RANDOM, SYSTEMATIC, AND COMMON CAUSE FAILURE: HOW DO YOU MANAGE THEM?

Angela Summers, Ph.D., P.E., President, and Michela Gentile, Ph.D., Design Consultant,
SIS-Tech Solutions, LP

"Failure Conundrum," Michela Gentile, PhD. And Angela Summers, PhD, PE, 8th Annual Symposium, Mary Kay O'Connor Process Safety Center "Beyond Regulatory Compliance: Making Safety Second Nature," Texas A&M University, College Station, Texas, October 24-27, 2005.

"Failure Conundrum," Michela Gentile, PhD. And Angela Summers, PhD, PE, American Institute of Chemical Engineers 2006 Spring National Meeting, Orlando, Florida, April 26-27, 2006.

"Random, Systematic, and Common Cause Failure: How do you manage them?" Process Safety Progress, December 2006.

ABSTRACT

A safety instrumented system (SIS) may fail to operate as desired when one or more of its devices fail due to random, systematic, and common cause events. IEC 61511 (ANSI/ISA 84.00.01-2004) stresses the importance of minimizing the propagation of device failure into system failure through design, operating, inspection, and maintenance practices. To fully understand the lifecycle requirements, it is first necessary to understand the types of failures and their potential effects on the SIS. Although several technical standards and other specialized literature address the topic, it is still a "fuzzy" matter, subject to misunderstanding and discussion.

IEC 61511 Clause 11.9 requires that the SIL be verified using quantitative analysis, such as reliability block diagrams, fault tree analysis, and Markov modeling. This analysis includes only those dangerous failures that are random in nature. Common cause failures may or may not be included in the verification calculation dependent on whether they exhibit random or systematic behavior. Any personnel assigned responsibility for verifying the SIL should understand each failure type and the strategies that can be used against it. Consequently, this paper provides an overview of random, systematic, and common cause failures and clarifies the differences in their management within IEC 61511.

INTRODUCTION

Modern control systems are designed to keep the process within specified parameters considered acceptable for normal and safe operation. However, when the process excursions are larger than the allowed range of variation, hazardous conditions may occur. If the situation is not addressed, it can propagate into a hazardous event with potential consequence to human life, environment, and plant assets.

The risk associated with such a process excursion may be reduced using a safety instrumented system (SIS).

During a hazard and risk analysis, risk reduction is allocated to safety functions, which are used to reduce the process risk below the owner/operator risk criteria. When a safety instrumented function (SIF) is identified, the allocated risk reduction is related to its safety integrity level (SIL). The risk reduction and SIL establish a benchmark for the design and management practices used throughout the SIF life. This benchmark serves two purposes: 1) it defines the minimum performance that should be achieved with regard to random failures and 2) it determines the rigor of the protective management system required to reduce the potential for systematic errors.

A Hazard and Risk Analysis (H&RA) team identifies hazardous events that pose catastrophic consequences. One such event might be loss of vessel level allowing high pressure gas to flow to downstream equipment which is not rated for the pressure. A Safety Instrumented Function (SIF) may be specified to reduce the risk of this event. The SIF detects low level and responds by closing an outlet block valve to prevent blow-by. Three redundant level transmitters are specified to detect a low level condition. Other level devices are used by the basic process control system to monitor and control the vessel level. When any two of the three level transmitters detect low level (two-out-of-three, 2oo3), the Safety Instrumented System (SIS) closes the outlet block valve. The SIF can still work if one level transmitter fails dangerously; however, if two transmitters fail dangerously, the SIF will fail to close the valve, allowing the level to be lost in the vessel with potential catastrophic consequences.

The dangerous failure of two level transmitters may be unlikely if the devices are designed, installed, operated, inspected, and maintained according to the specifications. However, when the device specifications are violated, multiple simultaneous failures can occur. This is known as a common cause failure (CCF). For example, the vent rate from the actuators of two redundant valves was not considered when calculating the valves response time, so the valves close too slowly for the process safety time. Redundant transmitters are installed using clean service process connections when the process contains solids. Or, inspection is not performed to detect pluggage of the vent ports on the solenoid operated valves.

