss of Gross H2 Flow Event Descriptions

SMEFH	Decrease in gross H2 flow found to be caused by distortions in the HPFTP turnaround manifold.
OF	The drop in H2 flow is detected by the flow meter in the low pressure fuel duct and transmitted to the Controller which commands the FPOV to open thereby increasing the power generated to drive the HPFTP and pumping more H2. Increasing the O2 delivered to the FPB will cause a high mixture ratio condition. Failing to open the FPOV will cause a loss of fuel to both preburners.

Table 5.5. High Mixture Ratio in FPB Event Descriptions

SMEMF	A high mixture ratio in the FPB may be caused by a control loop failure which erroneously increases O2 to the FPB or an FPOV malfunction with the same effect. Increasing the O2 to the FPB due to a loss of gross H2 flow (SMEFH event tree) will also cause a high mixture ratio in the FPB.
OR	A high mixture ratio should trigger the HPFTP high discharge temperature redline. Failure to detect the redline condition will cause catastrophic failure of the FPB due to high thermal stresses.
EH	Once an emergency shutdown is commanded the processes for satisfying the command must be performed successfully to avoid catastrophic failure. The most critical functions are closing the OPOV and purging the OPB.

Table 5.6. Loss of Fuel to Both Preburners Event Descriptions

SMEPB	A loss of fuel to both preburners may be caused by a control loop failure due to an erroneous flow meter calibration constant or by leakage of H2 from the nozzle coolant channels. Failing to increase the O2 to the FPB in the event of a loss of gross H2 flow (SMEFH event tree) will also cause a decrease in fuel to the preburners..
TR	Loss of fuel to the preburners will result in a high mixture ratio in both preburners, this should trigger one of the turbopump high discharge temperature redlines. Failure to detect both redline conditions will cause catastrophic failure of the FPB or OPB due to high thermal stresses.
EH	Once an emergency shutdown is commanded the processes for satisfying the command must be performed successfully to avoid catastrophic failure. The most critical functions are closing the OPOV and purging the OPB.

Table 5.7. HPFTP Coolant Liner Overpressure Event Descriptions

SMELO	An increase in HPFTP coolant liner pressure has occurred a number of times, most recently on STS-55, although the pressure increase has not been sufficient to cause a redline condition.
OP	A significant overpressure condition should be detected by the associated pressure sensor and a redline exceedance command generated. Failure to detect the redline condition will result in the buckling of the turnaround duct and subsequent LOV.
EH	Once an emergency shutdown is commanded the processes for satisfying the command must be performed successfully to avoid catastrophic failure. The most critical functions are closing the OPOV and purging the OPB.

Table 5.8. Failure to Maintain Proper SSME Propellant Valve Position Event Descriptions

SMEVP	SSME propellant valves must be maintained at ±10% of their commanded positions. Failure of any one valve to do so will result in a Servovalve Error Indication Interrupt (SEII). Upon the generation of a SEII the Controller de-energizes the fail-safe servo-switch in all five propellant valves.
HL	Once de-energized the servo-switch directs the hydraulic fluid such that the by-pass valve is actuated into its hydraulic lock-up position. Failure of any of the servo-switches to change positions will result in an emergency pneumatic shutdown command.
EP	Once an emergency shutdown is commanded the processes for satisfying the command must be performed successfully to avoid catastrophic failure. The most critical functions are closing the OPOV and purging the OPB.

Table 5.9. Hydraulic Lock-up Required Event Descriptions

SMEHL	Hydraulic lock-up is required in the event of a loss of hydraulic pressure or a commanded lock-up due to a significant deviation of one of the SSME propellant valves from its commanded positions (SMEVP event tree).
BL	The by-pass valve is spring loaded to move to the hydraulic lock-up position in the event of loss of hydraulic pressure or if the servo-switch de-energized. Failure of the by-pass valve to move to the lock-up position means the valve cannot be pneumatically closed since this requires the by-pass valve to move past the lock-up position. If the valve happens to be the OPOV then the engine cannot be shutdown and LOV is anticipated.
ND	Once the engine is in hydraulic lock-up the valves may drift due to vibration or other causes. If there is significant movement engine operation may be effected to the extent that a redline is exceeded in which case a pneumatic shutdown is commanded.
EP	Once an emergency shutdown is commanded the processes for satisfying the command must be performed successfully to avoid catastrophic failure. The most critical functions are closing the OPOV and purging the OPB.
ME	If no significant valve drift occurs, the engine continues to fire until a nominal MECO with one engine in hydraulic lock-up is called for. In such a case all three engines must be successfully shutdown.
PM	Once all three engines have been successfully shutdown, a propellant dump must be performed to evacuate the MPS plumbing of residual propellants. Failure to due so may result in overpressurization and rupture of the MPS plumbing. If critical components are effected then LOV may result.

Table 5.10. Structural Failure of Critical SSME Components Event Descriptions

SMEST	The strenuous conditions under which the SSME operates may cause structural failures despite the meticulous care taken to obviate such occurrences. Some of these failures are categorized as criticality 1 signifying that no protective or mitigation processes exist. The occurrence of such a structural failure leads directly to LOV.

Table 5.11. Simultaneous Dual SSME Shutdown Event Descriptions

SMEDS	Two engines can shutdown simultaneously if a shutdown is commanded in one engine and a redline is violated in a second engine before the MCC pressure has reached 30% of nominal in the first engine. This may be due to independent causes or a common cause dual APU failure which hydraulically locks up two engines thereby increasing the probability of redline exceedance due to valve drift.
BL	If the dual shutdown occurs prior to lift-off (ISRB ignition) an on-pad abort occurs. Dual shutdown after lift-off and prior to the droop 109 call is considered to lead to LOV due to a lack of an abort contingency.
AC	If the dual shutdown occurs after the droop 109 call an abort scenario ensues.

Table 5.12. Nominal MECO and Propellant Dump Event Descriptions

SMECD	If no events leading to LOV or abort scenarios occur between SSME ignition and the call for MECO all three engines must be successfully shutdown and the MPS purged of residual propellants.
MN	All three engines must be shutdown, of critical importance is the closing of all the OPOVs and the purging of all OPBs. Failure to perform either of these functions is considered an LOV event.
PD	Once all three engines have been successfully shutdown, a propellant dump must be performed to evacuate the MPS plumbing of residual propellants. Failure to due so may result in overpressurization and rupture of the MPS plumbing. If critical components are effected then LOV may result.

Table 5.13. Helium System Leakage Event Descriptions

SMELH	The helium system is designed to leak before rupture. Leakage may occur from the supply tanks or along any of the plumbing including the Pneumatic Control Assembly (PCA) in the SSME.
IL	Leakages in certain parts of the helium supply system may be isolated.
IH	The proper sequence of valves must be closed in order to isolate the leakage. This is accomplished by the crew which has control of the helium supply valve positions from the cockpit.
AH	Once the leak has been isolated a crosstie must be established from another leg of the helium system to maintain the helium supply to the engine effected by the leakage.
EM	If the leakage may not be isolated for some reason or a crosstie was unsuccessful the effected engine must be manually shutdown to avoid any adverse consequences due to insufficient helium. Failure to perform a manual shutdown under these circumstances is assumed to lead to LOV.
ME	If a crosstie proved to be successful the engine continues to fire until a nominal MECO is called for. In such a case all three engines must be successfully shutdown.
PM	Once all three engines have been successfully shutdown, a propellant dump must be performed to evacuate the MPS plumbing of residual propellants. Failure to due so may result in overpressurization and rupture of the MPS plumbing. If critical components are effected then LOV may result.

Table 5.14. Failure to Provide POGO Accumulator Charge Event Descriptions

SMEPG	Mechanical failures causing the loss of helium necessary to precharge the POGO accumulator.
PP	A pressure sensor is located in the helium precharge system to assure sufficient helium pressure is available for pogo charging. A low pressure condition will be cause for an emergency shutdown. Failure to detect the low pressure condition is considered to lead to LOV due to catastrophic start-up oscillations in the oxygen system.
EH	Once an emergency shutdown is commanded the processes for satisfying the command must be performed successfully to avoid catastrophic failure. The most critical functions are closing the OPOV and purging the OPB.

Table 5.15. Leakage of SSME/MPS Propellants Event Descriptions

SMELP	Propellant leakage causing fire or explosion in the aft-compartment. It is assumed that leakage of both oxygen and hydrogen is necessary to cause a fire or explosion. In addition it is also assumed that ignition always occurs given that both elements are present.

Table 5.16. Failure to Maintain Positions of MPS Propellant Supply Valves Event Descriptions

SMEPV	The two 17-inch disconnect valves and six prevalves must remain open during the duration of SSME operation. Closure of these valves while the SSME is operating will cause the turbopumps to overspeed and come apart due to sudden loss of pump load.

5.1.2.5. Data Analysis

Fault trees were developed for most of the events described in the previous section. The fault trees are shown in Appendix B.1. The data used in quantifying the basic events came from a variety of sources including analyses done by Rocketdyne (Ref. 32 & 33), previous PRA (Ref. 28 & 41) and generic data sources (Ref. 57 & 58). The exact source used for each basic event in the SSME/MPS model is shown in basic event listing in Volume III.

5.1.2.6. SSME/MPS Risk Contribution

One of the most insightful results of the SSME analysis was the dominant role which critical structural failures play in the risk of the SSME. Figure 5.6 shows the initial SSME initiator frequency distribution and the final LOV risk distribution. The initiators which induce a redline condition very rarely propagate to catastrophic failure because of the effective nature with which the Controller recognizes and acts upon the condition by shutting down the engine. The critical structural failures, which have no active protective features, end up dominating the SSME related risk even though they represent only 7.8% of the initiator frequencies. The contribution of the various SSME components to this risk is shown in Figure 5.7. Note that the HPOTP, HPFTP, and MCC account for 79% of the mean SSME risk. The failure modes constituting the risk of each of these components are shown in Figures 5.8, 5.9, and 5.10.

Figure 5.6. Propagation of SSME Specific Accident Initiating Events

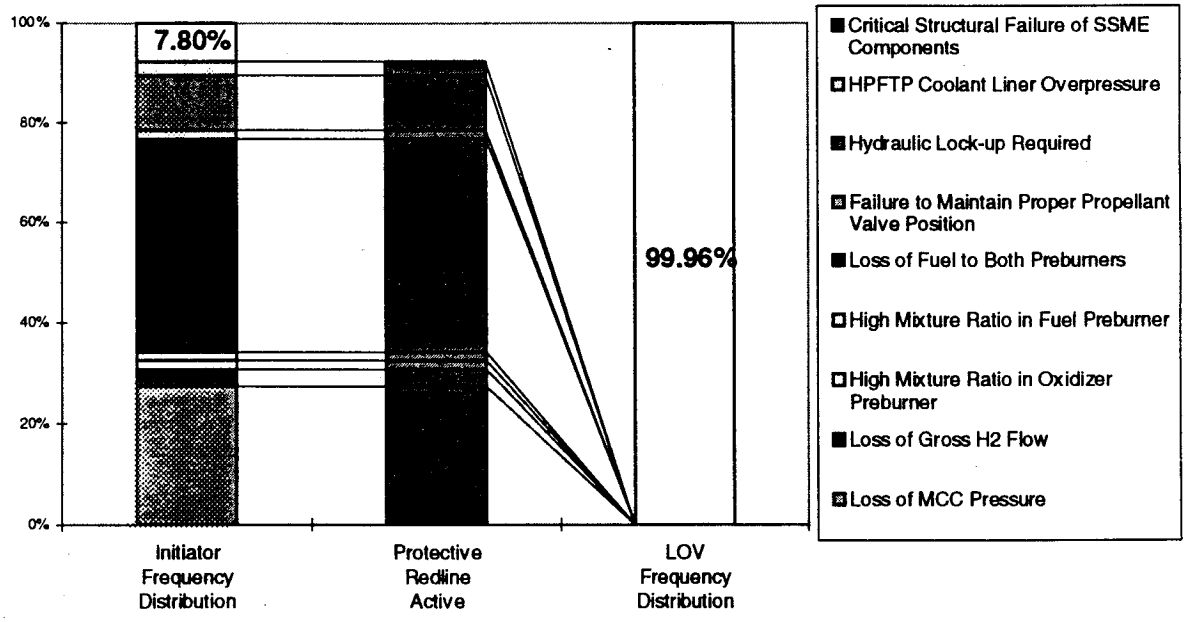

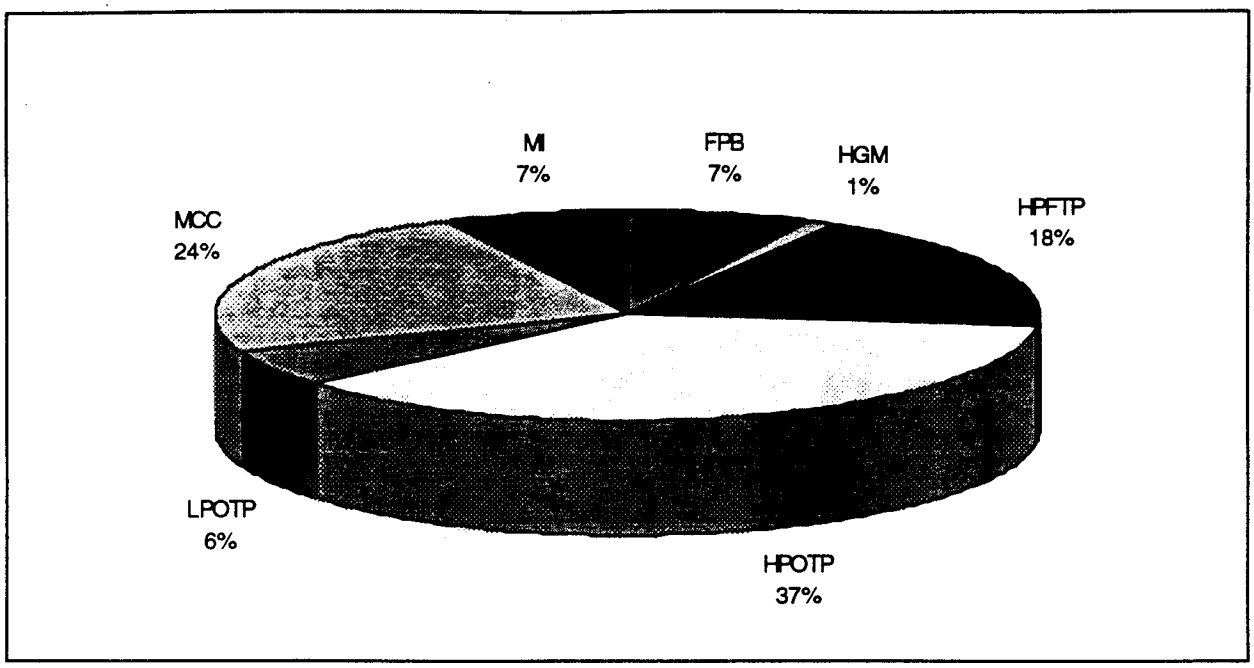

Figure 5.7. SSME Risk Contribution*

*Components which contribute less than 1% to the SSME critical structural failure risk are not shown: LPFTP, OPB, HEX, & Nozzle

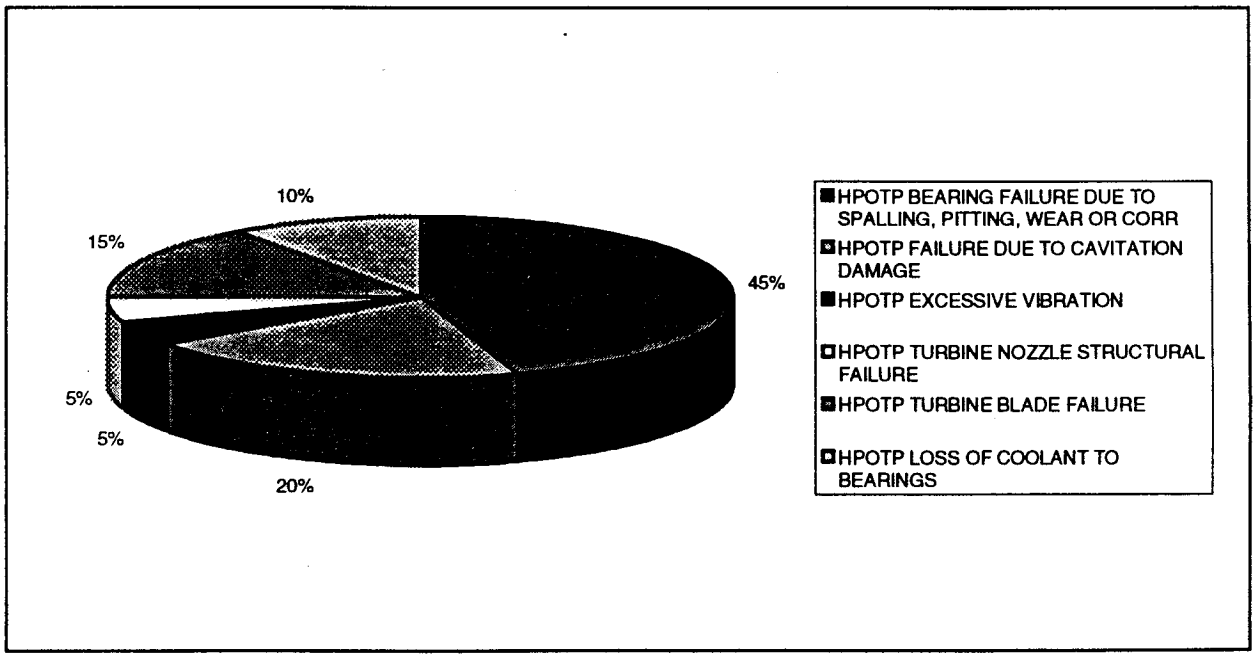

Figure 5.8. HPOTP Failure Mode Risk Contribution

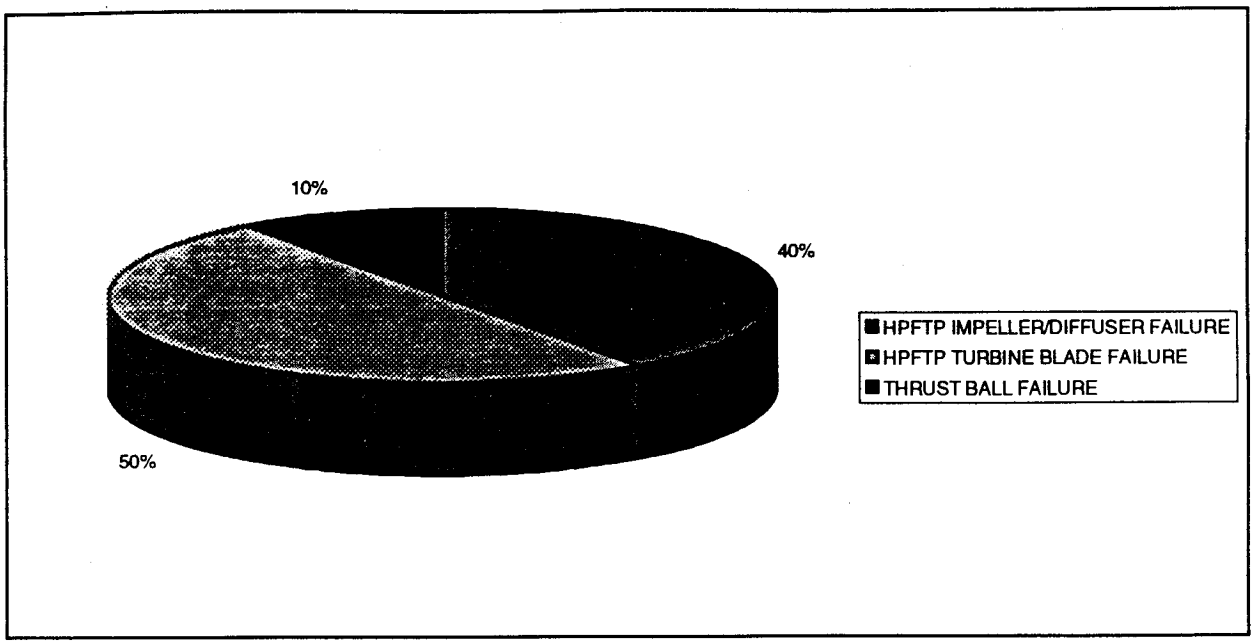

Figure 5.9. HPFTP Failure Mode Risk Contribution

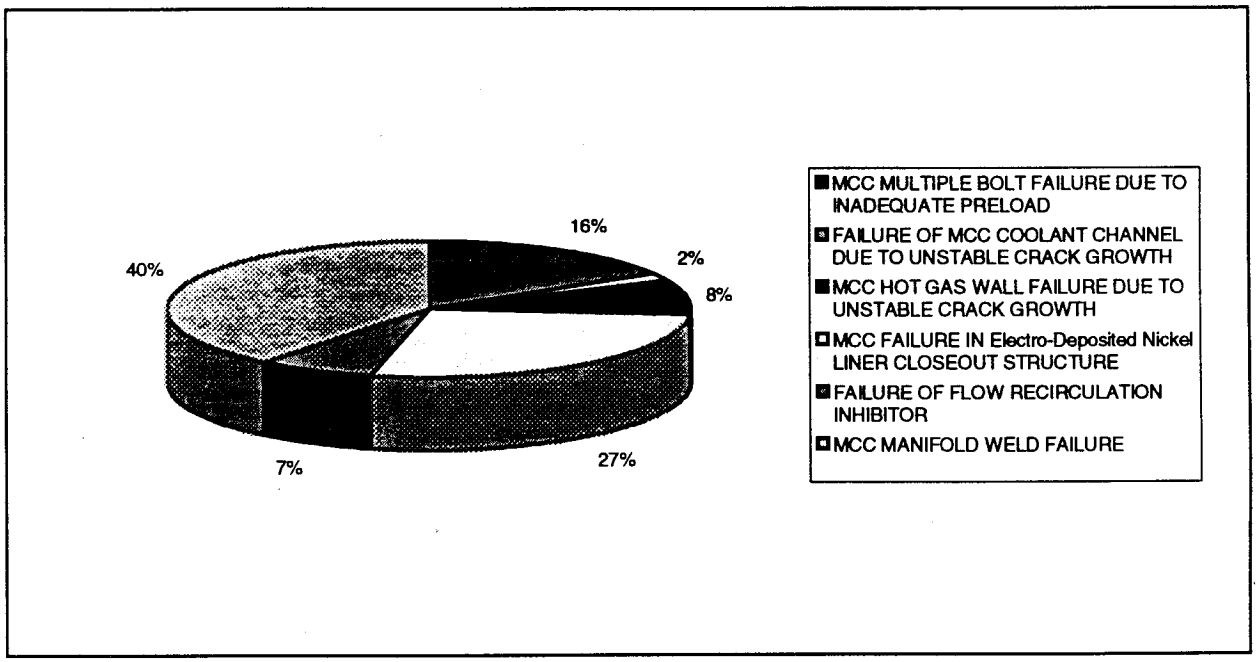

Figure 5.10. MCC Failure Mode Risk Contribution

Not surprisingly, from the discussion above, the next most significant contributor to SSME risk, although an extremely distant second, is the simultaneous shutdown of two engines. This incident accounts for 0.23% of the risk (1 in 155,000 missions) related to the SSME. Note however that approximately 50% of this risk is related to a common cause failure of two APU resulting in a loss of hydraulic pressure. The shutdown process at MECO was found to be extremely reliable; catastrophic failure was estimated at 1 in 375,000 missions.

The combined risk of the propellant management and helium supply system initiators is 1 in 46,000 missions of which 86% is accounted by failure to maintain the propellant supply valves open during SSME operation.

In conclusion, it is evident that the SSME/MPS risk is almost entirely attributable to critical structural failures. The most significant of the contributing failure modes are the HPOTP bearing failure, the HPFTP turbine blade failure and the MCC manifold weld failure. Although their quantitative contribution to risk has not been identified until now, these failure modes have long been recognized as important risk contributors by the engineers at Rocketdyne and NASA. This recognition has prompted requests for redesigned components and resulted in new hardware designs such as the Advanced Technology Developmental (ATD) HPOTP. This new pump design is currently undergoing certification for future flight. Additional development of other advanced technology SSME components has been suspended at this time.

5.1.2. Integrated Solid Rocket Booster

The integrated solid rocket booster (ISRB) refers to the combined solid rocket booster (SRB) and redesigned solid rocket motor (RSRM). The SRB consists of the structure, separation systems, operational flight instrumentation, recovery avionics, pyrotechnics, deceleration system, thrust vector control, and range safety destruct system. Each RSRM consists of its case, propellant, igniter, and nozzle.

5.1.2.1. Description

The two RSRM elements of the ISRBs provide the main thrust necessary to lift the Space Shuttle off the pad and up to an altitude of about 150,000 feet. In addition, the two SRBs carry the entire weight of the ET and Orbiter and transmit this static weight load through their structure to the mobile launcher platform. Each booster has a thrust (sea level) of approximately 3.3 million pounds at launch. They are ignited after the thrust level of the three SSMEs is verified. The two ISRB provide 71.4 percent of the thrust at lift-off and during the first stage of ascent.

