

6

The Role of Failure in Engineering Design

Consider the following scenario that describes the experience of two students working on the Peak-Performance Design Competition introduced in Chapter 2:

“The students had been working on their vehicle for almost a week. Having decided to enter the Peak-Performance Design Competition, they directed their design approach toward an articulated arm that would pick up the opposing vehicle and throw it off the track. They first outlined their design on paper and then tested it out on a computer-aided design (CAD) package available at one of their parent’s place of work. They built all the parts in the school machine shop, using their CAD drawings as a guide. A sketch of their design is shown in Figure 6.1. The students had just finished putting together all fifty eight machined pieces and were delighted to find that the arm worked perfectly on the very first try!”

6.1 HOW ENGINEERS LEARN FROM MISTAKES

Wouldn’t it be nice if engineering were so simple and fool-proof as the scenario depicted in the preceding, fictitious paragraph? In the real world, almost nothing works correctly the first time, and getting things to work perfectly almost always takes longer than expected. Fabricated parts

SECTIONS

- 6.1 How Engineers Learn from Mistakes
- 6.2 Classic Design Failures: Case Studies

OBJECTIVES

In this chapter, you will:

- Examine the role of failure in engineering design.
- Discuss classic design failures as case studies.
- Learn how to accept and utilize failure as part of the path to design success.




Figure 6.1. Hypothetical, impossible-to-build car design includes an articulated arm.

do not fit together, circuits have wiring errors, software modules have incompatibilities, and structural elements are incorrectly sized. Experienced engineers know that designs seldom work the first time and are never discouraged by initial failure.

Engineering tasks take longer than expected, because *failure* is an inevitable part of the design process. It's unreasonable to expect a new design to work the first time. When a device does not work as planned, it's a sure sign that something important has been overlooked. Perhaps two moving parts hit each other unexpectedly. Perhaps an electronic circuit doesn't work because stray effects were not included in the design model. A software program may fail because an unforeseen set of keystrokes leads to a logical dead end. A scale model of a bridge may reveal an overstressed support beam because a support pillar was omitted. A biomedical implant may be rejected because a tissue interaction was underestimated. Whatever the failure mode, it's better for a defect to appear *during* the design process than after it, when the device is in the field. Redesign after failure provides an important path to optimizing the final product and gives engineers needed time to correct deficiencies and work out bugs.*

A more realistic version of the paragraph that appeared at the start of this chapter might resemble the following:

"The students had been working on their Peak-Performance Design Competition vehicle for almost a month. Their first version of the car included an articulated arm designed to pick up the opposing vehicle and throw it off the track. After much trial and error, they succeeded in building an arm capable of lifting objects. At first, the students attempted to power the arm using the spring force from a single mousetrap, but after much testing, the students discovered that they had ignored frictional effects. Two mousetraps were required for adequate mechanical power. Eventually, they learned that powering the arm from a small electric motor and gear set provided much finer control of its movements. Their first attempt at assembly required that they return to the machine shop to redrill several holes they had put in the wrong place. Although the final version of the arm worked well on paper and also on a computer-aided design (CAD) package available at one parent's place of work, the arm failed miserably when mounted on the car. When an opposing vehicle was lifted into the air, the center of gravity fell outside the maximum allowed wheelbase of the vehicle, causing *both*

* The word "bug" was coined by Adm. Grace Hopper in 1945. In the early days of computers, logic gates were made from electromechanical relay switches rather than transistors. A moth flew inside one such computer, got stuck between two relay contacts, and prevented seemingly closed relay contacts from making a true connection. The computer malfunctioned because of the "bug."




Figure 6.2. More realistic battering ram design as offensive strategy.

vehicles to topple over and fall off the track. The students had to abandon their arm design and eventually settled on the battering ram of Figure 6.2 for their offensive strategy.”

6.2 CLASSIC DESIGN FAILURES: CASE STUDIES

The pages of engineering history are full of examples of design flaws that escaped detection in the design phase only to reveal themselves once the device was in actual use. Although all engineers are plagued by hidden design flaws from time to time, a few failure cases have become notorious because they affected many people, caused great property damage, or led to sweeping changes in engineering practice. In this section, we review several design failures from the annals of engineering lore. Each event involved the loss of human life or major destruction of property, and each was caused by engineering design failures. The mistakes were made by engineers who did the best they could, but had little prior experience or had major lapses in engineering judgement. After each incident, similar disasters were averted, because engineers were able to study the *causes* of the problems and establish new or revised engineering standards and guidelines. Studying these classic failures and the mistakes of the engineers who caused them will help you to avoid making similar mistakes in your own work.

