Software Engineering/Safety

SAFEWARE

SYSTEM SAFETY AND COMPUTERS

NANCY G. LEVESON

We are building systems today — and using computers to control them — that have the potential for large-scale destruction of life and environment. More than ever, software engineers and system developers, as well as their managers, must understand the issues and develop the skills needed to anticipate and prevent accidents before they occur. Professionals should not require a catastrophe to happen before taking action.

Addressing this need in her long awaited book, Nancy Leveson examines what is currently known about building safe electromechanical systems and looks at past accidents to see what practical lessons can be applied to new computer-controlled systems.

Safeware:

- Demonstrates the importance of integrating software safety efforts with system safety engineering
- Describes models of accidents and human error that underlie particular approaches to safety problems
- Presents the elements of a safeware program, including management, hazard analysis, requirements analysis, design for safety, design of the human-machine interface, and verification.

About the Author:

Nancy Leveson is Boeing Professor of Computer Science and Engineering at the University of Washington (and Adjunct Professor at the University of British Columbia). Dr. Leveson recently was awarded the Information System Award from the American Institute of Aeronautics and Astronautics, "For developing the field of software safety and for promoting responsible software and system engineering practices where life and liberty are at stake." She is the Editor-in-Chief of *IEEE Transactions on Software Engineering* and a member of the Board of Directors of the Computing Research Association, the National Research Council Commission on Engineering and Technical Systems, and the ACM Committee on Computers and Public Policy. She recently chaired a National Academy of Science study for NASA of the Space Shuttle software development process.

90000
9 780201 119725
ISBN 0-201-11972-2
ADDISON-WESLEY PUBLISHING COMPANY

SAFEWARE

LEVESON

ADDISON WESLEY

SAFEWARE

SYSTEM SAFETY AND COMPUTERS

NANCY G. LEVESON

A GUIDE TO PREVENTING ACCIDENTS AND LOSSES CAUSED BY TECHNOLOGY

Chapter

13

Hazard Analysis

The argument that the same risk was flown before without failure is often accepted as an argument for the safety of accepting it again. Because of this, obvious weaknesses are accepted again and again, sometimes without a sufficiently serious attempt to remedy them, or to delay a flight because of their continued presence.

—Richard Feynman
Personal Observations on Reliability of Shuttle[1]

Unfortunately, everyone had forgotten why the manuals came off the top of the mast and nobody realized that this was important.

Trevor Kletz
What Went Wrong?

Hazard analysis is at the heart of any effective safety program, providing visibility and coordination (see Figure 13.1). Information flows both to and from and back into the hazard analysis. The outward information, for example, helps designers perform trade studies and eliminate or mitigate hazards and helps quality assurance identify quality categories, acceptance tests, required inspections, and components that need special care. Any changes or additional information must flow back through the analyses so that (1) solutions or corrections can be integrated into the design and (2) the overall conceptual system model and approach to safety is maintained.

Although hazard analysis alone cannot ensure safety, it is a necessary first

[1] Appendix F of the Rogers Commission report (1986).

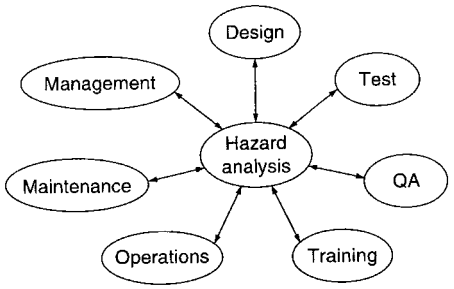

FIGURE 13.1
Hazard analysis provides visibility and coordination.

step before hazards can be eliminated or controlled through design or operational procedures. Simply knowing that a hazard exists may provide sufficient information to eliminate or control it, even without in-depth analyses of its causes. Often, general safeguards can be provided even if little is known about the hazard's precursors. A larger number of options for elimination and control usually exist, however, if more is known about the hazard and the conditions and events leading to it. Condition-specific or event-specific safeguards are frequently more effective and less costly in terms of the tradeoffs required.

Hazard analysis is not just performed at the beginning of a project or at fixed stages, but should be continuous throughout the life of the system, with increasing depth and extent as more information is obtained about the system design. As the project progresses, the uses of hazard analysis will vary—such as identifying hazards, testing basic assumptions and various scenarios about the system operation, specifying operational and maintenance tasks, planning training programs, evaluating potential changes, and evaluating the assumptions and foundations of the models as the system is used and feedback is obtained. In an operational system, analyses and their assumptions act as preconditions for safe operation and as constraints on management and operational procedures.

Just as the purpose of and the information available for hazard analyses change with time, so do the analysis requirements and the appropriate analysis methods. There is no one perfect method for every goal, but many that can and should be used. The various techniques provide formalisms for systematizing knowledge, draw attention to gaps in knowledge, help prevent important considerations from being missed, and aid in reasoning about systems and determining where improvements are most likely to be effective.