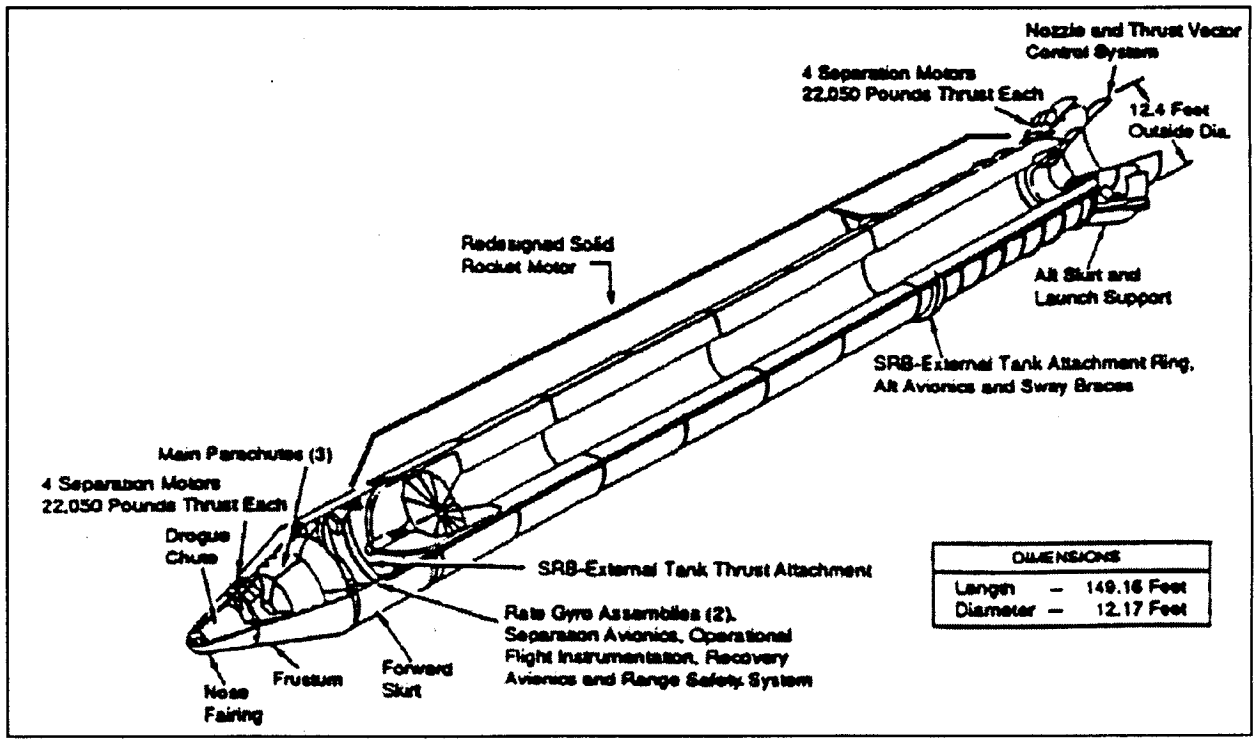

Figure 5.11. ISRB Diagram

Each booster is attached to the ET at the SRB's aft frame by two lateral sway braces and a diagonal attachment. The forward end of each SRB is attached to the ET at the forward end of the SRB's

forward skirt. On the launch pad, each booster also is attached to the mobile launcher platform at the aft skirt by frangible bolts which are fractured by small explosives at lift-off. The propellant mixture in each RSRM consists of an ammonium perchlorate (oxidizer, 69.6 percent by weight), aluminum (fuel, 16 percent), iron oxide (a catalyst, 0.4 percent), a polymer (a binder that holds the mixture together, 12.04 percent), and an epoxy curing agent (1.96 percent). The propellant is an 11-point star-shaped perforation in the forward motor segment and a double-truncated-cone perforation in each of the aft segments and aft closure. This configuration provides high thrust at ignition and then reduces the thrust by approximately a third 50 seconds after lift-off to prevent over-stressing the vehicle during maximum dynamic pressure.

The cone-shaped aft skirt transmits the aft loads between the SRB and the mobile launcher platform. The four aft separation motors are mounted on the skirt. The aft section contains avionics, a thrust vector control system that consists of two auxiliary power units and hydraulic pumps, hydraulic systems, and a nozzle extension jettison system. The forward section of each booster contains avionics, a sequencer, forward separation motors, a nose cone separation system, drogue and main parachutes, a recovery beacon, a recovery light, a parachute camera on selected flights, and a range safety system.

Each SRB has two integrated electronic assemblies, one forward and one aft. After burnout, the forward assembly initiates the release of the nose cap and frustrum and turns on the recovery aids. The aft assembly, mounted in the ET/SRB attach ring, connects with the forward assembly and the Orbiter avionics systems for SRB ignition commands and nozzle thrust vector control. Each integrated electronic assembly has a multiplexer/demultiplexer, which sends or receives more than one message, signal, or unit of information on a single channel.

Eight booster separation motors (four in the nose frustum and four in the aft skirt) of each SRB burn for 1.02 seconds at SRB separation from the ET. Each solid rocket separation motor is 31.1 inches long and 12.8 inches in diameter.

Location aids are provided for each SRB, frustum/drogue chutes, and main parachutes. These include a transmitter, antenna, strobe/converter, battery, and salt water switch electronics. The location aids are designed for a minimum operating life of 72 hours and when refurbished are considered usable up to 20 times. The flashing light is an exception. It has an operating life of 280 hours. The battery is used only once.

The SRB nose caps and nozzle extensions are not recovered.

The recovery crew retrieves the SRBs, frustum/drogue chutes, and main parachutes. The nozzles are plugged, the solid rocket motors are dewatered, and the SRBs are towed back to the launch site. Each booster is removed from the water, and its component are disassembled and washed with fresh and deionized water to limit salt water corrosion. The motor segments, igniter, and nozzle are shipped back to the manufacturer for refurbishment.

Each SRB has four hold-down posts that fit into corresponding support posts on the mobile launcher

platform. Hold-down bolts hold the SRB and launcher platform posts together. Each bolt has a nut at each end, but only the top nut is frangible. The top nut contains two NASA standard detonators (NSDs), which are ignited at solid rocket motor ignition commands.

When the two NSDs are ignited at each hold-down, the hold-down bolts travels downward because of the release of tension in the bolt (pretensioned before launch), NSD gas pressure, and gravity. The bolt is topped by the stud deceleration stand, which contains sand. The frangible nut is captured in a blast container.

The solid rocket motor ignition commands are issued by the orbiter's computers through the master events controllers to the hold-down pyrotechnic initiator controllers (PICs) on the mobile launcher platform. They provide the ignition to the hold-down NSDs. The launch procession system monitors the SRB hold-down PICs for low voltage during the last 16 seconds before launch. PIC low voltage will initiate a launch hold.

SRB ignition can occur only when a manual lock pin from each SRB safe and arm device has been removed. The ground crew removes the pin during prelaunch activities. At T minus 5 minutes, the SRB safe and arm device is rotated to the arm position. The solid rocket motor ignition commands are issued when the three SSMEs are at or above 90-percent rated thrust, no SSME fail and/or SRB ignition PIC low voltage is indicated, and there are no holds from the launch processing system.

The solid rocket motor ignition commands are sent by the orbiter computers through the master events controllers (MECs) to the safe and arm device NSDs in each SRB. A PIC single-channel capacitor discharge device controls the firing of each pyrotechnic device. Three signals must be present simultaneously for the PIC to generate the pyro firing output. These signals, arm, fire 1 and fire 2, originate in the GPCs and are transmitted to the MECs. The MECs reformat them to 28-volt dc signals for the PICs. The arm signal charges the PIC capacitor to 40 volts dc (minimum of 20 volts dc).

The fire 1 and 2 commands cause the redundant NSDs to fire through a thin barrier seal down a flame tunnel. This ignites a pyro booster charge, which is retained in the safe and arm device behind a perforated plate. The booster charge ignites the propellant in the igniter initiator, and combustion products of this propellant ignite the solid rocket motor initiator, which fires down the length of the solid rocket motor igniting the solid rocket motor propellant.

5.1.2.2. Success Criteria

The ISRB participates in providing three top-level STS functions; these are failure to provide proper propulsion, failure to contain energetic gas or debris, and failure to maintain proper configuration (i.e. TVC). Unlike the SSME, the ISRBs do not act as redundant units for each other in the event of a loss of propulsion or thrust vectoring capability. A failure of either ISRB to provide these functions was considered to lead to LOV.

The RSRMs have been found to experience chamber pressure spikes which tend to cause temporary increases in thrust. This issue was the topic of analysis in the first task of this project. During the course of that analysis it was determined that 124,000 lb thrust imbalance or greater would most likely produce shear stresses on the Shuttle stack sufficient enough to compromise the integrity of the ET structure.

The thrust vector control system for the ISRB must continue to function throughout the first stage of ascent to avoid LOV. Given the enormous thrust produced by each RSRM, any failure to control the orientation of that thrust is considered to lead to irrecoverable rotation of the Shuttle vehicle.

The ISRB must be released from the launch platform precisely at ISRB ignition, any failure which results in an improper holddown release has the potential of causing catastrophic damage to the ISRB aft skirt. For the purposes of this study holddown failures are assumed to be catastrophic occurrences.

During the first 128 seconds of ascent the thrust provided by the ISRB is critical to mission safety, however once the RSRMs have depleted their propellant their ejection from the vehicle is just as critical. Failure to separate the ISRBs from the ET at the proper time will induce adverse aerodynamic forces leading to vehicle breakup. A similar consequence will ensue if the recovery devices in the SRB nose release prematurely.

5.1.2.3. Initiating Events

Any external hot gas leakage was assumed to lead directly to LOV. This is admittedly a conservative assumption but past history supports this stance. Four types of mechanisms for gas leakage were identified:

- Leakage of RSRM Joints
- RSRM Nozzle Rupture
- RSRM Pressure Vessel Rupture
- SRB Structural Failure

Sustained or large transient thrust deviations (i.e. pressure spikes) can lead to high stress conditions or vehicle yaw rates which can cause a LOV. The following initiators were identified for these scenarios:

- RSRM Wrong Thrust
- No or Late Ignition of 1 SRB/RSRM

Configurational failures are malfunctions involving changes in orientation or physical connections between components. In this respect all configurational failures are due to SRB malfunctions and

these were identified as:

- SRB Thrust Vector Control System Failure
- SRB No, Late or Improper Holddown
- SRB Holddown Premature Release
- SRB Fails to Separate
- SRB Recovery Device Premature Release
- Premature Separation

These initiators do not evoke a system response but instead lead directly to a LOV. Their occurrence is not protected against by an active closed loop feedback control but instead rely on the success of more or less passive design features to obviate LOV. Protection from failure by the use of passive design features relies on the reliability of the components which provide the function without failure. There is no active monitoring system which can mitigate the accident sequence. Fault tolerance is gained by designing redundancy into a particular sub-component. For instance, RSRM joints are protected against leaking by a seal package with a number of different seals, although the failure of one seal is not necessarily catastrophic no active intervention can be is taken to reduce the stress upon the remaining seals. For this reason avoidance of a leakage is dependent upon the reliability of all the seals in the joint, the passive design feature of the joint seal system, and the guarantee of the integrity of the seal by preflight test and inspection.

The premature separation event was not considered incredible but it requires the active failure of multiply redundant electronic systems or a human error by the RSO. The RSO failure was out of scope and the multiple active failures have a frequency in the 10^{-6} range. Common cause failures are minimized by the active nature of the command loop and by separate arm and fire functions. Given these conditions the premature separation initiator was considered a non-significant risk contributor from this top-level analysis and not explicitly modeled.

5.1.2.4. Fault Trees

Since all of the initiators lead directly to LOV and no active failure responses are possible the event trees which were originally developed were not utilized. Instead the FESDs were converted directly into fault tree format. The fault trees contain the minimum set of sub-component failures, known as cutsets, which can cause the initiator to occur and thus lead to LOV. The FESDs from which the fault trees were derived were developed as discussed in section 3.3. The fault trees are shown in appendix B.2.

5.1.2.5. Data Analysis

Probability distributions for the failure rate of the sub-components included in the fault trees are determined from examining historical evidence. This poses a problem in the case of the ISRB because it is not routinely hot-fire tested; the vast majority of the data available is flight related. When there

is insufficient direct experience for a particular event to estimate its failure rate, it is necessary to make an estimate based on the rate of occurrences of similar events in similar environments. All US solid rocket experience was recieved as possibly relevant to establish a surrogate data set for the RSRM. From this data set the following solid rocket systems were analyzed for the purposes of serving as surrogate failure data sources for the ISRB:

- Castor IV
- Minuteman III Stage 1
- Poseidon C3
- Titan SRM

These solid rockets were selected because their size and construction were considered to be most similiar to the RSRM. However the analysts were well aware that even in these cases there remained significant differences between this set and the RSRM. Differences in physical characteristics and safety factors were accounted for and the aggregated surrogate failure rate was Bayesian updated with the RSRM specific data. This method was specifically applied to the determination of pressure vessel structural and thermal failure.

Leakage of RSRM Joints

Leakage of hot gas from any of the RSRM joints (igniter-to-case joint, igniter internal joints, case joints, nozzle joints) is considered an LOV event. The leakage of joints was likewise amended with surrogate data, in this case leak check data was used to supplement hot fire results. Because conditions cannot be considered to be identical leak check successes were only partially credited towards a hotfire success. Three criteria used in determining the amount of credit to give for a particular leak check:

1. Magnitude of pressure applied to the seal
2. Direction of pressure applied to the seal
3. Motor gap dynamics

A leak potentiality factor was assigned to the leak check according to how well it represented actual hotfire conditions; measured by the similarity of the three criteria above. The leak check was counted as 90% if all the criteria were met, 60% for two criteria and 20% if only one criterion was met. The hotfire equivalent failure rate was bayesian updated with actual hotfire data to arrive at the estimated failure rate. Since the leak checks were mostly successful, their inclusion in the estimate tended to decrease the failure rate of joint leakage as will be shown in section 6.3.

RSRM Wrong Thrust

Thrust deviations were attributed to two causes: slag accumulation and inhomogeneous iron oxide. Slag accumulation was concluded to be the most likely cause of the pressure spikes in the RSRM

observed during a number of Shuttle missions[9]. It was thought that the thrust transients could possibly cause the forward ET load bearing connection to fracture. If this were true then pressure spiking would be a considerable risk contribution and would change the magnitude of the ISRB contribution to the Shuttle significantly. However the PSA team was assured that NASA MSFC had constructed a sophisticated 3D two-phase computational fluid dynamics model and had performed simulations which indicated that it was physically impossible for the slag accumulation and ejection to develop to a degree which would induce maximum thrust transients sufficient for such an event. Therefore although the possibility of thrust transients due to slag accumulation (basic event APSLAG) is rather high at 1 in 33 missions, the possibility of the slag accumulation causing a catastrophic failure (basic event LOV_APSLAG) was considered as physically noncredible and a probability of zero was assigned to that event. There was no data available on the frequency and effects of inhomogeneous iron oxide in the propellant thus an estimate of 1 in 10,000 missions was assigned to both the occurrence and possibility of catastrophic consequences giving a probability of LOV due to inhomogeneous oxide of 1 in 100,000,000 missions.

No or Late Ignition of 1 SRB/RSRM

No or late ignition of one RSRM would cause the Shuttle to tip over on the launch pad, obviously a LOV consequence. This could be caused by two possibilities, either the propellant fails to ignite or the igniter does not function. The possibility of the propellant not igniting is an issue with little data and much debate, conversations with system experts set the mean probability at a conservative value of 1 in 1000 ignition attempts. The malfunctioning of the igniter may be due to failures of NSIs and PICs constituting the ignition mechanism. The various component failure modes were supplied by USBI.

SRB Thrust Vectoring Control System Failure

As mentioned in section 5.1.2.2., the thrust vectoring for the ISRB must operate throughout the first stage of ascent (lift-off to ISRB separation). There are four possible causes of a TVC malfunction which will cause a LOV; any one of four actuators fail, failure to supply hydraulic pressure to any actuator, gimbal joint failure, or failure to supply electrical power. The probability of the gimbal joint failing was obtained from Ref. 57. Hydraulic pressure is supplied by two redundant HPU which are similar in both design and operation to the Orbiter system hydraulic system which will be discussed in section 5.1.3.1.2.. Therefore the failure probability of an HPU is evaluated as the product of the failure rate of the Orbiter system and the operational time of an HPU (128 sec). Given the level of independence between the hydraulic systems, common cause HPU failures are assumed to account for 1 in 100 failures. Start up failures are not considered since the HPUs are in steady state operation at SSME ignition (beginning of risk profile). The failure rates for the components constituting the actuator systems were obtained from a generic database (Ref. 57) because they are not specific to the Shuttle. Common cause failures of the actuator components are assumed to account for 10% of all failures. Failures of the electrical system will be discussed in section 5.2.1.

---

[9]The most significant RSRM chamber pressure excursions were experienced in consecutive missions, STS-54 and STS-55, with deviations of approximately 13 psi.

No, Late or Improper Holddown

No, late or improper holddown release resulting in a catastrophic consequence may be due to two main causes: a bolt fails to release or two or more bolts are hung up. The bolt release mechanism includes the same pyrotechnic elements which constitute the RSRM ignition system. Therefore the quantification of the basic events is likewise performed using the same data as was used for the ignition system. The probability have having a holddown bolt hang up was estimated from incidences documented in PRACA records. From the frequency noted in the records a probability of having one bolt hang up of 1 in 260 missions was determined. Considering the permutations of two independent hang ups given four bolts leads to a probability of 1 in 11,300 missions for a catastrophic multiple hang up. In addition another scenario was included involving catastrophic damage to the ET or Orbiter due to holddown fragment debris impingement. The probability of such an occurrence was obtained from a USBI analysis and was estimated as 1 in 1 million missions. Premature release is believed to be a low probability event by system experts, a conservative estimate of 1 in 625,000 missions was made based upon CDF failure estimates.

SRB Fails to Separate

Separation of the ISRBs necessitates the successful operation of pyrotechnic and control function components. The failure rates of the NSI pressure cartridges were assumed equivalent to the NSIs in the ignition system. The PIC failure rates were estimated at three times the USBI estimates a conservative measure. The GPC and MDM failure modes resulting in a separation malfunction were quantified using previous failures found within PRACA. Other electrical components of the control system were standard and thus a generic databases were used to obtain their failure rates (Ref. 64). Failure of the MEC to generate an arm signal is estimated to occur at a frequency of 1 in 100,000 missions from a USBI analysis.

5.1.2.6. ISRB Initiator Risk Contribution

The final mean risk contribution of significant contributors (>1% of ISRB risk) is shown in Figure 5.12. Hot gas leaks continue to be the major risk contributors, accounting for 31% of the ISRB risk, despite the additional successful evidence provided by the leak checks. The next most risk significant events are SRB No, Late or Improper Holddown Release and No or Late Ignition of One ISRB with respective contributions of 21% and 18%. Note that this indicates that approximately 39% of the ISRB risk is concentrated at the moment of lift-off.

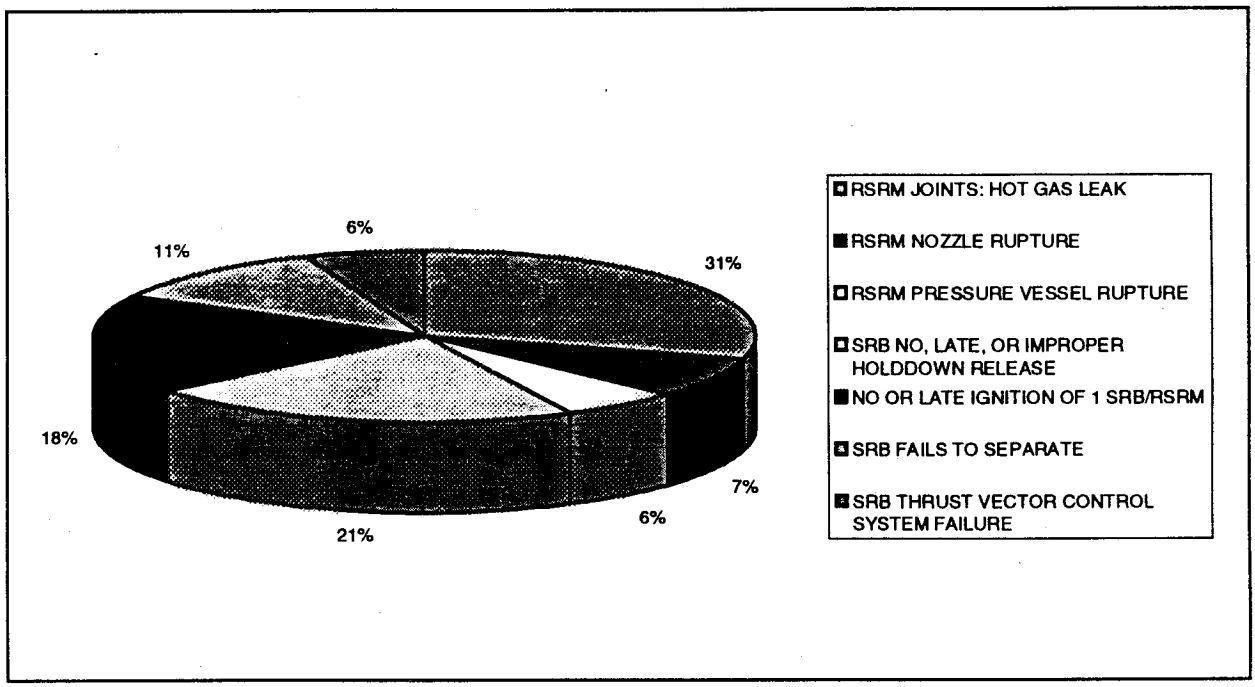

Figure 5.12. ISRB Risk Contribution

5.1.3. Orbiter Auxiliary Power Units

The Orbiter has three independent hydraulic systems, which are used to operate the aerosurfaces, direct the main engines, deploy and steer the landing gear, apply the brakes and retract the external tank/umbilical plate when the external tank separates from the Orbiter.

The power for the hydraulics is supplied by three identical auxiliary power units (APUs). The APUs convert the chemical energy in liquid hydrazine into mechanical shaft power to drive the hydraulic system main pumps, which in turn power the hydraulic systems.

To ensure a high level of reliability, the hydraulic systems are coupled together to supply redundant power to the various sub-systems. Table 5.16 gives a summary of how the hydraulic systems are used to provide redundancy.

Table 5.16. Summary of Hydraulic System Redundancy

Functions	Primary	Standby #1	Standby #2
NLG Uplk & Strut	1		
NLG Steering	1	2	3
RMG Uplk	1		
RMG Strut	1		
RMG Indbd Bks	1,2	3	
RMG Outbd Bks	1,2	3	
LMG Uplk	1		
LMG Strut	1		
LMG Inbd Bks	1,2	3	
LMG Outbd Bks	1,2	3	
L Outbd Elvn	3	1	2
L Inbd Elvn	2	1	3
R Outbd Elvn	2	1	3
R Inbd Elvn	3	1	2
Body Flap	1,2,3		
RUD/SPBK Logic	1	2	3
RUD Motors	1,2,3		
SPBK Motors	1,2,3		

5.1.3.1. Description

The Space Shuttle Orbiter has three independent hydraulic systems similar to those found on large aircraft. These hydraulic systems are used to actuate the Orbiter aero-surfaces, throttle and gimbal the Orbiter main engines, deploy and steer the landing gear, apply the landing gear brakes, and retract the external tank/umbilical plates when the external tank separates from the Orbiter. Figure 5.13 provides a schematic of the APU, Hydraulic and Water Spray Boiler assemblies.
Power for the Orbiter hydraulic systems is provided by three identical APUs, one for each hydraulic system. These APUs and their controllers are mounted on the forward bulkhead of the Orbiter aft compartment, as shown in Figure 5.14, and generate power by means of a catalytic reaction of liquid hydrazine.

The APUs are operated by the Orbiter flight crew, using flight deck controls and displays. The APUs cannot be controlled by ground command uplink. However, extensive telemetry on APU status is available to Space Shuttle ground controllers.

In a typical flight, the three APUs are started 5 minutes before lift-off and operate throughout the launch phase. They are shut down after the Orbital Maneuvering System (OMS) orbit insertion burn when hydraulic power is no longer required. The APUs are restarted for the deorbit burn and entry, and are shut down shortly after landing. In addition, one APU is usually run briefly the day before de-orbit to support a checkout of the Orbiter flight control system.

While the APUs are operating, they obtain lube oil cooling from three separate water spray boilers, one for each APU. During the inactive period on orbit, APU fluids are maintained within desired temperature ranges by thermostatically controlled heaters.

The APU is designed to achieve a high output of power in a compact package. It accomplishes this by means of a catalytic reaction of liquid hydrazine. This reaction produces a high velocity flow of hot gas, which is used to spin a turbine. A speed reduction gearbox transmits the power of the spinning turbine to the associated Orbiter main hydraulic pump.

The hydrazine fuel supply is stored in a 28-inch diameter titanium fuel tank and is pressurized with nitrogen during servicing. The gas pressure provides start capability through the fuel pump bypass valve until the fuel pump is running, and acts against the tank diaphragm to positively expel fuel to the APU. The fixed-displacement APU fuel pump provides a constant flow of hydrazine to the Gas Generator Valve Module (GGVM) after the initial bootstrap start. Approximately 325 lbs. of fuel is loaded into each fuel tank for a typical mission.

