# Organizational Aspects of Engineering System Safety: The Case of Offshore Platforms

M. ELISABETH PATÉ-CORNELL

**Organizational errors are often at the root of failures of critical engineering systems. Yet, when searching for risk management strategies, engineers tend to focus on technical solutions, in part because of the way risks and failures are analyzed. Probabilistic risk analysis allows assessment of the safety of a complex system by relating its failure probability to the performance of its components and operators. In this article, some organizational aspects are introduced to this analysis in an effort to describe the link between the probability of component failures and relevant features of the organization. Probabilities are used to analyze occurrences of organizational errors and their effects on system safety. Coarse estimates of the benefits of certain organizational improvements can then be derived. For jacket-type offshore platforms, improving the design review can provide substantial reliability gains, and the corresponding expense is about two orders of magnitude below the cost of achieving the same result by adding steel to structures.**

THE CHALLENGER, CHERNOBYL, THREE MILE ISLAND, AND the *Exxon Valdez* accidents (among others) have shaken the public's confidence in the safety of technology and stimulated national and international inquiries about the very nature of such events. After each of them, the consensus was that something should be done to prevent a recurrence. Eliminating a technology that does not seem to be managed properly may be tempting, but often it is not even an option. If we decide to live with the risk, we should understand what went wrong so that we do not let the same failure happen again and we should understand what else could go wrong so that we prevent accidents. Corporations tend to blame human errors or technical mishaps for catastrophic failures of engineering systems and treat them as bad luck. Yet, in many cases, the root of the problem is in the organization, even if the eventual failure can be traced back to a specific component or operator (*1–3*). Accidents come basically in two forms: those that are either totally unpredictable or so rare that one can reasonably decide to live with the risk, and those that are essentially self-inflicted, often through management practices that are bound to generate errors and defects with a much higher probability than generally estimated. Even though the distinction is sometimes fuzzy, the former can be attributed to bad luck and little can be done about them, whereas the latter are the result of organizational factors that often can be improved. At the root of the Challenger accident, for example, was an accumulation of organizational problems (*4*) that included miscommunication of technical uncertainties, failure to use information from past near-misses, and an error of judgment in balancing conflicting requirements of safety and schedule. The National Aeronautics and Space Administration (NASA) and its contractors had allowed the shuttle to fly several times below full capacity; yet, no accident had happened. It took a low temperature as an initiating event to cause the technical O-ring failure that proved fatal to flight 51-L (*5*).

Studies of such failure stories are informative, but provide only a narrow glimpse of a large number of potential failure scenarios. A systematic analysis is required to put these results in perspective and to learn from past experiences, which often involve few total failures, if any at all, but many partial failures and near-misses. Probabilistic risk analysis (PRA) is one of such techniques that was developed primarily in the nuclear power industry (*6, 7*). Portions of the oil industry now use PRA models to assess the reliability of offshore platforms (*8*). These analyses focus mainly on the probability that a platform fails because of extreme loads, such as excessive wave heights beyond the chosen design criteria. Provided that these criteria were reasonable in the first place, this particular type of failure can be attributed to bad luck. More often, as I discuss in this article, accidents result from organizational errors that decrease the platform's capacity and are rooted in the way the companies operate. In this study, I use probabilities to link some organizational factors to the performance of the components and jacket-type offshore platforms as an illustration of the method (*9*). The data include probabilities of errors and error detection, and probabilities of failure of the basic components (foundation, jacket, and deck) conditional on different error states. I obtained these probabilities from one expert (*10*). His assessments are based on his experience in the oil industry and on different data sets providing statistics about failure types and failure causes for a large class of structures (*11–14*).

## Organizational Errors and System Reliability

PRA models relate the probability of failure of a system to the probabilities of initiating events, human errors, and failures of the components. Initiating events that trigger system failures include accidents (such as fires) and overloads (such as excessive waves heights), in which case failure occurs when these loads exceed the capacity. There has been some effort to include in this analysis the possibility of management errors (*15*), for example, to estimate the effects of recurrent design and construction problems on the seismic capacity of nuclear reactors (*16*). More generally, one can expect a variety of management errors (that may not have been observed so far) caused by organizational factors, such as excessive time pressures or failure to monitor hazard signals. These errors increase the probabilities of component failures either by increasing the probability of some initiating events or by decreasing the system's

The author is in the Department of Industrial Engineering and Engineering Management, Stanford University, Stanford, CA 94305.

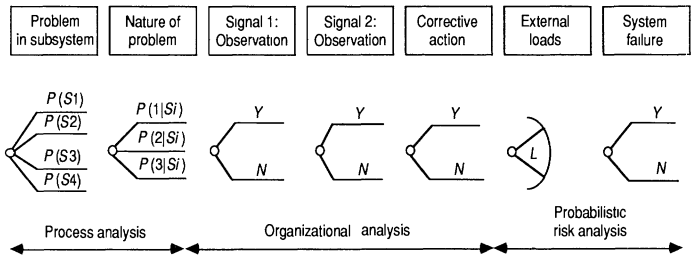


**Fig. 1.** Event tree showing the structure of the generalized reliability model. $P(Si)$, probability of a problem in subsystem $i$ (for example, foundation); $P(j|Si)$, probability of problem $j$ conditional on problem in subsystem $i$; Signal 1 or 2, observation of problem at each step of review or inspection; L, external loads (random variable).

capacity. I have thus extended the analysis to include explicitly the resulting variations of the component failure probabilities (Fig. 1). The first step is a study of the major components during the different phases of the system's life in which the focus is on organizational features. The objective of this analysis is to identify potential errors and to assess their probabilities of occurrences and the effectiveness of quality control. The second step is an evaluation of the effects of undetected errors on the failure probabilities of the components. The third step involves the use of PRA to compute the overall probability of failure for a spectrum of error scenarios.

*Information and incentives.* Reduction of a system's capacity can often be linked to wrong decisions. Some bad decisions may be simple human errors, but more often they are caused or encouraged by rules and goals set by the corporation. Two key elements of collective decision-making are information flow (who knows what and when) and incentive structure (how are individuals compensated and what are their objectives) (*17–19*). One wants to ensure that information is available when needed, but also that the organization is not swamped in a mass of irrelevant signals. First, one needs to determine which facts and data are actually observed and recorded in practice. Second, one needs to assess how information is communicated and, in particular, whether uncertainty and reservations are acknowledged and transmitted. Some organizations include redundant information channels (formal or informal) and mechanisms for full communication. Others tend to misrepresent or ignore relevant information.