The failure examples to follow all had dire consequences. Each occurred once the product was in use, long after the initial design, test, and evaluation phases. It's always better for problems to show up *before* the product has gone to market. Design problems can be corrected easily during testing, burn-in, and system evaluation. If a design flaw shows up in a product or system that has already been delivered for use, the consequences are far more serious. As you read the examples of this section, you might conclude that the causes of these failures in the field should have been obvious, and that failure to avoid them was the result of some engineer's carelessness. Indeed, it's relatively easy to play “Monday-morning quarterback” and analyze the cause of a failure *after* it has occurred. But as any experienced engineer will tell you, spotting a hidden flaw during the test phase is not always easy when a device or system is complex and has many parts or subsystems that interact in complicated ways. Even simple devices can be prone to hidden design flaws that elude the test and evaluation stages. Indeed, one of the marks of a good engineer is the ability to ferret out flaws and errors *before* the product finds its way to the end user. You can help to strengthen your abilities with the important intuitive skill of flaw detection by becoming familiar with the classic failure incidents discussed in this section. If you are interested in learning more details about any of the case studies, you might consult one of the references listed at the end of the chapter.

Case 1: Tacoma Narrows Bridge

The Tacoma Narrows Bridge, built across Puget Sound in Tacoma, Washington in 1940, was the longest suspension bridge of its day. The design engineers copied the structure


of smaller, existing suspension bridges and simply built a longer one. As had been done with countless shorter spans, support trusses deep in the structure of the bridge's framework were omitted to make it more graceful and visually appealing. No calculations were done to prove the structural integrity of a longer bridge lacking internal support trusses. Because the tried-and-true design methods used on shorter spans had been well tested, the engineers assumed that these design methods would work on longer spans. On November 7, 1940, during a particularly windy day, the bridge started to undulate and twist, entering into the magnificent torsional motion shown in Figure 6.3. After several hours, the bridge crumbled as if it were made from dry clay; not a piece remained between the two main center spans.

What went wrong? The engineers responsible for building the bridge had relied on calculations made for smaller bridges, even though the assumptions behind those calculations did not apply to the longer span of the Tacoma Narrows Bridge. Had the engineers heeded some basic scientific intuition, they would have realized that three-dimensional structures cannot be directly scaled upward without limits.

Case 2: Hartford Civic Center

The Hartford Civic Center was the first of its kind. At the time of its construction in the mid 1970's, no similar building had been built before. Its roof was made from a space frame structure of interconnected rods and ball sockets, much like a child's construction toy. Rods numbering in the hundreds were interconnected in a visually appealing geodesic pattern like the one shown in Figure 6.4. Instead of performing detailed hand calculations, the design engineers relied on the latest computer models to compute the loading on each individual member of the roof structure. Recall that computers in those days were much more primitive than those we enjoy today. The PC had not yet been invented, and all work was performed on large mainframe computers.

Figure 6.3. The Tacoma Narrows Bridge in torsional vibration.






Figure 6.4. Geodesic, rod-and-ball socket construction.

On January 18, 1978, just a few hours after the center had been filled to capacity with thousands of people watching a basketball game, the roof collapsed under a heavy snow load, demolishing the building. Miraculously, no one was hurt in the collapse.

Why did the collapse occur? Some attribute the failure to the engineers who designed the civic center and chose not to rely on their basic judgement and intuition gleaned from years of construction practice. Instead, they relied on computer models of their new space frame design. These computer models had been written by programmers, not structural engineers, during the days when computer modeling was in its infancy. The programmers based their code algorithms on structural formulas from textbooks. Not one of the programmers had ever actually built a roof truss. All failed to include basic derating factors at the structural joints to account for the slight changes in layout (e.g., minor variations in angles, lengths, and torsion) that occur when a complex structure is actually built. The design engineers trusted the output of computer models that never had been fully tested on actual construction. Under normal roof load, many ball-and-socket joints were stressed beyond their calculated limits. The addition of a heavy snow load to the roof load proved too much for the structure to bear.