The APU turbine speed is controlled by the GGVM. The valve module consists of two flapper-type valves in series. The primary or modulating valve downstream of the pump is normally open and allows flow to the secondary or shutoff valve. The secondary valve is normally in by-pass, which directs hydrazine flow back to the pump inlet. In the powered state, it allows hydrazine flow to the

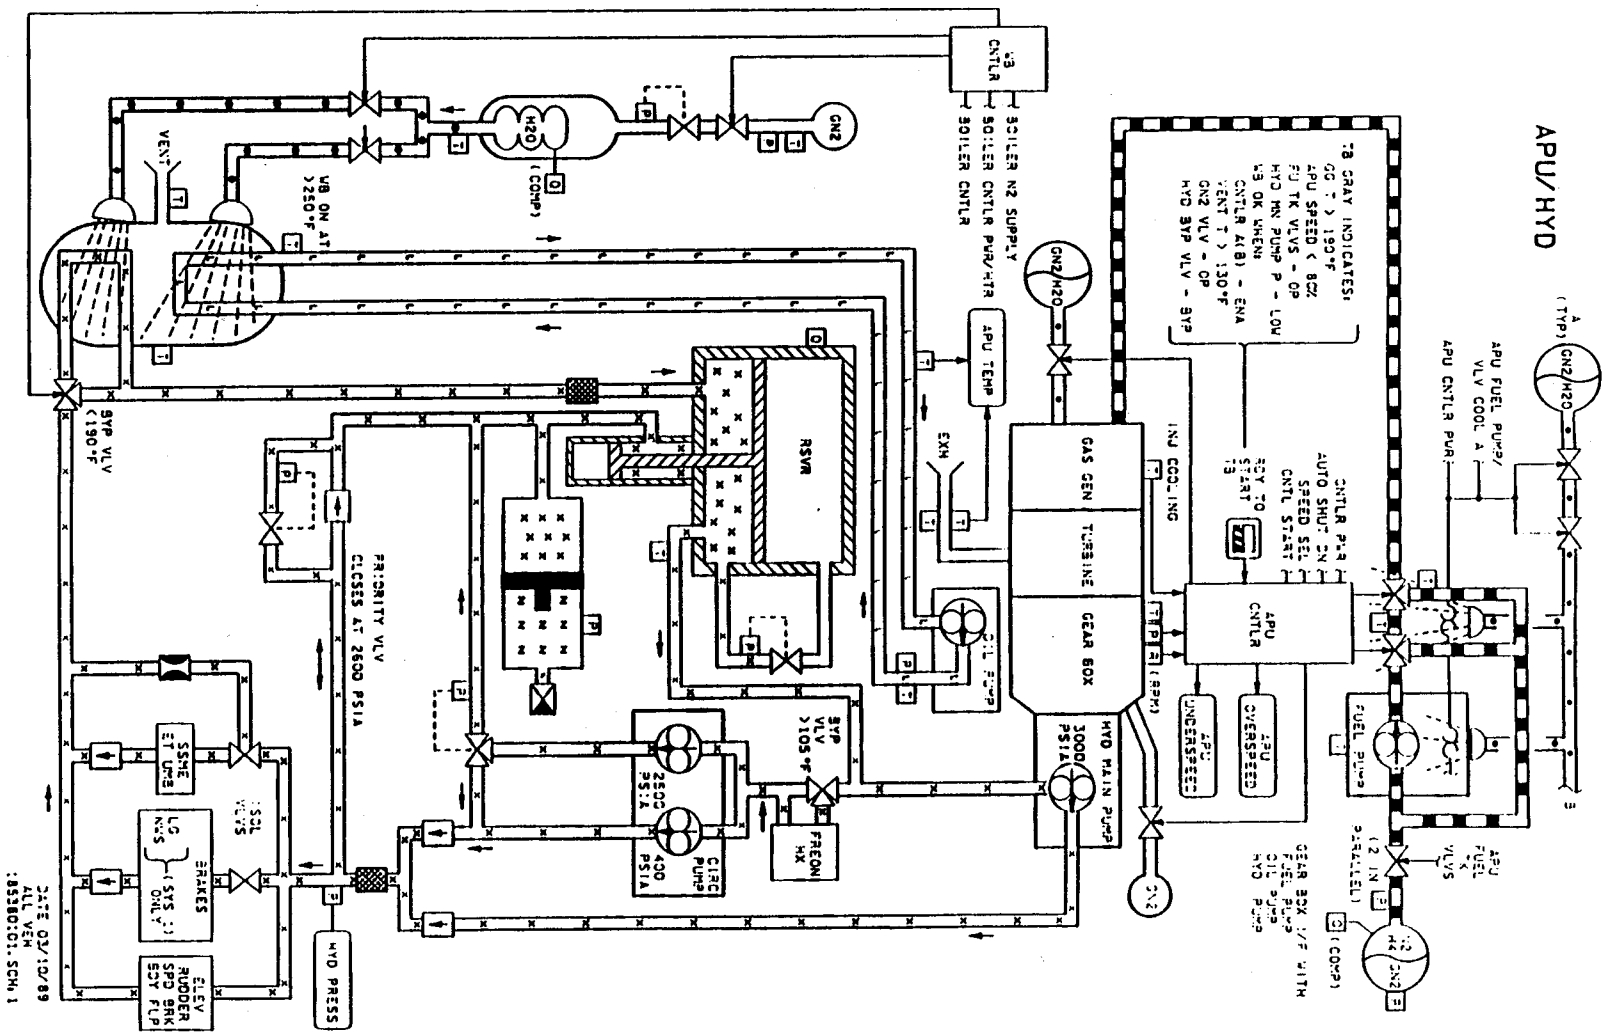

Figure 5.13. APU/HYD/WSB Schematic

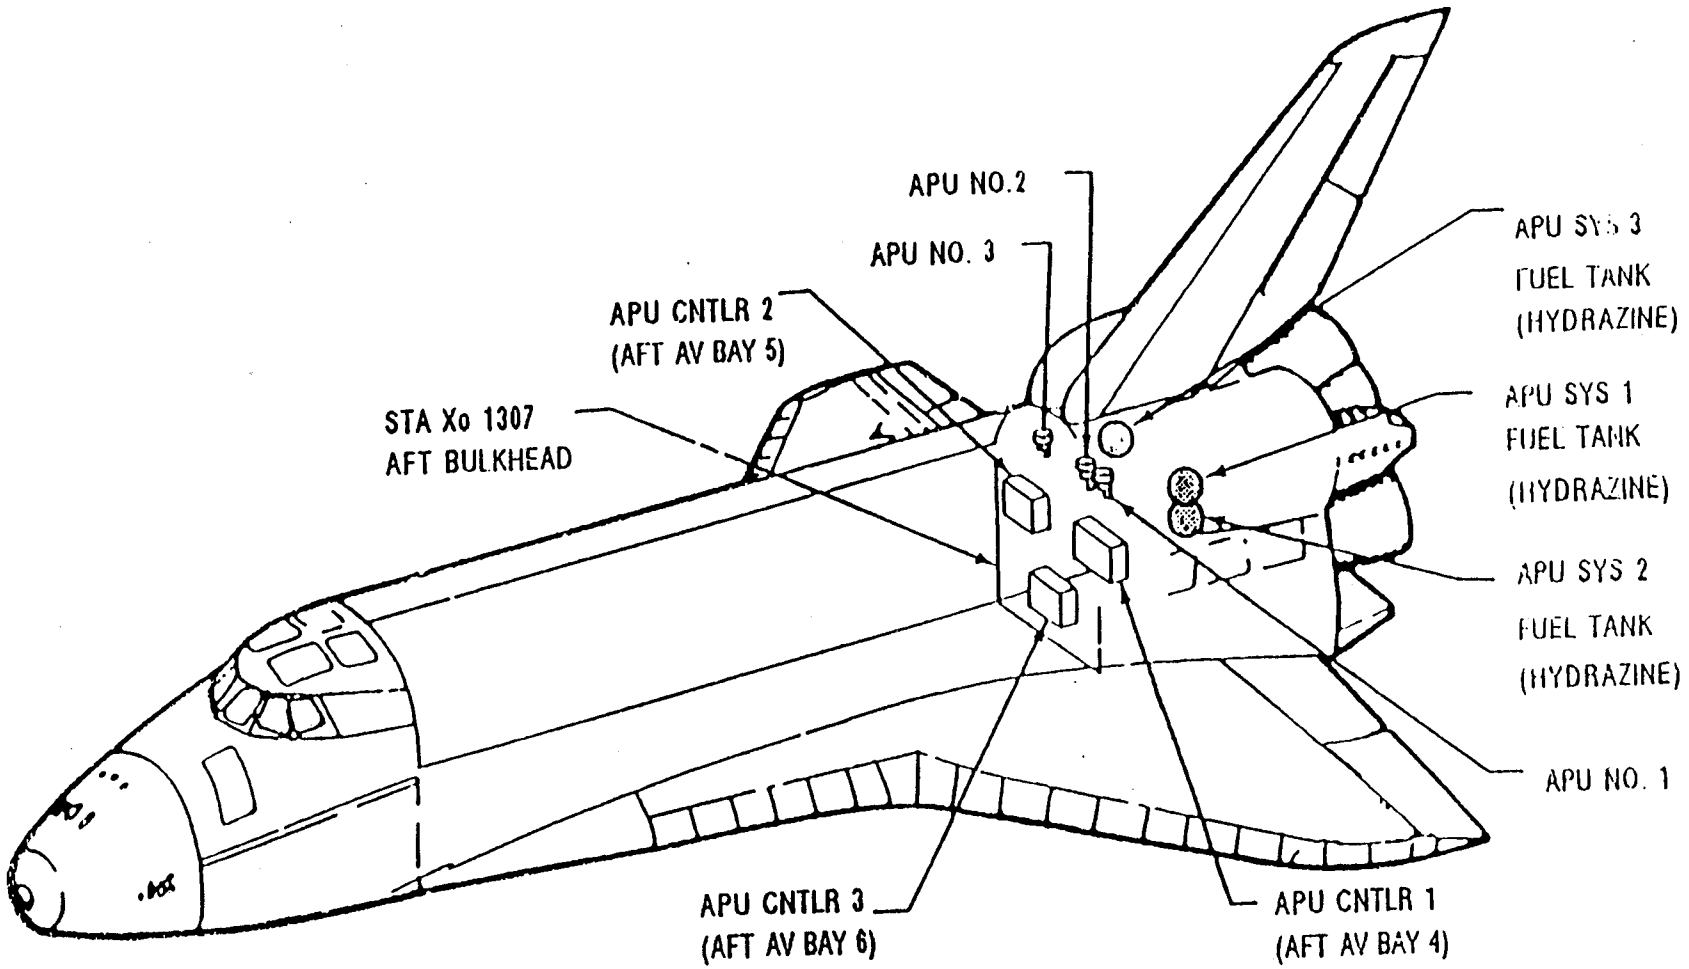

Figure 5.14. Auxiliary Power Unit Location

gas generator. The APU controller cycles the primary valve to maintain proper turbine speed (about 74,000 rpm). In the high speed mode, the controller cycles the secondary valve to maintain a speed of about 81,000 rpm. For safety, the primary valve will begin pulsing again to maintain a speed of about 83,000 rpm if the secondary valve fails open.

The gas generator (GG) is a pressure vessel containing a granular catalyst. Hydrazine flowing into the GG is decomposed by the catalyst, producing hot gases which are directed to the turbine assembly.

The dual-pass turbine assembly converts hot gas kinetic energy into mechanical shaft power at the desired speeds to operate the hydraulic pump, APU lube oil pump, and APU fuel pump.
The speed-reducing gearbox contains gears, bearings, seals, and a scavenger lubrication system. The gearbox is pressurized with nitrogen to prevent vaporization of the lubricant. A lube oil pump circulates the lube oil to the hydraulic system water boiler for cooling. The gearbox has a make-up pressurization system consisting of a small GN2 bottle and a solenoid shutoff valve actuated by the controller.

The APU electronic controller provides turbine speed control based on rotational speed sensors, logic for APU startup and shutdown, signal conditioning, gas generator catalyst bed heater control, gearbox make-up pressure control, and malfunction detection capability (flight crew alert signals to the Orbiter caution and warning system). Each controller is located remotely from its respective APU. One is located in each of the three aft avionics bays.

The exhaust duct assembly directs the APU exhaust products overboard through an exit at the upper aft fuselage skin. Exhaust duct assemblies 1 and 2 are located on the port side and duct 3 is on the starboard side of the aft fuselage at the base of the vertical stabilizer.

All APU fluid components (pumps, valves, lines) are equipped with thermostat-controlled heaters to maintain fluid temperatures in proper ranges during the APU quiescent period on orbit and pre-launch. Heaters are also used to maintain the gas generator bed at a proper temperature for APU start-up.

A single water tank with lines to all three APUs is provided to cool the gas generator injector should an APU restart be required before the gas generator can cool naturally. Control is via the APU controller. Starting a hot APU without this cooling risks detonation of the APU.

5.1.3.1.1. Water Spray Boiler

A water spray boiler (WSB) system provides cooling of both the APU gearbox oil and the orbiter hydraulic fluid. The system consists of three identical, independent water spray boilers; one for each APU and hydraulic system. Each WSB cools the corresponding APU lube oil system and hydraulic system by spraying water onto their lines; as the water boils off, the lube oil and hydraulic fluid are cooled. The steam that boils off in each water spray boiler exits through its own exhaust duct.

Water to cool the two heat exchangers is held in a bellows-type storage tank pressurized by GN2. Cooling of the oil and hydraulic fluid is effected by controlling the flow of water into the heat exchangers, as well as controlling the flow of hydraulic fluid through the exchanger. There are redundant controllers and temperature sensors for controlling the WSB.

5.1.3.1.2. Hydraulic System

The hydraulic system is located behind the Orbiter aft bulkhead 1307. Hydraulic lines branch to all of the Orbiter systems which require hydraulic power. As stated above the hydraulic system supplies power to operate the aerosurfaces, operate and direct the main engines, operate the landing gear and operate the umbilical system.

The components of each hydraulic system are:

1. Hydraulic main pump
2. Reservoir and accumulator
3. Circulation pump
4. Hydraulic/Freon heat exchanger
5. Heaters
6. Various valves, piping

The hydraulic system drives a number of different actuators. The types of actuators depends on the application.

The main pump provides 63 lb./hr. hydraulic flow at 3000 psia pressure to the hydraulic system. The pump is a variable speed pump driven by an APU. The main pump can be depressurized to 900 psia to reduce the torque at APU turbine start-up to reduce turbine spin-up time.
The hydraulic reservoir assures positive head pressure at the main pump and circulation pump inlets. The reservoir allows for thermal expansion and surges because of demand. The accumulator pressures the reservoir through a 40:1 differential area piston, and accumulator pressure is maintained by GN2.

In addition to the main pump, each hydraulic system has two circulation pumps. One is a high head low flow pump to re-pressurize the accumulator. The other is low head, high flow pump to circulate fluid through the Freon/heat exchanger to heat the hydraulic fluid. The operation of the circulation system is controlled by the General Purpose Computer while in orbit when the pump switch is in the GPC position.

As mentioned above, the hydraulic lines are warmed by hydraulic fluid passing through the Freon/hydraulic heat exchanger. The hydraulic lines in the various aerosurfaces are warmed by heaters. Each heated area has redundant heaters.

The hydraulic system interfaces with a number of sub-systems:

1. Body Flap System - The body flap is used during re-entry to adjust vehicle trim and limit hinge moment on the elevons. Hydraulic power is used to position the flap by means of three pilot operated control valves that are mechanically ganged together.

2. Rudder/Speed Brake - The vertical control surfaces consist of two sections; right and left hand. They operate together as a rudder and separately as a speed brake. This system receives power from one of the three hydraulic systems selected through a selector switch valve. Loss of the hydraulic system will be replaced by one of the other systems.

3. Elevon System - Each actuator of the elevon system is powered by one of three hydraulic systems through pressure-actuated switching valves. One is main and the other two are standby.

4. Brake System - The hydraulic system supplies pressure and flow to the main wheel brakes. A third system is connected to a switching valve on each brake as a standby.

5. Nosewheel Steering System - This system has been changed since the earlier flights. The Nosewheel Steering and deployment is supplied with hydraulic fluid via a switch valve in case of loss of pressure. As a protection against loss of complete pressure in the hydraulic system, an inhibit to switch over to system 2 is included in the valve control.

6. ET Umbilical Actuator - The actuator for the umbilical is operated hydraulically.

7. SSME Thrust Vector Control System - This is covered separately.

8. Main Engine Hydraulic Control Valves - Each of the three hydraulic systems provides power to the five valves on each of the main engines. The valves are:

- main fuel valves
- chamber coolant valves
- oxidizer preburner oxidizer valves
- main oxidizer valves
- fuel preburner oxidizer valves

5.1.3.2. Operation

There are three operational phases, namely Ascent through Orbital Insertion, Orbital Operations and Deorbit and Entry.

5.1.3.2.1. Ascent

As far as the APUs are concerned the ascent phase starts with the powering of the WSB controllers some 4 hours before launch. The WSB water tanks are pressurized at 1 hour and 10 minutes before launch, or T-1 hour and 10 minutes. Putting the WSB controllers into the **ON** mode activates heaters on the water tank, boiler and steam vent to ensure that the WSB is ready for launch.

At launch minus 6 minutes, the APUs are started. The pilot starts the APU pre-start sequence by activating the controllers and depressurizing the main hydraulic pumps to reduce the starting torque on the APU turbine. The pilot opens the APU fuel tank valves and looks for ready to start indications (gray talkbacks-annunciators) on the R2 panel, for all three APUs.

At T-5 minutes, the pilot starts all three APUs by putting the **APU CNTL** switches (R2 panel) in the **START/RUN** position and checks that the hydraulic pressure reaches 900 psi. If this is so, the pilot pressurizes the main pump and checks if the pressure reaches 3,000 psi on the gauges. The pressure must reach 2,800 psi by T-4 minutes or the mission will be aborted via the automatic launch sequencer.

The APUs are operated throughout the ascent phase and continue to be operated until the orbit insertion phase burn. While on the ground and during the first part of ascent, the tube bundles in the WSB are immersed in water. This boiler water precharge boils off about 8 minutes after launch and the WSB then enters spray mode. The hydraulic fluid usually does not heat up enough during ascent to require spray cooling. Once the main engine purge, dump and stow has been done the APUs and the WSBs are shutdown. The APU fuel pump and fuel valve module are cooled by running the 'A' cooling system, see Figure 5.13. Once orbit is achieved, the APU gas generators/fuel pump heaters are turned off. Heat soaking back will keep the GG bed warm enough for the next few hours. The heaters are reactivated 6 hours after lift-off. APU water and fuel line heaters are activated as the APU cools down to prevent freezing of the lines.

5.1.3.2.2. Orbital Operations

At 2 hours after lift-off the WSB steam vent heaters are turned on for 1.5 hours to eliminate ice from the WSB steam vents. Two and a half hours after lift-off, the APU fuel pump/fuel valve cooling is switched from the 'A' system to the 'B' system to avoid over heating the isolation solenoid. At 4 hours the APU fuel/valve cooling is shutdown and at 6 hours GG/fuel pump heaters are turned on as mentioned above.

On the day before deorbiting, one APU is started in order to have hydraulic pressure to check out the flight control system; i.e., to move the aerosurfaces. The associated WSB controller is activated, although the APU does not run long enough to require WSB operation. Fuel pump/valve cooling is activated. Landing gear isolation valves are closed before the APU is started and the isolation valves are reopened after the APU is shutdown.

5.1.3.2.3. Deorbiting and Entry

The WSB steam vent heaters are started 2.5 hours before deorbiting burn to prepare the WSB for operation during entry, at the same time the landing gear isolation valves are closed and circulation valves turned off.

Forty-five minutes before deorbit, the WSB water tanks are pressurized, the APU controllers activated and main hydraulic pumps set to **LO** pressure (controls on panel R2). The pilot opens APU tank valves and checks the status of grey talkbacks on panel R2. Five minutes prior to the reentry start sequence, one APU is started to ensure that one is available during descent. The hydraulic pump is left in **LO** pressure operation. This APU operates through deorbit burn. At 13 minutes to reaching 400,000 ft., the other two APUs are started and all hydraulic pumps are pressurized (**NORM**). Several operations involving the hydraulic system are carried out to ensure it is functional before the approach and landing.

After touchdown high flow tests are carried out on the APUs and hydraulic pumps; after this, the systems are shutdown.

5.1.3.3. Success Criteria

The APUs have been qualified to land with 2 out of 3 APUs operating. Pilots are trained on the simulator for single APU landings. Discussion with Don Williams of SAIC indicated that the entry and landing flight envelopes for single APU landings is not too dissimilar from landings with two or three APUs. There may be a combination of hydraulic demands and cross wind conditions that would make a single APU landing unsuccessful. The assessment team used the following success criteria.

1. Two and three APUs available upon landing - success
2. Single APU available upon landing - successful 80% to 100% of landings. This was represented by the event Unsuccessful Single APU Landing in the model. The event was quantified using a uniform distribution having a range of 0% to 20%.

Thrust vector control and aerosurface control use hydraulic actuators. Thrust vector control actuators have a switch valve connected to two APU/HYD systems. Loss of a single APU has no affect on thrust vector control. There are one pitch and one yaw actuator on each SSME. Loss of two APUs fails both TVC actuators on a single engine. Thrust aerosurface control actuators generally use all three APUs. Loss of two APUs maintains aerosurface control, but at 50% rate of

movement. (See Table 5.16 for more detail).

APU/HYDs power the hydraulic SSME propulsion valves during ascent. Each engine has a dedicated APU/HYD. Loss of an APU or hydraulic system during ascent sends that engine into hydraulic lockup. If hydraulic lockup occurs during the Thrust Bucket, an RTLS abort is called for in the Flight Rules. Other aborts depend on the potential for achieving orbit given a loss of APU or hydraulic system.

The WSB provides cooling to the APU lube oil that cools and lubes the gearbox and to the hydraulic fluid. Flight rules provide for declaring an APU lost and shutting it down if the WSB cooling fails and with lube oil temperature exceeding 250°F. Loss of lube oil cooling could fail bearings in the gearbox causing an APU underspeed shutdown. Such a failure would not be recoverable for a descent. Flight Rules are in place to help prevent that from happening. Flight records (in-flight anomaly descriptions and PRACA records) indicate that an APU has been shut down four times, once before MECO, because of lack of WSB cooling. An APU may be started at TAEM (instead of EI-13 or TIG-5) during descent if it is determined that the WSB can not support a longer run time.

Flight Rules provide for a MDF (minimum duration flight) if an APU fails or is declared lost. Flight Rules provide for a PLS (early primary landing site) landing if two APUs fail.

A leaking hydraulic system will be assessed for its ability to support landing before it is declared lost. It will be put into low pressure mode to preserve fluid. Such an assessment may be done during FCS checkout if the leak is known. There has been one incident of a serious hydraulic fluid leak in which the system's ability to support landing was in doubt. It was able to support to perform adequately. Our study treated this incident as a near miss. In other words it is partially credited as a failure.

A detected and confirmed hydrazine leak in an APU (other than into the pump seal catch bottle) should cause the APU to be declared lost in accordance with Flight Rules. Although PRACA records and In-Flight Anomaly Reports indicate six such leaks, no APU has been declared lost because of a hydrazine leak.

5.1.3.4. Initiating Events

The objective of this activity was to identify initiating event categories of scenarios triggered by malfunctions, perturbations, or failures of the orbiter APU, hydraulics, and thrust vector control. The identified events are such that all significant failure events that could potentially lead to LOC/V may be accounted for as part of the initiating event or part of subsequent scenario pivotal events. The initiating event categories were developed such that the pivotal events and end states are identical for all events within the category. Some of the initiating event categories may be broken down into their constituent failure modes using a fault tree later in future activities. Each event in a fault tree when taken individually should have the same pivotal events and end states as the category. On the other hand, no other initiating event category should have the same combination of pivotal events and end states. That is, some other initiating event categories may have some of the same pivotal events but

the combination of pivotal events and end states should be unique to each initiating event category. In a future activity, some pivotal events may be analyzed using fault trees as well.

Acute Catastrophic Failure

This category pertains to energetic failures that cause either immediate loss of orbiter structural integrity or immediate loss of equipment required to safely attain orbit or return to earth. Among the three subsystems analyzed, herein, the APU is the only credible source of such failures. The most risk significant sources of acute catastrophic failures are listed as turbine wheel overspeed, defective turbine wheel that fails, exhaust gas leak and large unisolatable hydrazine leak. The first two may lead to shrapnel that damages flight critical equipment or structures such as the large LO2 and LH2 lines, hydraulic lines, avionics boxes, other APUs and OMS bulkhead. Exhaust gas leaks may introduce more gas than can be removed through the aft compartment vents causing rupture of the orbiter. A large unisolatable leak may cause a large enough fire to catastrophically damage wire harnesses, APUs, hydraulic lines etc. such that there is insufficient flight control or landing control.

External Hydrazine Leaks

This category pertains to hydrazine leaks into the aft compartment. Hydrazine is damaging because of its combustible properties, ability to react with Kapton and insulation, and energetic disassociation properties in the absence of oxygen. This category includes leaks caused by material failures, under cooling or overheating that causes line ruptures, internal leaks within valves that cause overpressurization of valve operator coil cavity, or assembly errors. A fire in a single APU has the potential of propagating to an adjoining APU or flight critical equipment. Although it is theoretically possible to have a fire during ascent because of the uncontrolled level of N2 inerting in the aft compartment, entry and landing phases pose the most risk from this initiating event. Although helium inerting is effective down to about 70,000 feet altitude, during entry and descent, the low flammability limit for hydrazine is not reached until the atmospheric partial pressure of oxygen reaches about 4.7% which is between an altitude of 50,000 feet to 60,000 feet.