For many hazards and systems, analysis may consist merely of comparing

the design with various standards and codes that have been developed over time to deal with known hazards. Standards and codes provide a means of encoding historical information derived from accidents and near accidents so that safety reviews and analyses need not start from scratch each time. They help us encapsulate and learn from experience.

Nevertheless, as new technology is developed and systems are scaled up in size, new hazards arise and the possibility of introducing hazards increases. Systems with new and complex designs may require sophisticated, formal, documented analytical procedures.

In some countries or industries, the particular hazards to be analyzed and the depth and types of analysis may be established by regulatory authorities or by legislation. In others, these decisions must be made by the safety engineers early in project development. This chapter presents the information used in making such decisions.

13.1 The Hazard Analysis Process

The process chosen depends on the goals or purpose of the hazard analysis. Different goals will require very different processes. Once the goals are defined, the steps to be taken can be determined.

13.1.1 Goals of Hazard Analysis

The goals of safety analysis are related to three general tasks [323, 327]:

1. **Development**: the examination of a new system to identify and assess potential hazards and eliminate or control them.
2. **Operational management**: the examination of an existing system to identify and assess hazards in order to improve the level of safety, to formulate a safety management policy, to train personnel, and to increase motivation for efficiency and safety of operation.
3. **Certification**: the examination of a planned or existing system to demonstrate its level of safety and to be accepted by the authorities or the public.

The first two tasks have a common goal of *making the system safer*, while the third has the goal of convincing management or government that an existing design or system *is safe*. These three broad tasks can be divided into subtasks:

- **Development** and **operational management**
  1. Identify hazards that singly or in combination could cause an accident.
  2. Show that specific hazards are not present and that safeguards are not needed.
  3. Determine the possible damaging effects resulting from system hazards.

4. Evaluate the causal factors related to the hazards:
   (a) Determine how hazards could occur, their nature, and their possible consequences.
   (b) Examine the interrelationships among causal factors.
5. Identify safety design criteria, safety devices, or procedures that will eliminate or minimize and control the identified hazards.
6. Find ways to avoid or eliminate specific hazards.
7. Determine how to control hazards that cannot be eliminated and how to incorporate these controls into the design.
8. Evaluate the adequacy of the hazard controls.
9. Provide information for quality assurance—quality categories, required acceptance tests and inspections, and items needing special care.
10. Evaluate planned modifications.
11. Investigate accidents and near-miss reports. Determine whether they have validity and, if they do, the cause of the problem.

- **Certification**
  1. Demonstrate the level of safety achieved by the design.
  2. Evaluate the threat to society from the hazards that cannot be eliminated or avoided.

Each of these subtasks implies very different types of analysis. The goals and subtasks need to be clarified and documented at the start of any hazard analysis process.

13.1.2 Qualitative versus Quantitative Analyses

Safety analysis ranges from relatively simple qualitative methods to advanced quantitative methods in which numerical values for risk are derived. Even if quantitative methods are used, qualitative analyses must precede them—hazards and their causal factors must be identified before numerical values can be assigned to them. Thus, the quality of the quantitative analysis depends on how good the qualitative one was.

Often, the analysis can stop with the qualitative aspects. Knowing the causal factors and using best-guess estimates of relative hazard rankings are adequate for most purposes during the system development phase, when few accurate numerical values are available anyway. The most effective and accurate quantitative analyses use operational data from a working system to determine achieved reliability and to identify trends and changes that might affect safety. Quantitative analysis may also be useful in comparing the reliability of alternative features of an emerging design that uses standard parts with established failure rates.

The use of probabilistic risk assessment has been debated widely. The arguments in its favor are usually based on the technique's ability to provide input for decision making. Identifying hazards alone does not determine how funds should be allocated to reduce hazards. Comparative probability data can be useful in such

decisions. Decision making in the certification of plant designs is also eased by the use of such numbers.

However, many important factors—design deficiencies, for example—cannot be easily or reasonably quantified. Because equipment failures are the most easily assessed probabilistically, many quantitative risk assessments have been criticized for placing more emphasis on these failures than on less easily predicted and quantified factors such as design errors, construction deficiencies, operator actions, maintenance errors, and management deficiencies. For example, some probabilistic assessments have emphasized failure probabilities of devices that are in the range of 10^{-7} or 10^{-8} while ignoring installation errors or maintenance errors of those same devices with probabilities in the range of 10^{-2} or 10^{-3}. Quantifying only what can easily be quantified does not provide a realistic estimate of risk. In the space program, where probabilistic risk assessment based on fault tree analysis and failure modes and effects analysis was used extensively, almost 35 percent of the actual in-flight malfunctions were not identified by the method as "credible" [205].

Care also has to be taken that quantification does not divert attention away from risk reduction measures. The danger exists that system safety analysts and managers will become so enamored with the statistics that simpler and more meaningful engineering processes are ignored.

