White Paper on Approaches to Safety Engineering*

Nancy Leveson

April 23, 2003

> *A life without adventure is likely to be unsatisfying, but a life in which adventure is allowed to take whatever form it will, is likely to be short.*
>
> — Bertrand Russell

This white paper lays out some foundational information about different approaches to safety: how various industries differ in their approaches to safety engineering, and a comparison of three general approaches to safety (system safety, industrial safety engineering, and reliability engineering). An attempt is made to lay out the properties of industries and systems that make one approach more appropriate than another.

How do industries differ in their approaches to safety engineering?

While the concern about industrial safety dates back to at least the turn of the century (and before in some countries and industries) and individual efforts to design safe products and systems also goes back in time, rigorous and defined approaches to safety engineering mostly arose after World War II, when the AEC (and later the NRC) were engaged in a public debate about the safety of nuclear power; civil aviation was trying to convince a skeptical public to fly; the chemical industry was coping with larger plants, increasingly lethal chemicals, and heightened societal concern about pollution; and the DoD was developing ballistic missile systems and increasingly dangerous weapons. Each of these parallel efforts resulted in very different engineering approaches, mostly because the problems they needed to solve were different.

Commercial Aircraft

> *The [FAA] administrator was interviewed for a documentary film on the [Paris DC-10] accident. He was asked how he could still consider the infamous baggage door safe, given the door failure proven in the Paris accident and the precursor accident at Windsor, Ontario. The Administrator replied—and not facetiously either—'Of course it is safe, we certified it.'*
>
> — C.O. Miller
> *A Comparison of Military and*
> *Civilian Approaches to Aviation Safety*

In commercial aircraft, safety is assured by first identifying hazards and then performing a fault hazard analysis—the hazards are traced to the aircraft components and each is assigned a reliability target such that the aircraft as a whole will reach the FAA failure rate requirements. Then the components are designed and manufactured using these allocated reliability rates. The rates are

*Parts of this paper are taken from my book, *Safeware*, published in 1995 by Addison-Wesley.

assured by using a high degree of single element integrity, fail-safe design (using redundancy and other design approaches to handle single or multiple component failures), and careful *fly-fix-fly* approaches where designs are modified to prevent previous causes of accidents. One reason for the success of this approach is that commercial aircraft designs do not change dramatically over time and learning from the past is therefore very effective. The commercial aircraft industry is very conservative in their design approaches and does not usually push the technology envelope. When new technology has been introduced, such as fly-by-wire and glass cockpits in aviation, increased accident rates have resulted on these high-tech aircraft and the accident mechanisms have changed (e.g., pilots are making different types of errors). Another characteristic of the commercial aircraft industry that affects safety is that it is very tightly regulated.

Nuclear Power

> *"To me that is probably one of the most significant learnings of the whole accident [TMI] the degree to which the inadequacies of that experience feedback loop ... significantly contributed to making us and the plant vulnerable to this accident."*
>
> — Herman Dieckamp
> President of the utility at TMI

Although the terminology differs between countries, *design basis accidents* for nuclear power plants in the U.S. define the set of disturbances against which nuclear power plants are evaluated. Licensing is based on the identification and control of hazards under normal circumstances, and the use of shutdown systems to handle abnormal circumstances. Safety assurance is based on the use of multiple, independent barriers (*defense in depth*), a high degree of single element integrity, and the provision that no single failure of any active component will disable any barrier. With this defense-in-depth approach to safety, an accident requires a disturbance in the process, a protection system that fails, and inadequate or failing physical barriers. These events are assumed to be statistically independent because of differences in their underlying physical principles: A very low calculated probability of an accident can be obtained as a result of this independence assumption. The substitution of software for physical devices invalidates this assumption, which has slowed down the introduction of computers (although it has increased in the last few years). Again, the industry is tightly regulated by the government.

There are many similarities between the commercial aircraft and nuclear power approaches to safety: Both take a very conservative approach to introducing new technology and designs. Both concentrate on component failure as the cause of accidents. The fly-fix-fly approach to learning from experience for both is very effective because basic designs change extremely slowly over time. Because their approaches to safety are very reliability oriented, they both tend to rely heavily on redundancy. A significant difference is that a nuclear power plant can rely on shutdown as a primary safety measure, whereas shutting down systems on a plane is not always an option (i.e., the fail-safe states are different).