Risk management decisions typically strike a balance between different objectives, such as safety and productivity. In most cases, people behave according to actual rewards if they believe that their actions are likely to be noticed. Incentive effects can then be formally analyzed with a rational decision model (*20*), which requires assessment of a distribution of risk attitudes among operators and allows derivation of the probability of outcomes of binary decisions. In other cases, rationality is a poor assumption (*21*). Factors such as anger or fatigue may change entirely the risk attitudes, and behavioral theories provide a better approach to risk-taking (*22–24*). Another solution is thus to avoid the unwarranted hypothesis of rationality and to assess directly, based on experience and results of behavioral studies, the probability of what people choose to do.

Incentives, information flow, and resource allocation are, in turn, rooted in the structure, the procedures, and the culture of the organization (*25*), all of which are critical in system safety (*26*). For example, an organizational structure that is strictly compartmentalized with no feedback and no horizontal connections is likely to cause information gaps, inconsistencies, and inefficiencies in risk management. Even if one relies on safety factors to balance safety requirements across the different branches and components, the result may be a wide spectrum of failure probabilities. Furthermore, the marginal gain of reliability for an additional dollar invested in each of the subsystems may also vary widely. The resources could thus be reallocated among subsystems to achieve a higher level of overall safety. Integration functions across the different branches are needed to ensure optimality of resource allocation and reliability of the interfaces. In addition to organizational structure, many procedures affect system reliability—for example, hiring and training practices, reporting of incidents, or maintenance and inspection schedules. Finally, beyond structure and procedures, the more elusive factor of organizational culture affects system reliability insofar as it influences the true incentives and the effectiveness of communications. Up to a point, management can shape organizational culture through incentives and rules. Yet the actual culture may promote behaviors that do not seem to fit the official system but are rewarded in an informal way.

*Gross errors and errors of judgment.* An organization, through its structure, procedures, and culture, can commit systematic errors that affect part or all of an engineering system and act as common causes of failure. Concerns about human factors in the management of technology have increased with the complexity of engineering systems (*27–29*). Several taxonomies of human errors have been proposed, often to understand the psychological roots of human actions (*30*) and to classify data about human errors. For example, Rasmussen (*31*) links human malfunctions to an analysis of an operator's mental activity. Reason (*31, 32*) distinguishes managerial factors from individual actions and proposes a "generic error-modeling system" involving skill-based, rule-based, and knowledge-based behaviors, whereas Norman's analysis relies on the difference between slips and mistakes that are the results of intentional actions (*33*).

The taxonomy developed in this study focuses on individual decisions with specific consideration of risk and uncertainty. This approach allows linkage of errors to management problems of information and incentives. The primary distinction is between errors of judgment (in situations involving risks) and gross errors. In theory, gross errors and errors of judgment require different kinds of methods of analysis to capture the effects of uncertainties. In practice, the two types of errors often call for different management strategies. Errors are classified further into large categories that permit a rough assessment of probabilities of occurrence and detection (Fig. 2).

Gross errors are those about which there is no controversy or ambiguity ($2 + 2 = 6$) regardless of their severity. Presumably, the decisions would be reversed if they were reexamined and everyone would agree, including the individual who made the error in the first place. Gross errors are further divided among (i) communication problems (caused either by structural problems or defective procedures); (ii) cognitive problems due to fundamental ignorance, use of wrong models, or accidental slips (caused, for example, by inadequate training or excessive demands on workers); and (iii) errors due to physiological or psychological limitations, such as seasickness or sheer stupidity, that may be caused by inadequate hiring practices or poor job design.

Errors of judgment involve ambiguous or incomplete information, and therefore risk, in decision-making. They may be caused by wrong treatment of uncertainties or an attitude toward risk that does not correspond to that of the corporation. Cognitive errors of reasoning under uncertainty include known biases such as jumping to conclusions on the basis of too small a sample (*34*) or neglecting dissonant information (*35*). Errors of judgment may also occur when defective procedures and inappropriate organizational structures cause a rational individual to make decisions that are at odds with the global objectives of the organization. Unreasonable production goals and time constraints are common examples. Feedback to upper management is thus needed to ensure that they understand the consequences of the constraints that they set. Once the costs of

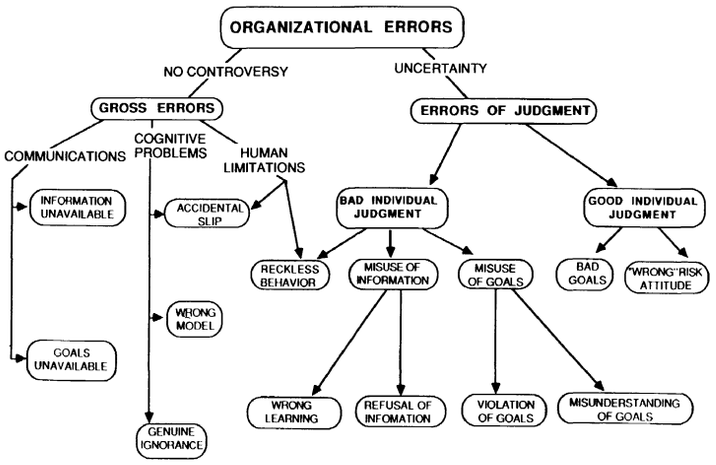


**Fig. 2.** A taxonomy of organizational errors.

requirements are understood, the choice of goals cannot be separated from the design of incentives, rewards, and guidelines about how to manage tradeoffs. One approach is the principal-agent paradigm (*36*) in which the goal is to design a reward system that leads individuals to act according to what management would want under various circumstances. These circumstances, however, are often difficult to predict.

## Failures of Jacket-Type Offshore Platforms

*Case histories*. Past failures of offshore platforms can often be linked to production pressures (*37*). In the four stories presented below (*38*), failures can be traced prima facie to human errors, but the problems were actually rooted deeper in the organization. In 1969, a platform slid in the mud of the Gulf of Mexico because the design of the foundation was completed before the results of soil tests were known. The design was completed early to stay on schedule, even though initial evaluations indicated the potential for mudslide. The risk was so obvious that the lead engineer had refused to sign the final construction drawings. The fundamental cause of the failure was the poor timing of the tests and the practice of parallel processing even when earlier information indicated high uncertainties.