Follensbee cites several recent commercial aircraft accidents caused by events that had been calculated to have a probability of 10^{-9} or less (one was calculated as 10^{-12}) using accepted techniques. In all of these cases, incorrect assumptions about the behavior of the pilots or the equipment led to underestimated failure or risk figures [86]. In several cases, the need for compliance with standard aircraft fail-safe design standards—which might have prevented these accidents—was judged unnecessary based upon the calculations [86].

The Union of Concerned Scientists and others have warned against a disproportionate emphasis on meeting predicted quantitative levels of safety rather than on consideration of technical problems and implementation of engineering fixes [3]. The same may be true for human error. Hornick, past president of the Human Factors Society, has suggested that "the general nuclear power community is couching a cavalier attitude towards human factors in the (false?) comfort of risk-assessment statistics" [127, p.114].

Quantitative risk assessment of a completed design sometimes is required by certification agencies or used in public arguments about the safety of a controversial technology such as nuclear power. The emphasis of this chapter (and this book), however, is on qualitative rather than quantitative hazard analysis.

13.1.3 Role and Qualifications of the Analyst

In general, the role of the safety analyst should be to generate alternatives as well as to eliminate them. A U.S. Air Force acquisition handbook suggests that pointing out new alternatives may be much more valuable to a program than exhaustive

analysis of given approaches [5]. The early selection of a single approach can be detrimental to the overall safety program when there are less expensive alternatives that yield equal or better safety or there are alternatives that require fewer tradeoffs with other system goals. Therefore, the job of the system safety analyst is to dig deeply into the design and suggest designs that will yield eventual system operation that is satisfactory from both a safety and performance standpoint.

Successful hazard analysis requires an understanding of the system under consideration. Hazard analysis does not remove the need for engineering expertise and judgment: Standardized analysis approaches and terminology simply help to clarify the problems in order to provide a context for expert judgment and to enhance interdisciplinary communication. The process usually requires a team of people with a wide variety of knowledge and skill. System safety engineers provide expertise in safety analysis, while other team members bring expertise in specific engineering disciplines along with alternative viewpoints. The approaches to solving a problem suggested by system engineers, subsystem engineers, reliability engineers, safety engineers, application experts, operators, and management may be entirely different, but all contribute significantly to finding a satisfactory solution [5]. When humans and computers play important roles in system operation, operators and software engineers should also be involved in the hazard analysis and in the design of hazard reduction measures.

13.1.4 General Features of an Effective Hazard Analysis Process

The hazard analysis process is both *continual* and *iterative* (see Figure 13.2). Hazard identification and analysis begin at the conceptual stage of the project and continue through decommissioning. Starting early is imperative if safety considerations are to be incorporated into trade studies and early design decisions, when hazards can be most effectively and cheaply handled. Planning for hazard elimination and control should begin as soon as safety problems are uncovered—preferably before unsafe features become firmly embedded in the design.

The forms of analysis will be different as the system matures, but all are part of a single analysis process. Each stage of analysis acts as a baseline upon which later steps build. The stages reflect the quality of information available, with analysis depth and breadth increasing as more information is obtained. In the early project phases, hazard resolution may involve simply getting more information about the hazard or generating alternative design solutions. As the project progresses and the design is elaborated, more detailed analyses and tests may uncover new hazards or eliminate old ones from consideration.

If a hazard cannot be resolved at a particular stage, follow-up evaluation and review may be necessary. Organizational controls, such as audit trails and tracking systems, must be installed to ensure that this follow-up occurs.

The operational safety achieved depends on the accuracy of the assumptions and models underlying the design process. The system must be monitored to ensure (1) that it is constructed, operated, and maintained in the manner assumed

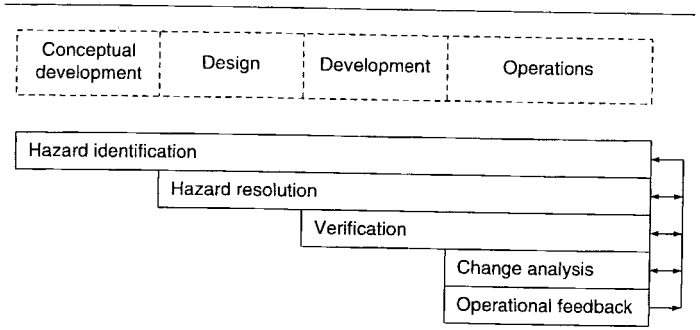

FIGURE 13.2
The hazard analysis process is continual and iterative.

by the designers; (2) that the models and assumptions used during initial decision making and design were correct; and (3) that the models and assumptions are not violated by changes in the system, such as workarounds or unauthorized changes in procedures, or by changes in its environment. Operational feedback on trends, incidents, and accidents should trigger reanalysis.