The Chemical Process Industry

> *My company has had a safety program for 150 years. The program was instituted as a result of a French law requiring an explosives manufacturer to live on the premises with his family.*
>
> — Crawford Greenwalt
> Former president of Dupont[1]

[1]Quoted in William Johnson, *MORT Safety Assurance Systems*, Marcel Dekker, Inc., 1980.

The chemical industry differs from the prior two in that it is not government regulated to the same extent. Instead, the process industry approach to safety was by insurance needs—the term commonly used in the industry, *loss prevention*, reflects these origins. Loss, in this case, refers to the financial loss of damaged plant, third party claims, and lost production.Frank P. Lees, Loss Prevention in the Process Industries, Vol. 1 and 2, Butterworths, 1980.. The three traditional hazards in the industry—fire, explosion, and toxic release—have remained virtually unchanged in their nature for many years. Design and operating procedures to eliminate or control these hazards have evolved and been incorporated into codes and standards. As the chemical and petrochemical industries began to grow in complexity, size, and the use of new technology, there has been a concomitant increase in the consequences of accidents and environmental concerns and major accidents have increased political pressure for legislation to enact controls on the industry.

Although the chemical industry does use some standard reliability and hazard analysis techniques, the unique aspects of the application have led to the development of industry-specific techniques (e.g., the Dow and Mond Indexes and HAZOP). Hazard analysis on chemical plants is often done late in the design process or on existing plant or on designs where the only alternative for controlling hazards without costly design changes or retrofits is to add on protective devices. There have been some, e.g., Trevor Kletz, who have argued strongly for moving the industry toward a system safety approach, but component reliability and protection systems (called *safety systems* in the industry) are still the main emphasis.

Defense and Military Aviation

At the same time as civil aviation and the process industries (including nuclear power) were developing their approaches to safety after World War II, system safety was developing in the defense industry. Although many of the basic concepts of system safety, such as anticipating hazards and accidents and building in safety, predate the post–World War II period, much of the early development of system safety as a separate discipline began with flight engineers immediately after World War II. The Air Force had long had problems with aircraft accidents. For example, from 1952 to 1966, it lost 7715 aircraft, in which 8547 persons, including 3822 pilots, were killed[2] Most of these accidents were blamed on pilots. Many industry flight engineers, however, did not believe the cause was so simple: They argued that safety must be designed and built into aircraft just as are performance, stability, and structural integrity.

Seminars were conducted by the Flight Safety Foundation, headed by Jerome Lederer (who would later head the NASA Apollo safety program), that brought together engineering, operations, and management personnel. It was in 1954, at one of these seminars, that the term "system safety" may have first been used—in a paper by one of the aviation safety pioneers, C.O. Miller, titled "Applying Lessons Learned from Accident Investigations to Design Through a Systems Safety Concept." Around the same time, the Air Force began holding symposiums that fostered a professional approach to safety in propulsion, electrical, flight control, and other aircraft subsystems, but they did not at that time treat safety as a system problem.

When the Air Force began to develop intercontinental ballistic missiles (ICBMs), there were no pilots to blame for accidents, yet the liquid-propellant missiles blew up frequently and with devastating results. The Department of Defense and the Atomic Energy Commission were also facing the problems of building and handling nuclear weapons and finding it necessary to establish rigid controls and requirements on nuclear materials and weapon design.

System safety itself arose out of these ballistic missile programs. In the fifties, when the Atlas and Titan ICBMs were being developed, intense political pressure was focused on building a nuclear warhead with delivery capability as a deterrent to nuclear war. On these first missile projects,

[2]Willie Hammer, *Product Safety Management and Engineering*, Prentice Hall, 1972.

system safety was not identified and assigned as a specific responsibility. Instead, as was usual at the time, each designer, manager, and engineer was assigned responsibility for safety. These projects, however, involved advanced technology and much greater complexity than had previously been attempted, and the drawbacks of the standard approach to safety became clear when many interface problems went unnoticed until it was too late.