Case 3: Space Shuttle Challenger

The NASA Space Shuttle Challenger blew up during launch on a cold day in January 1986 at Cape Kennedy (Canaveral) in Florida. Thousands witnessed the explosion as it happened. (See Figure 6.5). Hundreds of millions watched news tapes of the event for weeks afterwards. After months of investigation, NASA traced the problem to a set of O-rings used to seal sections of the multisegmented booster rockets. The seals were never designed to be operated in cold weather, and on that particular day, it was about 28°F (−2°C; a very cold day for Florida). The frozen O-rings were either too stiff to properly seal the sections of the booster rocket or became brittle and cracked due to the unusually cold temperatures. Flames spewed from an open seal during acceleration and ignited an adjacent fuel tank. The entire spacecraft blew up, killing all seven astronauts on board, including a high school teacher. It was the worst space disaster in U.S. history.

In using O-rings to seal adjacent cylindrical surfaces, such as those depicted in Figure 6.6, the engineers had relied on a standard design technique for rockets. The Challenger's booster rockets, however, were much larger than any on which O-rings had been used before. This factor, combined with the unusually cold temperature, brought the seal to its limit, and it failed.





Figure 6.5. The Space Shuttle Challenger explodes during launch. (Photo courtesy of RJS Associates.)

There was, however, another dimension to the failure. *Why* had the booster been built in multiple sections, requiring O-rings in the first place? The answer is complex, but the cause was largely attributable to one factor: The decision to build a multi-section booster was, in part, *political*. Had engineering common sense been the sole factor, the boosters would have been built in one piece without O-rings. Joints are notoriously weak spots, and a solid body is almost always stronger than a comparable

Figure 6.6. Schematic depiction of o-ring seals.



one assembled from sections. The manufacturing technology existed to build large, one-piece rockets of appropriate size. But a senator from Utah lobbied heavily to have the contract for constructing the booster rockets awarded to a company in his state. It was not physically possible to transport a large, one-piece booster rocket all the way from Utah to Florida over existing rail lines. Trucks were too small, and no ships were available that could sail to land-locked Utah, which lies in the middle of the United States. The decision by NASA to award the contract to the Utah company resulted in a multisection, O-ring-sealed booster rocket whose smaller pieces would easily be shipped by rail or truck.

Some say the catastrophe resulted from a lack of ethics on the part of the design engineers who suspected the O-ring design of having potential problems. Some say it was the fault of NASA for succumbing to political pressure from Congress, its ultimate funding source. Others say it was just an unusual convergence of circumstances, since neither the Utah senator nor the design engineers knowingly advocated for a substandard product. The sectioned booster had worked flawlessly on many previous shuttle flights that had not been launched in subfreezing temperatures. Still others say that by putting more weight on a political element of the project, rather than on pure engineering concerns, the engineers were forced into a less-than-desirable design concept that had never before been attempted on something so large.

Case 4: Kansas City Hyatt

If you've ever been inside a Hyatt hotel, you know that their internal architectures are very unique. The typical Hyatt hotel has cantilevered floors that form an inner trapezoidal atrium, and the walkways and halls are open, inviting structures. There's nothing quite like the inside of a Hyatt. In the case of the Kansas City Hyatt, first opened in 1981, the design included a two-layer, open-air walkway that spanned the entire lobby in midair, from one balcony to another. During a party that took place not long after the hotel opened, the walkway was filled with people dancing in time to the music. The weight and rhythm of the load of people, perhaps in resonance with the walkway, caused it to collapse suddenly. Over one hundred people died, and the event will be remembered forever in the history of hotel management. Although the hotel eventually reopened, to this day the walkway has never been rebuilt.

The collapse of the Hyatt walkway is a classic example of failure due to lack of construction experience. In this case, however, the error originated during the *design* phase, not the construction phase. In order to explain how the walkway collapsed, consider the sketch of the skeletal frame of the walkway, as specified by the design engineer, shown here in Figure 6.7.

Each box beam was to be held up by a separate nut threaded onto a suspended steel rod. The rated load for each nut-to-beam joint was intended to be above the maximum weight encountered during the time of the accident. What's wrong with this picture? The problem is that the structure as specified was not a realistic structure to build. The design called for the walkway's two decks to be hung from the ceiling by a single rod at each support point. The rods were made from smooth steel having no threads. Threading reduces the diameter of a rod, so it's impossible to get a nut to the middle of a rod unless the rod is threaded for at least half its length. In order to construct the walkway as specified, each rod would have to be threaded along about 20 feet of its length, and numerous of rods were needed for the long span of the walkway. Even with an electric threading machine, it would have taken days to thread all the needed rods. The contractor who actually built the walkway proposed a modification to the construction so that only the very ends of the rods would have to be threaded. The modification is illustrated in Figure 6.8.