For many years, the process industry has addressed SIS integrity and reliability issues through the use of prescriptive design and maintenance practices, which are sufficiently conservative to assure that the risk reduction can be met in a wide variety of process applications. These internal practices are frequently criticized for what is perceived as “over-design,” because the practices seem to go beyond the minimum requirements of IEC 61511(1) or other good engineering practices. However, it should be recognized that any good engineering practice documents general practices that are acceptable to participating stakeholders. A specific application may necessitate more rigor, including additional fault tolerance, diagnostics, testing, etc. For more guidance related to IEC 61511, refer to ISA TR84.00.04(5).

Project and plant personnel are under pressure to “optimize” anything and everything, while adhering to good engineering practices. This “optimization” is generally intended to make the design as simple and cost effective as possible. However, when taken to extreme, optimization can reduce the design safety margin to a minimum and can result in the heavy reliance on devices whose failures are not well known. As less and less safety margin is built into the design, greater importance is placed on the precision of the predictive failure calculation. Thus, the device’s failure rate can become a critical design

parameter rather than simply a performance benchmark. Excessive focus on minimal cost can result in the SIF being designed and operated with less tolerance for failure.

Consequently, while this paper presents a discussion of how to classify and manage failures, it must be recognized that the calculations are not a means to an end. Safe operation is achieved by designing and implementing SISs that take into account a wide variety of site specific criteria. The primary advantage of internal practices is that, rather than turning every SIS design into an exercise in optimization, the risk management philosophy is clearly laid out with a sufficient safety margin to address uncertainty. The SIL is then verified using estimates of the random dangerous failure rates of the devices that comprise the SIS. Long term performance is monitored to see if the actual failures exceed expectations.

A "BATHTUB" VIEWPOINT

The failure behavior of a population of hypothetical devices over their lifecycle (i.e., from production to disposal) is commonly represented by the "bathtub" curve, as shown in Figure 1A. This plot shows how the overall failure rate of a device changes with time. The initial failure rate of the hypothetical device is driven by its "burn-in" or "infant mortality" rate, which declines rapidly. The middle flat section of the curve represents the "useful life" of the device and is typically characterized by a constant failure rate. The SIF performance calculation is based on this middle section. The last part represents the "wear out" or "end-of-life" failure rate and is characterized by an increasing failure frequency.

As indicated by Smith(2), the Bathtub Curve is in reality composed by three overlapping curves, one for each of the three sections, Figure 1B.

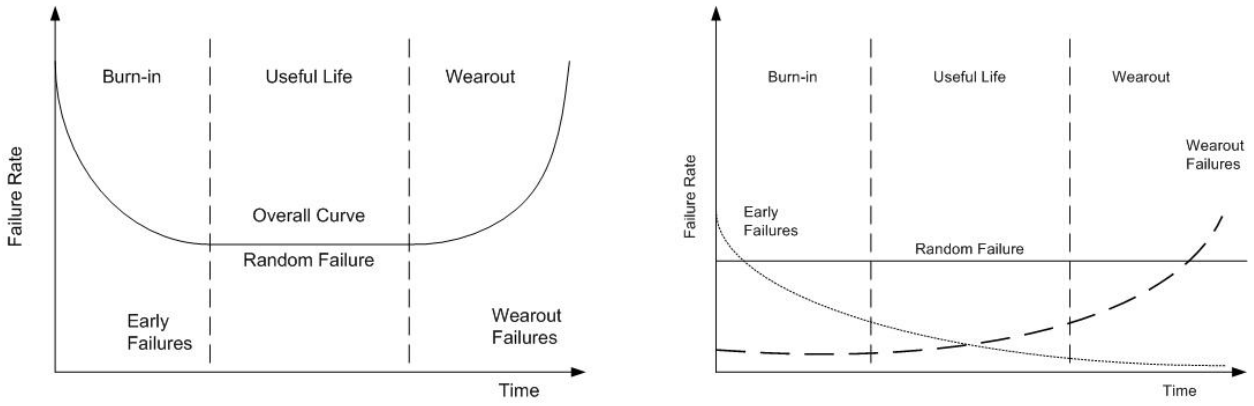

Figure 1: Overall bathtub curve (A) and components of the bathtub curve (B).