Shutdown Events

This category pertains to events that lead to the inability of either an APU or hydraulic system to provide sufficient hydraulic pressure. Both independent and common cause events are included in this category. Many events in this category are simply component failures for which the component ceases to function and recovery is not possible. However, some APU and hydraulic malfunctions are covered by flight rules. These provide for temporary shutdown of the affected subsystem and the potential recovery or restart of the subsystem should it be needed for landing. Both types of events are included in this category. An example of a common cause event is plugging of two or more APU turbine bearing oil filters by a waxy substance during the same mission. This category excludes catastrophic and hydrazine/hydraulic fluid leakage events, in the other initiating event categories, that could damage equipment external to the APU/hydraulic system. This category of events could potentially lead to loss of vehicle if at least two APU or hydraulic systems were incapacitated at any time except during the thrust bucket. Within the thrust bucket failure to provide hydraulic pressure from a single system is considered an initiating event for an RTLS or TAL abort.

Events within this initiating event category include:

i failure to function of active components such as pumps, valves, controllers, water spray boiler, cooling, and heating.
ii excessive leakage into the fuel pump drain cavity
iii events that cause premature depletion of APU fuel
iv events leading to permanent loss of hydraulic pressure such as loss of N2 overpressure or hydraulic fluid.
v events that prevent hydraulic pressure from being transferred to the actuators (e.g., plugging on hydraulic feed side, closure of actuator control valves)

High Hydraulic Pressure

This category pertains to events that lead to providing excessive hydraulic pressure to actuators. This category includes events such as premature valve motion that allows premature deployment of the landing gear or excessive pressure on the brakes during landing. It also includes events that lead to high sustained hydraulic pressure such as plugging of the hydraulic return line and failure of hydraulic pressure reduction during APU startup. Other potential events are:

i EXCESSIVE OR PREMATURE PRESSURE ON LANDING GEAR AND BRAKES
- PREMATURE LANDING GEAR DEPLOY
- UNCOMMANDED BRAKE PRESSURE

ii CLOGGING OR INADVERTENT CLOSURE OF FILTERS OR VALVES IN RETURN LINES OR EXCESSIVE SUSTAINED MAIN HYDRAULIC PUMP PRESSURE (E.G., FROM SPOOL/SLEEVE BINDING)
- MAINTAINS HIGH PRESSURE AT SSME VALVES => PNEUMATIC SYSTEM CAN NOT OVERCOME HIGH HYDRAULIC PRESSURE => POSSIBLE OXYGEN RICH MIXTURE DURING ATTEMPTED SHUTDOWN
- LANDING GEAR DOES NOT DEPLOY PROPERLY & UNCOMMANDED BRAKE PRESSURE

iii PREMATURE OPEN OF LANDING GEAR CONTROL VALVE

5.1.3.5. Event Sequences

The scenarios are modeled in two phases: ascent and descent. Failures that might occur during orbit are treated within 'failure to start' events during descent. The end states of the ascent model are used as initiating events of the descent model. This is necessary in order to 1) provide separate risk contributions for ascent, 2) determine the length of the mission (e.g., full duration, minimum duration flight, or next primary landing site), and 3) determine the number of and the condition of APUs that are available for descent.

The fault trees presented are for entire scenarios (that is, they are already linked). Each fault tree is identified with the scenario number of its associated event tree. We found this to be necessary in order to accurately model common cause and other dependencies among events. It was not necessary to model the APUs down to the component level within the fault trees. It was necessary to develop

expressions for some of the fault tree basic events in order to properly use the PRACA and IFA data. This is part of our data analysis effort.

In essence, the basic events to be entered into CAFTA™ are at the subsystem level. This was necessary for two reasons. First, in order to accurately capture common cause APU failures and common cause APU hydrazine leaks, we used the MGL method. This method requires equations that are not easily (if at all) treatable within CAFTA™. We then backfitted the fault tree event probabilities such that the correct answer, via the MGL method, would be reproduced by CAFTA™.

Second, we used a Bayesian analysis for APU/HYD/WSB failure probabilities. We developed prior distributions for *individual* APU/HYD/WSB run and start failures by the logical combinations of probabilities of their APU components. Generally, we used the 1987 McDonnell Douglas, Shuttle PRA Proof of Concept Study as the resource for the APU prior probabilities. Because that study did not include the hydraulic and water spray boiler subsystems, we used current applicable generic data for the hydraulic and WSB subsystem. Combining the probabilities of the components gave us a good prior distribution for an individual APU/HYD/WSB subsystem. We developed separate priors for ascent and descent. We updated the ascent prior distribution with the data - 4 WSB induced APU/HYD failures in 63 missions. We updated the descent prior distribution with the data - 1 near miss APU/HYD failure in 63 missions. Therefore, because our updated probabilities used subsystem level information, only subsystem level basic events were appropriate for CAFTA input.

As a final comment. This analysis used the terminology APU1, APU2, APU3 or Unit 1, Unit 2, and Unit 3. This does *not* refer to specific APU units on the orbiter. The meaning is as follows:

- APU 1 = 1st APU/HYD/WSB to leak (or fail)
- APU 2 = 2nd APU/HYD/WSB to leak (or fail) given 1 other is leaking (or has failed)
- APU 3 = Another APU/HYD/WSB to leak (or fail) given either 1 or 2 others are leaking (or have failed). (The definition is situation dependent)

APU/HYD Hub Breakup and Overspeed

This sequence of events is initiated by an APU/HYD turbine overspeed sufficient to cause hub disassembly, or a hub defect that allows hub disassembly. Given an overspeed condition, if the hub is OK, then the unit would be shutdown without the possibility of causing adjacent damage. This is an OK condition. Otherwise, a hub breakup would pose a serious threat to surrounding equipment. If the breakup is contained within the APU/HYD, then adjacent damage is averted, and an OK condition still exists. If projectiles are released within the aft compartment, then the concern is not only for the other APU/HYD units, but also critical flight equipment necessary for a successful flight.

Regardless of whether or not the remaining APU/HYD units have been damaged by the turbine breakup, a LOV would occur if critical flight equipment is damaged. With the critical flight equipment working, successful operation of the remaining APU/HYD units is necessary for a successful mission. If both are working, then an OK condition exists. If a second unit fails, then a single APU/HYD reentry, TAEM and landing would be necessary for an OK condition. A LOV would occur if the third APU/HYD unit fails.

The initiating and pivotal events of this sequence are described in Table 5.17. The turbine overspeed or hub failure event can occur if both primary and secondary fuel control valves fail in the open position while the APU is operating and the overspeed trip fails to close the secondary valve or it has been disabled. Closure of the fuel tank isolation valves following an overspeed trip would not prevent turbine overspeed or subsequent breakup. Calculations have shown that the hydrazine quantity downstream of the isolation valves may be sufficient to allow the turbine to reach breakup speed.

Mechanical, electrical and controller causes of turbine overspeed are all possible. Turbine overspeed implies that the APU, or some other control item such as a control circuit, has failed. Given a state of turbine overspeed, the next question is whether or not the resulting shrapnel and hydrazine cause another APU or other critical flight component to fail. There are also independent APU/HYD system failures to consider. Occurrence of this event after launch and in the absence of other failures leads to a PLS reentry unless it occurs in the thrust bucket, in which case an intact abort would occur.

The possibility of two APU/HYD units failing independently in the same mission from turbine overspeed is not modeled because the frequency of this sequence is much smaller than the frequency of sequences leading to a loss of vehicle/crew that involves one turbine overspeed followed by other failures. This event sequence diagram is for an entire flight from launch to touchdown.

Table 5.17. APU/HYD Hub Breakup and Overspeed Event Descriptions

EVENT	EXPLANATION	CONDITIONS FOR FAILURE
APU/HYD Turbine Overspeed or Hub Failure	Turbine speed greater than normal operating range or hub fails to support the turbine.	Both fuel control valves fail open. Failure of pressure relief valve. Failure or disability of speed control. Hub failure may be due to turbine overspeed or physical defect.
Hub OK	No hub failure.	Overspeed condition or physical defects. (1)
Contained within the APU	The turbine pieces are contained within the APU; i.e., no missiles.	Retention ring and turbine housing do not prevent turbine pieces from becoming missiles. (2)
Flight critical equipment OK	Aft compartment equipment critical to the flight, such as avionics, fuel lines, oxidizer lines, hydraulic lines, hydraulic tanks/accumulators, lube oil lines, OMS tanks, wire harnesses, controllers, etc. are all operational.	Missiles may penetrate lines, bulkheads, avionics boxes, controllers , wires (3)
Other APU/HYD units OK	The other APU/HYD units are not damaged by the missiles.	Missile damage, see above.
Third APU/HYD unit OK	If other APU/HYD units are damaged, is it only one, or both remaining units are damaged.	Missile damage, see above.
Single APU/HYD unit landing successful	If two APU/HYD units are disabled, then a single unit reentry, TAEM and landing is needed for a successful flight.	The landing sequence may not be successful due to the lack of hydraulic pressure to control the various control surfaces, particularly if weather conditions are not good.

1) Overspeed tests demonstrate that unnotched hubs do not come apart below 135% speed.
2) Retention ring designed for less than 135% speed.
3) Previous calculations (Shuttle PRA Proof of Concept Study) show that hub fragments have sufficient energy for damage.

Large Exhaust Gas or Hydrazine Leak

With a large gas leak, the first concern is over the critical flight equipment. Gas is exhausted in the temperature range between 800°F and 1100°F, and the impact of these hot gases on electronic equipment, fuel or fluid bearing APU/HYD lines may cause their failure. If some of the critical flight equipment is affected by the leak and fails, then the mission would end with a LOV. A second concern is over structural integrity of the orbiter's aft compartment. When a large gas leak occurs, over-pressurization may overstress the aft bulkheads, in which case a LOV would also occur. A third concern is ignition of the exhaust gas owing to non-catalyzed hydrazine in the exhaust gas. This event occurred during landing of STS - 51 (Sept. 12 - Sept. 22, 1993). A scenario would included overheating of the exhaust duct, causing escape of hot gas and burning hydrazine into the aft compartment. This scenario is unlikely to lead to a LOV and was not explicitly modeled.

With an unisolatable hydrazine leak, concerns over both the possibility of fire or explosion if ignition sources are present, and the corrosive affects of hydrazine exist. Hydrazine is very flammable, particularly if it comes into contact with a porous material. Hydrazine is also very corrosive and can attack Kapton wire insulation upon contact.

The initiating and pivotal events of this sequence are described in Table 5.18.

The first damage scenario in this sequence is the possible damage of critical flight equipment due to leaking hydrazine or exhaust gases in the aft compartment.

This second damage scenario is overpressurization of the aft compartment due either to a large hydrazine leak or the accumulation of exhaust gas in the compartment. Prior to reentry, the vent doors are closed at the Software Major Mode (MM 304) transition (EI-5 minutes). Gas accumulation can begin at this point until the vent doors open at approximately Mach 2.4. Only 0.3 PSID pressure is required to cause structural failure to the aft compartment.

Table 5.18. Large Gas or Hydrazine Leak Event Descriptions

EVENT	EXPLANATION	CONDITIONS FOR FAILURE
Large gas or hydrazine leak	A large exhaust gas leak or unisolatable hydrazine leak.	Exhaust leaks may be due to failures in the exhaust duct or turbine housing. Unisolatable hydrazine leaks may occur from any part of the APU, including valves, lines, couplings, fittings, fuel tank, etc..
All flight critical equipment OK	Equipment critical to the flight, such as avionics, fuel lines, oxidizer lines, hydraulic lines, hydraulic tanks/accumulators, lube oil lines, OMS tanks, wire harnesses, controllers, etc., are all operational.	Missiles may penetrate lines, bulkheads, avionics boxes, controllers, wires.
Aft compartment of main orbiter maintains structural integrity	Only 0.3 PSID pressure is required to cause structural failure to the aft compartment.	Overpressurization may be due to the hot exhaust gases leaking into the aft compartment, or a large hydrazine fire within the compartment.

Calculations have shown that a leak rate of approximately 10 percent of the total exhaust gas flow, starting at MM 304, is required to cause damage to the aft compartment structure before the vent doors open. Overpressurization may occur based on the size of the leak and the available time. However, hydrazine is also hazardous because of its flammability and corrosive properties. A large unisolated leak could potentially lead to massive wire stripping or a large fire in the aft compartment.

This event sequence groups the severe exhaust gas leak into the aft compartment with unisolatable hydrazine leaks. This is a conservative treatment that simplifies the model, and is acceptable, because the low frequency of occurrence of severe exhaust gas leaks that are severe enough to overpressurize the aft compartment or fail flight critical equipment. Separately categorizing these events would insignificantly change the estimate risk

APU/HYD Hydrazine Leaks During Ascent

This event sequence diagram starts with the event that at least one APU/HYD unit is leaking. There are many possible end states of this event sequence diagram. Table 5.19 lists and describes the various end states.

As described previously, the end states carry information about aborts, failures and leaks from the ascent model to the descent model. The first pivotal event is to determine the number of APU/HYDs leaking. Either one APU/HYD unit leaks or all three leak.

The sequence of events then depends on whether or not the leaks are detected and confirmed by ground control. If one APU/HYD is leaking, and the leak is detected and confirmed, than the Flight Rules dictate that the leaking APU would be shutdown post-MECO. The next pivotal event is whether or not the leaking APU/HYD unit is recoverable. If the leaking APU/HYD unit is recoverable, then the resulting end state for the reentry sequence depends on the number of other APU/HYD failures. If the leaking APU/HYD unit is the only failure, then the orbiter can safely land with the remaining two units with an MDFR condition. If one APU/HYD unit fails, then a PLSRU condition exists. If both non-leaking APU/HYDs fail, then a more serious PLSR2U condition exists.

Table 5.19. End State Definitions

END STATES	DESCRIPTION
IL0	Undetected or unconfirmed leak in APU. Full duration flight.
ILT	Undetected or unconfirmed leak in all three APUs. Full duration flight.
LOV	A situation in which the vehicle is presumed to be lost before wheelstop.
MDFR	Minimum duration flight declared with one unrecoverable APU or hydraulic subsystem.
MDFU	Minimum duration flight declared with one unrecoverable APU or hydraulic subsystem.
MDFRU	Minimum duration flight declared with one unconfirmed leaking APU and one unrecoverable APU or hydraulic subsystem.
MDF2RU	Minimum duration flight declared with two unconfirmed leaking APUs and one unrecoverable APU or hydraulic subsystem.
PLSRU	Primary landing site declared with a recoverable (confirmed leak) APU and an unrecoverable APU or hydraulic subsystem.
PLSR2U	Primary landing site declared with a recoverable (confirmed leak) APU and two unrecoverable APU or hydraulic subsystems.
PLS2RU	Primary landing site declared with two recoverable (confirmed leak) APUs and an unrecoverable APU or hydraulic subsystem.
PLS2U	Primary landing site declared with two unrecoverable APU or hydraulic subsystems.
PLS3R	Primary landing site declared with three confirmed leaking APUs.

If the leaking APU/HYD unit has an unrecoverable failure and is shut down permanently, then the orbiter is down to two fully functional APU/HYDs. If neither of the functional units fail, then a MDFU condition exists. However, if one unit fails, then a PLS2U condition exists, and if both units fail, then a LOV condition results.

A non-detected leak would not result in a crew or ground initiated shutdown. If all of the APU/HYDs, including the leaking one, survive the ascent with no problems, then an IL0 condition exists. The leaking hydrazine may affect the other APU/HYDs without failing the leaking unit, in which case if one APU/HYD unit fails, a MDFRU condition exists. If both remaining units fail, a PLSR2U condition exists.

The leaking unit may fail eventually. If it is the only failure, then a MDFU condition exists. If a second unit fails, then a PLS2U condition exists, while given that if the third unit fails a LOV condition exists.

A more serious situation exists if all three APU/HYDs are leaking. If this situation is detected and confirmed, then the Flight Rules dictate that only one of the three would be shutdown post-MECO. If all units operate and no failures occur, then a PLS3R condition exists. If one operating unit fails, then a PLS2RU condition exists, while if two units fail, a PLSR2U condition exists. When all three leaking units fail, a LOV condition exists.

If the shutdown APU/HYD subsequently fails, and cannot be recovered, then only two leaking units remain. If the remaining units survive, then a PLS2RU condition exists. When one remaining unit fails, then a PLSR2U condition exists, while if both remaining units fail, then a LOV condition exists.

Without detection, all three leaking units would be left in normal operating mode. If all three survive the ascent, then a ILT condition exists. If one unit fails, a MDF2RU condition exists, or if two units fail a PLSR2U situation results. If all three units fail during the ascent, then a LOV condition exists.

This model accounts for the conditional probability of APU failure owing to independent and common cause failures, as well as leaks. Table 5.20 lists the various pivotal events with descriptions and possible failure modes.

Hydrazine leaks may develop in the APU/HYD units during ascent. A detected and confirmed leak during ascent is indicated for crew shutdown by flight rules after MECO. The leaking APU/HYD may or may not cause an unrecoverable failure. If the leaking APU/HYD does not fail, then the event sequence diagram (ESD) questions whether or not either of the non-leaking APU/HYD units have failed. A nonrecoverable failure owing to hydrazine leaks would occur if:

- hydrazine causes electrical shorts or open circuits owing to its chemical information with wire insulation
- hydrazine causes a fire that damages the APU
- remaining fuel or pressure can not support the descent and landing

Nitrogen inerting of the aft compartment greatly reduces the conditional probability of fire during ascent.

Flight Rules guide actions regarding failed or leaking APUs. This model assumes that Flight Rules are followed.

Table 5.20. APU/HYD Hydrazine Leaks During Ascent Event Descriptions

EVENT	EXPLANATION	CONDITIONS FOR FAILURE
Hydrazine leak	An unisolatable hydrazine leak from or one of the APU/HYD units exists.	Unisolatable hydrazine leaks may occur from any part of the APU, including valves, lines, couplings, fittings, fuel tank, etc.
Remaining APU/HYD units do not leak	Either one APU/HYD unit leaks, or all three leak.	The condition of all three APU/HYD units leaking is probably due to a common cause.
Leak is detected an confirmed	Leaks need to be detected and confirmed by ground control before crew actions are taken.	Leaks may be detected through the unexplained loss of fuel tank pressure, unexplained cooling of the fuel tanks, lines, and valves , or high gearbox pressure.
Leaking APU/HYD unit OK	The leaking APU/HYD unit is recoverable.	A failure would be either independent/ dependent or leakage induced.
Other APU/HYD units OK	More than one APU/HYD unit is damaged.	A failure would be either independent/ dependent or leakage induced.
Third APU/HYD unit OK (one APU/HYD unit is leaking)	If other APU/HYD units are damaged, is it only one, or are both remaining units damaged.	If all three units have failures, they are either independent, common cause or leakage induced.
All APU/HYD units OK	All three APU/HYD units leak and are recoverable.	A failure would be either independent/ dependent or leakage induced.
Second APU/HYD units OK	If one APU/HYD unit is damaged, then are two or three damaged.	A failure would be either independent/ dependent or leakage induced.
Third APU/HYD unit OK (three APU/HYD units are leaking)	If other APU/HYD units are damaged, is it only one, or are both remaining units damaged.	The three failures are independent, common cause, own leak induced, or leakage from other unit induced.

This event sequence diagram makes several assumptions. First, it is assumed for modeling simplicity that if more than one APU/HYD unit leaks, all three units leak. This was done for the sake of simplifying the models. It introduces a minor nonconservatism on the probability of leaks because the probability of common cause contributions of 2 units leaking is not included. The error in the leakage probability is about 11 %. However, the overall LOV probability owing to leaks is still conservatively assessed by this simplification. This is because the conditional probability of failure of an APU owing to its own leak is about 20 times higher than the conditional probability of an APU failure owing to the leak of another APU. The second assumption deals with detecting and confirming the isolatable leaks. If all three APU/HYDs are leaking, then it is assumed that if one leak is detected, then ground control would be more cautious with the other units, and would detect and confirm those leaks also.

All APU/HYD Units OK Without A Hydrazine Leak During Ascent

This event sequence diagram begins with a single APU/HYD unit failure without any hydrazine leaks. The loss of one unit, according to the Flight Rules, suggests a minimum duration flight (MDF). If a second unit suffers an unrecoverable shutdown, then the Flight Rules suggest a landing at the next available primary landing site (PLS). If the third unit fails, then a LOV situation occurs.
Table 5.21 lists the various pivotal events with descriptions and possible failure modes.

This event includes equipment failures of any of the APU/HYDs that cause a shutdown of that particular system. For example, any underspeed or overspeed condition would cause an APU/HYD shutdown. Other failures may include the failure to start the APU, failures of the pump or turbine, the fuel line is clogged, etc. The failure of more than one APU/HYD unit would lead to a PLS situation. This event sequence diagram assumes that the failure of the first APU/HYD unit was not caused by a hydrazine leak.

Table 5.21. At Least One APU/HYD Unit Fails Without a Hydrazine Leak During Ascent Event Descriptions

EVENT	EXPLANATION	CONDITIONS FOR FAILURE
APU Fails	One APU/HYD unit has an unrecoverable failure with no leaks.	This type of a failure results from some independent failure of the APU or hydraulic subsystem.
Remaining APU/HYD Units OK	Neither of the two remaining APU/HYD units have failures	Independent failures would cause at least one unit to fail, a common cause failure could affect both units.
Third APU/HYD unit OK	A second unit has failed, and this pivotal event addresses whether or not the third unit has failed.	If both remaining units fail, this could be due either to two independent failures or a common cause failure.

Initially Operational APU/HYD Units During Reentry, TAEM And Landing
This event sequence diagram represents the various paths to either a successful landing or LOV situation during reentry, TAEM and landing. The term "successful" or "OK" would be used to describe anything but a LOV situation, although in reality there could be different levels of success. This sequence starts essentially with no initiating events. All three APU/HYD units are fully operational before reentry.

Consider the possibility that during reentry, TAEM or landing that isolatable hydrazine leaks may develop in the APU/HYD units. If no leaks develop, than the only real concern is how many, if any, of the APU/HYD units would have unrecoverable failures. If no failures occur, then success, or an OK state, would result. Even if one unit fails, an OK state would also result since a two APU/HYD unit landing is considered a success. When two units fail, success depends on whether or not a single APU/HYD unit landing can be performed. If so, then an OK state; if not, then a LOV occurs.

When a single APU/HYD unit begins leaking, then detection becomes a concern. If the leak is detected and confirmed, then the Flight Rules dictate that the unit would be shut down. If no other units fail, then an OK situation would result. If either, or both, of the non-leaking units fails, then a restart of the unit that was shutdown would be attempted. A restart of any APU/HYD system includes both the starting procedure itself and the unit running after the starting procedure. If one APU/HYD unit fails with a successful restart, then two units are operational, and the landing is successful. If both non-leaking units fail, and the leaking unit is successfully restarted, then the next pivotal event is whether or not the previously described single APU landing can be performed. An OK situation would result if the landing is successful, otherwise a LOV occurs.

Any failure with the inability to restart the shutdown APU/HYD unit is serious. If one unit has failed, then a single unit landing would be needed for an OK state. If both non-leaking units fail, then all three units have unrecoverable failures, and a LOV would occur.

When a single APU/HYD leak is not detected, it would not be shutdown, and a similar condition exists to that of a no leak situation. The model, however, considers the conditional probability that a leaking unit will subsequently fail or fail another unit. The subsequent APU/HYD unit failures would dictate an OK or LOV condition, depending strictly on the number of APU/HYD unit failures, and if the single APU/HYD unit reentry, TAEM and landing is successful if needed. If only one unit fails, then an OK would result. If two units fail, a single APU landing would be needed for a OK state. Obviously, if all three units fail, a LOV condition exists.

When all three APU/HYD units are leaking and are detected and confirmed by ground control, then the Flight Rules suggest that only one unit would be shutdown. If no other failures occur, then an OK state would result. If one or more of the remaining units fail, then a restart would be attempted on the unit that was shutdown. Without a successful restart, then one unit failing leads to an attempted single APU/HYD landing situation, but if both remaining units fail, a LOV occurs. If only one unit has failed, and the restart is successful, then an OK situation results. If

both of the working units fail with the successful restart, than a single APU/HYD landing would also be needed for success.

Some pivotal events used in this event sequence have been described in other diagrams. The pivotal events that have not been previously described are listed in Table 5.22.

All APU/HYD units are fully operational to begin this phase. The initiating events within the sequence can be APU/HYD unit failures, a single APU/HYD unit leak or three APU/HYD units leaking.

Table 5.22. Fully Operational APU/HYD Units During Reentry, TAEM and Landing Event Descriptions

EVENT	EXPLANATION	CONDITIONS FOR FAILURE
All APU/HYD units have integrity (no leaks)	If no units leak, then all APU/HYD units have integrity.	Both independent and common cause failures are included.
Successful restart/run of shutdown APU/HYD unit (not a hot restart)	Once a unit is shutdown, there may be a failure that either prevents it from restarting, or once it has started, a failure occurs and the unit becomes unrecoverable.	These failures may be independent, own leak induced, or leakage from another unit induced. A failure may either prevent a restart or cause an unrecoverable shutdown after the restart sequence.

PLSRU State During Reentry, TAEM And Landing

This sequence of events begins with the unrecoverable loss of one APU/HYD unit, another that is leaking hydrazine, and the third being fully operational. The first concern during the descent is whether or not the leaking unit can be restarted. Without the leaking unit, the orbiter would need to make a successful one APU/HYD unit landing for an OK state. Otherwise, a LOV would occur. If the leaking unit is restarted, then the next pivotal event is the number of failures that may occur during descent.