In 1979, a mobile drilling rig originally designed for the Gulf of Mexico was to be sited in Lower Cook Inlet in Alaska. The results of an analysis had indicated that failure risks were high as a result of severe weather conditions during the Alaska winter. Reluctantly, management decided to delay the siting by several months. Later on, when another risk analysis indicated a high probability of scour (erosion of the sea floor by currents) around the rig's footing, the results were ignored. Scour occurred and two divers were killed during subsequent placement of scour protection. The unit was then towed toward the California coast and sank in the Aleutian trench (because a door was left open) during the towing. The insurer incurred the cost. The fundamental error was an error of judgment similar to the classic refusal to disregard sunk costs (costs that have been incurred in the past): because of earlier delays, relevant information regarding the risk of scour was ignored.

In another instance, a platform was constructed without a foundation because, in a lowest-bid procedure, the job was awarded to a contractor who simply drove piles into the ground instead of drilling and grouting, in spite of soil test results that indicated that the soil was brittle. During a pull-out test on a well conductor, the engineers discovered that there was no skin or shaft resistance. The platform had to be derated and fixed at high cost. In this case, a gross error occurred under the incentives created by the lowest-bid practice.

In yet another case, a steel jacket was designed to be towed to the drilling site and launched from a barge. Buoyancy tanks were placed at the upper face of the top end of the jacket; but when the jacket was launched, it rotated because of high momentum. The tanks were ineffective in slowing the movement, and the jacket embedded upside down. The initial error was the result of ignorance and the use of a wrong model by the engineer who conceived the system. The incident, however, was kept secret, and nothing was learned from it. Two years later someone else made the same mistake at a different site, at which point the lesson was finally absorbed by the industry. This case is an example of gross error compounded by inadequate learning.

These four examples are a small sample of organizational errors that have led to platform failures. For risk management purposes, it is neither necessary nor desirable to anticipate all detailed sequences of events that might lead to an accident. Instead, the different types of organizational errors and element failures are structured into broad classes of scenarios for the assessment of their contributions to the probability of platform failure.

*Errors and risks*. A jacket-type offshore platform consists of a steel tower (the jacket), anchored to the sea floor by its foundation, and supporting the deck (or top side) on which oil and gas production takes place (Fig. 3). Oil and gas are brought to the top side by a group of pipes (risers) connecting the deck and the wells. Probabilistic risk analysis models are based on the probability that the loads (for instance, wave heights) exceed the platform's capacity, or that a severe accident such as a fire or a blowout causes a platform failure (event $F$). An initiating event is an event that triggers an accident sequence—for example, a wave that exceeds the jacket's capacity or an earthquake that, in turn, triggers a blowout that causes failure of the foundation. As initiating events, they are mutually exclusive: only one of them starts the accident sequence. A catastrophic platform failure can start by failure of the foundation ($O$), failure of the jacket ($A$), or failure of the deck ($E$). These initiating failures are also (by definition) mutually exclusive and constitute the basic events of the PRA model in its simplest form. $Y_i$ is an initiating event (index $i$), such as an excessive wave load; $\gamma$ is a particular level

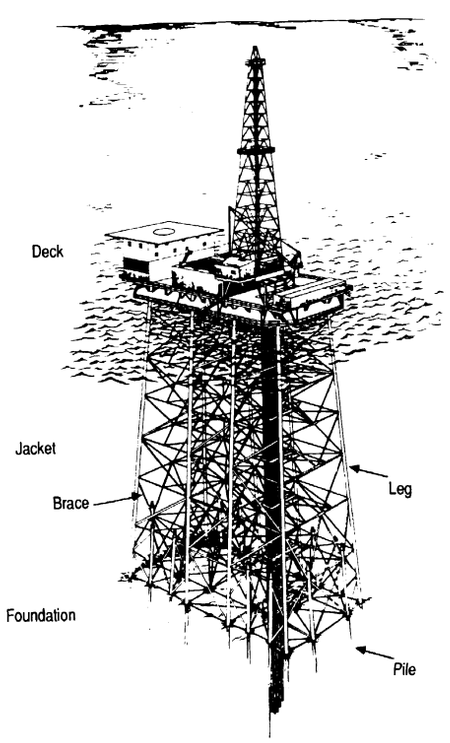


**Fig. 3.** Example of jacket-type offshore platform. [Adapted from (*57*), courtesy R. G. Bea]

**Table 1.** Probabilities of detection of design errors by the successive reviewers given that the errors were not detected earlier, and probabilities of error correction given detection. Illustrative values are given.

| Type of error | Lead engineer | Engineering manager | Constructor | Corrective action |
|---|---|---|---|---|
| Gross error | | | | |
| High severity | 0.45 | 0.8 | 0.7 | 0.9 |
| Low severity | 0.20 | 0.65 | 0.4 | 0.6 |
| Error of judgment | | | | |
| High severity | 0.20 | 0.5 | 0.1 | 0.6 |
| Low severity | 0.05 | 0.3 | 0.01 | 0.2 |

**Table 2.** Events and random variables for the analysis of the design review process.

| Variable | Definition |
|---|---|
| $e_{t,s}$ | Error of type $t$ (for example, gross errors) and severity level $s$. |
| $e_s$ | Error of severity level $s$ (including no error: severity 0). |
| $E_i$ | Initial error state: occurrence of errors in the process (random variable). |
| $E_f$ | Final error state after review and correction (random variable). |
| $D_j$ | Detection of an error at step $j$ of the review process. |
| $P(Dj)$ | Probability of error detection at step $j$ given that is was not detected before. |
| $C$ | Error correction. |

of severity of $Y_i$ (for example, a wave height); $f_{Yi}(y)$ is the probability density function of the annual maximum value of $Y_i$; $P(X)$ is the annual probability of failure $X$ of the different components ($X = O$, $A$, or $E$); and $P(X|y)$ is the probability of possible failures $X$ conditional on $y$. A simple model yielding the marginal annual probability of platform failure is

$$P(F) = P(O) + P(A) + P(E) \quad (1)$$

$$P(X) = \sum [\int f_{Yi}(y) \times P(X|y)\, dy] \quad (2)$$

for $X = O$, $A$, or $E$.