Failures during the burn-in period may be caused by manufacturing flaws in the components that comprise the device. Manufacturers affect the shape of this curve when they perform burn-in and function testing prior to release from production. Infant mortality also occurs as a result of device handling and installation. Devices can be damaged during shipment, unpacking, storage, transport to work site, and installation. Many early failures are caused by rough handling of the device, poor installation practices, and sloppy construction practices. Materials used for installation activities, such as paint, pipe dope, insulation, and small pieces of a wielding rod, have been shown to cause devices to fail by getting into places where they are not supposed to be. Rigorous inspection, commissioning and validation activities are necessary to identify and correct these failures.

The transition between the burn-in and useful life curves (Figure 1A) occurs when the failure rate becomes constant. Failures during the useful life of the device occur mainly due to random events that increase the stress on the device beyond its physical limits. When the wear-out failure rate becomes dominant over the useful life failure rate, the overall failure rate increases ending the useful life of the device.

The wear out period is characterized by an increased slope that is related to the device type. Programmable devices tend to have a very sharp increase in the failure rate due to the large number of aging components. Electromechanical devices tend to have a more gradual increase in failure rate as they age, and the slope of their failure rate curve is heavily affected by how well their mechanical components are inspected and maintained. Preventive maintenance has been shown to extend the useful life of electromechanical devices. A lack of inspection and preventive maintenance has been cited as a primary cause of early failure. As with infant mortality, not all wear-out failures occur after the end of the useful life; these failures can happen during any stage of the device lifecycle and as time passes its rate of occurrence increases.

A device should undergo burn-in and function testing to detect as many infant mortalities as reasonably possible. The device should also be replaced when maintenance records demonstrate that it has reached its wear-out phase. When infant mortality and end-of-life issues are addressed within the device installation, commissioning and maintenance plans, it is assumed that the device failure rate is constant. The useful life section of Figure 1A illustrates the constant failure rate that can be achieved through an effective inspection and maintenance program.

The constant failure rate assumption is a fundamental one in the analysis of the SIF performance; however, it has been demonstrated that this region does not exist for all devices(3). For example, some devices, such as block valves, are characterized by dominant failure mechanisms which are a function of time (i.e., lower failure rate at early stages of useful life, high failure rate at wear-out stages). In this case, specific statistical distributions, such as Weibull may be considered, as well as the random failure models.

The random failure rate is characteristic to the device when it is operated in accordance with its specification. The failure rate is estimated based on specific operational and maintenance conditions implicit in collected data. Reducing inspection and maintenance rigor will affect a device's failure rate. Operating the device outside of the manufacturer's specification for the device may damage or stress it, thereby causing early failure. Unexpected process application impact, such as chemical attack, corrosion, and deposition, or external environmental impact, such as electromagnetic interference, vibration, and heat, can also shorten its useful life.

SYSTEMATIC, RANDOM, AND COMMON CAUSE FAILURES

An SIF uses many different devices to execute a safety function intended to reduce the risk of an identified hazardous event. The SIF performance depends on the characteristics of each SIF device (e.g., individual failure rates), the properties of the system itself, and the interactions among its components (e.g., voting architecture, common cause failures). An implicit assumption made during SIF design and its performance verification is that the devices are in their useful life period and that the SIF is installed and managed according to its design and operating basis

The SIF probability of failure is a function of the random failure rate of its devices (i.e., field sensors, final elements, logic solvers, and other events) and the design basis parameters, such as redundancy, test interval, and diagnostic coverage. Other systematic, non-random contributions affect the observed performance of the SIF. Systematic and random faults can cause the failure of an individual device or the simultaneous failure of multiple devices. The individual device failure is easily assessed using probabilistic techniques that are described in many publications, such as Smith(1), Lees(2), and ISA TR84.00.02 (6).

Independent failures of devices are tracked and analyzed as part of a protective management system. The goal of a failure tracking program is to identify root causes and data trends, and to develop means to reduce the potential for failure reoccurrence. The analysis may identify systematic errors and common cause failures. When a common cause failure occurs, the root causes tend to be installation- and/or application-specific. Systematic errors occur because some human error broke through the administrative barriers that were supposed to detect and correct it. Therefore, systematic errors tend to be perceived as a local personnel issue. Regardless of the nature of the failure, dangerous failure reports and identified trends should be communicated to appropriate personnel, because as personnel better understand how devices can fail, the better they can work to prevent failure.

Failures are managed within the IEC 61511 work process using different strategies depending on whether they are random or systematic. Consequently, the following sections briefly discuss random, systematic, and common cause failures, so that the strategies for their management can be better understood.