With both remaining APU/HYD units operating successfully during descent, the flight would end with an OK state. If one unit has an unrecoverable failure, then a single APU/HYD unit landing is required for an OK state to occur. If both remaining units suffer unrecoverable failures, then the mission would end with a LOV state.

This event begins with one unit fully operational, one unit unrecoverable, and one unit recoverable, but leaking. Unrecoverable failures and leaking units have been discussed previously. With the one APU/HYD unit unrecoverable, both remaining units are needed. One unit is leaking hydrazine, and the leak may lead to the unit failing itself, or the other unit failing due to hydrazine exposure. Both

remaining units may fail to either independent failures, due to the hydrazine leak, or due to a common cause failure.

PLSR2U State During Reentry, TAEM And Landing

This sequence of events is relatively simple. Two APU/HYD units have suffered unrecoverable errors. If the third unit does not suffer any unrecoverable errors, and a single APU/HYD reentry, TAEM and landing can be performed, then an OK state would result. However, if the unit suffers an unrecoverable failure, or the single unit reentry, TAEM and landing is unsuccessful, then a LOV state would occur.

This event is initiated by the unrecoverable loss of two APU/HYD units. Unrecoverable failures have been discussed previously. The third unit is leaking hydrazine, but is recoverable. Any failure of the third unit would be catastrophic, regardless of whether the failure was caused by the leak itself or some other, independent failure.

PLS3R State During Reentry, TAEM And Landing

This sequence begins with the state that all three APU/HYD units have detected and confirmed leaks, but are recoverable, and one unit is shutdown according to the Flight Rules. The next pivotal event in this sequence of events is how many of the two remaining APU/HYD units suffer unrecoverable failures. If both units are successful through landing, then an OK situation would result. If one unit fails, then success would result only if a single APU/HYD unit landing can be accomplished, otherwise a LOV would occur. If the two remaining APU/HYD units fail before the orbiter has landed, then all three are lost, and a LOV would occur.

This sequence of events is initiated by all three APU/HYD units leaking hydrazine in the aft compartment. This event sequence has been described previously.

An assumption is made that this initiating event is of low probability, so a simple conservative event sequence diagram can be used to model the possible event outcomes without adversely affecting the overall success probability.

5.1.3.6. Data Analysis

The development of probability distributions for the fault trees is done using Bayesian updating methods. Prior probability distributions for failure rates are taken from the 1987 APU/HPU study (Ref. 56 - see Volume V, Appendix C.6), NPRD-95 (Ref. 58), IEEE Std. 500 (Ref. 64), WASH 1400 (Ref. 65), Shuttle experience and expert judgment. System level priors for the entire APU/HYD/WSB system (failure to start and failure to run distributions) are developed using component data mostly from the 1987 study. Bayesian updating was done at the system level using data found in the in-flight anomaly list (IFAS), PRACA reports, and Post Flight Mission Safety Evaluation Reports.

A method known as the Multiple Greek Letter (MGL) method was implimented to estimate the fraction of component failures attributed to dual or triple common cause failures of the APUs. A detailed algebraic development of the MLG method, along with other APU data analysis information, is contained information is provided in Appendix B.3.

5.1.3.7. APU Initiator Risk Contribution

The top-level APU significant risk contributors (>1% of APU risk) may be grouped into failures leading to loss of hydraulic pressure during ascent or entry and descent and the risk due to APU turbine hub failure as shown in Figure 5.15. Notice that the APU risk is dominated by the possibility of losing hydraulic pressure during entry or descent. The failure modes contributing to the risk of this scenario are shown in Figure 5.16. The major causes of failure were found to be common cause failure of the APU (52%) and a single APU leakage induced failure (36%) during entry and descent.

The hydrazine leakage of one APU may propagate to cause the entire loss of function of the APU leaking as well as the other two APUs. The leakage of hydrazine from one APU may affect the operation of the other APU due to the hydrazine igniting in the aft compartment or corroding the sub-components of the other APU systems. Common cause failures of the APU are prevalent due to the confined space which is shared by all three APU. This proximity makes the potential of an adverse condition in the orbiter aft-compartment taking out multiple APU a significant possibility.

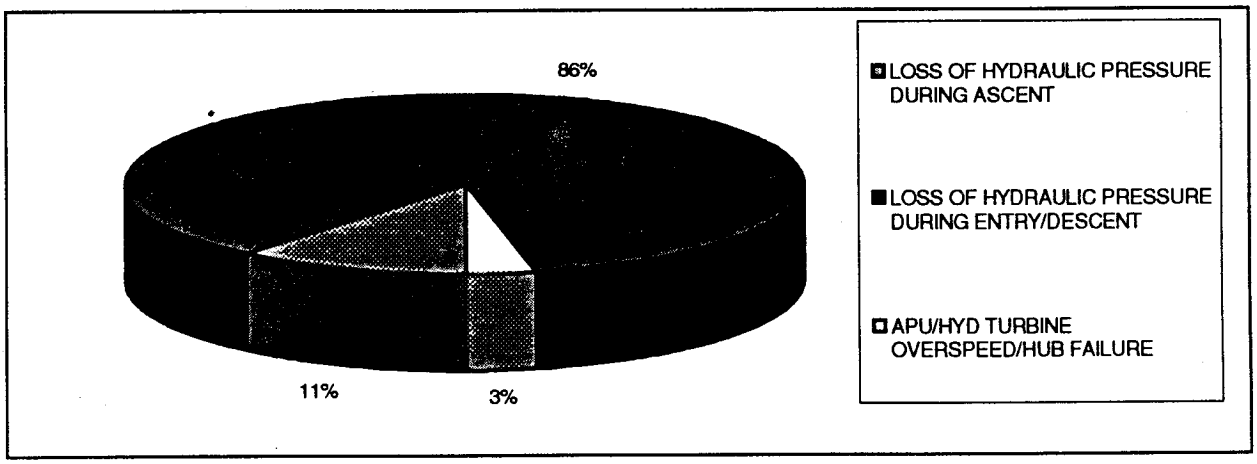

5.15. APU Risk Contribution

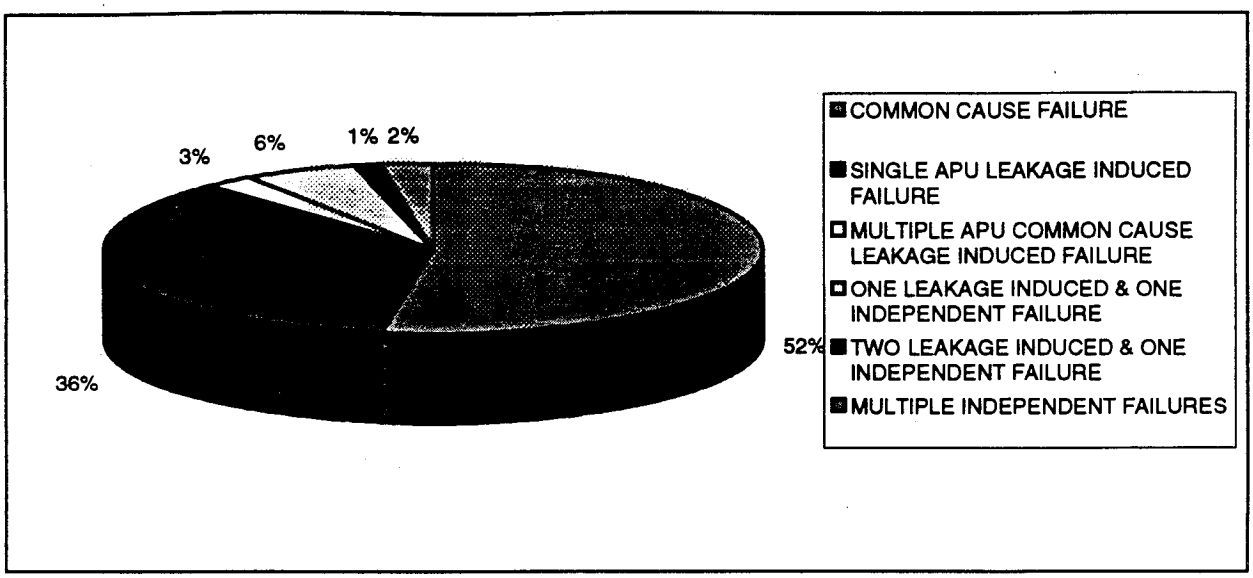

5.16. APU Entry/Descent Failure Mode Risk Contribution

5.1.4. Orbital Maneuvering System and Reaction Control System

The Orbital Maneuvering System provides the necessary thrust for orbit insertion, orbit circularization, orbit transfer, rendezvous, deorbit, abort to orbit and abort once around. It can also provide up to 1,000 lbs of propellant to the aft reaction control system.

The Reaction Control System provides propulsive forces from a collection of jet thrusters to control the motion of the orbiter. The selective firing of individual jets or specific combination of jets provides thrust for:

- Attitude control
- Rotational maneuvers (pitch, yaw and roll)
- Small velocity changes along the orbiter axes (translation maneuvers).

5.1.4.1. Description

The OMS is located in two independent pods in the aft end of the orbiter on either side of the vertical tail. Each OMS pod contains one OMS engine, a fuel tank, an oxidizer tank, and a helium tank, along with propellant feedlines and other supporting equipment.

Each OMS module consists of :

- One high pressure gaseous helium storage tank
- Pressure regulation system
- Fuel and oxidizer tanks
- Propellant distribution system
- Thermal control system
- Engine (injector plate, thrust chamber, nozzle)
- Thrust Vector Control
- Bipropellant Valve Assembly and Nitrogen System

The RCS consists of three separate vehicle modules: forward, left aft and right aft. The forward module is located in the forward fuselage nose area and the aft modules are located with the orbital maneuvering system (OMS) in the OMS/RCS pods.

Each RCS module consists of :

- Two high pressure gaseous helium storage tanks
- Pressure regulation and relief systems
- Fuel and oxidizer tanks
- Propellant distribution system
- Thermal control system
- Reaction Control jets (the forward module has 14 primary and 2 vernier engines, each aft module contains 12 primary and 2 vernier engines)

All primary thrusters contain instrumentation for chamber pressure.

5.1.4.2. Phase Operations

Ascent
Generally, there are two OMS thrusting periods:

During the first OMS thrusting period, both OMS engines are used to raise the orbiter to a predetermined elliptical orbit. Vehicle attitude is maintained by gimbaling (swiveling) the OMS engines. RCS attitude control may be used if OMS gimbal rate or gimbal limits are exceeded.

The second OMS thrusting period takes place near the apogee of the orbit established by the first OMS thrusting period. It uses both OMS engines and is used to circularize the predetermined orbit for the mission.

Some missions have only one OMS period and are called direct insertion. Such missions replace the first OMS period with a 5-feet-per-second RCS translation maneuver to facilitate the dump of liquid propellants from the main propulsion system. The only OMS period is then used to achieve orbit insertion.

During ascent, the RCS is used for rotational control during mated coast with the external tank. It is then also used toprovide -Z translation at external tank separation, using all 10 down (-Z) primary jets. This is the only RCS translational maneuver done automatically. The RCS is also used during ascent to maneuver to OMS burn 1 and burn 2 attitudes and to trim residual velocities post-OMS burn 1, post-OMS burn 2, and post-deorbit burn, if required.

On-orbit
Additional OMS thrusting periods using one or both OMS engines are performed on orbit according to the mission's requirements to modify the orbit for rendezvous, payload deployment or transfer to another orbit.

During orbit, the RCS provides for attitude control, including pointing and attitude hold, and rendezvous maneuvers.

The OMS is interconnected to the RCS in three cases: during ascent aborts, when the OMS engines and RCS jets are all firing to dump OMS propellant (the jets are used to help complete the dump more quickly); during ascent if two main engines have failed and RCS jets are fired to maintain roll control (OMS propellant is used so that RCS propellant can be saved); and during certain times in orbit when OMS propellant is fed to the RCS jets so that RCS propellant can be conserved. The reason for special concern about conserving RCS propellant is that during entry when RCS jets are used for control, it is not possible to feed OMS propellant to the jets because the g forces are perpendicular to the OMS tank axis. It is essential that enough propellant remain in the RCS tanks to provide control through entry.

Deorbit
The two OMS engines are later used to deorbit, in which the engines fire in the general direction of the velocity vector, thus decreasing orbital velocity and lowering the perigee attitude to nearly zero.

The crew requires more data for the deorbit burn than for other burns because of special procedures in failure situations. The decisions about continuing or stopping the deorbit burn in case of failures depend on how far the burn has progressed. In general, if a serious failure occurs during a deorbit burn and the current perigee is above 80 n. mi., the burn is terminated and the problem is analyzed. The vehicle can remain in orbit for several revolutions and the deorbit. If the perigee altitude is below 80 n. mi. when a failure occurs, the burn is continued by whatever means are available because the orbit would be too unstable to stop at that point.

During entry, the RCS provides for center of gravity management through the forward propellant dump. Specific functions are decreasing landing weight, controlling the center of gravity, and reducing OMS/RCS tank structural loads at touchdown. Also during entry, yaw, roll and pitch control is provided by the aft left/right (Y) and up/down (Z) jets.

Note:
The aft RCS plus X jets can be used to complete any planned OMS thrusting period.

Aborts
The RCS is also used during off-nominal situations. These include single-engine roll control in the case of loss of two SSMEs on ascent. The OMS-to-RCS interconnect is automatically commanded, and the RCS jets control the roll. If the OMS gimballing system is not performing adequately to control vehicle attitude during an OMS burn, an RCS wraparound is used. The RCS is also used manually if the OMS fails prematurely.

Constraints and Limitations
The minimum altitude for an OMS engine burn is 70,000 feet. Below this altitude the pressure difference between the inside and the outside of the nozzle could cause it to collapse.

The maximum number of starts that can be supported by the nitrogen system is 17.

The minimum allowable Pc during an OMS burn is 80 percent of nominal. Below 80 percent there may be unacceptable cooling of the engine and imbalances in the mixture ratio which could cause engine damage. If the engine is being operated in the blowdown mode (helium press/vapor isolation valves closed), the minimum allowable Pc is reduced to 72 percent.

The OMS propellant quantity should be less than 22 percent of the lift-off amount to remain within the structural limit for landing weight.

5.1.4.3. OMS Initiators

Orbiter Explosion due to OMS/RCS explosive failures

- Critical Structural Failure of OMS
  - i Structural failure of MMH propellant containers (tanks, pipes, valves)
  - ii Structural failure of N2O4 propellant containers (tanks, pipes, valves)
  - iii Structural failure of He press. system containers (tanks, pipes, valves)
  - iv Heater fails on, causing propellant auto-decomposition and overpressure
  - v Thruster burn-through during OMS operation; combustion instability

These initiating causes are such that they cause the loss of vehicle directly, that is, once the initiator has occured, the sequence goes to an LOV with no pivotal events in between. However, the frequency of such failure modes is low when compared to accident sequences for other portions of the Shuttle, so that they are not considered in the study.

OMS/RCS functional failures (non-explosive)

- Failure to feed MMH or N2O4 propellant to the thruster chamber
  - i Small external leak in MMH or N2O4 tank, pipe or valve depleting all propellant
  - ii MMH or N2O4 thruster inlet valve fails open depleting all propellant
  - iii MMH propellant freezes (Heater failure)
  - iv Control failure depletes all propellant (Interface with Avionics System)
  - v Human error depletes all propellant
  - vi External leak in He tank, pipe or valves depleting all He
  - vii Pressure relief valve fails open

OMS/RCS configurational failures

- Failure to maintain proper valve configuration
  - i Thruster valves fail to open on demand
  - ii Both thruster inlet propellant valves fail to close simultaneously
  - iii Human error

The OMS/RCS combination was assessed to be a non-significant risk contributor in this study for the following reasons:

1. Low system pressures and "leak before rupture" design makes the possibility of an explosive rupture highly improbable.
2. System malfunctions occur at portions of the trajectory which are less sensitive to loss of propulsion leading to abort rather than LOV events. Moreover some or all of the impulse lost due to OMS malfunctions may be provided by the RCS.
3. Successful system operation is critical during some abort scenarios which as noted before are out of scope.

5.1.5. Orbiter Thermal Protection System

5.1.5.1. Description

The thermal protection system (TPS) consists of various materials applied to the outer structural skin of the orbiter to maintain the skin within acceptable temperatures, primarily during the entry phase of the mission. The orbiter's outer structural skin is constructed primarily of aluminum and graphite epoxy.

During entry, the TPS materials protect the orbiter outer skin from temperatures above . In addition, they are reusable for 100 missions with refurbishment and maintenance. These materials perform in temperature ranges from minus 250° F in the cold soak of space to entry temperatures that reach nearly 3000° F. The TPS also sustains the forces induced by deflections of the orbiter airframe as it responds to the various external environments. Because the TPS is installed on the outside of the orbiter skin, it establishes the aerodynamics over the vehicle in addition to acting as the heat sink.

5.1.5.2. Success Criteria

As described above the TPS primary function is to maintain the orbiter skin below 350° F during entry. However a LOV does not ensue unless a burnthrough of the orbiter skin causes hot re-entry gases to impinge upon a critical flight system or weaken the orbiter structure to a point which makes it unable to withstand aerodynamic loads or landing impulses.

5.1.5.3. TPS Study Integration

This project included a literature search for previous quantitative risk related studies performed on various Shuttle systems to preclude an unnecessary reduplication of effort. A study involving the safety of the TPS was uncovered which was deemed to meet the requirements established for the Shuttle PRA and thereby justified its inclusion in the PRA model. The study is titled *Safety of the Thermal Protection System of the Space Shuttle Orbiter: Quantitative Analysis and Organizational Factors*; the report is presented in its entirety in Appendix C.4.

The following is a recapitulation of the principal ideas and results from the fore mentioned report:

There are three initiating events considered which may lead to burthrough and subsequent LOV:

- Initial debris impact on one tile only
- Initial debris impact on several tiles
- Debonding caused by factors other than debris impact

Data on TPS damage was obtained from PRACA, TIPS, and PCASS. It was determined that much of the severe damage was caused by insulation from the cone area of the right SRB, prior to STS-27R the major source of debris was found to be from portions of SOFI insulation from the ET. Note that

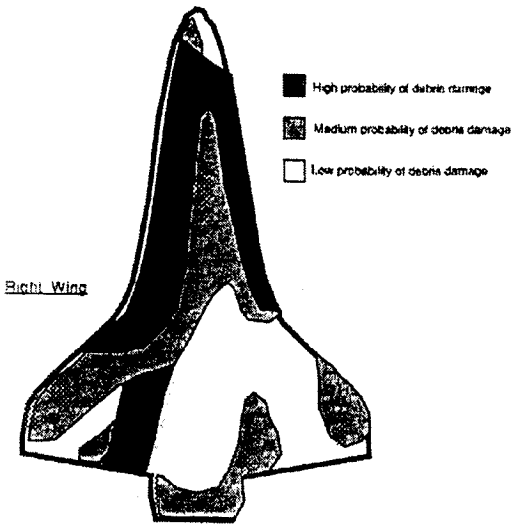

Figure 5.17. Debris Impact Profile

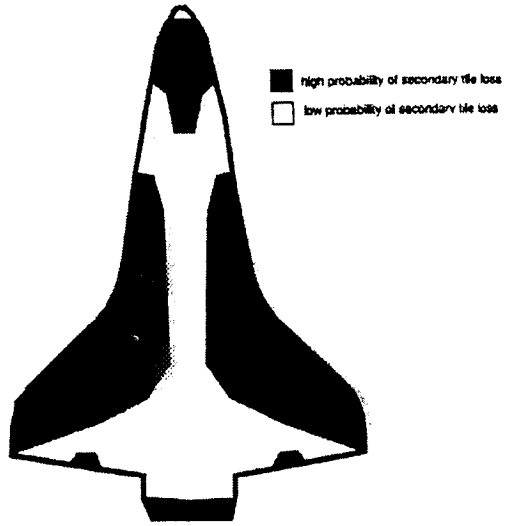

Figure 5.18. Secondary Tile Loss Profile

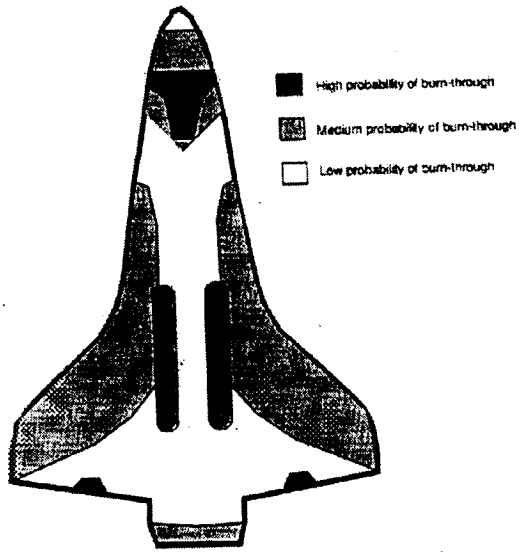

Figure 5.19. Burnthrough Profile

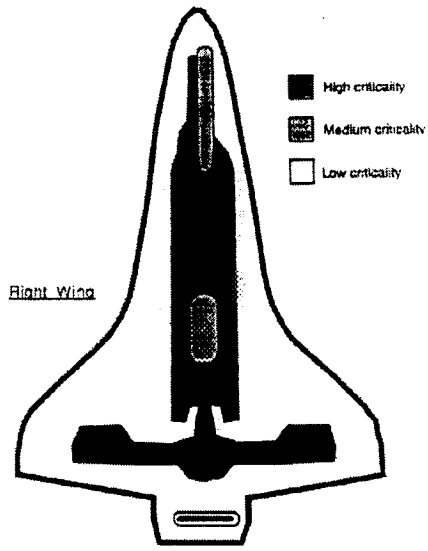

Figure 5.20. Criticality Profile

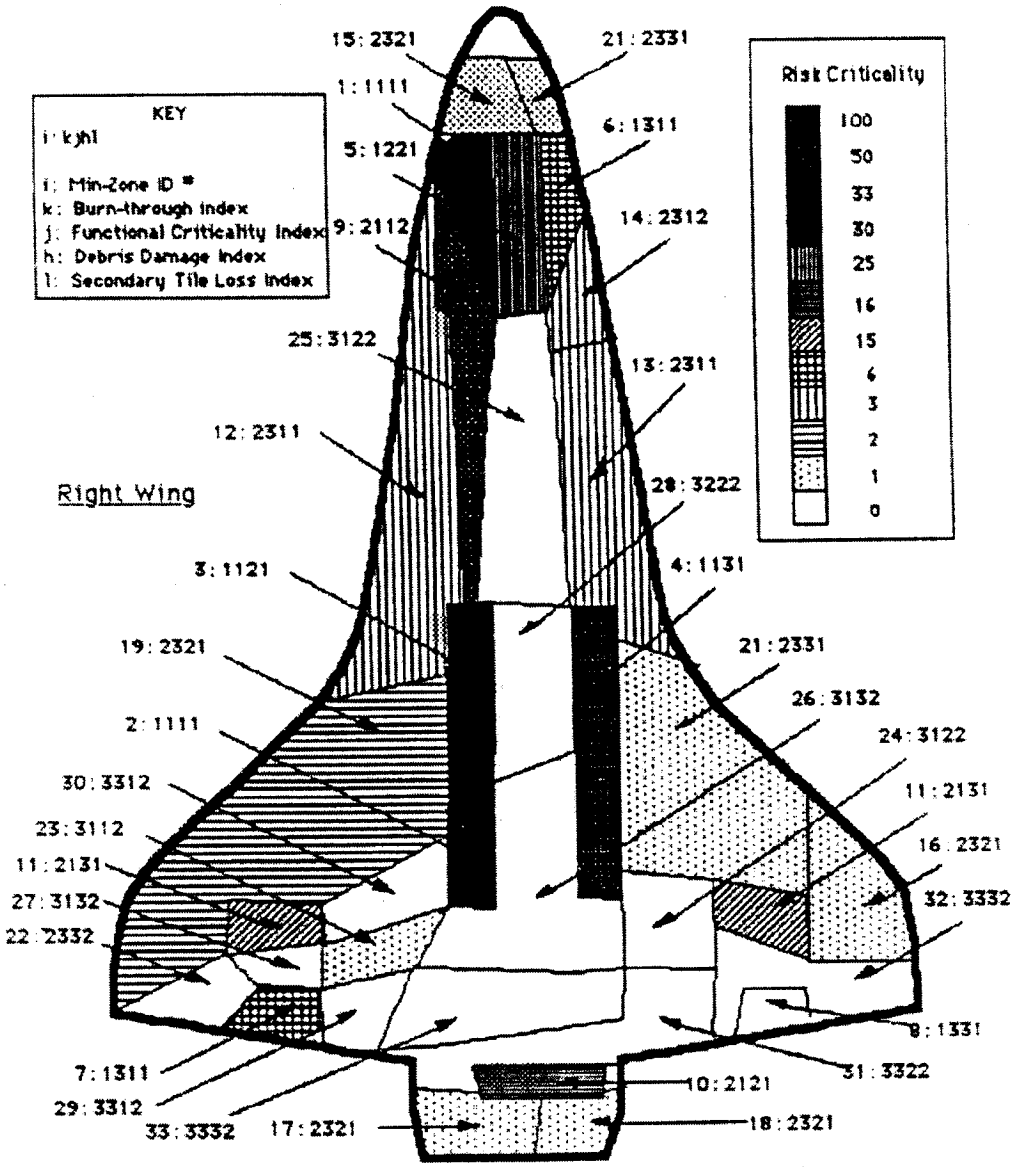

Figure 5.21. TPS Min-Zone Partitioning

given this information, the risk of the TPS may be partially allocated to the SRB however, as was done for the SSME-APU interface, the risk is allotted to the system which exacerbates a benign event into a catastrophic occurrence. The relative density of debris damage is shown in Figure 5.17, the biasing is of course due to the fact that most of the debris originates from the right SRB. Based on history it was proposed that a 1 in 2,000 probability exists of a large hit causing a tile to lose its insulating capability. The loss of one tile may lead to the loss of adjacent tiles during re-entry. Figure 5.18 illustrates the locations which have a higher likelihood of secondary tile loss. Once the tile or tiles have failed the susceptibility to burnthrough varies depending on where the damage has occurred. Figure 5.19 shows the relative probability of burthrough for various locations on the orbiter (refer to Appendix C.4: Table 4 for actual probabilistic estimates). Given that a burnthrough has occurred, the degree of criticality varies upon the equipment directly beneath the damaged area. The relative criticality of burnthrough is shown in figure 5.20.