Figure 6.7. Kansas City Hyatt walkway support structure as designed.

The problem with this modification is that the nut (A) at the lower end of the upper rod now had to support the weight of *both* walkways. A good analogy would be two mountain climbers hanging onto a rope. If both grabbed the rope simultaneously, but independently, the rope could hold their weight. If the lower climber grabbed the ankles of the upper climber instead of the rope, however, the upper climber's hands would have to hold the weight of *two* climbers. Under the full, or maybe excessive, load conditions of that day, the weight on nut (A) of the Hyatt walkway was just too much, and the joint gave way. Once the joint on one rod failed, the complete collapse of the rest of the joints and the entire walkway quickly followed.

Some attributed the fatal flaw to the inexperience of the design engineer who specified single rods requiring 20 feet of threading. Others blamed it on lack of experience. Yet others blamed both the junior engineer, who signed off on the modifications presented by the construction crew at the construction site, and the senior engineer,




Figure 6.8. Kansas City Hyatt walkway support structure as actually built.

who should have communicated to the junior engineer the critical nature of the rod structure as specified. Perhaps both engineers lacked seasoning—the process of getting their hands dirty on real construction problems as a way of gaining a feeling for how things are made in the real world.

Regardless of who was at fault, the design also left little room for *safety margins*. It's common practice in structural design to leave *at least* a factor-of-two safety margin between the calculated maximum load and the expected maximum load on a structure. The safety margin allows for inaccuracies in load calculations due to approximation, random variations in material strengths, and small errors in fabrication. Had the walkway included a safety margin of a factor of two or more, the doubly stressed joint on the walkway might not have collapsed, even given its modified construction. The design engineers specified a walkway structure that was possible, but not practical, to build. The construction supervisor, unaware of the structural implications, but wishing to see the job to completion, ordered a small, seemingly innocent, but ultimately fatal, change in the construction method. Had but one of the design engineers ever spent time working on a construction site, this shortcoming might have been discovered. Errors such as the one that occurred at the Kansas City Hyatt can be prevented by including workers from all phases of construction in the design process, ensuring adequate communication between all levels of employees, and adding far more than minimal safety margins where public safety is at risk.

Case 5: Three Mile Island

Three Mile Island was a large nuclear power plant in Pennsylvania (See Figure 6.9.) It was the sight of the worst nuclear accident in the United States and nearly comparable to the total meltdown at Chernobyl, Ukraine. Fortunately, the incident at Three Mile Island resulted in only a near miss at a meltdown, but it also led to the permanent shutting down and trashing of a billion-dollar electric power plant and significant loss of electrical generation capacity on the power grid in the eastern United States.

On the day of the accident, a pressure buildup occurred inside the reactor vessel. It was normal procedure to open a relief valve in such situations to reduce the pressure to safe levels. The valve in question was held closed by a spring and was opened by applying voltage to an electromagnetic actuator. The designer of the electrical control system had made one critical mistake. As suggested by the schematic diagram shown in Figure 6.10, indicator lights in the control room lit up when power was applied to or removed from the valve actuator coil, but the control panel gave no indication about the *actual* position of the valve. After a pressure-relief operation, the valve at Three Mile Island became stuck in the open position. Although the actuation voltage had been turned off and lights in the control room indicated the valve to be closed, it was actually stuck open. The mechanical spring responsible for closing the valve did not have enough force to overcome the sticking force. While the operators, believing the valve to be closed, tried to diagnose the problem, coolant leaked from the vessel for almost two hours. Had the operators known that the valve was open, they could have closed it manually or taken other corrective measures. In the panic that followed, however, the operators continually believed their control-panel indicator lights and thought that the valve was closed. Eventually the problem was contained, but not before a rupture nearly occurred in the vessel. Such an event would have resulted in a complete core meltdown and spewed radioactive gas into the atmosphere. Even so, damage to the reactor core was so severe that the plant was permanently shut down. It has never reopened.

The valve actuation system at Three Mile Island was designed with a poor human-machine interface. The ultimate test of such a system, of course, would be during an emergency when the need for absolutely accurate information would be critical. The operators assumed that the information they were receiving was accurate, while in reality it was not. The power plant's control panel provided the key information by inference, rather than by direct confirmation. A better design would have been one that

Figure 6.9. Three Mile Island power plant.