If a change is proposed to a baseline or completed design, or an unplanned change is detected during operation, it must be analyzed for its potential effect on safety. The process involves reviewing previously generated analyses to identify the impact of the proposed change, updating the analyses and documentation to reflect the changes, and identifying new hazards and hazard causes. The reanalysis must start at the highest level of system design at which the change becomes visible and show that the change does not create a new hazard, affect a hazard that has already been resolved, or increase the severity level of a currently existing hazard.

13.1.5 Steps in the Process

A hazard analysis consists of the following steps:

1. Definition of objectives.
2. Definition of scope.
3. Definition and description of the system, system boundaries, and information to be used in the analysis.
4. Identification of hazards.
5. Collection of data (such as historical data, related standards and codes of practice, scientific tests and experimental results).

6. Qualitative ranking of hazards based on their potential effects (immediate or protracted) and perhaps their likelihood (qualitative or quantitative).
7. Identification of causal factors.
8. Identification of preventive or corrective measures and general design criteria and controls.
9. Evaluation of preventive or corrective measures, including estimates of cost. Relative cost rankings may be adequate.
10. Verification that controls have been implemented correctly and are effective.
11. Quantification of selected, unresolved hazards, including probability of occurrence, economic impact, potential losses, and costs of preventive or corrective measures.
12. Quantification of residual risk.
13. Feedback and evaluation of operational experience.

Each step requires documentation of the results and of any underlying assumptions and models. The purpose of the hazard analysis process is to use the results as a reference for judgment [274] in design, maintenance, and management decisions. Accomplishing this goal requires an explicit formulation of the models, premises, and assumptions underlying the safety analysis and the design features used to eliminate or control hazards. There must also be an explicit description of the assumptions about work procedures and organizational structures that will constrain management's freedom of action.

Not all of the steps may need to be performed for every system and for every hazard. For standard designs with well-established risk mitigation features (perhaps included in standards and codes), only the first 5 steps (plus step 13) may be needed. For new designs, usually the first 10 steps are necessary. Step 11 may be needed when risk control measures require tradeoffs with critical functions or constraints (such as weight or space), or hazards are expensive to resolve. Step 12 may be required for certification or licensing of systems by government agencies when the potential consequences of a hazard are catastrophic. Completely eliminating or controlling hazards may not be possible or practical due to political factors, lack of time, potential cost, or the magnitude or nature of the hazard; quantification of the residual risk provides an estimate of the risk assumed in operating the completed system. Step 13 always needs to be done.

Because, as has been widely noted, what can go wrong probably will, all foreseeable uses and modes of the system over its entire lifetime need to be examined. At various times in its operational life, a system will be exposed to different environments, processes, conditions, and loads, and the effects will differ according to when the stress or condition occurs.

Hazard analysis can be divided into three basic functions: (1) identifying hazards, (2) identifying and evaluating the hazard causal factors, and (3) evaluating risk. Although names such as *Preliminary Hazard Analysis* and *Subsystem Hazard Analysis* are often applied to these functions, these are merely phases, as noted earlier, in a continuing and iterative process rather than disjoint sets of analyses.

13.1.6 Hazard Identification

Hazard identification starts early in the concept formation stages of the project, and hazard lists are continually updated with new hazards and information about previously identified hazards throughout the entire lifetime of the system.

Hazard identification in the earliest stages of a program, often called Preliminary Hazard Analysis (PHA), involves

1. Determining what hazards might exist during operation of the system and their relative magnitude.
2. Developing guidelines, specifications, and criteria to be followed in system design.
3. Initiating actions for the control of particular hazards.
4. Identifying management and technical responsibilities for action and risk acceptance and assuring that effective control is exercised over the hazards.
5. Determining the magnitude and complexity of the safety problems in the program (how much management and engineering attention is required to minimize and control hazards).

The output of the hazard identification process is used in developing system safety requirements, preparing performance and design specifications, test planning, preparing operational instructions, and management planning. The results serve as a framework or baseline for later analyses and as a checklist to ensure that management and technical responsibilities for safety tasks are carried out.

The first step is to identify the system, system boundaries, and the limits of resolution, as described in earlier chapters. If multiple subsystems or systems are involved, the boundaries must be consistent. For simple or well-understood products or systems, the hazards may already be known and the analyst can skip to the next step.

For some special systems regulated by government agencies, hazards or hazard categories may be mandated. For example, the U.S. Department of Defense (DoD) identifies four hazards (see Section 12.2.2) that must be considered when constructing nuclear weapon systems. Special processes, such as Nuclear Safety Cross Check Analysis (NSCCA) and Software Nuclear Safety Analysis (SNSA) may be required on such programs to ensure that the software cannot contribute to any of these four hazards. NSCCA and SNSA are not specific analysis techniques but rather rigorous independent verification and validation procedures; they are described briefly in Chapter 18.

In most systems, however, the hazards are not immediately known or subject to government mandate and need to be determined. Childs warns that the analyst should not just list every conceivable hazard, including those without any clear-cut relationship to the system being studied [51].