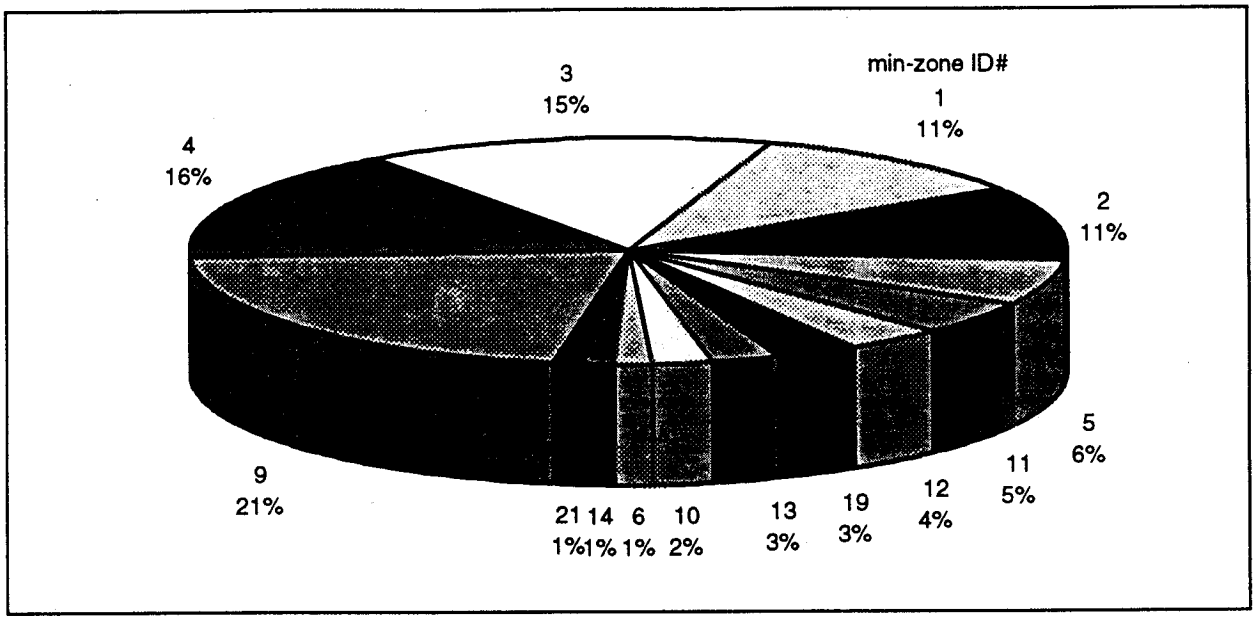

Figure 5.22. TPS Min-Zones Risk Contribution

By overlaying the functional criticality, burn-through, debris damage, and secondary tile loss areas, 33 min-zones were established (Figure 5.19). Table 5.23 shows the final numerical results of the study. The index is determined by combining the relative measure of probability for each of the failure contributing phenomena with 1 being the highest probability and 3 the lowest. The first digit refers to functional criticality with the remaining respective digits referring to the relative probability of burn-through, debris damage, and secondary tile loss areas for the min-zone shown.

The catastrophic failure of each min-zone was included as a basic event in the integrated PRA model along with the respective probability given in the TPS study. The risk contribution of significant min-zones to TPS failure is shown in figure 5.22.

				x 1E-4	x 1E-4	x 1E-4
ID #	INDEX	LOCATION	#TILES	DEBRIS	DEBOND	TOTAL
1	1111	Right side TPS, under crew	156	0.87	0.36	1.23
2	1111	Right side near main ldg gear (af	156	0.87	0.36	1.23
3	1121	Right side near main ldg gear (fw	676	0.13	1.62	1.75
4	1131	Left side near main ldg gear	780	0.00	1.87	1.87
5	1211	Centerline under crew	364	0.51	0.22	0.73
6	1311	Left side TPS, under crew	312	0.11	0.04	0.15
7	1331	Center of right elevon	104	0.04	0.01	0.05
8	2112	Center of left elevon	401	0.00	0.00	0.00
9	2121	Right side TPS, fwd mid edge	624	1.73	0.75	2.48
10	2131	Center of body flap	208	0.02	0.24	0.26
11	2311	Left wing TPS, center	468	0.00	0.56	0.56
12	2311	Right side TPS, mid edge	1664	0.30	0.13	0.43
13	2312	Left side TPS, mid edge	1196	0.21	0.08	0.29
14	2321	Left side TPS, fwd mid edge	572	0.10	0.04	0.14
15	2321	Right side TPS, nose	277	0.01	0.02	0.03
16	2321	Left wing TPS, center	832	0.01	0.06	0.07
17	2321	Right side TPS, body flap	104	0.00	0.01	0.01
18	2321	Left side TPS, body flap	104	0.00	0.01	0.01
19	2321	Right wing	2132	0.18	0.16	0.34
20	2321	Left side nose	312	0.00	0.02	0.02
21	2321	Left wing TPS, fwd	1768	0.00	0.13	0.13
22	2332	Right elevon TPS, outboard	312	0.00	0.02	0.02
23	3112	Right wing TPS, center	364	0.01	0.01	0.02
24	3122	Left wing TPS, center	468	0.00	0.01	0.01
25	3122	Payload bay TPS, fwd	1664	0.00	0.02	0.02
26	3132	Payload bay TPS, aft	1976	0.00	0.02	0.02
27	3132	Right wing TPS, center	468	0.00	0.01	0.01
28	3222	Payload bay TPS, mid	520	0.00	0.00	0.00
29	3312	Right elevon TPS, in board	312	0.00	0.00	0.00
30	3312	Right wing TPS, center	416	0.00	0.00	0.00
31	3322	Left elevon in/ center body flap	728	0.00	0.00	0.00
32	3332	Left elevon TPS, outboard	572	0.00	0.00	0.00
33	3332	Center TPS, aft	1040	0.00	0.00	0.00
			Totals	5.09	6.79	11.88

Table 5.23. TPS Min-Zone Catastrophic Failure Probabilities

5.1.6. Risk Contribution of Orbiter Components

The components of the Orbiter were analyzed independently, however, their risk contribution to the Orbiter may aggregated and compared. The risk contribution of the analyzed Orbiter components is shown in Figure 5.23. The APUs and TPS have been shown to dominate the Orbiter risk which implies that most of the Orbiter risk (91%) is realized during entry and descent.

Figure 5.23. Orbiter Risk Contribution

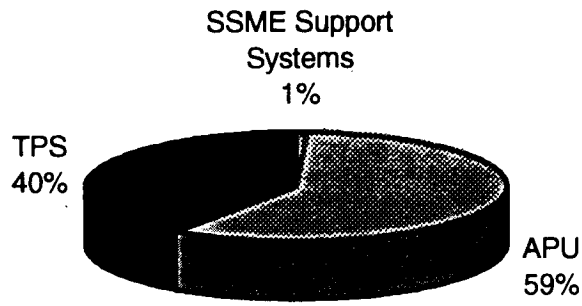

5.2. Supporting Systems

5.2.1. Orbiter Electric Power

5.2.1.1. Description

The electrical power system (EPS) consists of the equipment and reactants that produce electrical power for distribution throughout the orbiter vehicle, and fulfill all the orbiter external tank, solid rocket booster, and payload power requirements when not connected to ground support equipment. The EPS operates during all flight phases. For nominal operations, very little flight crew interaction is required by the EPS.

The EPS is functionally divided into three subsystems: power reactants storage and distribution, three fuel cell power plants (fuel cells), and electrical power distribution and control.

Through a chemical reaction, the three fuel cells generate all 28-volt direct-current electrical power for the vehicle from launch minus 3 minutes and 30 seconds through landing rollout. Prior to that electrical power is provided by ground power supplies and the onboard fuel cells.

Power is controlled and distributed by assemblies located in the forward, mid, and aft sections of the orbiter. Each assembly is a housing for electrical components such as remote switching devices, buses, resistors, diodes, and fuses. Each assembly usually contains a power bus or buses and remote switching devices for distributing bus power to subsystems located in its area.

The power reactants storage and distribution (PRSD) system stores the reactants (cryogenic hydrogen and oxygen) and supplies them to the three fuel cells that generate the electrical power for the vehicle during all mission phases. In addition, the subsystem supplies cryogenic oxygen to the environmental control and life support system for crew cabin pressurization. They hydrogen and oxygen are stored in tanks at cryogenic temperatures (-285° F for liquid oxygen and -420° F for liquid hydrogen) and supercritical pressures (above 731 psia for oxygen and above 188 psia for hydrogen).

The PRSD system components are located in the orbiter midbody underneath the payload bay or on a payload bay pallet for 10+ day missions in Extended Duration Orbiter (EDO) vehicles. The system stores the reactants hydrogen and oxygen in double-walled, thermally insulated spherical tanks with a vacuum annulus between the inner pressure vessel and outer tank shell. Each tank has heaters to add energy to the reactants during depletion to control pressure. Each tank is capable of measuring quantity remaining.

The tanks are grouped in sets of one hydrogen and one oxygen tank. The number of tank sets installed depends on the specific mission requirement and vehicle. Up to five tank sets can be installed in the midfuselage under the payload bay liner of OVs 102, 104 and 105 (four sets maximum for OV 103). Up to four additional tank sets can be flown on the EDO pallet in the payload bay of these vehicles.

5.2.1.2. Success Criteria

The electrical power system supplies power to all systems on the Shuttle vehicle and therefore the success criteria is linked to the power requirements to operate the supported systems within their safety success criteria. During ascent power must be supplied to all the Shuttle elements whereas only the orbiter requires power during the other phases since it is the only Shuttle element left. For this reason, two out of three fuel cells must be operational during ascent but the orbiter can orbit, de-orbit, descend, and land with the capacity of one fuel cell.

5.2.1.3. Data Analysis

The probability of failure for the EPS components was estimated from generic databases. The components of the EPS did not vary significantly from those utilized in other capacities and therefore the generic data offered the best combination of applicability, degree of failure, and sample size. The estimated frequency of identified electrical system accident sequences are shown in Appendix B.4.

5.2.2. Orbiter Environmental Control and Life Support (ECLSS)

5.2.2.1. Description

The ECLSS maintains the orbiter's thermal stability and provides a pressurized, habitable environment for the crew and onboard avionics. The ECLSS also manages the storage and disposal of water and crew waste.

ECLSS is functionally divided into four systems:

Pressure control system, which maintains the crew compartment at 14.7 psia with a breathable mixture of oxygen and nitrogen. Nitrogen is also used to pressurize the supply and waste water tanks.

Atmospheric revitalization system, which uses air circulation and water coolant loops to remove heat, control humidity, and clean and purify cabin air.

Active thermal control system, which consists of two freon loops that collect waste heat from orbiter systems and transfer the heat overboard.

Supply and waste water system, which stores water produced by the fuel cells for drinking, personal hygiene, and orbiter cooling. The waste water system stores crew liquid waste and waste water from the humidity separator. The system also has the capability to dump supply and waste water overboard.

5.2.2.2. Risk Implications

ECLSS was assessed as a non-contributing component to Shuttle LOV. The system's function is primarily one of crew maintenance and although the loss of the crew is a highly undesirable event the focus of this study is upon the preservation of the vehicle. It may be argued that the crew is an integral part of the Shuttle and is instrumental in providing for its safety but in reality the crew's function is one of mission success not vehicle safety. The crew becomes involved in some accident sequences of the helium supply system during ascent and the APU during descent and landing but the probability of a concurrent failure of ECLSS incapacitating the crew during either of these phases and the occurrence of an initiator requiring human intervention is considered negligible.

In most instances LOV implies loss of the crew however the converse is not considered to be an accurate statement. The loss of the crew while the vehicle is in orbit does not necessarily mean that the vehicle is lost, the vehicle is said to be stranded and contingencies for vehicle retrieval are certainly possible. With the sustained human presence once the International Space Station is operational the reaction time for such a contingency could be just a few days.

The presence of humans in the Shuttle are the best defense against its failure. The rich array of sensory information possible with the human body should make most ECLSS problems identifiable and correctable before they escalate to fatal proportions.

5.2.3. General Purpose Computer and Data Management

5.2.3.1. Description

The Data Processing System (DPS), consisting of various hardware components and self-contained software, provides the entire Space Shuttle vehicle with computerized monitoring and control. DPS functions include:

- Support the guidance, navigation, and control of the vehicle, including calculations of trajectories, SSME burn data, and vehicle attitude control data.
- Monitor and control vehicle subsystems, such as the electrical power system and the environmental control and life support system.
- Process vehicle data for the flight crew and for transmission to the ground, and allow ground control of some vehicle systems via transmitted commands.
- Check data transmission errors and crew control input errors; support annunciation of vehicle system failures and out-of-tolerance system conditions.
- Support payloads with crew/software interface for activation, deployment, deactivation, and retrieval.

- Process rendezvous, tracking, and data transmissions between payloads and the ground.

The DPS hardware consists of five general-purpose computers (GPCs), two mass memory units (MMUs) for large-volume bulk storage, and a network of serial digital data buses to accommodate the data traffic between the GPCs and vehicle systems. The DPS also includes 19 orbiter and four SRB multiplexers/demultiplexers (MDMs) to convert and format data from the various vehicle systems, three SSME interface units to command the SSMEs, four multifunction CRT display systems used by the flight crew to monitor and control the vehicle and payload systems, two data bus isolation amplifiers to interface with the ground support equipment/launch processing system and the SRBs, two master events controllers, and a master timing unit.

DPS software accommodates almost every aspect of Space Shuttle operations, including orbiter checkout, prelaunch and final countdown for launch, turnaround activities, control and monitoring during launch, ascent, on-orbit activities, entry, and landing, and aborts or other contingency mission phases. A multicomputer mode is used for the critical phases of the mission, such as launch, ascent, orbit, entry, landing, and aborts.

5.2.3.2. Risk Implications

The initial strategy of analysis for the DPS was to study the system as an independent entity but this proved to be ineffective for practical purposes. Instead the system was analyzed on a "need to basis" as its functions became necessary to meet the success criteria of other systems. The hardware of the DPS was not a risk significant contributing factor, electrical components in general are relatively high reliability components. In addition, redundancy is easily designed into computer systems, making the system extremely fault tolerant. The software of the DPS was a contributor to the frequency of some initiating events however accident sequences initiated by a software problem included other software related monitoring parameters which were effective in identifying the problem and mitigating it. Therefore no independent DPS model exists, instead the critical DPS functional failures are included in the fault trees of the front-line systems.

5.2.4. SSME Thrust Vector Control

5.2.4.1. Description

The six main engine TVC actuators receive hydraulic pressure from the three main propulsion system (MPS) isolation valves. There are two actuators for each SSME: one for controlling pitch, the other for controlling yaw. Each actuator receives hydraulic pressure from 2 out of the 3 hydraulic power systems; one configured as the primary system, one as the secondary system was shown in Table 5.16. The hydraulic systems are each powered by an auxiliary power unit (APU), which drives a hydraulic pump. Each actuator has a switching valve which switches to secondary pressure should the primary system pressure drop.

The MPS thrust vector control command flow starts in the GPCs, in which the flight control system (FCS) generates the position commands, and terminates at the TVC actuators, which gimbal the SSMEs in response to commands. All the position commands are issued to the MPS subsystem operating program (SOP). This program processes and transmits commands to their corresponding flight aft multiplexors. The commands are then separated and distributed to ascent thrust vector control (ATVC) channels, which generate equivalent analog voltages for each command issued. These voltages are then sent to the actuators, commanding the hydraulic actuator ram to extend or retract, thus gimbaling the main engine to which it is fastened.

The SSME actuators change each main engine's thrust vector direction as needed during the flight sequence. The three pitch actuators gimbal the engines up or down from the null position; the three yaw actuators gimbal the engines left or right from the installed position. Each actuator consists of a control module and an actuator ram.

The ram is fastened to the orbiter thrust structure by the actuator ram tail stock and to the gimbal bearing, which is on the engine attachment truss, by the actuator ram rod end. The aft end of the ram is attached to the mechanical feedback linkage. The actuator ram position is controlled by a primary hydraulic imbalance created by the power valve, which in turn is controlled by the force summed output of the secondary hydraulic pressure generated by four electro-hydraulic servovalves. Each servovalve receives identical commands from its own flight control system channel. These commands are in the form of current to torque motors which direct hydraulic pressure to the servovalve spools.

There are four flight control channels, each with an ATVC driver. Each main engine actuator contains four independent, two-stage servovalves, which receive signals from the drivers. Each ATVC driver provides a signal to one servovalve in each actuator. Each servovalve controls the power spool in concert with the other servovalves. This method is called *force-summed majority voting*. Because each servovalve receives identical commands, a single servovalve delivering the incorrect pressure to the power spool cannot dominate the correct action of the other three servovalves in gimbaling an SSME to the correct position.

The entire servo feedback loop is closed by mechanical linkages inside the actuator control module and does not rely on electrical feedback loops to the ATVCs or GPCs. Servoloop closure is provided by the mechanical position feedback linkage which transmits position from the actuator ram movement to the servovalve flappers, assuring a stable equilibrium point for a given constant command current input to the servovalve.

5.2.4.2. Success Criteria

The failure of any one pitch actuator, one yaw actuator, or one gimbal joint resulting in a loss to vector one SSME is not considered a catastrophic incident. Based on information in the *Operational Flight Rules* (Ref. 16), the other two SSME should have enough authority to overcome the one failure. Therefore a LOV is assumed to occur if two out of three SSME cannot be vectored.

5.2.4.4. Fault Tree Model & Data Analysis

A fault tree model was developed which satisfied the success criteria in the previous section. The model for the TVC is shown on page 551 of the integrated model. The model is developed such that any combination of two failures in two separate SSME which causes a loss of thrust vectoring control leads to an LOV event. The probability of such an occurrence in one flight was estimated at approximately 1 in 10 million. Note that this estimate is only for TVC component hardware failures, functional failure of the SSME TVC because of a loss of hydraulic pressure are attributed to the APU driven hydraulic system.

The SSME TVC is composed of components which have similiar counterparts in other industrial applications. For this reason it was found that the generic databases (Ref. 57 & 58) were sufficient sources of data to provide failure estimates for the TVC components.

5.3. Terminal Phase Risk

Experience indicated that human error might be a significant contributor to the risk of the Shuttle landing terminal phase. Therefore it was agreed that if the terminal phase were addressed directly, human error would become an important issue. Unfortunately human error risk models are not as mature as hardware risk models and the controversial analysis might detract from the rest of the study. Since the terminal phase was believed to be a low risk contributor to overall mission risk, a comparative study was conducted to bound the risk without a direct analysis.

The landing risk for commercial transport aircraft and military multiple crew member non-fighter aircraft were investigated as a basis of comparison to the risk of landing the Shuttle. The basic methodology involved comparing the Shuttle design and operational aspects to the comparative cases and identify issues which would lead to expected higher or lower Shuttle landing risk.

The following data was utilized in the comparative study:

- National Transportation Safety Board: Commercial Landing Mishaps, 1983-1993
- US NAVY P-3 Landings: Pilot Error Mishaps, Land Base, 1981-1994
- US NAVY E-2 Landings: Pilot Error Mishaps, Land Base, 1981-1994; Carrier, 1981-1994
- US NAVY C130 Landings: Pilot Error Mishaps, Land Base, 1981-1994

There are factors which should make the Shuttle landing safer than the studied incidences. The Shuttle pilots are all experienced test pilot school graduates and the crews train together for at least 12 months. Shuttle pilots are required to have over 1000 simulated approaches using the Shuttle Training Aircraft prior to flying the Shuttle. In addition an approach/landing outside the certified envelope is precluded by the flight rules.

On the other hand, there are certain issues which could make the case for the Shuttle worse. The Shuttle is committed to land one hour prior to touchdown. In addition the Shuttle has no go-around capability, lacks autonomous navigation, and has a limited divert capability. The crew does not train

on the actual flight vehicle and the landing speed is 9% below certification.

In the final analysis the probability of a landing incident for large commercial air carriers was found to be 1 in 100,000 landings, the Shuttle has experienced 7 incidences in 62 landings yielding an estimate of 1 in 9 landings. This seems to suggest that the Shuttle is 11,000 times more likely of having a landing incident than a large commercial air carrier. A study done by Rockwell International assumed a landing and rollout risk of 3% of overall risk. To be conservative SAIC applied this 3% criteria to the Galileo estimate giving a mean landing risk estimate of approximately 1 in 2,400 missions which falls within the bounds established by the comparative study.

5.4. Systems Analysis Summary Results

The results of the systems analysis discussed in this section are summarized in Table 5.24. The top-level Shuttle estimate is an aggregation of the element mean risk estimates. The mean number of missions between failures is simply the inverse of the probability of having a LOV incident in one mission. The percent contribution to the top-level Shuttle risk estimate is shown for those initiators or events which contribute more than one tenth of 1% of the total risk.

Table 5.24. Systems Analysis Summary Results

STS Element	Initiator Code	Initiator Description	Mean LOV Frequency (per mission)	Mean # of Missions Between Failures	% Contrib to STS Risk
STS			**7.66E-03**	**131**	
SSME			**2.87E-03**	**348**	**37%**
	SMECD	FAILURE TO SHUTDOWN ENGINE(S) AND DUMP PROPELLANTS	2.67E-06	374532	
	SMEDS	SIMULTANEOUS DUAL SSME PREMATURE SHUTDOWN	6.46E-06	154799	0.1%
	SMEFH	LOSS OF GROSS H2 FLOW	1.89E-07	5291005	
	SMEFO	LOSS OF MCC PRESSURE	1.48E-08	67567568	
	SMEHL	HYDRAULIC LOCKUP REQUIRED	5.79E-07	1727116	
	SMELO	HPFTP COOLANT LINER OVERPRESSURE	1.51E-07	6622517	
	SMELP	PROPELLANT MANAGEMENT SYSTEM AND/OR SSME COMBUSTIBLE LEAKAGE	1.55E-07	6451613	
	SMEMF	HIGH MIXTURE RATIO IN FUEL PREBURNER	9.48E-08	10548523	
	SMEMO	HIGH MIXTURE RATIO IN OXIDIZER PREBURNER	9.48E-08	10548523	
	SMEPB	LOSS OF FUEL TO BOTH PREBURNERS	1.84E-08	54347826	
	SMEVP	FAILURE TO MAINTAIN PROPELLANT VALVE POSITION; SEII GENERATED	9.08E-08	11013216	
	SMEST	FAILURE TO MAINTAIN CRITICAL STRUCTURAL INTEGRITY OF SSME	2.86E-03	350	37.3%
		LOV due to FPB structural failure	2.01E-04	4968	2.6%
		LOV due to HEX structural failure	4.20E-08	23809524	
		LOV due to HGM structural failure	3.00E-05	33333	0.4%
		LOV due to HPFTP structural failure	5.02E-04	1990	6.6%
		LOV due to HPOTP structural failure	1.00E-03	995	13.1%
		LOV due to LPFTP structural failure	4.20E-08	23809524	
		LOV due to LPOTP structural failure	1.51E-04	6623	2.0%
		LOV due to MCC structural failure	6.45E-04	1550	8.4%
		LOV due to MI structural failure	2.01E-04	4967	2.6%
		LOV due to NOZZLE structural failure	4.20E-08	23809524	
		LOV due to OPB structural failure	1.26E-07	7936508	
Orbiter			**2.95E-03**	**339**	**39%**
		SSME Support Systems	**2.20E-05**	**45375**	**0.3%**
	SMETV	SSME THRUST VECTOR CONTROL SYSTEM FAILURE	1.07E-07	9345794	
	SMEPV	FAILURE TO MAINTAIN PROPELLANT SUPPLY SYSTEM VALVE POSITIONS	1.89E-05	52910	0.2%
	SMELH	HELIUM SYSTEM LEAKAGE	2.94E-06	340136	
	SMEPG	FAILURE TO PROVIDE HELIUM POGO CHARGE	9.15E-08	10928962	
		APU	**1.74E-03**	**576**	**22.7%**
	APUAOK	LOSS OF HYDRAULIC PRESSURE DURING ASCENT	1.92E-04	5208	2.5%
	APUEL0	ASCENT INITIATED APU HYDRAZINE LEAKAGE DURING ENTRY/DECENT	1.64E-06	609756	
	APUELL	LARGE GAS/HYD LEAK ADVERSELY EFFECTS AFT COMPARTMENT ENVIROMENT	5.32E-06	187970	0.1%
	APUELT	ASCENT INITIATED 3 APU HYDRAZINE LEAKAGE DURING ENTRY/DECENT	1.25E-07	8000000	
	APUEOK	LOSS OF HYDRAULIC PRESSURE DURING ENTRY/DESCENT	1.48E-03	676	19.3%
	APUETU	APU/HYD TURBINE OVERSPEED/HUB FAILURE	5.65E-05	17699	0.7%
	APULK	APU HYDRAZINE LEAKAGE DURING ASCENT	3.86E-08	25906736	
		TPS	**1.19E-03**	**840**	**15.5%**
ISRB			**1.26E-03**	**794**	**16%**
		RSRM	**5.59E-04**	**1788**	**7.3%**
	RSRHGLK	RSRM JOINTS: HOT GAS LEAK	3.98E-04	2513	5.2%
	RSRNZRUP	RSRM NOZZLE RUPTURE	8.90E-05	11236	1.2%
	RSRPVRUP	RSRM PRESSURE VESSEL RUPTURE	7.22E-05	13850	0.9%
	RSRWRTHR	RSRM WRONG THRUST	1.00E-08	100000000	
		SRB	**6.99E-04**	**1431**	**9.1%**
	SRBNOHLDN	SRB NO, LATE, OR IMPROPER HOLDDOWN RELEASE	2.58E-04	3876	3.4%
	SRBNOIGN	NO OR LATE IGNITION OF 1 SRB/RSRM	2.22E-04	4505	2.9%
	SRBNOSEP	SRB FAILS TO SEPARATE	1.39E-04	7194	1.8%
	SRBPREMHD	SRB HOLDDOWN: PREMATURE RELEASE	1.60E-06	625000	
	SRBRECPREM	SRB RECOVERY DEVICE: PREMATURE RELEASE	6.00E-06	166667	0.1%
	SRBSTR	SRB STRUCTURAL FAILURES	1.00E-06	1000000	
	SRBTV	SRB THRUST VECTOR CONTROL SYSTEM FAILURE	7.13E-05	14025	0.9%
ET			**1.92E-04**	**5200**	**3%**
LANDING			**4.11E-04**	**2433**	**5%**

6.0. Model Evaluation & Results

6.0 Model Evaluation & Results

6.1 Representation and Propagation of Uncertainties

6.1.1. Application of Uncertainty Bounds

The data analysis effort yields a point estimate which would imply that the actual failure rate is known with 100% certainty. This would only be possible if an infinite numbers of actual missions had been completed, in other words, a failure rate can never be known with 100% certainty. However, as experience is gained in the operation of a system, the additional knowledge gained tends to decrease uncertainty. Therefore an uncertainty distribution is a measure of the current state of knowledge on the particular failure in question. Here, the Loss of the Shuttle Vehicle (LOV).