Figure 6.10. Valve indicator system as actually designed.

included an independent sensor that unambiguously verified the true position of the valve, as suggested by the diagram of Figure 6.11.

Case 6: USS Vincennes

The Vincennes was a U.S. missile cruiser stationed in the Persian Gulf during the Iran-Iraq war. On July 3, 1988, while patrolling the Persian Gulf, the Vincennes received two IFF (Identification: Friend or Foe) signals on its Aegis air-defense system. Aegis was the Navy's complex, billion-dollar, state-of-the-art information-processing system that displayed more information than any one operator could possibly hope to digest. Information saturation was commonplace among operators of the Aegis system. The Vincennes had received two IFF signals, one for a civilian plane and the other for a military plane. Under the pressure of anticipating a possible attack, the overstimulated operator misread the cluttered radar display and concluded that only one airplane was approaching the Vincennes. Repeated attempts to reach the nonexistent warplane by radio failed. The captain concluded that his ship was under attack and made the split-second decision to have the airplane shot down. Two hundred and ninety civilians died needlessly.

Figure 6.11. Valve indicator system as it should have been designed and built.



What caused this catastrophic outcome? Was it bad military judgement? Was it an operating error? Were the engineers who designed the system at fault? The Navy officially attributed the accident to “operator error” by an enlisted sailor, but in some circles the blame was placed on the engineers who had designed the system. Under the stress of possible attack and deluged with information, the operator simply could not cope with an ill-conceived human-machine interface designed by engineers. Critical information, being needed most during crisis situations, should have been uncluttered and easy to interpret. The complex display of the Aegis system was an example of something that was designed just because it was technically possible. It resulted in a human-machine interface that became a weak link in the system.

Case 7: Hubble Telescope

The Hubble is an orbiting telescope that was put into space at a cost of over a billion dollars. Unaffected by the distortion experienced by ground-based telescopes due to atmospheric turbulence, the Hubble has provided spectacular photos of space and has made possible numerous astronomical discoveries. Yet the Hubble telescope did not escape design flaws. Of the many problems that plagued the Hubble during its first few years, the most famous was its improperly fabricated mirrors. They were distorted and had to be corrected by the installation of an adaptive optic mirror that compensated for aberrations. The repairs were carried out by a NASA Space Shuttle crew. Although this particular flaw is the one most often associated with the Hubble, it was attributed to sloppy mirror fabrication rather than to a design error. Another, less-well-known design error more closely illustrates the lessons of this chapter. The Hubble’s solar panels were deployed in the environment of space, where they were subjected to alternate heating and cooling as the telescope moved in and out of the earth’s shadow. The resulting expansion and contraction cycles caused the solar panels to flap like the wings of a bird. Attempts to compensate for the unexpected motion by the spacecraft’s computer-controlled stabilizing program led to a positive feedback effect which only made the problem worse. Had the design engineers anticipated the environment in which the telescope was to be operated, they could have compensated for the heating and cooling cycles and avoided the problem. This example illustrates that it’s difficult to anticipate all the conditions under which a device or system may be operated. Nevertheless, extremes in operating environment often are responsible for engineering failures. Engineers must compensate for this problem by testing and *retesting* devices under different temperatures, load conditions, operating environments, and weather conditions. Whenever possible (though obviously not possible in the case of the Hubble), a system should be developed and tested in as many different environmental conditions as possible if a chance exists that those conditions will be encountered in the field.

Case 8: De Haviland Comet

The De Haviland Comet was the first commercial passenger jet aircraft. A British design, the Comet enjoyed many months of trouble-free flying in the 1950’s until several went down in unexplained crashes. Investigations of the wreckages suggested that the fuselages of these planes had ripped apart in mid flight. For years, the engineers assigned the task of determining the cause of the crashes were baffled. What, short of an explosion, could have caused the fuselage of an aircraft to blow apart in flight? No evidence of sabotage was found at any of the wreckage sites. After some time, the cause of the crashes was discovered. No one had foreseen the effects of the numerous pressurization and depressurization cycles that were an inevitable consequence of takeoffs and landings. Before jet aircraft, lower altitude airplanes were not routinely operated under pressure. Higher altitude jet travel brought with it the need to pressurize the cabin. In the case of the Comet, the locations of the rivets holding in the windows developed

fatigue cracks, which, after many pressurization and depressurization cycles, grew into large, full-blown cracks in the fuselage. This mode of failure is depicted in Figure 6.12.