Within 18 months after the fleet of 71 Atlas F missiles became operational, four blew up in their silos during operational testing. The missiles also had an extremely low launch success rate. An Air Force manual describes several of these accidents:

> An ICBM silo was destroyed because the counterweights, used to balance the silo elevator on the way up and down in the silo, were designed with consideration only to raising a fueled missile to the surface for firing. There was no consideration that, when you were not firing in anger, you had to bring the fueled missile back down to defuel. The first operation with a fueled missile was nearly successful. The drive mechanism held it for all but the last five feet when gravity took over and the missile dropped back. Very suddenly, the 40-foot diameter silo was altered to about 100-foot diameter.
>
> During operational tests on another silo, the decision was made to continue a test against the safety engineer's advice when all indications were that, because of high oxygen concentrations in the silo, a catastrophe was imminent. The resulting fire destroyed a missile and caused extensive silo damage. In another accident, five people were killed when a single-point failure in a hydraulic system caused a 120-ton door to fall.
>
> Launch failures were caused by reversed gyros, reversed electrical plugs, bypass of procedural steps, and by management decisions to continue, in spite of contrary indications, because of schedule pressures (from the Air Force *System Safety Handbook for Acquisition Managers*, Air Force Space Division, January 1984.)

Not only were the losses themselves costly, but the resulting investigations detected serious safety deficiencies in the system that would require extensive modifications to correct. In fact, the cost of the modifications would have been so high that a decision was made to retire the entire weapon system and accelerate deployment of the Minuteman missile system.[3]

When the early aerospace accidents were investigated, it became apparent that the causes of a large percentage of them could be traced to deficiencies in design, operations, and management. The previous "fly–fix–fly" approach was clearly not adequate. In this approach, investigations were conducted to reconstruct the causes of accidents, action was taken to prevent or minimize the recurrence of accidents with the same cause, and eventually these preventive actions were incorporated into standards, codes of practice, and regulations. Although the fly–fix–fly approach was effective in reducing the repetition of accidents with identical causes, it became clear to the Department of Defense (DoD), and later to others, that it was too costly and, in the case of nuclear weapons, unacceptable. to prevent accidents before they occur the first time.[4] This realization led to the adoption of system safety approaches to try to prevent accidents before they occur the first time.

[3]William P. Rogers, Introduction to System Safety Engineering, John Wiley & Sons, 1971

[4]This lesson seems to require repeated relearning. As an example, consider the placing of a Milstar satellite in an incorrect and unusable orbit by a Titan IV B-32/Centaur launch in 1999. As happened with Challenger, the Titan program office decided that system safety, system engineering, and mission assurance could be reduced during operations because it was felt they were no longer needed or the resources were needed more elsewhere. The risk analysis that was done was not based on the steps critical to mission success but instead only considered the problems that had occurred in previous launches. One of the critical factors in the loss (software constant generation) was considered to be low risk because there had been no previous problems with it. The loss—$800 million for the satellite and an additional $433 million for the launcher—was one of the most costly unmanned losses in the history of Cape Canaveral Launch Operations.

The first military specification on system safety was published by the Air Force (Ballistic Systems Division) in 1962, and the Minuteman ICBM became the first weapon system to have a contractual, formal, disciplined system safety program. From that time on, system safety received increasing attention, especially in Air Force missile programs where testing was limited and accident consequences serious. The Army soon adopted system safety programs because of the many personnel it was losing in helicopter accidents, and the Navy followed suit. In 1966, the DoD issued a single directive requiring system safety programs on all development or modification contracts. The one exception to the use of this approach has been the Navy submarine nuclear reactor program, which has always followed a different approach to safety due to its unique requirements and history, although system safety as specified by MIL-STD-882 is used in the non-nuclear aspects of Navy submarine engineering.

At first, there were few techniques that could be used on these complex defense systems. But, step by step, the specialized safety engineering and operational safety practices that had evolved over the years were integrated with scientific, technical, and management techniques that were newly developed or adapted from other activities. Particular emphasis was placed on hazard analysis techniques, such as fault trees, which were first developed to cope with complex programs such as Minuteman.