Uncertainty distributions were defined for the basic events which contributed the top 99.9% of the risk as determined from the propagation of the point estimates through model. This accounted for approximately 20% of the basic events, saving both computation time and labor (defining uncertainty distributions for the other 80% of basic events is unjustifiable given the negligible effect on the final result).

There are two basic sources of uncertainty:

- Model Development Uncertainty
- Data Analysis Uncertainty

Model uncertainty concerns the degree of confidence placed in the accuracy of the representation of system operation in the PRA model. Given the breadth of information supplied by NASA and the Shuttle contractors concerning the operation of the Shuttle, this aspect of uncertainty was not considered the driving factor of uncertainty. Data related uncertainty is influenced by both the amount of data available and how it was modified to estimate the equivalent flight failures. The statistical uncertainty (that uncertainty due to the finite nature of a data set) may be determined from the amount of data available. The uncertainty contributed by the modification of the data involves considering both the degree of applicability and degree of failure realized. This can be done by performing a detailed analysis of the phenomenological effects which drive the mechanism of failure. This approach was beyond the scope of this study, instead conservative uncertainty estimates were made using previous studies and expert judgment as a guide. Mathematically the uncertainty is represented by the error factor of the distribution, which as discussed above was determined by considering the attributes of the data set and the method of alteration. The final error factors (for the top 99.9% risk contributing basic events) are shown in Volume III: Appendix A.2.

All uncertainty distributions were assumed to be lognormally distributed with the point estimates serving as the mean and an error factor representing the level of uncertainty in the mean estimate. The lognormal distribution is a "natural" distribution for describing data which can vary by orders of magnitude. If the failure rate is expressed as 10^{e}, where e is some exponent, then describing the data

as having a lognormal distribution is equivalent to describing the exponent, *e*, as having a normal distribution. The positive skewness of the lognormal distribution also lends itself to the general reliability-associated behavior of assessed data by accounting for less likely but large deviations such as abnormally high failure rates due to period slips in quality control.

6.1.2. Evaluation of Top-level Uncertainty Distributions

The first step in obtaining a top-level result is integrating the various system models into a unified PRA framework. The boolean format chosen as the fundamental structure of the integrated model was a fault tree framework. Therefore all accident sequences leading to LOV were converted into fault trees. Only the sequences with an LOV consequence were converted because the model will be used to determine the probability of LOV. The conversion process involves combining the top events in the event trees which contribute to LOV with and gates. Some of these top events have associated fault trees which were included in the integrated model. Each of the LOV sequences were grouped under the initiating event of the event tree. The initiating events are grouped under the Shuttle element functional failures they contribute to. At this point a fault tree exists for each Shuttle element and these are linked to form the overall Shuttle risk model. The model is shown in Appendix A.

The reason for defining uncertainty distributions at the basic event level is to obtain an uncertainty bound for the top-level risk. This is accomplished by propagating the basic event uncertainty distributions through the model. CAFTA™, a microcomputer-based fault tree analysis workstation, was used to automate the quantification process. The integrated Shuttle fault tree model is input into the fault tree editor module of the program and the mean and error factors for the basic events are entered into the reliability database editor. The output is a list of minimal cutsets and associated probabilities. The minimal cutsets are the combinations of basic events which cause a catastrophic

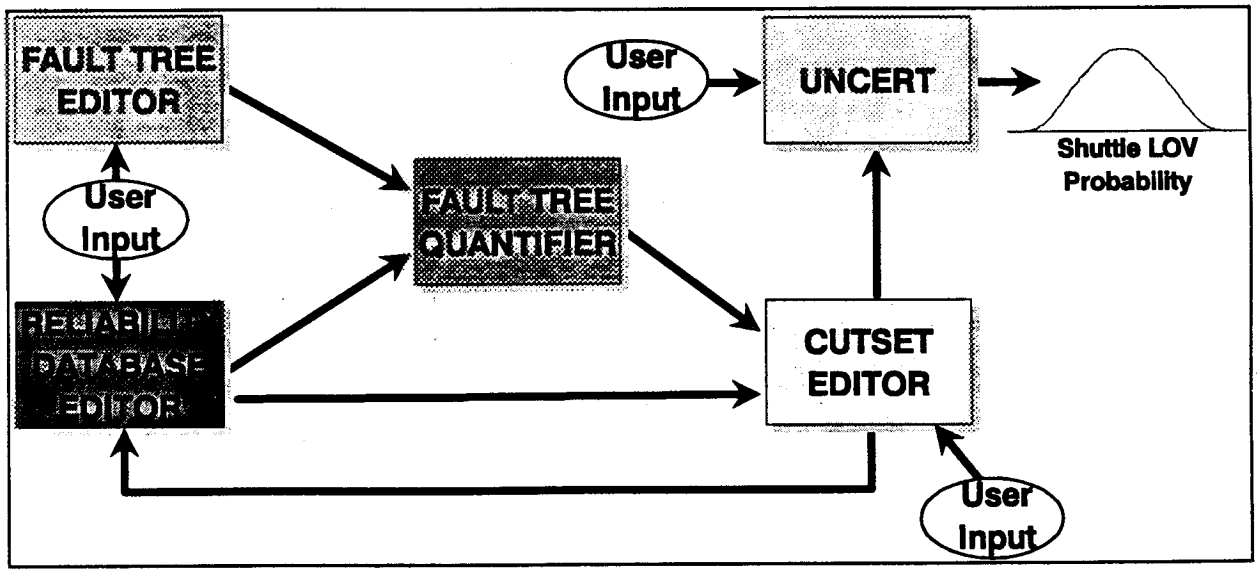

Figure 6.1. Information Flow between Analysis Computer Modules

failure. The basic event probabilities are aggregated to give the cutset probability which in turn are aggregated to obtain the top-level point risk estimate.

The uncertainty analysis was conducted using another routine called UNCERT™. UNCERT™ uses the cutsets from CAFTA™ as inputs along with the basic event mean and error factor estimates. The routine uses a Monte Carlo method to combine the basic event distributions to obtain uncertainty distributions for predefined top events. Uncertainty distributions were obtained in this fashion for the Shuttle vehicle as a whole and for each of the individual elements. The ET and landing uncertainty distributions were not determined from a Monte Carlo method since they are basic events in the current model.

6.2 Base Case Risk Evaluation

The base case model is the model which has resulted from the analysis as it has been described in the preceding chapters, in the next section variations to this model will be assessed to show the effect or sensitivity of certain assumptions or methodology to the final result. The final results of the base case model are shown in Figure 6.2 and tabulated in Table 6.1. The uncertainty bars show the range, with 90% confidence of the risk or LOV contributing probability of each Shuttle element. The median has the largest associated probability of being the actual risk but in actuality the risk may lie anywhere in the uncertainty range which increases as the confidence interval is increased. Although the mean or median risk of one element is larger than another, if the two uncertainty distributions overlap there is a probability that the element with the smaller mean estimate could actually be the higher risk contributor. For instance, in the case of the Orbiter and SSME, which have practically identical distributions, either one has a 50% chance of being the higher risk contributor. On the other hand the degree of overlap between the ISRB and SSME implies that there is approximately a 25% chance that the ISRB is a higher risk contributor than the SSME.

Table 6.1. Summary of PRA Results: Estimated Loss-of-Vehicle Frequency

	5th Percentile	Median	Mean	95th Percentile
STS	1/230	1/145	1/131	1/76
Orbiter	1/758	1/397	1/330	1/169
SSME	1/820	1/410	1/348	1/172
ISRB	1/2591	1/1152	1/775	1/322
ET	1/86207	1/11223	1/5208	1/1460
LANDING	1/141528	1/9435	1/2433	1/629

Figure 6.2. LOV Risk Uncertainty Distributions for Total Shuttle Mission

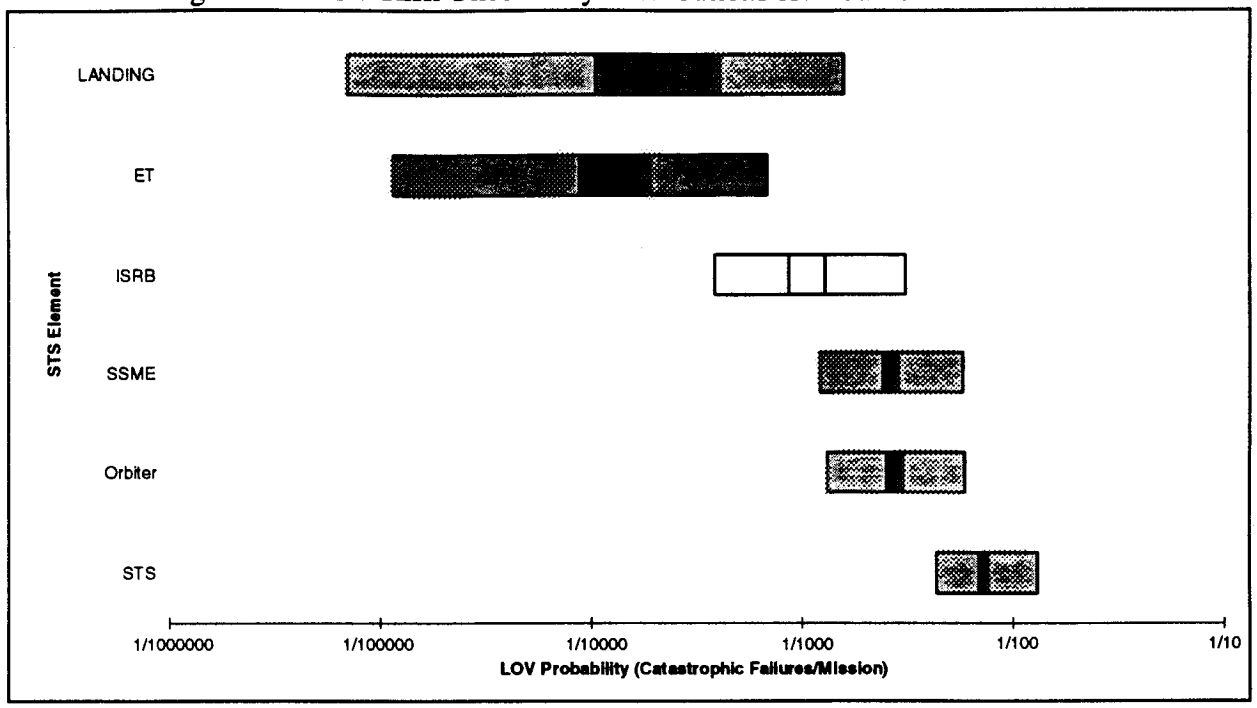

Figure 6.3. Distribution of Mean Loss-of-Vehicle Risk Among Shuttle Vehicle Elements.

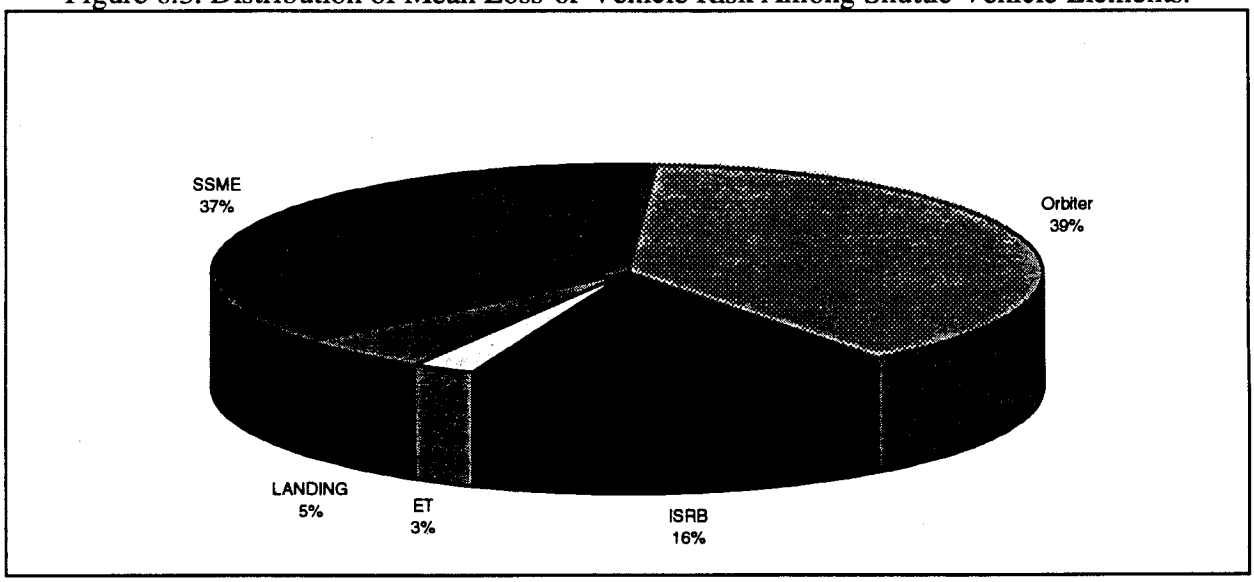

The degree of uncertainty for any component tends to decrease up to a point as more flight experience is gained. The uncertainty for particular items such as landing may be reduced by a more detailed study. Another method of reducing uncertainty is a properly defined testing philosophy geared to collecting data on phenomena which is not well understood at the moment or cannot be

studied during flight (e.g. landing gear tire reliability).

Table 6.2. Comparison of Current Shuttle PRA Ascent Results and Previous Studies

	5th Percentile	Median	Mean	95th Percentile
STS PRA	1/429	1/248	1/219	1/118
PRA Phase 1 Study	1/223	1/90	1/73	1/31
Galileo Study	1/350	1/78	1/55	1/18

Figure 6.4. Comparison of Current PRA Ascent Results with Previous Risk Studies

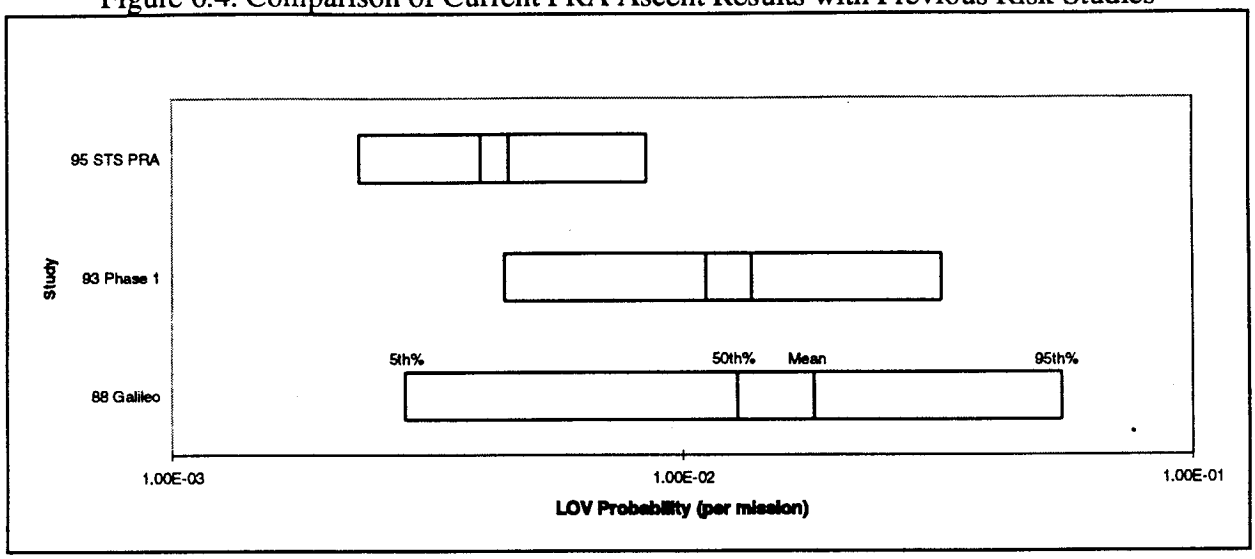

Uncertainty distributions may be developed for various portions of the mission. This involves identifying and isolating functional failures which are phase specific and propagating the basic event uncertainties through those particular portions of the model. This was done for the ascent phase which has been the focus of past studies such as the Galileo study which was updated in Phase 1 of this project. The comparison of the Galileo and updated Phase 1 results are compared with the ascent specific results of this PRA in Figure 6.4. The first striking feature is the relatively smaller uncertainty which is not surprising considering the more detailed analysis of the current study when compared to the top-level Galileo study. The other differentiating attribute is the lower mean risk displayed by the PRA results. One of the main reasons for this difference is the additional credit given to the ISRB for leak checks performed prior to flight. Another is the number of successful flights which have occurred since the Galileo study in 1988. In previous studies only actual hotfire exposure was considered in estimating the reliability of RSRM seals. Fortunately most of the seals are not exposed to combustion gases during flight therefore flight success alone does not allow credit to be given to redundant sealing function and thus a conservative estimate of their reliability was made in the Galileo study. During the PRA, Thiokol supplied SAIC with additional data (shown in Appendix B.2) concerning leak checks of the various seals in the RSRM. Although success of a leak check does not

guarantee that the seal will not leak during flight it does give some indication of seal integrity which can be used to give some credit towards seal reliability. The accounting of partial credit was discussed in section 5.1.2.5., the result was an overall increase in the reliability estimate of the RSRM with an associated drop in risk, as can be seen in the element risk comparison in Figure 6.5. The magnitude of the effect the inclusion of the leak check data will be one of the sensitivity cases in the next section.

Figure 6.5. Comparison of Ascent Risk Uncertainty Distributions for Shuttle Elements

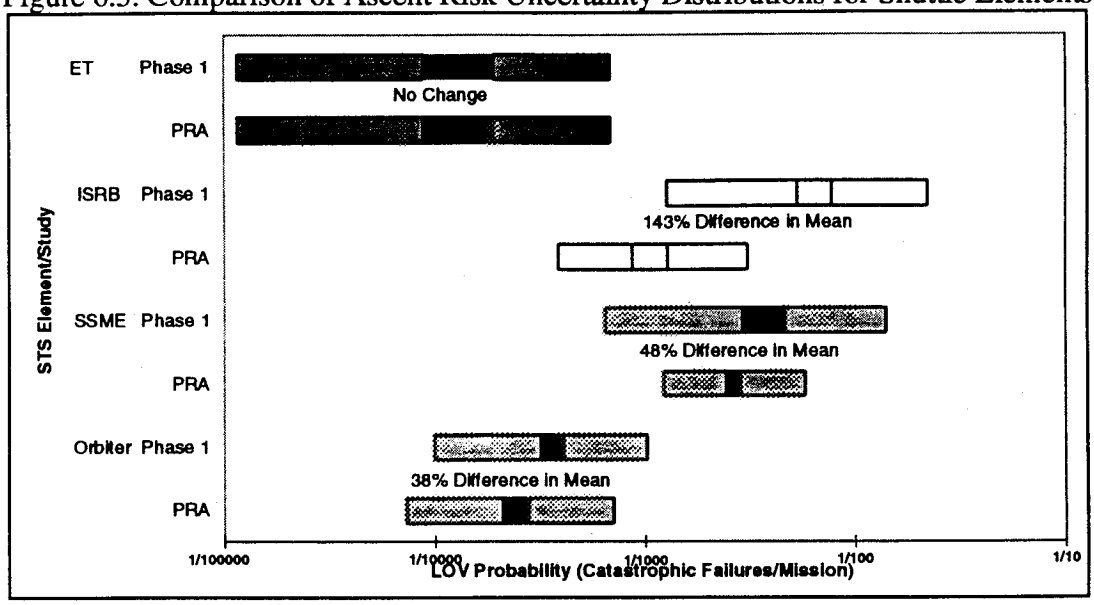

An estimate of the risk of orbit and re-entry/descent may be conducted in a similar fashion. The in-orbit risk was attributed to latent hydrazine leakage which deflagrates during re-entry however the estimate showed relatively little risk compared to ascent and descent. The risk for ascent and descent as well as the respective contribution to each is shown in Figure 6.6.

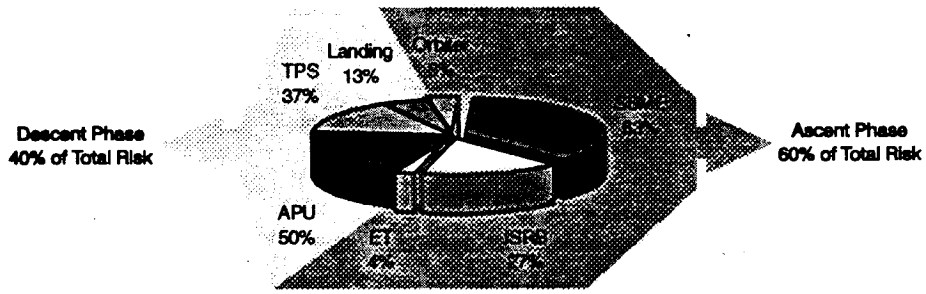

Phase Description	
Ascent	SSME Ignition Through APU Shutdown
Orbit*	APU Shutdown to De-orbit Burn
Descent	De-orbit Burn Through Wheel Stop at Landing

*Orbit phase did not contain any significant LOV risk drivers

Figure 6.6. Mission Phase and Relative Element Risk Contribution

The risk may be broken down to even smaller mission intervals to show how the functional and environmental variations during the mission effect the degree of risk encountered by the Shuttle (Figure 6.7). Note that the risk is instantaneously attributed to the point at which catastrophic failure is realized. For instance, in the case of the TPS, the damage which causes the tiles to fail during re-entry may have been inflicted during ascent at ISRB separation. However the risk of TPS failure is considered to occur once the environmental conditions challenge the effectiveness or vulnerability of the system. The bars illustrate the specific risk or risk per unit time for each defined interval. Notice how concentrated the risk is during the first 128 seconds of flight when most of the "energetic" active systems are operating such as the ISRB and SSME. The overall results show the Orbiter as the major risk contributor but this risk is distributed along the entire mission while that of the ISRB and SSME is concentrated during the fleeting moments of ascent. This is a conclusive demonstration of the intuitive concern which most members of the Shuttle community express about the relative risk of ascent. A surprising result of the study was that the risk of descent may be on the same order of magnitude as ascent, as shown in Figure 6.6, but the short time under which the risk is realized during ascent makes it much more noticeable.

Figure 6.7. Relative Risk Versus Mission Intervals: Linear Risk Scale.