A few techniques for identifying hazards have been developed. They are essentially structured ways to stimulate a group of people to apply their personal knowledge to the task—for example, by raising a series of *what-if* questions guided by a model of the system.

Most hazard identification involves less structured procedures, but various activities can be helpful:

- Review pertinent historical safety experience, lessons learned, trouble reports, and accident and incident files.
- Use published lists and checklists of hazards (Table 13.1 is an example from [232]). Investigate standards and codes of practice; these often reflect known hazards that have caused accidents in the past.
- Examine the basic energy sources, energy flows, and high-energy items in the system and provisions for their control.
- Consider hazardous materials such as fuels, propellants, lasers, explosives, toxic substances, and pressure systems.
- Look at potential interface problems such as material incompatibilities, possibilities for inadvertent activation, contamination, and adverse environmental scenarios.
- Examine hazard analyses on previous systems.
- Review the mission and basic performance requirements, including the environments in which operations will take place. Look at all possible system uses, all modes of operation, all possible environments, and all times during operation. Stresses and effects of particular conditions will differ as a system is exposed to varying environments, processes, conditions, and loads.
- Take advantage of the general engineering experience and personal experience of safety and application experts. *Tiger team attacks*, which are essentially brainstorming sessions, are a standard security technique that might be useful for safety.
- Examine the human–machine interface and the interaction between the operators and the automated equipment.
- Look particularly at transition phases—changes in the system, changes in the technical and social environment, and changes between modes of operation in the system. Accidents often occur during non-routine operating modes: startup, restart, shutdown, testing, trials of new methods, breakdown, maintenance, repair, inspection, troubleshooting, modifications, changeovers, adjacent system change, nonstandard input, stresses (including budget, schedule, delays, and catch-up), and adverse conditions.
- Use scientific investigation of physical, chemical, and other properties of the system, which may involve theoretical studies and small-scale tests.
- Think through the entire process, step by step, anticipating what might go wrong, how to prepare for it, and what to do if the worst happens.

As hazards are identified, information about them needs to be recorded.

TABLE 13.1
Generic hazards for the Space Shuttle.

Generic hazards	Generic hazard type	
Contamination/corrosion	Chemical disassociation	Oxidation
	Chemical replacement/combination	Organic (fungus/bacterial, etc.)
	Moisture	Particulate
Electrical discharge/shock	External shock	Corona
	Internal shock	Short
	Static discharge	
Environment/weather	Fog	Radiation
	Fungus/bacterial	Sand/dust
	Lightning	Vacuum
	Precipitation (fog, rain, snow, sleet, hail)	Wind
		Temperature extremes
Fire/explosion	Chemical change (exothermic/endothermic)	Pressure release/implosion
	Fuel and oxydizer in presence of fuel and ignition source	High heat source
Impact/collision	Acceleration (including gravity)	Meteoroids/meteorites
	Detached equipment	Moving/rotating equipment
	Mechanical shock/vibration/acoustical	
Loss of habitable environment	Contamination	Toxicity
	High pressure	Low temperature
	Low oxygen pressure	High temperature
	Low pressure	
Pathological/physiological/psychological	Acceleration/shock/impact/vibration	Sharp edges
		Lack of sleep
	Atmospheric pressure (high, low, rapid change)	Visibility (glare, window/helmet fogging)
	Humidity	Temperature
	Illness	Excessive workload
	Noise	
Radiation	Electromagnetic	Thermal/infrared
	Ionizing	Ultraviolet
Temperature extremes	High	
	Low	
	Variations	

Often tabular forms are used. The most effective and efficient way to record hazard information is to use one form, filling it in as the analysis progresses through the various stages. At the beginning, only parts of the form will be completed, but by the end of the project development, all of the information should be available. An example of such a form is shown in Figure 12.11.

The form may include some or all of the following information:

- System, subsystem, unit (equipment grouping where the potential hazard exists).
- Hazard description.
- Hazard cause.
- Possible effects on the system and the environment.
- Category (hazard level).
- Corrective or preventive measures (compensation and control), possible safeguards, recommended action, and design criteria.
- Operational phase when hazardous.
- Organizations responsible for ensuring that safeguards are provided for the specific hazard.
- Verification methods (tests, demonstrations, analysis, inspection) to verify that the hazard is effectively controlled.
- Other proposed and necessary actions.
- Remarks and status of the hazard resolution process. The hazard is closed when it has been verified that the recommended actions have been implemented and are effective.

Hazard Level. The *hazard category* or *level* is defined by likelihood and severity and often specified in the form of a matrix to aid in prioritization. Example matrices are shown in Figures 12.3 and 12.5. Since the depth of the analysis usually depends upon the severity of the hazard, the worst-case consequences must be determined early. The evaluation of susceptibility to a hazard should consider duration and exposure—how many people are exposed for how long. The warning time (the interval between identification of the problem and the occurrence of injury or damage) may also be important.