Random failures

Random failures are physical failures brought on by excessive stress on the device. Random failures can occur at any time during the device life and are generally well understood by the time a device achieves user-approved status. The occurrence of these failures does not follow a pattern, but instead occur randomly in time. Stress factors are the result of a variety of causes, such as abnormal process conditions, the presence of an adverse microenvironment (e.g., corrosion produced by chemical impurities or metallurgy), or normal device wear and tear. Low frequency atmospheric events (e.g. snow in Houston, Texas, USA) can also be considered random events.

The bathtub curve for an individual device is developed by counting the failures that occur in a population of identical (or sufficiently similar) devices over a certain period of time. When the devices are in their useful life period, the collected data set is used to calculate the failure rate using statistical techniques. Collected data typically includes some systematic failures, which contribute to the observed failure rate. For example, it may take multiple failure reports before it is recognized that the instrument air quality is causing the failure. Figure 2A and 2B illustrate how the observed failure rate is higher due to the presence of systematic events that are exhibited randomly. These failures are illustrated in Figure 2B using additional failure rate lines. Since the intent of the quantitative analysis is to predict SIF performance, the systematic failures should be tracked and their inherent presence considered when estimating the random failure rate. As more information is collected, trends can be identified and used to minimize random and systematic failures in new or modified designs.

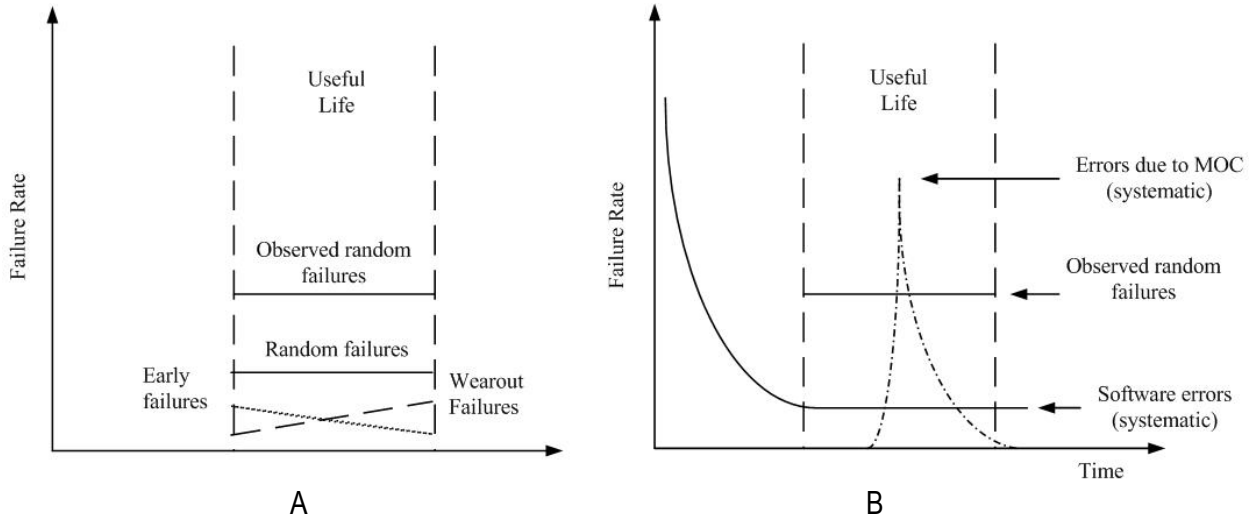

Figure 2: Hypothetical device's random failure rate is constant (A), however other non-random failure sources affect the observed failure rate of the device (B).

Random failures can be detected by diagnostics internal to the device, externally-configured diagnostics, and proof tests. Redundant, fault tolerant subsystems are often used to reduce the probability that a single failure will cause the SIF to fail. Redundant subsystems also provide the potential for diagnostic coverage, where a diagnostic algorithm is executed at a specified interval to detect certain device failures. The test interval is generally chosen based on maintenance history, manufacturer's recommendations, good engineering practice, insurance requirements, regulatory requirements, and what is necessary to achieve the required integrity and reliability. As the test interval gets longer, there is an increased probability that multiple devices within a subsystem have malfunctioned prior to fault detection.