Different analytical methods are used for the computation of the failure probabilities of each subsystem. For example, the probability of failure of a jacket is computed on the basis of identification of the sequences of member failures leading to jacket failure and study of the reallocation of internal forces after the failure of each member (*39*). Typical annual results of a PRA for an offshore structure are $P(F) \approx 2 \times 10^{-3}$, with $P(E) \approx 10^{-3}$ (about 50%), $P(A) \approx 6 \times 10^{-4}$ (about 30%), and $P(O) \approx 4 \times 10^{-4}$ (about 20%).

The capacity of each subsystem and, therefore, its annual failure probability, varies with the occurrence of errors in phases of design, construction, or operation. For the jacket, examples of errors in the design phase include ill-conceived configuration; in the construction phase, the use of the wrong steel (thus, lower yield stress than assumed in the design); and in both phases, wrong sizing of the members, resulting in decreased resistances. Errors in the operation phase, such as failure to maintain the structure and prevent corrosion, can also result in a decrease of the members' resistances, with probabilistic dependencies induced by common maintenance procedures. Each error scenario can be linked to its effects on the inputs of the technical analysis (for example, probability distribution of cross sections of members and yield stresses of the steel). The PRA model for the corresponding subsystem yields the resulting failure probabilities. In this study, the analysis is more global, and the effects of errors in the different subsystems are assessed directly through expert opinion.

*Design process and errors.* In the decision to design and construct a new platform, corporate management generally fixes a target production level, the site, the general type of platform, the schedule, and the budget. An engineering development group chooses the actual platform type and configuration. The engineering design group decides on the details of the configuration, design parameters, and inspection and maintenance requirements within the limits set by management and development. Other agents include the contractors and outside participants such as the regulators, public interest groups, and other oil companies who have an interest in the image and performance of the industry as a whole. The incentive system is dominated by corporate goals, generally set in a rigid manner with limited and filtered feedback to upper management. Each level can ask confirmation and clarification regarding particular decisions, but there are strong incentives to stick to goals and constraints as set. There is relatively little individual penalty for technical failures, which are rare in any case; sanctions are mainly setbacks in the careers of key personnel involved. There is, however, high individual penalty in the long run for not reaching corporate targets. Furthermore, the custom of awarding jobs to the lowest bidder is often viewed as the most efficient way to satisfy cost constraints, but can yield poor-quality work that decreases system safety.

The design process is divided among specific steps, such as preliminary configuration and sizing of platform elements. At each step, errors can occur in the different categories of the taxonomy presented above. For simplicity, I have limited the analysis to the distinction between gross errors and errors of judgment and two levels of error severity. As a first approximation, I assume that in each phase, no more than one error occurs in each of the subsystems (or that the probability of two or more errors is negligible compared to that of one error only). Examples of design errors include: for the foundation, the use of a wrong formula for the calculation of pile capacity (gross error) or the underestimation of soil stiffness because of reliance on inaccurate laboratory soil tests (error of judgment); for the jacket, omission of the tidal currents in the calculation of design forces (gross error); and for the deck, a mistake in load computations (gross error).

The design review process is sequential (Table 1). The first review is performed by the lead engineer, typically competent and knowledgeable, but not necessarily someone who has had long experience in the field. This person is thus more likely to detect gross errors than errors of judgment. The second review is performed by the engineering manager, generally someone experienced who may not check the detail of all computations but can detect errors of judgment more easily than the lead engineer. Finally, the constructor may detect an error when actually doing the work in the field. At

**Table 3.** Annual probability of failure of an offshore platform. Contribution of gross errors and errors of judgment are considred in the design phase only.

| Design errors | Subsystem | | | Whole structure | |
|---|---|---|---|---|---|
| | Foundation | Jacket | Deck | Failure probability | Percent |
| High severity | $3.6 \times 10^{-4}$ | $10^{-4}$ | $2.2 \times 10^{-5}$ | $4.8 \times 10^{-4}$ | 25% |
| Gross errors | $7.2 \times 10^{-5}$ | $5 \times 10^{-5}$ | $1.8 \times 10^{-5}$ | $1.4 \times 10^{-4}$ | 7% |
| Errors of judgment | $2.9 \times 10^{-4}$ | $5 \times 10^{-5}$ | $4.4 \times 10^{-6}$ | $3.4 \times 10^{-4}$ | 18% |
| Low severity | $2.4 \times 10^{-5}$ | $2 \times 10^{-4}$ | $5.2 \times 10^{-5}$ | $2.9 \times 10^{-4}$ | 15% |
| Gross errors | $4.8 \times 10^{-6}$ | $10^{-4}$ | $4.2 \times 10^{-5}$ | $1.5 \times 10^{-4}$ | 8% |
| Errors of judgment | $1.9 \times 10^{-5}$ | $10^{-4}$ | $10^{-5}$ | $1.3 \times 10^{-4}$ | 7% |
| No design error | $3.2 \times 10^{-5}$ | $2.9 \times 10^{-4}$ | $9.8 \times 10^{-4}$ | $1.3 \times 10^{-3}$ | 60% |
| Total | $4.1 \times 10^{-4}$ | $6.1 \times 10^{-4}$ | $10^{-3}$ | $2.2 \times 10^{-3}$ | |

**Table 4.** Probability of failure of offshore platforms given the possibility of cumulation of errors in design, construction, and operation; the error severity is the dominant (most severe) one in each error combination.

| Error severity | Subsystem | | | Whole structure | Percent |
|---|---|---|---|---|---|
| | Foundation | Jacket | Deck | | |
| High severity | $3.9 \times 10^{-4}$ | $4.1 \times 10^{-4}$ | $8.0 \times 10^{-4}$ | $1.6 \times 10^{-3}$ | 79% |
| Low severity | $1.7 \times 10^{-5}$ | $1.6 \times 10^{-4}$ | $1.7 \times 10^{-4}$ | $3.5 \times 10^{-4}$ | 17% |
| No error | $7.2 \times 10^{-6}$ | $3.1 \times 10^{-5}$ | $4.1 \times 10^{-5}$ | $7.9 \times 10^{-5}$ | 4% |
| Total | $4.1 \times 10^{-4}$ | $6.1 \times 10^{-4}$ | $1.0 \times 10^{-3}$ | $2.0 \times 10^{-3}$ | 100% |

this late stage, it will be easier to correct gross errors than errors of judgment about which the constructor, in spite of his experience, may have little to say. Although the review process is similar for the different subsystems (foundation, deck, and jacket), the details of the procedures vary. For the foundation, a significant part of the review may involve questioning the assumptions, whereas for the rest of the structure, it may focus more on the verification of the analysis and the computations.