Hazard severity categories reflect worst possible consequences. The categories used are specific to the industry and sometimes the system. For example, the DoD (MIL-STD-882B: System Safety Program Requirements) uses these categories:

Category I: Catastrophic; may cause death or system loss.

Category II: Critical; may cause severe injury, severe occupational illness, or major system damage.

Category III: Marginal; may cause minor injury, minor occupational illness, or minor system damage.

Category IV: Negligible; will not result in injury, occupational illness, or system damage.

In contrast, a NASA document (NHB 5300.4) lists NASA hazard categories as:

Category 1: Loss of life or vehicle (includes loss or injury to the public).

Category 2: Loss of mission (includes postlaunch abort and launch delay sufficient to cause mission scrub).

Category 3: All others.

As a final example, a Department of Energy standard (DOE 5481.1) for nuclear systems defines three categories of hazard severity:

High: Hazards with potential for major onsite or offsite impacts to people or the environment.

Moderate: Hazards that present considerable potential onsite impacts to people or environment but at most only minor offsite impacts.

Low: Hazards that present minor onsite and negligible offsite impacts to people or the environment.

The other component of hazard level, likelihood, is commonly divided into discrete categories, such as

Frequent: Likely to occur frequently to an individual item, continuously experienced throughout the fleet or inventory.

Probable: Will occur several times during the life of an individual item, frequently throughout the fleet or inventory.

Occasional: Likely to occur sometime during the life of an individual item, several times throughout the fleet or inventory.

Remote: Unlikely to occur but possible during the life of an individual item; unlikely but reasonably expected to occur in a fleet or inventory.

Improbable: Extremely unlikely to occur to an individual item; possible for a fleet or inventory.

Physically Impossible: Cannot occur to an item or in a fleet or inventory.

Quantitative probability assessment, if used, is stated in terms of likelihood of occurrence of the hazard per unit of time, events, population, items, or activity, such as 10^{-7} per year.

Design Criteria. These broad concepts state *what* has to be achieved, leaving the designer free to use ingenuity in deciding *how* the goal may best be achieved. A design criterion for a collision avoidance system, for example, is that maneuvers should be avoided that require the objects to cross paths. A typical criterion for a pressure system is that all pressure tanks have a relief valve of sufficient size to reduce the tank pressure when the pressure exceeds a specific amount above

TABLE 13.2
The relationship between hazards and design criteria for the doors in a rapid transit system.

Hazard	Design criterion
Train starts with door open.	Train must not be capable of moving with any door open.
Door opens while train is in motion.	Doors must remain closed while train is in motion.
Door opens while improperly aligned with station platform.	Door must be capable of opening only after train is stopped and properly aligned with platform unless emergency exists (see below).
Door closes while someone is in doorway.	Door areas must be clear before door closing begins.
Door that closes on an obstruction does not reopen, or reopened door does not reclose.	An obstructed door must reopen to permit removal of obstruction and then automatically reclose.
Doors cannot be opened for emergency evacuation.	Means must be provided to open doors for emergency evacuation when the train is stopped anywhere.

Source: Adapted from Willie Hammer. *Product Safety Management and Engineering*. Prentice-Hall, Inc., Englewood Cliffs, N.J., 1980, page 207.

normal operating pressure. Table 13.2 shows the relationship between system hazards and system design criteria for the door control in a rapid transit system [108].

Design criteria are not requirements (which are much more specific in nature and apply to a particular design), but rather are used to derive the design requirements. Design criteria are general in nature and can be applied to many different systems and designs.

Operational Phase. Whether an accident actually occurs as a result of a hazard and the severity of the hazard may depend upon the operational phase—the system and environmental conditions—in which the hazard occurs. Therefore, the phase must also be documented for each hazard. Failure of a missile launch system during launch might simply leave the missile sitting on the pad, whereas failure right after launch might result in a fallback and total destruction. Similarly, loss of control of a missile immediately after liftoff might result in more damage (since the missile is filled with propellant) than loss of control far down range [106].

13.1.7 Hazard Causal Analysis

After the hazards are identified, the next step is to determine the causes and effects associated with each hazardous condition. The hazards can be analyzed in greater detail as specific system design features are elaborated during development. Information about the causes of hazards is helpful in specifying the safety requirements and design constraints needed to minimize or control hazards to an acceptable level.

Each hazard can have several potential causes, and each cause can have several potential consequences or effects. The factors considered and the process of identifying them will depend on the underlying accident model used (see Chapter 10), either consciously or subconsciously, by the analyst.

When examining complex systems, the analyst may need to consider conditions or events involving several components or the interfaces among components. For such systems, a model of some type and an analysis method defined for it are often useful in tracing the effects of conditions or events either backward or forward. Sensitivity analysis may be used to determine the most important factors.

Causal analysis can be divided into system (whole) and subsystem (unit or part) analyses. These two types of analysis are merely different parts of a total process and different ways of looking at the system. Both are usually performed, since each contributes important information.