Systematic failures

While random failures are caused mainly by physical factors, systematic failures are the direct consequence of device and SIF complexity. Every hardware and software component in a device is subject to failure due to design, specification, operating, maintenance, and installation errors. These mistakes immediately put the devices on the path to failure. As the complexity of the SIF increases, the potential for systematic errors increases due to the combination of failures. As complexity increases, the probability of detecting these errors decreases. Each device has many known opportunities for systematic error. With any new technology, there is the potential for many unknown (or as yet unidentified) failures. When issues associated with interconnectivity, communication, and support systems are added to the analysis, there are generally a large number of potential systematic failures. Due to the complex and non-random nature of systematic failures, it is difficult to predict or statistically analyze them.

Further, only a limited number of device failures and failure paths can be tested. When the failure patterns are not detected by the limited testing that is practically achievable, systematic failures can happen every time the specific set of conditions occurs. These conditions become an intrinsic part of the device and thus the SIF. Systematic errors are a major source of common cause failure, because of the potential to disable redundant devices.

Systematic failures include many types of errors, such as:

- Manufacturing defects, e.g., software and hardware errors built into the device by the manufacture.
- Specification mistakes, e.g. incorrect design basis and inaccurate software specification.
- Implementation errors, e.g., improper installation, incorrect programming, interface problems, and not following the safety manual for the SIS devices.
- Operation and maintenance, e.g., poor inspection, incomplete testing and improper bypassing.

Systematic errors related to manufacturing defects can be reduced through the use of diverse redundancy. Diversity typically involves the use of different technologies or different manufacturers. While device manufacturer errors can be addressed by diversity, this increases the SIF complexity. Incorrect specification, implementation, operation, and maintenance constitute root causes that are not solved by diversity and can actually increase if unnecessary complexity is introduced into the SIS design. The need for diversity should be balanced with the need for error free operation and maintenance.

Further, the perceived improvements gained by diversity are based on the assumption that different devices exhibit different failures and failure modes. In other words, it is less probable for all of them to fail simultaneously, if they are different. However, diversity only reduces the potential for common mode failures. Many common cause failures are not addressed by diversity. As an example of a systematic error, consider a specification mistake that occurs during the engineering, installation, commissioning and validation phase. This error may be as simple as not properly specifying the process application on the hardware specification form, leading to the purchase of a device with inappropriate materials of construction. While the device itself may fulfill all of the requirements necessary to gain user-approved status, the incorrect specification results in the premature failure of the device. Redundancy and voting do not address the specification error, because it is typical practice to purchase both devices using the same specification.

To illustrate further, let's assume that there is a subsystem composed of 2oo3 voting temperature transmitters installed in the same thermowell. The designer feels that this is acceptable, because the specification indicates that the process is clean. However, in reality, the process produces a polymer which rapidly fouls the thermowells, insulating them from the actual process. None of the temperature sensors detect the actual process condition. This mistake may result in the sensors not being able to detect high temperature within the process safety time.

Improper management of change applied during manufacture, specification, installation, operation and maintenance may generate systematic failures. For example, the latest version of the embedded software in a transmitter may change how the technician configures the device. If the maintenance procedure is not changed, errors may occur during commissioning or maintenance. For this reason, the protective management system incorporates configuration management as a specialized management of change process.

According to the data from a sample of 34 incidents published by the Health and Safety Executive(4) there is no single lifecycle phase to blame for SIS non-performance. Mistakes are made throughout the lifecycle. Systematic errors are best addressed through the implementation of a protective management system, which overlays a quality management system with a project development process. A rigorous system is required to decrease systematic errors and enhance safe and reliable operation. Each

verification, functional assessment, audit, and validation is aimed at reducing the probability of systematic error to a sufficiently low level.

The management system should define work processes, which seek to identify and correct human error. Internal guidelines and procedures should be developed to support the day-to-day work processes for project engineering and on-going plant operation and maintenance. Procedures also serve as a training tool and ensure consistent execution of required activities. As errors or failures are detected, their occurrence should be investigated, so that lessons can be learned and communicated to potentially affected personnel.

Common cause failures (CCF)

All common cause failures have the potential to reduce the SIF performance; however, they are addressed in different ways depending on the nature of the failure (e.g., systematic or random). Throughout the IEC 61511(1) lifecycle, it is recommended that the devices, systems, or protection layers be assessed for independence and the potential for common cause failure. The concepts of independence and common cause are interrelated. A lack of independence means that there is a potential for a common cause failure. Likewise, an identified common cause indicates a lack of independence and therefore some level of dependency.