To compute the probability that the system fails because of undetected errors, I used an analytical model based on events and random variables described in Table 2.

$$P(E_f = E_i = e_s) = \sum_t \{P(E_i = e_{t,s}) \times \prod_j [1 - P(D_j \mid E_i = e_{t,s})P(C \mid D_j, E_i = e_{t,s})]\} \quad (3)$$

$$P(X) = \sum_s P(X, e_s) = \sum_s P(E_f = e_s) \times P(X \mid e_s) \quad (4)$$

for $X = O$, $A$, or $E$.

Equation 3 yields the probability that an error of given severity remains at the end of the review process. It is the probability that an error of this severity initially occurs, is not detected, or is not corrected. The effect decreases in the system's capacity to withstand loads and is characterized by the probability of failure conditional on the final error state. Equation 4 yields the probability of failure of each subsystem as the sum, for all error severity levels, of the joint probabilities of failure and errors. The probability of failure of the whole platform is the sum of the probabilities of (initial) failure of the foundation, the jacket, and the deck (Eq. 1). Table 3 shows its allocation among error types and severity levels. Because a large fraction of accidents starts on the deck with operation problems, design errors account for only about 40% of the total failure probability.

*Error accumulations and compounded effects.* Errors in the design are thus only part of the story. Offshore platform accidents often occur because errors in different phases of the structure's lifetime introduce "resident pathogens" (*40*) that contribute to weakening the system. Design errors are therefore compounded by errors of construction and operation. During construction, a foundation pile sleeve may be only partially grouted, the wrong welding rods may be used on critical joints of the jacket, or deck sections may be installed in bad weather. During operation, drilling blowouts at the foundation may undermine the foundation and significantly reduce its capacity; someone may decide not to repair strength-degrading dents in braces of the jacket; or production may be maintained during a platform fire, causing an explosion of the production pipeline. I analyzed conjunctions of errors of design, construction, and operation and their effect on system reliability using a model similar to Eqs. 3 and 4. The data include probabilities of occurrence and detection of construction errors of different types and severity levels in each subsystem and annual probabilities of operation errors. Design errors, construction errors, and operation errors are assumed to be independent among themselves and across subsystems. The consequences of error scenarios for the different subsystems (foundation, jacket, and deck) are characterized by probabilities of failure conditional on each possible combination of errors of different types (for example, construction and operation) and different severity levels. These probabilities capture the synergies among errors—for example, the fact that low-severity design or construction errors weaken the system and worsen the effects of high-severity errors in the operation phase.

The results of the allocation of the probability of failure among subsystems according to the highest severity level in each error combination are shown in Table 4. Failure scenarios in which the error severity is low represent only about 20% of the overall failure probability. But when the total contribution to the overall failure probability is computed by difference of platform failure probabilities with and without low-severity errors, they actually account for about 50% of the failure probabilities of the foundation and jacket, 20% for the deck, and about 40% for the whole platform. Therefore, because of the synergistic effects mentioned above, low-severity errors are important contributors to the overall probability of system failure.

## Some Roots of Organizational Errors and Possible Improvements

Overlooking a construction defect that would take time to correct, continuing construction in bad weather, or delaying maintenance are examples of errors of judgment induced by production and schedule pressures. Without sufficient incentives to take reasonable safety measures, these may be rational individual decisions if there is an immediate cost of delay and no imminent threat of accident, even though they increase the probability of failure in the long term. Five particular management problems seem to be at the root of these errors: (i) time pressures, (ii) observation of warnings of deterioration and signals of malfunction, (iii) design of an incentive system to handle properly the tradeoffs between productivity and safety, (iv) learning in a changing environment where there are few incentives to disclose mistakes, and (v) communication and processing of uncertainties.

*Time pressures.* One of the key mechanisms that shapes engineering decisions in many organizations is the use of goals set by corporate management (*22*). Time pressures induced by schedule constraints have several effects. Up to a point, they may be stimulating for individuals as well as teams; but with too much pressure, people tend to cut corners, in particular on the critical path in construction and operation phases. For example, Carson shows how time pressures in the exploitation of the petroleum resources in the North Sea and inadequate safety regulations led to an unusually high rate of accidents (*37*). Under time pressures, parallel processing becomes an attractive option (for example, the oil industry practice of simultaneously conducting soil testing and preliminary foundation design). Given the costs of delays, parallel processing may be a rational decision if there is little uncertainty about the soil characteristics and if the design can be modified later at low cost. Otherwise, it may be a costly gamble because of the technical and financial risks involved. Furthermore, once the design has started, there may be an "escalation of commitment" (*41*), and it may be tempting to ignore or downplay the test results if they invalidate the efforts that have been put into the initial drafts. Time pressures not only increase the probability of errors, but also decrease the chances that they are detected by the regular procedures. Low-severity errors and errors of judgment are more likely to be accepted because of the time that it

would take to correct them. Time and schedule pressures are part of corporate life and some schedule is in order, but schedules must be adaptable to accommodate response to incidents. Feedback to management is therefore needed to ensure that the reward system does not penalize those who respond to avoid future failures, at the cost of immediate delays.

*Missed signals of deterioration.* Many accidents and failures, technical or not, involve missed warnings. In the management of engineering systems, the problem can be generally attributed to inadequate inspection and maintenance procedures, failure to record performance trends, or incentives to ignore the warnings. An example, for marine structures, is missing signs of severe corrosion. Setting an appropriate warning system often involves selecting signals to be observed, trend analysis to monitor component deterioration, and the choice of an alert threshold that strikes a balance between reacting too late (or not at all) and intervening too often (*42*). The choice of an optimal inspection and maintenance system (for instance, on schedule or on demand) can be based on a stochastic model that involves the probabilities of detection, false alerts, and missed signals (*43*). Such a model requires identification of classes of accident scenarios, the corresponding deterioration rates, the lead time required for action, and rates of human response to signals given past warning patterns. Even when signals of malfunction are observed, identification of the corresponding malfunction may be difficult in complex situations and may require the help of appropriate decision-support software (*44, 45*).