System Hazard Analysis. This analysis considers the system as a whole and identifies how the system operation, the interfaces between the system components, and the interfaces between the system and its operators can contribute to hazards. As stated earlier, segmenting the system operation into sequences of events and actions (operational phases) by considering the mission and goals is often useful. The time at which a condition is critical varies with the system configuration, its location, and the time that an event occurs.

Subsystem Hazard Analysis. This type of analysis looks at individual subsystems and determines the effect of their operating or failure modes on system hazards. Typical subsystems include power, structural, control, sensor, operator, communications, propulsion, and environmental control. This type of analysis identifies the impact on overall safety of component failure modes, critical erroneous human inputs, and operating or failure modes of the subsystem related to performance, operational degradation, functional failure, unintended function, and inadvertent function (proper function but at the wrong time or in the wrong order).

Software Hazard Analysis. Software is just like any other component, especially various types of control or monitoring components, and thus will be included in the system hazard analysis as well as the focus of a subsystem hazard analysis (sometimes called *Software Hazard Analysis*). The computer's behavior

and its interfaces to the rest of the system must be evaluated for potential contribution to system hazards. When such critical behavior is identified, it can be traced into the software design and code to identify parts of the software that require special design features or that need to be analyzed in depth.

Later, the software is evaluated to determine if the system design criteria have been satisfied in the software specifications and if the software implementation has impaired or degraded system safety or introduced new hazards. Because of the difficulty of these tasks, normally each part of the software development process is evaluated rather than just waiting until the software is complete.

All software, including commercial off-the-shelf software (COTS), needs to be analyzed to the level necessary to determine any impact or influence on the identified system hazards. In addition, changes to the software must be evaluated for potential impact on hazards.

Operational Hazard Analysis. A separate analysis of the risk controls that have been assigned to operational procedures, called *Operational Hazard Analysis*, may also be performed. Often this analysis is left until the design is complete. The process, however, is no different than the other analyses, and there is no good reason to separate it or to leave it until the end. Instead, the operational use of the system should be considered from the very beginning (like the other analyses) and should be an integral part of the system and subsystem analyses. Otherwise, hazardous conditions may be left for procedural control that are impossible, unwise, or could have been handled more easily and safely in the design.

In addition, the information uncovered in the analysis of operator procedures should be used in the human–machine interface design and thus needs to be available early. This approach implies a very different design process than is often applied: The operational procedures are defined in concert with the automated systems and do not consist merely of functions that the designers did not know how to automate.

Operational analysis, like the other subsystem analyses, looks at all the ways that the operators can contribute to system hazards if they follow defined procedures (the system operates as it was designed to operate) or if they do not. Typically, operational procedures are divided into phases or tasks (as is done for the other analyses), which are analyzed in detail for their potential contribution to hazards, including proper and improper sequencing of actions.

13.1.8 Risk Assessment and Acceptance Analysis

Once design and development are complete, the system design can be evaluated. The goal here is not to guide the design process but to evaluate the final product. The results may be used internally or for independent certification to determine the residual risk and whether the system is acceptable for use.

The acceptance analysis should include more than just estimates of the probability and consequences of hazards and accidents. Each hazard should be documented to show its potential causes, the implemented controls and tracing of the hazard into the detailed design, and the results of the verification efforts. Basically, this assessment is a description of the potential problem and what is being done about it. The documentation should allow the assessor or reviewer to follow the analyst's reasoning in order to check the correctness of the results [299]. For very critical systems, the informal reasoning may be augmented by formal (and perhaps mathematical) reasoning and argument. The assumptions underlying any formal models or methods used need to be verified.

The acceptance analysis may also include a probabilistic risk assessment, in which the probabilities and severity of accidents are evaluated in addition to the hazard level. A risk assessment alone is not adequate, but it can augment the other information provided. The accuracy of probabilistic risk assessment is quite controversial, however. Briefly, the major limitations lie in the inaccuracies of the available failure and human error data, in the consequence models (such as dispersion of heavy gas and explosions), and in the toxicity data. Some studies have shown widely varying results for parallel assessments performed on the same system by different analysts or analysis teams. The variations seem to be caused by different initial assumptions, different restrictions of the object under study, the particular failure data used, and different analyses of the hazard consequences [299].

Typically, worst-case effects of hazards are all that is needed for hazard analysis and assessment, and these are relatively easily determined. Risk assessment, however, requires elaborate quantitative analyses. Unfortunately, obtaining probabilistic data about the harmful consequences of some hazards and about human errors is quite difficult and perhaps the most error-prone part of any quantitative risk assessment.