The first system safety standard, MIL-STD-882 was issued in June 1969 and a system safety program became mandatory on all DoD-procured products and systems until Perry abolished all military standards in 1994 in favor of using commercial standards. Because there was no comparable commercial standard and there was a clear need for a system safety standard, MIL-STD-882 was readopted a few years ago, but in a watered-down form due to the anti-standard culture in the defense industry at the time.

The space program was the second major application area to apply system safety approaches in a disciplined fashion. Until the *Apollo 204* fire in 1967 at Cape Kennedy, in which three astronauts were killed, NASA safety efforts had focused on industrial worker safety. The accident alerted NASA, and they commissioned the General Electric Company at Daytona Beach, among others, to develop policies and procedures that became models for civilian aerospace safety activities.[5] Jerome Lederer was hired to head manned space-flight safety and, later, all NASA safety efforts. Through his leadership, an extensive program of system safety was set up for space projects, much of it patterned after the Air Force and DoD programs. Many of the same engineers and companies that had established formal system safety programs for DoD contracts also were involved in space programs, and the technology and management activities were transferred to this new application.

It is not surprising that defense and space should have chosen the same approach to safety—they share a lot of common features that differ with the other industries described: tremendously complex designs stretching the limits of current engineering techniques, requirements for new and innovative designs to meet mission goals, continual introduction of new and unproven technology, and limitations in the ability to test.

What is System Safety?

System safety uses systems theory and systems engineering approaches to prevent foreseeable accidents and to minimize the result of unforeseen ones. Losses in general, not just human death or injury, are considered. Such losses may include destruction of property, loss of mission, and environmental harm.

The primary concern of system safety is the management of hazards: their identification, evaluation, elimination, and control through analysis, design and management procedures. Mueller, in 1968, described the then new discipline of system safety engineering as "organized common sense"

[5] C.O. Miller, The Broader Lesson from Challenger, *Hazard Prevention*, 1986

.[6] It is a planned, disciplined, and systematic approach to identifying, analyzing, and controlling hazards throughout the life cycle of a system in order to prevent or reduce accidents.

System safety activities start in the earliest concept development stages of a project and continue through design, production, testing, operational use, and disposal. One aspect that distinguishes system safety from other approaches to safety is its primary emphasis on the early identification and classification of hazards so that corrective action can be taken to eliminate or minimize those hazards before final design decisions are made.

Although system safety is a relatively new discipline and still evolving, some general principles are constant throughout its various manifestations and distinguish it from other approaches to safety and risk management.

- *System safety emphasizes building in safety, not adding it on to a completed design:* Safety considerations must be part of the initial stage of concept development and requirements definition: From 70 to 90 percent of the design decisions that affect safety will be made in these early project phases.[7] The degree to which it is economically feasible to eliminate a hazard rather than to control it depends upon the stage in system development at which the hazard is identified and considered. Early integration of safety considerations into the system development process allows maximum safety with minimal negative impact. The alternative is to design the plant, identify the hazards, and then add on protective equipment to control the hazards when they occur—which is usually more expensive and less effective.

- *System safety deals with systems as a whole rather than with subsystems or components:* Safety is an emergent property of systems, not a component property. One of the principle responsibilities of system safety is to evaluate the interfaces between the system components and determine the effects of component interaction, where the set of components includes humans, machines, and the environment.

- *System safety takes a larger view of hazards than just failures:* Hazards are not always caused by failures, and all failures do not cause hazards. Serious accidents have occurred while system components were all functioning exactly as specified—that is, without failure. If failures only are considered in a safety analysis, many potential accidents will be missed. In addition, the engineering approaches to preventing failures (increasing reliability) and preventing hazards (increasing safety) are different and sometimes conflict.

- *System safety emphasizes analysis rather than past experience and standards:* Standards and codes of practice incorporate experience and knowledge about how to reduce hazards, usually accumulated over long periods of time and resulting from previous mistakes. While such standards and learning from experience are essential in all aspects of engineering, including safety, the pace of change today does not always allow for such experience to accumulate and for proven designs to be used. System safety analysis attempts to anticipate and prevent accidents and near-accidents *before* they occur.

- *System safety emphasizes qualitative rather than quantitative approaches:* System safety places major emphasis on identifying hazards as early as possible in the design stage and then designing to eliminate or control those hazards. At these early stages, quantitative information usually does not exist. Although such quantitative information would be useful in prioritizing hazards, subjective judgments about the likelihood of a hazard are usually adequate and all that is possible at the time that design decisions must be made.