*Productivity versus safety.* In the long term, reliability contributes to productivity, in particular for expensive systems or for those whose failures can have disastrous consequences. Yet the daily management of critical systems often involves short-term tradeoffs between productivity and safety. Preferences appear to be driven more by costs than by rewards (*46*), hence the willingness to take risks to stay on schedule. In the oil industry, it may thus be tempting to delay maintenance or pursue operation under severe environmental conditions or at reduced system capacity. Corporations concerned with long-term results must thus provide adequate training, a safety culture, reasonable production goals, and appropriate guidelines and incentives for handling incidents that can degenerate into catastrophes.

The problem of designing incentives for balancing productivity and safety seldom has a clear-cut solution (*47, 48*). External control and regulation or an independent safety office have a limited effectiveness when the controlling organization depends on the industry it is regulating for critical information (*49*). In practice, when production and safety are inseparable, at least two management approaches can be considered. The first is to issue strict and detailed guidelines, a strategy that demands that most problematic situations be foreseen when designing these guidelines. This approach may reduce the probability of serious errors of judgment, but it sacrifices flexibility. The second solution is for management to leave the decision to competent operators, to emphasize the need for both productivity and safety, and to give operators the responsibility to use their judgment in balancing the two. This strategy has the advantage of making the operators directly responsible, but it leaves more room for major errors of judgment. The competence and the experience of the decision makers on the spot are therefore critical and require that they learn from past corporate experience.

*Learning.* Clearly, learning has occurred in the oil industry in the past 50 years, as demonstrated by a marked decrease in the annual failure rate of offshore platforms (*50*). Yet the transfer of experience has been slowed down both by some management procedures and by the recent history of the oil industry. Management by goals in high-pressure industries encourages an image of super performance and creates a tendency to cover up past mistakes (*38*). In such an environment, learning is difficult for the individual and the corporation. Furthermore, promotions and transfers in the oil industry often occur so fast that people do not have the time to observe the effects of their past actions. These transfers and incentives to remove evidence of past errors make it still more difficult for the engineer who inherits a problem to understand what happened and to learn from it. Therefore, not only does the person who made the mistake miss the opportunity to learn, but so does the organization.

Transfer of experience from one environment to another may also become an organizational issue. Experience gathered in the Gulf of Mexico is to some extent relevant, but not sufficient to capture all the problems that may arise on the Alaska North Slope. In extreme cases, the experience of senior management is not even available, because they have left the corporation. Corporate learning requires formal or informal mechanisms to observe, record, and retrieve past collective experience, including mistakes (*51*). When the appearance of performance is essential to personnel evaluation, protection of identity may be a necessary condition for the disclosure of past errors. More generally, learning in an innovative environment often involves gradual resolution of uncertainties on the basis of new evidence. The formal use of probabilities may reduce the effect of errors of logic in the updating of incomplete information. Learning from past accidents and near misses allows the organization not only to avoid the repetition of gross errors, but also to improve the understanding of potential hazards and to calibrate better its collective judgment.

*Uncertainties.* Recognition, communication, and management of uncertainties are major issues in many engineering fields. In the oil industry, where the environment is often poorly known and highly variable, uncertainties are inescapable (*52*) in exploration decisions, in the development of new offshore structure technologies, and in the choice of design parameters. A common strategy is to try to eliminate uncertainties from decisions, sometimes simply by redefining the problem (*53*). When this cannot be done, incentives and culture often lead to denial, biases in the interpretation of conflicting evidence, and overconfidence in either the most likely or the most favorable hypothesis. Unless there is a clear danger of failure, the natural tendency, in the communication of incomplete information, is to tell people what they prefer to hear. As the information travels along a hierarchical path, qualifiers and restrictions are thus likely to be dropped, particularly when statistics have not been gathered and when uncertainties can only be characterized by engineering judgment. An accurate description of the state of knowledge may never reach the decision maker, thus increasing the probability of error. Therefore, incentives to disclose all relevant information need to be integrated in the reward structure to balance this tendency toward optimism and wishful thinking, in particular when accidents are rare and when, most of the time, there is no technical feedback.

The root of the problem is sometimes outside of the corporation and in the legal system. If there is evidence that an adverse situation of low probability has been examined and judged unlikely enough to be ignored, the corporation may be punished for having recognized its possibility. The incentives are thus to make sure that no trace of it can be found in the paper trail. Yet, for the organization's sake and that of society, it is better off dealing explicitly with undesirable and unlikely prospects when safety is at stake.

*Some organizational improvements.* Inspection alone is not the most efficient way to provide quality and reliability. The above analysis suggests other organizational improvements—for example, (i) ensuring the effectiveness of learning mechanisms (as a complement to personnel selection and training) by carefully maintaining corporate memory and updating databases; (ii) using the concepts and vocabulary of probability in the management process to improve communications as well as decision-making; (iii) adjusting schedul-

ing procedures to include uncertainties and the possibility of delays in the different tasks of the program; (iv) improving feedback mechanisms within the corporation to make managers more aware of the consequences of the goals that they set; and (v) having project engineers check that technical changes do not compromise system safety. The model presented above can be used to quantify the corresponding safety gains as a function of the corresponding decrease of the probability of errors or increase of their probability of detection, or both.

For example, in the U.S. oil industry, the design review process for offshore platforms can be improved by the intervention of a certified verifying authority. A high-quality verification process could decrease significantly the probabilities of undetected gross errors and errors of judgment for an additional cost of about $100,000. I computed the corresponding safety gains on the basis of expert assessments of the reduction of the probability of undetected errors (*9*). The considered improvements of the design review could reduce approximately by a factor of 2 the probability of failure of the foundation where most of the uncertainties and the difficulties are encountered. The improved design review would reduce by about 20% the probability of failure of the jacket, but would improve little the reliability of the deck, which is more susceptible to operations errors than design errors. Altogether, an independent review process would therefore decrease by about half the contribution of design errors to the overall failure probability (which was found to be about 40%), thus decreasing by about 20% the probability of platform failure. To achieve such safety benefits, engineers tend to consider first structural reinforcements such as increasing the number and the strength of jacket members. For a structure that costs about $400 million, the cost of reducing the failure probability by 20% is assessed by the oil companies at about $9 million (*38*). For the same benefit, the cost of the structural solution is thus roughly two orders of magnitude above the cost of the proposed improvement of the design review. Yet, the technical solution is generally preferred because it seems to bring sure benefits if one makes the implicit assumption that the work will be done as planned.