Assessing Harmful Consequences. Probabilistic risk assessments have been performed primarily on chemical or radiation hazards. The main physical effects of plant or equipment failure arise from escaping gases or liquids catching fire or exploding [126], or from toxicity. To assess the magnitude of an accident, the analyst first needs to determine (1) how much material is likely to escape, (2) what is likely to happen to it over time and distance (the physical consequences), and (3) the effect on people. Chemical engineering models exist to calculate some of these factors, but the accuracy of the models for some factors, such as gas dispersion, is poor. Considerable simplifying assumptions are needed to make the models manageable since the effects of a release depend on the physical state of the released material, its release rate, natural topography, intervening structure, atmospheric conditions, homogeneity of the gas cloud, ignition sources, and so on. The translation of structural damage into human casualties "is so speculative that in practice it can be no more than a statistical assumption, which in any given case may be orders of magnitude wrong" [126].

The accuracy of consequence assessment not involving release of material, such as the probability that two planes violating minimum separation standards

actually collide, depends on the hazard and system involved. Usually, the number of uncontrolled variables makes prediction very difficult.

Assessing Human Error. A second controversial aspect of probabilistic risk assessment is the accuracy of human error data. Most systems do not provide enough data on human error to be useful for probabilistic modeling. The alternatives are to use laboratory studies or to use numbers collected over a long time and over many types of systems.

The difficulty in extrapolating from laboratory studies stems from the significant differences between the laboratory and industrial settings [269]. Laboratory tasks tend to have (1) a well-defined goal, (2) stable requirements, (3) specific instructions, (4) artificial and low-valued payoffs, and (5) the subject controlled by the task. In contrast,

> In "real" tasks only a (sometimes vague) overall performance criterion is given and the detailed goal structure must be inferred by the operator. . . . The task may vary as the demands of the system vary in real time. Operating conditions and the system itself are liable to change. Costs and benefits may have enormous values. There is a hierarchy of performance goals. The operator is usually highly trained, and largely controls the task, being allowed to use what strategies he will. Risk is occurred in ways which can never be simulated in the laboratory [269].

The other alternative, using historical data from a large number of systems, has two main drawbacks: (1) the data and tasks from one system may not apply to a different system, and (2) data collection is often biased, incomplete, or inaccurate. Case studies and incidents do not represent all the errors that operators make, but merely those that are reported. Unreported errors tend to be those that the operator is able to correct before damage occurs. Monitoring is one way to collect data, but human behavior may be abnormal if the person being monitored is aware of the monitoring.

Human error data also suffers from the difficulty of classifying errors, such as determining whether an error was an operating error or a design error. The variety of classification schemes makes use of data in a different environment unreliable.

Given the discussion of the relationship between human behavior, psychological mechanisms, and task characteristics in Chapter 10, it seems misleading to collect empirical data about human errors without also noting subtle differences in the environment and system design when those errors occurred. There can be wide variation in the environmental situations and physical aspects of the tasks, including stress factors such as noise, temperature, emotional stress, and vibration. For example, collecting probabilistic data about humans misreading a particular type of dial in poor lighting and then applying the "probability of misreading a dial" to systems with a different dial design and better lighting may be unjustified; noting in the database all conditions under which every error was made is probably impossible, and the number of instances is not large enough for

such differences to become unimportant. Errors on particular tasks may also depend on the other tasks the operator is performing. Thus, the context of the task is extremely important, further limiting the situations in which historical data applies.

Besides the problems in collecting it, human numerical error rate data also suffers from various other kinds of uncertainty. Human performance exhibits considerable natural variability based on skill, experience, and personal characteristics [189]. Not only do people differ in their innate capabilities, but the performance of any one individual will vary over time. Some of these variations are unpredictable, while others seem to be circadian. Performance variations over a 24-hour day arise from fluctuations in the work situation and also from modifications in human capabilities [276]. Historically, safety and productivity are low at night. The fact that we are a diurnal species may explain why many of the major industrial accidents involving human error have occurred at night [85]. Variability in performance may also arise from interactions with an unstable environment, from stress, and from interactions with other workers.

Specific methods for assessing human error rates and including them in hazard analysis and risk assessment have been proposed despite these problems. They are discussed in the next chapter.

13.2 Types of System Models

Every hazard analysis requires some type of model of the system, which may range from a fuzzy idea in the analyst's mind to a complex and carefully specified mathematical model. The model may also range from a high-level abstraction to a low-level and detailed prototype. Nevertheless, information about the system must exist in some form, and that constitutes the system model upon which the analysis is performed.

A model is a representation of a system that can be manipulated in order to obtain information about the system itself. Models can be categorized along different dimensions [50]:

- *Material* models (which represent a complex system by another physical system that is simpler yet similar in important respects) versus *symbolic* or *formal* models (which represent the structural properties of a system in terms of assertions or logical statements).
- *Dynamic* models (where the features of the model vary perceptibly with time) versus *static* models.
- *Stochastic* models (containing intrinsic probabilistic or random elements that affect the outcome or response of the model) versus *deterministic* models.
- *Iconic* models (those that pictorially or visually represent aspects of the system), *analog* models (those that employ one set of properties to represent