[6] Jerome Lederer. How Far Have We Come? A Look Back at the Leading Edge of System Safety Eighteen Years Ago, *Hazard Prevention*, 1986

[7] William Johnson. *MORT Safety Assurance Systems*, Marcel Dekker, 1980

- *Recognition of tradeoffs and conflicts:* Nothing is absolutely safe, and safety is not the only, and is rarely the primary, goal in building systems. Most of the time, safety acts as a constraint on the possible system designs and may conflict with other design goals such as operational effectiveness, performance, ease of use, time, and cost. System safety techniques and approaches focus on providing information for decision making about risk management tradeoffs.

- *System safety is more than just system engineering:* System safety engineering is an important part of system safety, but the concerns of system safety extend beyond the traditional boundaries of engineering. In 1968, Jerome Lederer, then the director of the NASA Manned Flight Safety Program for Apollo wrote:

  > System safety covers the total spectrum of risk management. It goes *beyond the hardware* and associated procedures of system safety engineering. It involves: attitudes and motivation of designers and production people, employee/management rapport, the relation of industrial associations among themselves and with government, human factors in supervision and quality control, documentation on the interfaces of industrial and public safety with design and operations, the interest and attitudes of top management, the effects of the legal system on accident investigations and exchange of information, the certification of critical workers, political considerations, resources, public sentiment and many other non-technical but vital influences on the attainment of an acceptable level of risk control. These non-technical aspects of system safety cannot be ignored.

Using these general principles, system safety attempts to manage hazards through analysis, design, and management procedures. Key activities include top-down system hazard analyses (starting in the early concept design stage to eliminate or control hazards and continuing during the life of the system to evaluate changes in the system or the environment), documenting and tracking hazards and their resolution (establishing audit trails); designing to eliminate or control hazards and minimize damage, maintaining safety information systems and documentation; and establishing reporting and information channels.

What is the difference between Industrial Safety and System Safety?

Industrial, or occupational has traditionally focused primarily on controlling injuries to employees on the job. The industrial safety engineer usually is dealing with a fixed manufacturing design and hazards that have existed for a long time, many of which are accepted as necessary for operations. More emphasis is often placed on teaching employees to work within this environment than on removing the hazards.

Industrial safety engineers collect data during the operational life of the system and eliminate or control unacceptable hazards if possible or practical. When accidents occur, they are investigated and action is taken to reduce the likelihood of a recurrence—either changing the plant or changing employee work rules and training. The hazards associated with high-energy or dangerous processes are usually controlled either (1) by disturbance control algorithms implemented by operators or an automated control system or (2) by transferring the plant to a safe state using a separate protection system.

Safety reviews and audits are conducted by industrial safety divisions within the company or by safety committees to ensure that unsafe conditions in the plant are corrected and that employees are following the work rules specified in manuals and instructions. Lessons learned from accidents are incorporated into design standards, and much of the emphasis in the design of new plants and work

rules is on implementing these standards. Often, the standards are enforced by the government through occupational safety and health legislation.

In contrast, system safety is concerned primarily with new systems. The concept of loss is treated much more broadly: Relevant losses may include injury to nonemployees; damage to equipment, property, or the environment; and loss of mission. As has been seen, instead of making changes as a result of operational experience with the system, system safety attempts to identify potential hazards before the system is designed, to define and incorporate safety design criteria, and to build safety into the design before the system becomes operational. Although standards are used in system safety, they usually are process rather than product standards—reliance on design or product standards is often inadequate for new types of systems, and more emphasis is placed on upfront analysis and designing for safety.

There have been attempts to incorporate system safety techniques and approaches into traditional industrial safety programs, especially when new plants and processes are being built. Although system safety techniques are considered "overkill" for many industrial safety problems, larger plants and increasingly dangerous processes have raised concern about injuries to people outside the plant and about pollution and have made system safety approaches more relevant. Furthermore, with the increase in size and cost of plant equipment, changes and retrofits to increase safety are costly and may require discontinuing operations for a period of time.