*Limitations of the quantitative approach.* Although the use of probability concepts seems to have gained acceptance in the U.S. oil industry, managers, in general, do not like them as a descriptor of risk (*54*). Even if one accepts in theory the framework of probabilistic reasoning, the same objections that have been raised in the past against probabilities in technical PRA (*55, 56*) can be raised a fortiori when they are used to describe management problems (for example, that it is impossible to ensure that all failure modes, failure causes, and errors have been identified; or, that with little information, probabilities are "soft," therefore, lack credibility and can be manipulated). It is difficult, in particular, to assess probabilities of human errors. The data are often scarce, personalities and situations vary widely, and human behaviors in general seem to be less amenable to probabilistic evaluation than the performance of mechanical components. Yet, past experience provides some information. If some errors are more frequent than others, probability allows setting priorities among corrective measures based on frequencies and consequences of these errors. The results, however, may be coarse because of uncertainties in the inputs. In all instances—human as well as technical performances—an element of subjectivity is unavoidable when probability is used, whether in the encoding of expert opinions or simply in the transfer of statistical data to a particular case.

## Conclusions

Quantification of risks is not necessary if the system is simple, if the solutions are fairly obvious, or if there is no significant resource constraint and, therefore, no need to set priorities. In many cases, however, resource constraints are inescapable, and hazardous technologies will remain because society values their benefits. Probabilistic risk analysis is one of the sources of information that can provide guidance to improve management practices. If the current models are extended to include organizational factors, the analysis can support organizational as well as technical organizational improvements.

### REFERENCES AND NOTES


1. C. Perrow, *Normal Accidents* (Basic Books, New York, 1984).
2. B. M. Turner, *Admin. Sci. Q.* **21**, 378 (1976).
3. ———, *Man-Made Disasters* (Wykeham, London, 1978).
4. D. Vaughn, *Admin. Sci. Q.* **35**, 225 (1990).
5. *Report of the Presidential Commission on the Space Shuttle Challenger Accident* (U.S. Government Printing Office, Washington, DC, 1986).
6. E. J. Henley and H. Kumamoto, *Reliability Engineering and Risk Assessment* (Prentice-Hall, Englewood Cliffs, NJ, 1981).
7. *Report WASH-1400 (NUREG-75/014)* [U.S. Nuclear Regulatory Commission (USNRC), Washington, DC, 1975].
8. *Application of Risk Analysis to Offshore Oil and Gas Operations* (U.S. Department of Commerce, National Bureau of Standards, Spec. Publ. 695, Washington, DC, 1985).
9. M. E. Paté-Cornell and R. Bea, *Organizational Aspects of Reliability Management: Design, Construction, and Operation of Offshore Platforms*, Res. Rep. No. 89-1 (Department of Industrial Engineering and Engineering Management, Stanford Univ., Stanford, CA, 1989).
10. Because of the generic nature of the example, the opinion of one expert only was solicited (R. G. Bea from the Department of Naval Architecture, University of California, Berkeley). The implications of his assessments for the probabilities of different types of platform failures were found to be consistent with statistics of worldwide accidents. The study of a particular structure in a given location requires a specific reliability model based on formal analysis of the local loads, and the opinions of several experts should be solicited if necessary.
11. *Rationalization of Safety and Serviceability Factors in Structural Codes* (Construction Industry Research and Information Association, Rep. 63, London, 1977).
12. *Comparative Safety Evaluation of Arrangements for Accommodating Personnel Offshore* (Offshore Certification Bureau, Rep. OTN-88-175, London, 1988).
13. *Risk Assessment Report of the Norwegian Offshore Petroleum Activities* (Royal Norwegian Council for Scientific and Industrial Research, Oslo, Norway, 1979).
14. Veritec: The Worldwide Offshore Accident Data Bank (WOAD 85) (Veritec, Oslo, Norway, 1985).
15. J. S. Wu, G. E. Apostolakis, D. Okrent, in *Proceedings of the 1989 Annual Meeting of the Society for Risk Analysis*, 1985 (Plenum, New York, 1991), p. 429.
16. M. Gonzalez-Cuesta and D. Okrent, *Nucl. Eng. Des.* **97**, 89 (1986).
17. J. G. March and H. A. Simon, *Organizations* (Wiley, New York, 1958).
18. K. J. Arrow, *Decision and Organization* (North-Holland, Amsterdam, 1972).
19. H. A. Simon, *Administrative Behavior* (Free Press, New York, 1976).
20. H. Raiffa, *Decision Analysis* (Addison-Wesley, Reading, MA, 1968).
21. H. A. Simon, *Models of Bounded Rationality* (MIT Press, Cambridge, MA, 1982).
22. R. M. Cyert and J. G. March, *A Behavioral Theory of the Firm* (Prentice-Hall, Englewood Cliffs, NJ, 1963).
23. B. Fischhoff, S. Lichtenstein, P. Slovic, S. Derby, R. Keeney, *Acceptable Risk* (Cambridge Univ. Press, New York, 1981).
24. P. Slovic, B. Fischhoff, S. Lichtenstein, *Annu. Rev. Psychol.* **28**, 1 (1977).
25. K. E. Weick, *Calif. Manage. Rev.* **29**, 112 (1987).
26. T. R. La Porte, *High Reliability Organization Project* (University of California, Department of Political Science, Research Report, Berkeley, California, 1988).
27. L. P. Goodstein, H. B. Andersen, S. E. Olsen, Eds., *Tasks, Errors and Mental Models* (Taylor and Francis, London, 1988).
28. J. T. Reason, *Reliabil. Eng. Syst. Saf.* **22**, 137 (1988).
29. ———, *Human Error* (Cambridge Univ. Press, New York, 1990).
30. D. D. Woods, E. M. Roth, H. Pople, *Reliabil. Eng. Syst. Saf.* **22**, 169 (1988).
31. J. Rasmussen, K. Duncan, J. Leplat, Eds., *New Technology and Human Error* (Wiley, New York, 1987).
32. J. T. Reason, "The human contribution to organizational accidents," paper presented at the Second World Bank Workshop on Safety Control and Risk Management (World Bank, Washington, DC, 1989).
33. D. A. Norman, *Psychol Rev.* **88**, 1 (1981).
34. D. Kahneman, P. Slovic, A. Tversky, *Judgment Under Uncertainty: Heuristics and Biases* (Cambridge Univ. Press, New York, 1982).
35. E. L. Janis, *Victims of Group Think* (Houghton Mifflin, Boston, MA, 1972).
36. D. Levinthal, *Econ. Behav. Organ.* **9**, 153 (1988).
37. W. G. Carson, *The Other Price of Britain's Oil: Safety and Control in the North Sea* (Rutgers Univ. Press, New Brunswick, NJ, 1982).
38. R. G. Bea, personal communication.
39. H. Nordal, C. A. Cornell, A. Karamchandani, in *Proceedings of the Marine Structural Reliability Symposium* (Society of Naval Architecture and Marine Engineering, Arlington, VA, 1987), pp. 193–216.
40. J. T. Reason, "Resident pathogens and risk management," paper presented at the First World Bank Workshop on Safety Control and Risk Management (World Bank, Washington, DC, 1988).
41. B. Staw, *Acad. Manage. Rev.* **6**, 577 (1981).
42. M. E. Paté-Cornell, *Risk Anal.* **6**, 223 (1986).