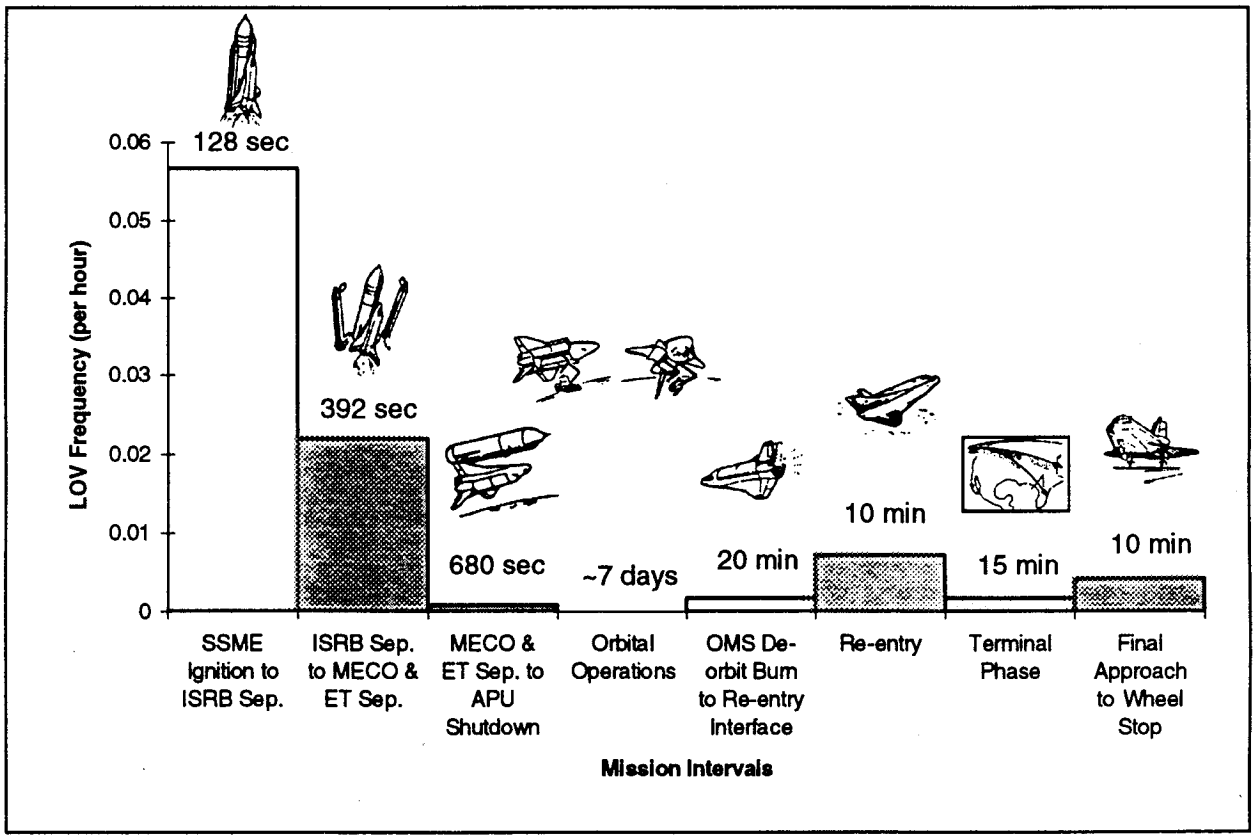

The cutsets generated from the model quantification may be analyzed and grouped to obtain a ranking for the risk significant accident sequences. The top 20 risk contributing accident sequences are shown in Table 6.3 along with their mean contribution to the total Shuttle risk. Examining the results

Table 6.3. Summary of Top 20 Risk-Contributing Accident Sequences.

Rank	Accident Description	Mean LOV Prob	Mean Percent Contrib. to Total Risk
1	SSME HPOTP Bearing Failure Due To Spalling, Pitting, Wear Or Corrosion	4.52E-04	5.89%
2	Two Leakage Induced Orbiter APU Failures During Re-entry/Descent and Failure To Land Using One APU	4.28E-04	5.57%
3	Two Orbiter APUs Fail To Start Or Run During Re-entry/Descent Due to Common Cause Failure and Failure To Land Using One APU	3.99E-04	5.20%
4	All Three Orbiter APUs Fail To Start Or Run During Re-entry/Descent Due to Common Cause Failure	3.43E-04	4.47%
5	SSME MCC Manifold Weld Failure	2.53E-04	3.29%
6	SSME HPFTP Turbine Blade Failure	2.51E-04	3.27%
7	Catastrophic Failure Of Right Side TPS, Fwd Mid Edge (624 Tiles)	2.48E-04	3.23%
8	Common Cause Failure of ISRB Igniter Joint S&A Primary and Secondary Gasket Seals	2.10E-04	2.73%
9	SSME HPOTP Failure Due To Cavitation Damage	2.01E-04	2.62%
10	SSME HPFTP Impeller/Diffuser Failure	2.01E-04	2.62%
11	Propellant Fails To Ignite In One Of The ISRBs	2.00E-04	2.60%
12	All Three Orbiter APUs Fail To Run During Ascent Due to Common Cause Failure	1.92E-04	2.50%
13	Catastrophic Failure Of Left Side Near Main Landing Gear TPS (780 Tiles)	1.87E-04	2.43%
14	Two or more ISRB Holddown Studs Hang-up	1.78E-04	2.31%
15	Failure In SSME MCC EDNi Liner Closeout Structure	1.76E-04	2.29%
16	Catastrophic Failure Of Forward Right Side Near Main Landing Gear TPS (676 Tiles)	1.75E-04	2.28%
17	SSME MI Lox Post Structural Failure	1.51E-04	1.97%
18	Structural Failure Of SSME LPOTP	1.51E-04	1.97%
19	SSME HPOTP Turbine Blade Failure	1.51E-04	1.97%
20	SSME FPB Faceplate Failure Due To Erosion	1.51E-04	1.97%

shows that this implies that the first two accident sequences alone contribute over 10% of the entire Shuttle risk. These sequences are considered to be the risk drivers of the system and the improvement of the associated components have the most potential of positively impacting Shuttle safety. Assessing the cost effectiveness of these improvements using the current PRA model will be discussed in the next section. Table 6.4. shows the contribution of risk driving element components to as a percentage of the risk contribution of the top 10 and top 20 ranked accident sequences. Note that the Orbiter APUs and the SSME turbomachinery have the most significant impact from a potential risk reduction perspective. That is these components have a combination of both criticality and failure possibility which warrants further investigation and potential mitigation through redesign.

Table 6.4. Risk Summary Statistics of Most Significant Accident Sequences
(Sequences shown in Table 6.3)

	Top 10 Accident Seq.	Top 20 Accident Seq.		Top 10 Accident Seq.	Top 20 Accident Seq.
Percent of Total Risk	38.88%	61.17%			
Orbiter	47.49%	41.98%	Auxiliary Power Units	39.18%	28.99%
			Thermal Protection System	8.31%	12.99%
SSME	45.48%	45.51%	Turbomachinery	37.01%	29.95%
			Combustion Devices	8.47%	15.56%
ISRB	7.03%	12.51%	Redesigned Solid Rocket Motor	7.03%	8.73%
			Solid Rocket Booster	-	3.78%

6.3 Sensitivity Analysis and Redesign Effectiveness Measures

The base case of the PRA model has yielded some interesting results but its real utility is realized by introducing proposed variations to demonstrate the effect of the applied methodology, assumptions and proposed changes to Shuttle design features. The primary variations to the model were made to demonstrate the effect of two issues raised during the development of this model. The first is the inclusion of leak check data into the ISRB estimate and the second is the consideration of common cause failures in the analysis of the orbiter APU. Excluding each independently and in unison, the changes in the associated elements were as shown in Table 6.5. Note that the variations are well within the predetermined uncertainty distributions for the base case.

Table 6.5. Estimated Loss-of-Vehicle Frequency for Base and Sensitivity Cases.

Sensitivity Case Description		STS	Orbiter	SSME	ISRB	ET	Landing
Base Case	Mean Mission	1	1	1	1	1	1
As Described in Report	LOV Probability	131	330	348	775	5208	2433
	Risk Percentage		39%	37%	17%	2%	5%
Sensitivity Case 1	Mean Mission	1	1	1	1	1	1
Additional Credit for Successful	LOV Probability	106	330	348	337	5208	2433
ISRB Leak Checks Not Considered	Risk Percentage		32%	30%	31%	2%	4%
Sensitivity Case 2	Mean Mission	1	1	1	1	1	1
Common Cause Failures not	LOV Probability	150	508	348	775	5208	2433
Considered for APUs	Risk Percentage		29%	43%	19%	3%	6%
Sensitivity Case 3	Mean Mission	1	1	1	1	1	1
Both ISRB Leak Checks and APU	LOV Probability	119	508	348	337	5208	2433
Common Cause Failures Neglected	Risk Percentage		23%	34%	35%	2%	5%

The model may also be used to assess the risk related cost effectiveness of proposed design changes or improvements. For a system such as the Shuttle which has a proven performance record, modular design improvements are the key to increasing safety and reducing costs. The model may be used to weigh the cost of the proposed changes against the expected potential loss of not making the design modification. For instance, APU risk due to hydrazine leakage was found to contribute to about 10% of the overall Space Shuttle risk, which could be argued makes this failure mode a candidate for a redesign effort. Taking a hypothetical simplistic case, a design modification is proposed for the hydrazine piping which testing shows reduces the probability of leakage by 75% from the current estimate. Propagating this estimate through the PRA model yields a 6.5% decrease in Shuttle risk. If the cost of losing a Shuttle is assessed at $5 billion and it is assumed that the Shuttle will be the primary launch vehicle for the indefinite future the 6.5% change in risk translates into a $325 million cost of protection against expected loss of the vehicle. Therefore if the design change is justifiable from a risk perspective if it may be made for less than $325 million. Conversely, design improvement objectives may be established for various cost estimates.

7.0. Conclusions & Recommendations

7.0 Conclusions and Recommendations

7.1. Risk Insights Based on PRA Results

The results of the PRA indicate that the Shuttle has been demonstrated to be by far the most reliable and least risky of all launch vehicles in the world (see Figure 7.1). However the distinctive advantage of the Shuttle as a returnable and reusable vehicle makes even this comparison fall short of the Shuttle's clear dominate position with respect to other vehicles. Despite this dramatic improvement of the Shuttle over the current inventory of launch vehicles Shuttle LOV risk continues to be substantial. The probability of having a catastrophic failure during a nominal flight was assessed to be, with 90% confidence, between 1 in 76 and 1 in 230 per mission. This implies that if the Shuttle is flown until the year 2030 with an average of 7 missions per year (245 missions), the risk of the occurrence of at least one more catastrophic failure is substantial. Flying the Shuttle until 2015 at the same launch rate corresponds to a 50-50 chance of a catastrophic failure occurring. Note that these risk estimates correspond to the current Shuttle design; changes in design or processing could substantially improve the reliability of certain components thereby reducing the risk to the Shuttle.

Figure 7.1. Risk Comparison of International Launch Vehicles

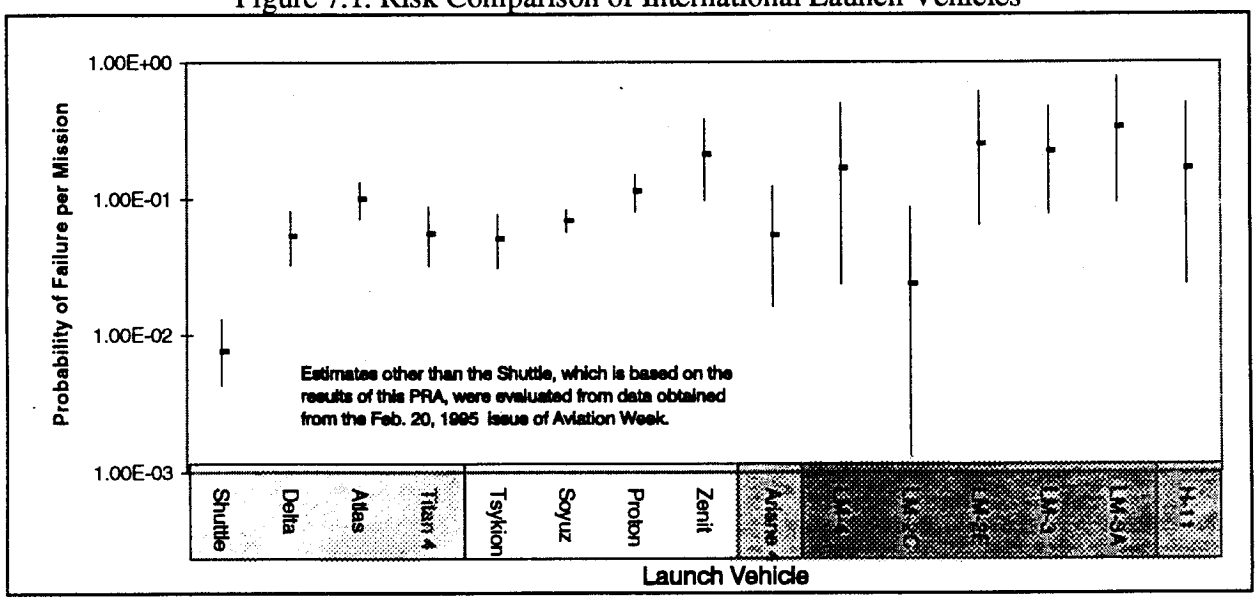

The redesign of the solid rocket boosters seems to have significantly reduced the risk due to the failure mode which caused the Challenger accident. However the ISRB still remains an operationally risky element of the Shuttle vehicle. Although the SAIC team was convinced that the Shuttle booster is the most reliable rocket motor of its type to be built, operating such powerful propulsion systems will always pose a challenge to the safety of a launch vehicle. This is substantiated by the fact that the ISRB risk rate (risk per unit time of operation) continues to be by far the highest of all the Shuttle elements.

The three SSME were shown to contribute a significant portion of the Shuttle risk. They account for 37% of the overall Shuttle flight risk even though they are active only during ascent. Practically all of the SSME risk is due to sudden catastrophic structural failure of one of the high energy components (HPOTP, HPFTP and MCC). The redlines which were established to shutdown the engine in the event of off-nominal operation were found to be extremely effective at accomplishing this task. However, an SSME shutdown leads to Shuttle operational conditions which may prove to be even more dangerous than continuing to fire the engine which was to be shutdown. Abort scenarios were not included in this study because of their second order impact. However the results of the study indicate that they should probably be considered in any extension of this study.

The risk of the Orbiter is dominated by failures of two of its main systems, the APU driven hydraulic system and the tiled thermal protection system (TPS). The APU system was found to be susceptible to common cause failures which resulted in multiple APU losses. Although the system was designed to be redundant the propensity for multiple failures negates the advantages of having back-up components. A significant amount of the common cause failures are due to hydrazine leakage. The TPS risk was found to be dominated by certain portions of the tiles which are susceptible to debris generated during separation of the right ISRB. Even though this damage occurs during ascent there is currently no opportunity for inspecting the tiles and repairing damaged ones before they are required during re-entry.

7.2. Design and Operations Recommendations

Further detailed study would be necessary to make effective recommendations for design and operational modifications but some salient safety issues may be discussed. The propensity of the APUs to leak hydrazine might be curtailed by improving the hydrazine plumbing or perhaps by eliminating hydrazine altogether and using electric powered APUs instead. Of course these options and any other risk reduction recommendations requires that their potential risk reduction benefits in terms of the potential loss protection justifies the associated cost development.

Risk reduction efforts on the ISRB would probably be best applied to the pyrotechnic related processes. A redesign may not be necessary but at the very least the failure modes discussed in this PRA should be studied further to ensure that an acceptable degree of reliability has been realized. On a related issue, efforts should be made to reduce the amount of debris which impinges upon the orbiter from detonation of the separation motors.

The SSME, not unlike the orbiter, also is a prime candidate for risk reduction efforts. Some of these might be cost effective from an operations stand point. The process of redesigning SSME components has already been initiated, a new HPOTP is currently being certified and is slated to be flown in mid-1995. As mentioned in previous sections the catastrophic failure modes of the SSME are primarily driven by single point structural failures. The utilization of advanced materials tailored towards eliminating the mechanisms which drive certain components to failure (e.g. crack initiation and propagation) would offer the most effective means of reducing SSME risk.

7.3. Recommendations for Continuing Risk Assessment and Management Work

The PRA model developed herein does not represent a complete Shuttle risk model. Nor does SAIC claim it to be. However it is SAIC's belief that the model has been developed to a stage which captures a significant portion of the Shuttle risk. Additional expansions would certainly be worth considering. For example although abort scenarios were identified they were not developed and therefore the associated potential risk can only be roughly estimated. For this reason the model has been developed to be a "living" model which may be modified and amended as deemed necessary to provide risk insights to a variety of management inquiries.

For example the model might be used to establish realistic cost objectives for redesigning the risk driving components. A simplified case was shown in section 6.3 for the orbiter APU. The cost estimates for any proposed design improvement could be tied to exact improvement objectives on a risk based criteria. This methodology will assure that limited resources are focused towards solving the problems which will have the most impact on safety.

The model may also be extended and modified to include turnaround processing and maintenance to illustrate the effect on operational risk. Such an analysis would provide a mechanism for ensuring that cutbacks in processing budgets do not significantly influence Shuttle safety. Extensions of this sort would allow processing tasks to be ranked according to their risk reduction worth and the cost incurred to perform the task. In this way management may quickly and concisely compare a task's overall worthiness in meeting future cost constraints and safety objectives.

The current study indicates that it would be useful to consider abort scenarios. The current reasonably high estimated probability of their occurrence (approximately 3 in 100 missions) warrants for their attention. The risk analysis of abort scenarios differs from the current PRA in that the time at which the initial event occurs is crucial to the criticality of the final consequence. The dynamic nature of this problem further increases the complexity of the analysis process in order to properly represent the true abort risk.

A part of the nominal mission risk, as well as abort risk, originates from landing related processes. Although this study did account for this risk, the associated uncertainty was found to be rather high. This may not be as much of an issue for a nominal flight as it would be for an abort scenario which would require Shuttle pilots and equipment to operate under less tolerant and more strenuous conditions. Therefore a more involved study of the landing process would offer more concise bounds on the related risk and provide insights and set the groundwork for the an analysis of abort scenarios.

In the near future the Shuttle will be utilized in constructing the International Space Station Alpha (ISSA) and will later dock with the ISSA for extended periods of time. These activities introduce processes which differ appreciably from today's nominal orbital operations and in effect introduce associated risks. One of the more obvious risks being the potential for problems during the docking maneuvers which involve two large space structures rendezvousing, precisely maneuvering in close proximity and docking to allow exchange of materials and personnel. Not unlike the propagation of accidents from one Shuttle system to another, attaching two complex systems together for extended

periods of time introduces interfacing risks which should be studied and understood.

Extending beyond the sphere of influence of the Shuttle system, the ISSA is a system which merits analysis independent of the Shuttle. The initial risk related activities which may impact the ISSA are those involved in the assembly process. There are direct risks involved in assembling the ISSA such as the unprecedented amount of EVA required and the manipulation of large construction materials in orbit and the risk of the ISSA maintaining favorable attitude. There are also indirect risks which include latent failures made during the assembly process which could later impact the operational phase (i.e. flawed mirror on Hubble Space Telescope). Both direct and indirect risk of the assembly process as well as the operational risk of the ISSA should be studied to ensure reliable and efficient service during its usable life.

In conclusion, just like any other tool the utilization of the PRA as a risk management tool is only limited by the ingenuity of its users. This document has shown what type of results a PRA may provide and how those results may be used to better allocate resources for the purposes of safety. Future space transportation systems and in-orbit facilities may be better served by conducting such an analysis as part of the design process.

[redacted]

References

1. STS Technical Manual, SSME Description and Operation (E41000/RSS-8559-1-1-1, Sept.1,1983)

2. STS Technical Manual, SSME Intermediate Level Maintenance (E41000/RSS-8559-1-1-1, Volume I, June15,1987)

3. STS Technical Manual, SSME Illustrated Parts Breakdown (E41000/RSS-8559-1-1-1, Volume III, Feb.14,1989)

4. SSME Phase 2 Predicted Engine Performance (OL 87RC06038)

5. The SSME Performance Book, 1st Edition (December 1992)

6. SSME Reliability Predictions with ATD HPOTP Phase II+ Powerhead and Large Throat MCC (A.L. Hallden, Dec. 1992)

7. Using Variable-Power Test Results to Estimate Engine Reliability (JSC-26244, Alan H. Feiveson, June 1993)

8. SSME Reliability (Bob Biggs, Feb. 1991)

9. Space Shuttle Main Engine Reliability Analysis (Faysal Safie, May 3, 1993)

10. Space Shuttle Main Propulsion Pressurization System PRA: Results Review (Lockheed, H.E. Smith, Dec. 17,1987)

11. Report of the Shuttle Processing Review Team (June 30, 1993)

12. SSME Contract End Item Engine Specification (CP320R0003B, May 10, 1973)

13. Space Shuttle Interface Control Document, Space Shuttle Orbiter Vehicle/Main Engine (ICD-13M15000, April 4, 1993)

14. STS Training Manual, SSME Orientation (Part A-Engine) (ME-110(A)RIR, Jan. 1991)

15. SSME Description and Operation, Appendix A (RSS 8559-2-1)

16. Space Shuttle Operational Flight Rules (JSC-12820, Mission Operations Directorate, Jan. 20, 1989)

17. SSME Operations (GSS SSME OPS), (JSC-19395)

18. SSME: The First Ten Years (Bob Biggs, Nov. 2, 1989)

19. Report of the SSME Assessment Team (Jan. 1993)

20. SSME FRR Technical Issues Database (Jan. 1, 1993)

21. SSME Test Configuration & Objective Summary (May 13, 1993)

22. Space Shuttle MPS Integrated Fault Tree Description (SAIC, Jan. 8, 1989)

23. Galileo RTG Risk Assessment Data Analysis, Final Report (SAIC, Sept. 30, 1988)

24. PRA of the Space Shuttle Phase 1:Space Shuttle Catastrophic Failure Frequency Final Report (J.Karns, SAIC,Aug.16, 1993)

25. SSME Weld Quality (Rockwell International, Mar. 1993)

26. Shuttle Integrated Risk Assessment Program, SSME Weld Risk and Quantified Data Development Program (SAIC, Sep. 1990)

27. Excerpts from Mission Safety Evaluation Reports from STS-28 Through STS-32

28. Space Shuttle Main Propulsion Pressurization System Probabilistic Risk Assessment, Final Report (Lockheed, Feb. 1988)

29. STS PRACA DATA: Problem Reporting and Corrective Action One-Line Descriptions (2 Books, NASA/MSFC)

30. SSME FMEA-CIL (Rocketdyne)

31. SSME Integrated Hazard Analysis (RSS-8545-20 Rev. G, Rocketdyne, Nov. 1991)

32. SSME Premature Cutoff Database (B.Biggs, Rocketdyne, September 1994)

33. SSME Redline Sensor Reliability Analysis Spreadsheet (B.Biggs, Rocketdyne, September 1994)

34. ACTS Database, SSME Operational Duration Data (MSFC, March 1993)

35. MPS Pathfinder Program C.A.R.s Set 2 (Oct. 31, 1990)

36. MPS Pathfinder Program FMEA/CILs

37. MPS Components Manual (Rockwell International)

38. Propellant Management System Selected CILs

39. Propellant Management System Component Database

40. Shuttle Integrated Risk Assessment: MPS-PMS Diagraph Modules

41. Shuttle Integrated Risk Assessment Program (SIRA): Main Propulsion System Propellant Management System

42. Reliability Failure Data (SAIC, April 19, 1991)

43. MPS BlueBook Vol II Part I & Vol III Part I

44. Preliminary SSME Integrated Hazards Analysis Volume I & II (Rockwell International)

45. Studies and Analysis of the SSME: SSME Failure Data Review (Battelle, Dec. 15, 1986)

46. Booster Systems Briefs (NASA/JSC, October 1, 1984)

47. Shuttle Flight Operations Manual, Volume 8A Main Propulsion Systems (NASA/JSC, October 1979)

48. Solid Rocket Boosters (NSTS Reference, June 1992)

49. Solid Rocket Motor Redesign (NSTS Reference, June 1992)

50. SRB Fact Book (MSFC SA 44-80-02, December 1984)

51. Techniques for Reliability Assessment & Growth in Large Solid Rocket Motor Development (D.K. Lloyd, AIAA Paper No. 67-435, July 1971)

52. Forecasting Reliability Growth (D.K. Lloyd, John Wiley & Sons, June 1990)

53. Trend Analysis for Large SRMs - a Program Level Approach (N.E. Babbitt III, AIAA 92-3357, July 1992)

54. Probabilistic Assessment for the Design of High Reliability Objects Such as Solid Propellant Grains (S. Nahon et. al., AIAA 92-3359, July 1992)

55. Elimination of Process-Induced Failure Modes as a Source of SRM Unreliability (B. E. Suta, R. J. Kunz, AIAA 90-2711, July 1990)

56. Space Shuttle Probabilistic Risk Assessment Proof-of-Concept Study, Vol. III, Auxiliary Power Unit and Hydraulic Power Unit Analysis Report (December 18, 1987)

57. Nonelectric Parts Reliability Data (ITT Research Institute, 1985)

58. Nonelectric Parts Reliability Data (Reliability Analysis Center, 1991)

59. Safety of the Thermal Protection System of the Space Shuttle Orbiter (Stanford University; Carnegie-Mellon University)

60. Reaction Control System Training Manual (RCS2102A, TD 340A, November 1992)

61. Orbital Maneuvering System, Orbiter Systems Training Manual (OMS 2102, JSC-19950, March 1984)

62. National Space Transportation System Reference, Volume 1- Systems and Facilities (NASA, June 1988)

63. Electrical Power System 2102 Training Manual (EPS 2102, JSC, February 1994)

64. IEEE Std 500 Reliability Data (IEE Std 500-1984)

65. Reactor Safety Study: An Assessment of Accident Risks in U.S. Commercial Nuclear Power Plants (WASH-1400:NUREG-75/014, United States Nuclear Regulatory Commision, October 1975)

66. Design Data Book for Space Shuttle Redesigned Solid Rocket Motor (TWR-16881, Thiokol, November 1989)