From the other side, system safety is increasingly considering issues that have been traditionally thought to be industrial safety concerns. In some cases, the neglect of these issues has caused serious losses:

> Over a period of two years, a contractor experienced 26 satellite damaging mishaps during manufacturing! Twice they hit it with a forklift. Twice more they hit it with a crane hook. Wrenches were dropped into the satellite. Makeshift workstands failed. It appeared as if there were forces bent on destroying the satellite before it got to the launch site. Investigation revealed that the System Safety activity never had addressed the manufacturing phase of the program because the phase was covered by existing industrial safety activities.[8]

In summary, industrial safety activities are designed to protect workers in an industrial environment; extensive standards are imposed by federal codes or regulations providing for a safe workplace. However, few, if any, of these codes apply to protection of the product being manufactured. With the relatively recent introduction of robots into the workplace environment and with long-lived engineering programs like the Space Shuttle that have substantial continuing complex engineering design activities, the traditional concerns of industrial safety and system safety have become more intertwined.

How do Reliability Engineering and System Safety differ?

In the years following World War II, the growth in military electronics gave rise to reliability engineering. Reliability was also important to NASA and our space efforts, as evidenced by the high failure rate of space missions in the late 1950s and early 1960s.

Reliability engineering is concerned primarily with failures and failure rate reduction. The reliability engineering approach to safety thus concentrates on failure as the cause of accidents. Reliability engineering uses a variety of techniques to minimize component failures and thereby the failures of complex systems caused by component failure, including parallel redundancy, standby sparing, built-in safety factors and margins, derating, screening, and timed replacements.

[8] *System Safety Handbook for the Acquisitions Manager*, Air Force Space Division, 1987

While these techniques are often effective in increasing reliability, they do not necessarily increase safety. In fact, their use under some conditions may actually reduce safety. For example, increasing the burst-pressure to working-pressure ratio of a tank often introduces new dangers of an explosion or chemical reaction in the event of a rupture. System safety hazard analyses look at these interactions and not just at failures or engineering strengths.

Reliability engineers often assume that reliability and safety are synonymous, but this assumption is true only in special cases. In general, safety has a broader scope than failures, and failures may not compromise safety. There is obviously an overlap between reliability and safety, but many accidents occur without any component failure—the individual components were operating exactly as specified or intended, that is, without failure. The opposite is also true—components may fail without a resulting accident.

Accidents may be caused by equipment operation outside the parameters and time limits upon which the reliability analyses are based. Therefore, a system may have high reliability and still have accidents. In addition, generalized probabilities and reliability analyses may not apply to specific, localized areas: The probability of a bird strike causing an aircraft accident, for example, is much higher at Midway Island than at most other places. Most important, accidents are often not the result of a simple combination of component failures.

Safety is an emergent property that arises at the system level when components are operating together. The events leading to an accident may be a complex combination of equipment failure, faulty maintenance, instrumentation and control problems, human actions, and design errors. Reliability analysis considers only the possibility of accidents related to failures; it does not investigate potential damage that could result from *successful* operation of the individual components.

Consider an accident that occurred in a batch chemical reactor in England. The design of this system is shown in Figure 1. The computer was responsible for controlling the flow of catalyst into the reactor and also the flow of water into the reflux condenser to cool off the reaction. Additionally, sensor inputs to the computer were supposed to warn of any problems in various parts of the plant. The programmers were told that if a fault occurred in the plant, they were to leave all controlled variables as they were and to sound an alarm. On one occasion, the computer received a signal telling it that there was a low oil level in a gearbox. The computer reacted as the requirements specified: It sounded an alarm and left the controls as they were. By coincidence, a catalyst had just been added to the reactor, and the computer had just started to increase the cooling-water flow to the reflux condenser; the flow was therefore kept at a low rate. The reactor overheated, the relief valve lifted, and the contents of the reactor were discharged into the atmosphere.

This accident involved a sequence of events, none of which was a component failure. The individual components worked as specified, but together they created a hazardous system state. Reliability uses a bottom-up approach (e.g., FMEA) to evaluate the effect of component failures on system function, while safety requires a top-down approach that evaluates how hazardous states can occur from a combination of both incorrect and correct component behavior, such as proper behavior of a component at an improper time or under the wrong environmental conditions.