Common cause failure is a term that is used to describe random and systematic events that cause multiple devices, systems, or layers to fail simultaneously. Another term is “common mode” failure, which describes the simultaneous failure of two devices in the same mode. Common mode failure is related to the use of identical devices in the redundant subsystem. For example, two redundant differential pressure sensors on a pipeline can be simultaneously disabled due to loss of signal (common mode failure) originated from diaphragm damage (failure cause) caused by water hammer. Common mode failure is a subset of common cause failure.

Diversity is often suggested as a means to eliminate common cause failure. However, common cause failure can impact identical and diverse devices. For example, process application or external environmental conditions can affect different technologies simultaneously when the conditions trigger each device’s failure mechanisms. These devices may eventually fail due to different reasons, but the abnormal process condition is root cause that started the failure propagation. The use of different technologies (i.e., diversity) does reduce the potential for common mode failure. In summary, diversity reduces the potential for dependent failure by minimizing common mode failure, but does not eliminate the potential for common cause failure.

The approach taken to manage common cause failure (CCF) is specific to the nature of the failure. Within the standard(1), two types of CCF are addressed: 1) single points of failure where one malfunctioning device causes an SIF failure and 2) single events that lead to multiple failures in a redundant subsystem. Figure 3 and Table 1 refer to these as “single point of failure” and “redundant,” respectively.

Single points of failure were discussed in random and systematic sections above. As shown in Figure 3 and Table 1, single points of failure can occur due to systematic (A) or random (B) events. Systematic failures (A) occur when human errors result in the violation or invalidation of design and

operating basis assumptions (e.g., process assumed to be clean but in reality is not). Random failures (B) can occur throughout the useful life of a device. These failures are managed using redundancy, diagnostics, and function testing.

As with single points of failure, redundant subsystems can fail due to systematic errors (C) in the device manufacture, specification, design, installation, and maintenance. These errors typically happen due to lack of knowledge, information, and training and are generally unknown to personnel. Test procedures may not identify these errors, since they are not expected. Systematic errors are difficult to test even if easily identified. For example, it is possible that the valve actuator is incorrectly specified, but how do you test to determine that the valve actuator will not close under emergency process conditions? Specification errors must be caught during the design and engineering phases using independent verification. Checklists can be used to identify CCF. The list of questions guides the engineer through the design aspects examining opportunities for CCF. Installation, commissioning, and maintenance errors are reduced by independent checks, verifications, and audits.

Random failures of redundant subsystems (D, E, and F) can be caused either by conditions that are inherent to the device (F) or inherent to the system (D and E). Random failures inherent to the device (F) are generally manufacturing defects which may include hardware and/or software failures. These failures are estimated using the beta factor method. The operating environment, the installation, and the interconnection to other systems (D and E) affect the device operation. This system-induced random failure can be divided into two categories depending on the availability of failure frequency data. If data is available (D) the failure can be modeled explicitly as an event. For example, if fault tree analysis is the selected analytical methodology, this type of CCF is treated as a basic event with its own failure rate. If data is not available, CCF can be addressed using the beta factor method (E).

The beta factor accounts for random events that cause the device to fail in the operating environment. The value of the beta-factor is selected based on engineering judgment. Many owner/operators use a beta factor between 0.1% and 5.0% when good engineering practices are applied in the design, installation, inspection, and maintenance practices. The beta factor can be substantially higher if good engineering practices are not followed.

Table 1: Methods used to address the different types of common cause failures and examples (see Figure 3)