Had the design engineers thought about the environment under which the finished product would be used, the problem could have been avoided. Content instead with laboratory stress tests that did not mimic the actual pressurization and depressurization cycles, the engineers were lulled into a false sense of security about the soundness of their design. This example of failure again underscores an important engineering lesson: Always test a design under the most realistic conditions possible. Always assume that environmental conditions will affect performance and reliability.


Preparing for Failure in Your Own Design

First-time designs often betray previously hidden flaws after an initial period of successful use. Design flaws eventually show up because the operating environment changes, a previously untried sequence of events occurs, or weak points in the design encounter repeated stress. Sometimes, failure occurs just because of plain old statistics. As the saying goes, “If something is bound to fail, it *will* fail sooner or later.” After a failure occurs, it’s the engineer’s job to determine the cause, fix what’s wrong, and begin tests again. At the same time, it’s up to the design engineer to identify as many of the bugs and weaknesses as possible during the early phases of the design process. Thorough testing and retesting under all sorts of operating conditions is essential. An unsuccessful first prototype presents an excellent opportunity to discover and weed out bugs before the final version of the device is put to market.

In the commercial sector, the rush to bring a product to market ahead of competitors puts pressure on the engineer to complete the test and evaluation phases as quickly as possible. For this reason, many consumer products, including automobiles, computers, and software, develop problems soon after they are released. If you purchase one of these items during the first year of issue, expect to find bugs and weaknesses that were not discovered on the factory floor.

Despite this admonition, you should be eager to apply your engineering design skills to new technology and innovation. If all engineers were content to stay with tried-and-true designs, technical progress would cease. Understanding when to stay within the

Figure 6.12. Stress cracks around the window rivets of the DeHaviland Comet.



bounds of a traditional design and when to move on to new creative frontiers requires experience, practice, and intuition. When you encounter failure in your own design projects, do not be discouraged. Recognize failure as an inevitable part of the design process, and use it to learn, discover, and expand your foundation of engineering knowledge.

KEY TERMS

Failure
Safety margin

Flaw

Bugs

REFERENCES

- J. FELD, KENNETH and L. Carper, *Construction Failure—Wiley Series of Practical Construction Guides*, New York: John Wiley & Sons, 1996.
- E. S. FERGUSON, "How Engineers Lose Touch," *Invention and Technology*. Vol. 8 (3), Winter 1993, pp.16–24.
- H. PETROSKI, *To Engineer is Human: The Role of Failure in Successful Design*. New York: Vintage Books, 1992.
- H. PETROSKI, *Design Paradigms: Case Histories of Error and Judgement in Engineering*. Cambridge: Cambridge Univ. Press, 1994.
- D. D. A. PIESOLD, *Civil Engineering Practice: Engineering Success by Analysis of Failure*. New York: McGraw-Hill, 1991.
- R. UHL, G. M. DAVIDSON, and K. A. ESAKLUL, *Handbook of Case Histories in Failure Analysis*. ASM International, 1992.
- P. VISWANADHAM and P. SINGH, *Failure Modes and Mechanisms in Electronic Packages*. Chapman & Hall, 1997.

Problems

1. Identify a product or system in your own experience that has failed. Write a short summary of the cause of the failure and how you might improve upon the design.

Look up and write a synopsis of the following classic engineering failure incidents:

2. Exxon Bayway refinery, Linden, New Jersey (1990)
3. General Electric rotary compressor refrigerators (1990)
4. Green Bank radio telescope (1989)
5. Union Carbide chemical leak, Bhopal, India (1984)
6. Korean Airlines Flight 007 (1983)
7. Interstate 95 Bridge, Mianus River, Connecticut (1983)
8. Alexander L. Kielland oil platform, North Sea (1980)
9. American Airlines DC-10 (1979)
10. Skylab (1979)
11. The New York City Power Blackout (1976)
12. The windows in the John Hancock Tower, Boston, Massachusetts (1976)
13. Big Ben, London (1976)
14. Bay Area Rapid Transit (BART) (1973)
15. Point Pleasant Bridge, Ohio River, Ohio-West Virginia (1967)
16. The Apollo 1 capsule fire (1967)
17. The Great Northeast Power Blackout (1965)
18. Quebec City bridge (1907)
19. The Johnstown flood (1889)
20. Liberty Bell, Philadelphia (1835)