Care needs to be taken when applying reliability assessment techniques to safety. Since accidents are not necessarily caused by events that can be measured this way, it should not be used as a measure of risk. Reliability assessment measures the probability of random failures—not the probability of hazards or accidents. Also, if a design error is found in a system, safety will be more effectively enhanced by removing the design error than by measuring it in order to convince someone that it will never cause an accident. In the case of the batch reactor, if the scenario that unfolded had been known and therefore could have been included in a system reliability assessment, then the engineers would simply have changed the design to require opening the water valve before the catalyst valve, rather than attempting to measure its probability. High reliability numbers do not guarantee safety, and safety need not require ultrahigh reliability.

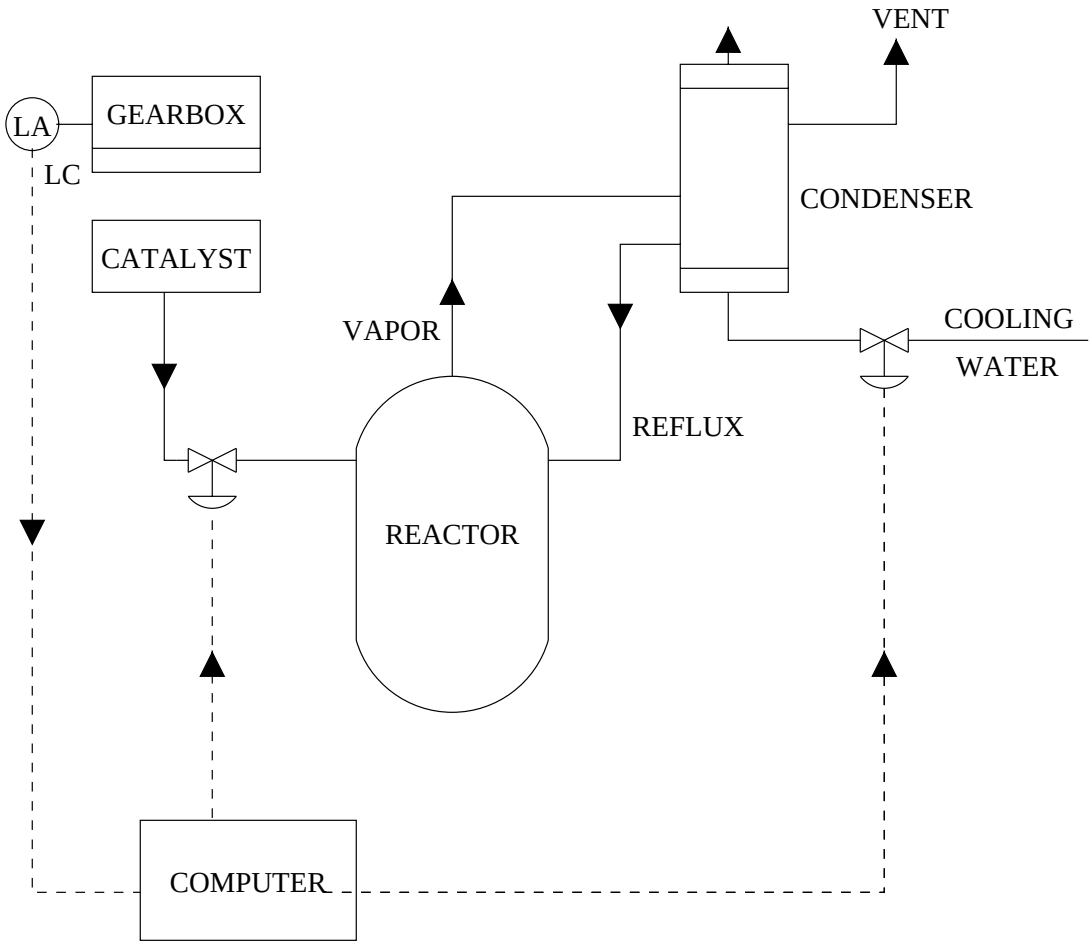

Figure 1: A chemical reactor design. (Source: Trevor Kletz, Human problems with computer control, *Plant/Operations Progress*, 1(4), October 1982. page 6.