43. ———, H. L. Lee, G. Tagaras, *Manage. Sci.* **33**, 1277 (1987).
44. W. B. Rouse, in *Human Detection and Diagnosis of System Failures*, J. Rasmussen and W. B. Rouse, Eds. (Plenum, New York, 1981), pp. 199–216.
45. R. E. Curry, *ibid.*, pp. 171–184.
46. C. A. Heimer, *Annu. Rev Sociol.* **14**, 491 (1988).
47. J. G. March and J. P. Olsen, *Ambiguity and Choice in Organizations* (Universitetsforlaget, Bergen, Norway, 1976).
48. J. G. March, *Decisions in Organizations* (Blackwell, New York, 1988).
49. J. Pfeffer and G. R. Salancik, *The External Control of Organizations: A Resource Dependence Perspective* (Harper and Row, New York, 1978).
50. R. G. Bea, *J Struct. Div. Am. Soc. Civ. Eng.* **106**, 1835 (1980).
51. B. Levitt and J. G. March, *Annu. Rev. Sociol.* **14**, 319 (1988).
52. A. L. Stinchcombe and C. A. Heimer, *Organization Theory and Project Management. Administering Uncertainty in Norwegian Offshore Oil* (Norwegian Univ. Press, Oslo, 1985).
53. A. Tversky and D. Kahneman, *Science* **211**, 453 (1981).
54. J. G. March and Z. Shapira, *Manage. Sci.* **33**, 1404 (1988).
55. *Report NUREG/CR-0400* (USNRC, Washington, DC, 1978).
56. W. R. Freudenburg, *Science* **242**, 44 (1988).
57. R. G. Bea, *Oceanology Int.* (1975), p. 11.

# Highly Parallel Computation

PETER J. DENNING AND WALTER F. TICHY

**Highly parallel computing architectures are the only means to achieve the computational rates demanded by advanced scientific problems. A decade of research has demonstrated the feasibility of such machines, and current research focuses on which architectures are best suited for particular classes of problems. The architectures designated as MIMD and SIMD have produced the best results to date; neither shows a decisive advantage for most near-homogeneous scientific problems. For scientific problems with many dissimilar parts, more speculative architectures such as neural networks or data flow may be needed.**

COMPUTATION HAS EMERGED AS AN IMPORTANT NEW method in science. It gives access to solutions of fundamental problems that pure analysis and pure experiment cannot reach. Aerospace engineers, for example, estimate that a complete numerical simulation of an aircraft in flight could be performed in a matter of hours on a supercomputer capable of sustaining at least 1 trillion floating point operations per second (teraflops, or tflops). Researchers in materials analysis, oil exploration, circuit design, visual recognition, high-energy physics, cosmology, earthquake prediction, atmospherics, oceanography, and other disciplines report that breakthroughs are likely with machines that can compute at a tflops rate.

The fastest workstations today operate at maximum speeds of slightly beyond 10 million flops (10 megaflops, or mflops). In contrast, the fastest supercomputers have peak rates in excess of 1 billion flops (gigaflops, or gflops)—for example, the NEC SX-2 is rated at 1.0 gflops and the Cray Y-MP at 2.7 gflops. Even faster computers are being designed: the four-processor NEC SX-3 (1990) will have a peak rate of 22 gflops and the Cray 4 (1992) 128 gflops. When recompiled for these machines, standard Fortran programs typically realize 10 to 20% of the peak rate. When algorithms are carefully redesigned for the machine architecture, they realize 70 to 90% of the peak rate (*1*). There is an obvious payoff in learning systematic ways to design algorithms for parallel machines.

Bell anticipates that machines capable of 1 tflops and containing thousands (or even millions) of processors will be available as early as 1995 (*2*). For example, IBM Research is developing the Vulcan machine, which will consist of 32,768 ($2^{15}$) 50-mflops processors, and Thinking Machines Corporation is considering a Connection Machine with over a million ($2^{20}$) processors. These supermachines may cost on the order of $50 million apiece. Bell anticipates that low-cost, single-processor, reduced instruction set chips with speeds on the order of 20 mflops will be common in workstations by 1995. It is clear that tflops machines will be multicomputers consisting of large numbers of processing elements (processor plus memory) connected by a high-speed message exchange network. Smaller multicomputers will proliferate in the next 5 years: we must learn to program them.

Speed-up is a common measure of the performance gain from a parallel processor. It is defined as the ratio of the time required to complete the job with one processor to the time required to complete the job with $N$ processors (*3*). Perfect speed-up, a factor of $N$, can be attained in one of two ways. In a machine where each piece of the work is permanently assigned to its own processor, perfect speed-up is attained only when the pieces are computationally equal and processors experience no significant delays in exchanging information. In a machine where work can be dynamically assigned to available processors, it is attained as long as the number of pieces of work ready for processing is at least $N$.

In discussing speed-up, it is important to distinguish between problem size and computational work. Problem size measures the number of elements in the data space, and computational work measures the number of operations required to complete the solution. For example, an $N \times N$ square matrix occupies $N^2$ storage locations, and it takes about $N^3$ operations to form the product of two of these matrices. If $N$ is doubled, the storage requirement will be multiplied by four and the computational work by eight. Conversely, if the number of processors is doubled, two matrices of dimension 26% larger than $N$ can be multiplied in the same amount of time. This has important consequences for multiprocessors: there

P. J. Denning is a research fellow of the Research Institute for Advanced Computer Science, National Aeronautics and Space Administration, Ames Research Center, Moffett Field, CA 94035. W. F. Tichy is in the Computer Science Department of the University of Karlsruhe, Karlsruhe, Germany.