CCF Type	METHOD	COMMENTS	EXAMPLES
SINGLE POINT FAILURE			
A	Lifecycle management	Caused by mistakes and puts devices immediately on failure path. Assumptions are violated and invalidated. Failures were unknown due to lack of technical knowledge. Difficult to model these errors.	Software error, poor installation, incorrect configuration, lack of inspection, maintenance leaves device in bypass.
B	Analytical (i.e., Single event in fault tree)	Device fails due to stressors in the operating environment, infant mortality, or wear out.	Typical random failure.
REDUNDANT ARCHITECTURE			
C	Lifecycle management	Caused by mistakes and puts devices immediately on failure path. Assumptions are violated and invalidated. Failures were unknown due to lack of technical knowledge. Difficult to model these errors.	Wrong device design, software bug, defective device due to manufacturer error. Wrong device specification due to engineering error, therefore device plugs because used in dirty environment believed to be clean.
D	Analytical (i.e., Single event in fault tree)	Caused by the installation and device interconnection. Explicit failure data available.	Single power supply, single instrument air supply, shared digital communication.
E	Analytical (i.e., β-factor)	Caused by the operating environment, installation and device interconnection No explicit failure data available therefore β-factor estimated on prior-use information and experience for similar devices.	Environmental stress on the device (e.g., abnormal corrosivity, vibration, pressure, temperature, solids, erosion, etc.) due to transient or abnormal operation.
F	Analytical (i.e., β-factor)	Inherent part of the design of the device. β-factor estimated on prior-use information and experience.	Device failure due to manufacturing errors, such as defective raw materials, poor quality production processes, and software bugs.

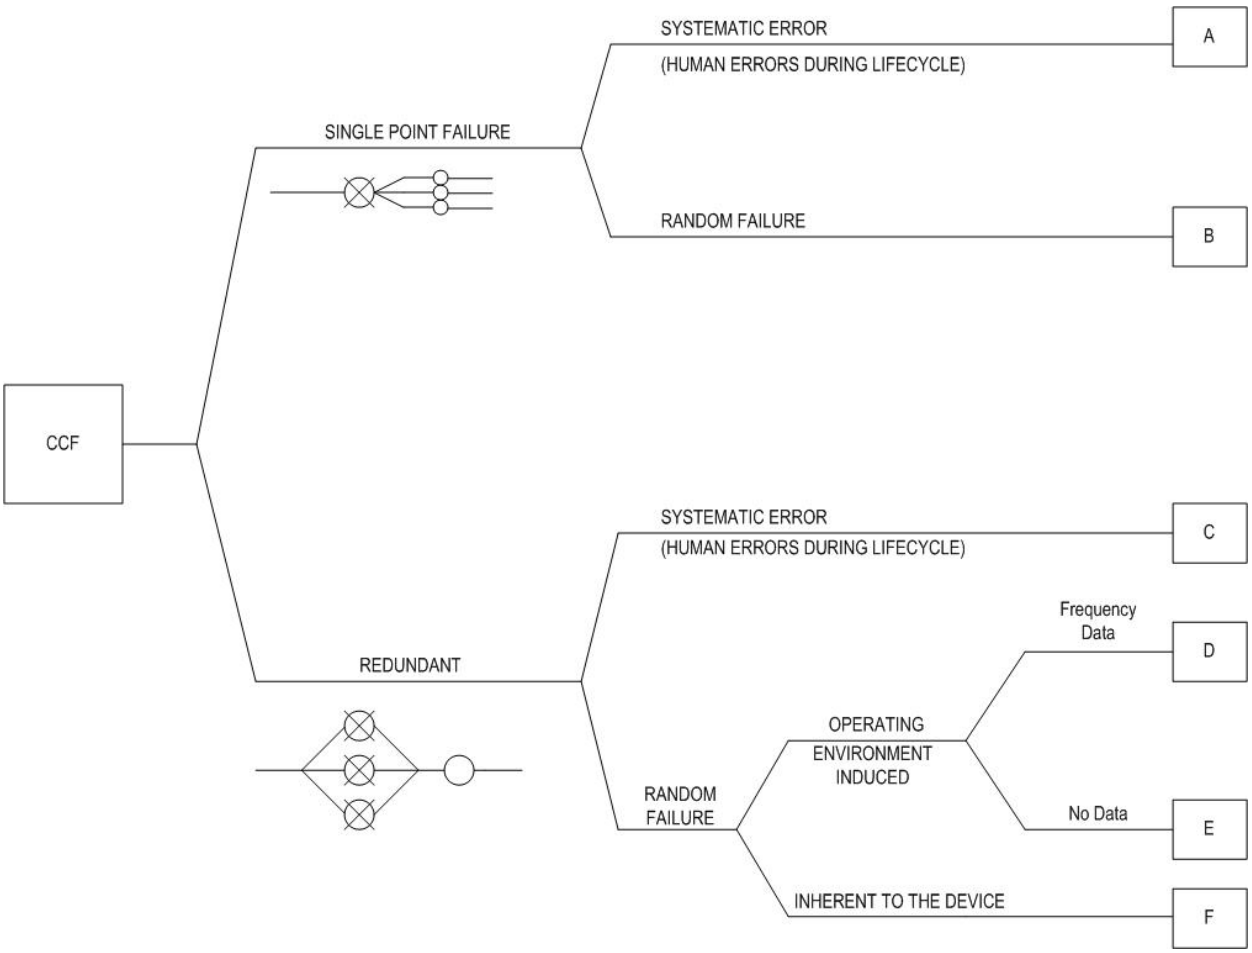

Figure 3: Taxonomy of common cause factors (CCF) and methods for their analysis (see Table 1).

CONCLUSIONS

It is important for process and control engineers to understand how devices fail in order to identify the optimal solutions for each situation. Device failures can be classified as random, systematic, and common cause. Random failures occur due to different types of random events and can be assumed statistically to occur at a constant failure rate. Random failures are easily modeled using probabilistic math, allowing the SIF performance to be estimated. Systematic errors are caused, or indirectly induced, by human error or unforeseeable complex conditions. Systematic failures are not random events and are addressed by the management system. Since they are not easily analyzed, most systematic errors are not included in the SIL verification calculations. Both, random and systematic events can induce common cause failure (CCF) in the form of single points of failure or the failure of redundant devices. The nature of the event determines the techniques necessary to reduce its potential and the analytical method used to predict it.

DEFINITIONS

CCF or Common Cause Failures: Failure of more than one component, device, or system due to the same cause.

H&RA or Hazard and Risk Analysis study: Identification of hazardous events, the analysis of the mechanisms by which these undesired events could occur, and the estimation of risk based on an engineering evaluation and mathematical techniques for combining estimates of event frequency and consequence severity.

Management System: A program or activity involving the application of management principles and analytical techniques to ensure the safe and reliable operation of process equipment.

MOC or Management of Change: Core attribute of a protection layer, where a system is used to identify, review, and approve all modifications to equipment, procedures, raw materials, processing conditions, other than "replacement in kind," prior to implementation.

SIF or Safety Instrumented Function: A safety function allocated to the safety instrumented system with a safety integrity level (SIL) necessary to achieve the desired risk reduction for an identified hazardous event.

SIL or Safety Integrity Level: Discrete level (one out of a possible four) used to specify the probability that a safety instrumented system will perform its required function under all stated conditions within a specified time period. Integrity level 4 has the highest level of expected performance, while integrity level 1 has the lowest.

SIS or Safety Instrumented System: Composed of a separate and independent combination of sensors, logic solvers, final elements, and support systems designed and managed to achieve a specified safety integrity level. An SIS may implement one or more safety instrumented functions (SIFs).

Voting: Specific configuration of hardware components within a subsystem. Voting is expressed as MooN (M out of N). "N" designates the total number of devices that are implemented in parallel; "M" designates the minimum number of devices out of N that are required to initiate shutdown conditions or to achieve a defined output action. This is also called a voting system or voting architecture.

2oo3: Out of a total of 3 redundant devices, 2 must work properly in order for the safety function to detect or respond to a hazardous condition.

REFERENCES

International Electrotechnical Commission, IEC 61511, Functional Safety: Safety Instrumented Systems for the Process Sector, Geneva, Switzerland (2003).

Smith D.J, "Reliability Maintainability and Risk," 5th edition, Butterworth-Heinemann Ltd., Oxford (1997).

Mannan, Sam, Lee's Loss Prevention in the Process Industries, Volumes 1-3, Elsevier Butterworth-Heinemann, UK (2005).

Health and Safety Executive, Out of Control – Why control systems go wrong and how to prevent failure, Second Edition, HSE Books, UK (2003).

Instrumentation, Systems, and Automation society, ISA TR84.00.04, Guidelines on the implementation of ANSI/ISA 84.00.01-2004 (ISA 61511 modified), Research Triangle Park, NC (2006).

Instrumentation, Systems, and Automation society, ISA TR84.00.02, Safety Instrumented Functions (SIF) - Safety Integrity Level (SIL) Evaluation Techniques, Research Triangle Park, NC (2002).

NOTE: This paper is adapted from a draft Center for Chemical Process Safety (CCPS) book, Guidelines for Safety and Reliable Instrumented Protective Systems, American Institute of Chemical Engineers, New York, New York (Expected publication 2006).