

Maximizing Safety Without Compromising Reliability

Artesyn Embedded Technologies

www.artesyn.com

October 2015

A programmable electronic system can be defined as functionally safe if it operates correctly and predictably, so that even in the event of failures it remains safe for persons and the environment. Such a system can be defined as reliable if it performs its function without failure for a specified period of time. These attributes can lead to conflicting requirements and very different designs.

For example, to achieve high levels of functional safety, one method is to compare two or more channels as a diagnostic so that if a difference is detected, the system enters a “fail-safe” state and stops delivering its prescribed service.

On the other hand, achieving high reliability also requires two or more channels. But in this case, upon failure in one channel, the secondary standby channel becomes active, and the system continues to deliver its prescribed service.

This paper provides more detail on some of the processes and steps that must be undertaken to make a safe system and also discusses how high reliability is achieved without compromising functional safety.

Designing for Reliability and Fail-Safe Operation

In a railway environment and in other industrial environments, fail-safe operation is commonly used to achieve high levels of functional safety. The rationale being that it is acceptable for the system under control to simply stop when a failure is detected, rather than to keep running and potentially causing harm. Achieving high levels of functional safety is easier and less costly to implement when the option for a fail-safe design is possible.

However, a safe shutdown can still be expensive for the operator. For example, consider the subsea oil and gas exploration environment where a failure to a system could cost the operator millions of dollars per day in lost revenue and maintenance costs. Thus, in addition to high levels of functional safety, such systems must also deliver high levels of reliability.

Safety must be designed into a system from the ground up, starting with the marketing requirements, through high level product requirements, and through all phases of the design, testing and qualifications of the product. This requires a significantly more structured and formal methodology to designing safe systems than would be typical of a system that only needs to meet reliability requirements.

This methodology includes performing hazard and risk analysis on all system elements to assure that all faults are detected or safe, and that the system remains in a safe state in the event of failure. Probabilistic analysis is performed to determine the probability of a hazardous failure, and whether or not the risks have been reduced to tolerable levels. This kind of analysis will drive the design in ways high reliability requirements cannot.

Explaining the Differences

To evaluate how safety and reliability correspond and inter-relate in an application environment that uses embedded computing, such as rail signaling, we must first clarify how these properties are defined.

Safety

Functional safety applies to systems that could cause injury or loss of human life if they fail, or could cause significant environmental damage. Flight-control systems, automotive drive-by-wire, nuclear

reactor management, oil and gas systems, or operating room heart/lung bypass machines are some of the applications that come to mind. However, other simpler devices as common as the power windows in your car are also safety-critical because some failures could cause injury. Safety integrity is the degree of freedom from unacceptable risk to persons or the environment. It is quantified by terms such as the probability of a hazardous failure (PFH) used in IEC 61508, or the tolerable hazard rate (THR) used in EN 50129. These measures are related to the unavailability of the safety functions. A safety related system implements the necessary safety functions to achieve the desired safety integrity level (SIL).

Fail-Safe

A fail-safe device is one that, in the event of a failure or other potentially hazardous event (such as a person entering a hazardous area), will respond by stopping the equipment under control in a way that will prevent harm to personnel or the environment. Rail interlocking systems are typically designed to be fail-safe so that a train cannot proceed over a route in the event of a failure. Industrial machines such as presses or assembly robots might use light curtains in a fail-safe system to stop machinery when a person enters the vicinity.

Reliability

Reliability is defined as the probability a component (or system) will survive its mission time. Reliability is measured in terms of a failure rate, which can be expressed in several different ways: in failures per hour, failures per billion hours (FIT or “failures-in-time”) or the inverse of the failure rate as mean-time-to-failure (MTTF). All components have a failure rate and designers have limited control over the aggregate failure rate of a computing system. But they do have control over how a failure would affect delivered service.

High Availability

Highly available systems are systems that continue to deliver their prescribed service in the face of scheduled events (e.g., maintenance) or unscheduled events (e.g., failure). Many systems are designed to recover from a failure by detecting the failed component and switching to an alternate standby system. These systems, although

sometimes called fault-tolerant, are more widely known as “high availability” systems.

Many high availability systems do not pose a safety threat in cases of failure and instead are designed to maximize ‘uptime’ and minimize ‘downtime’. High availability systems today strive to be up and running 99.999% of the time (the so-called ‘five-nines availability’) or more, equivalent to a total of about five minutes down time out of 8,760 hours per year.

A key difference between high availability and functional safety is that functional safety must guarantee, to a high degree of probability, that the system will not compromise safety when a demand event occurs, while high availability must guarantee to a high degree of probability that the system continues to operate in the event of a failure.

Fault Tolerance

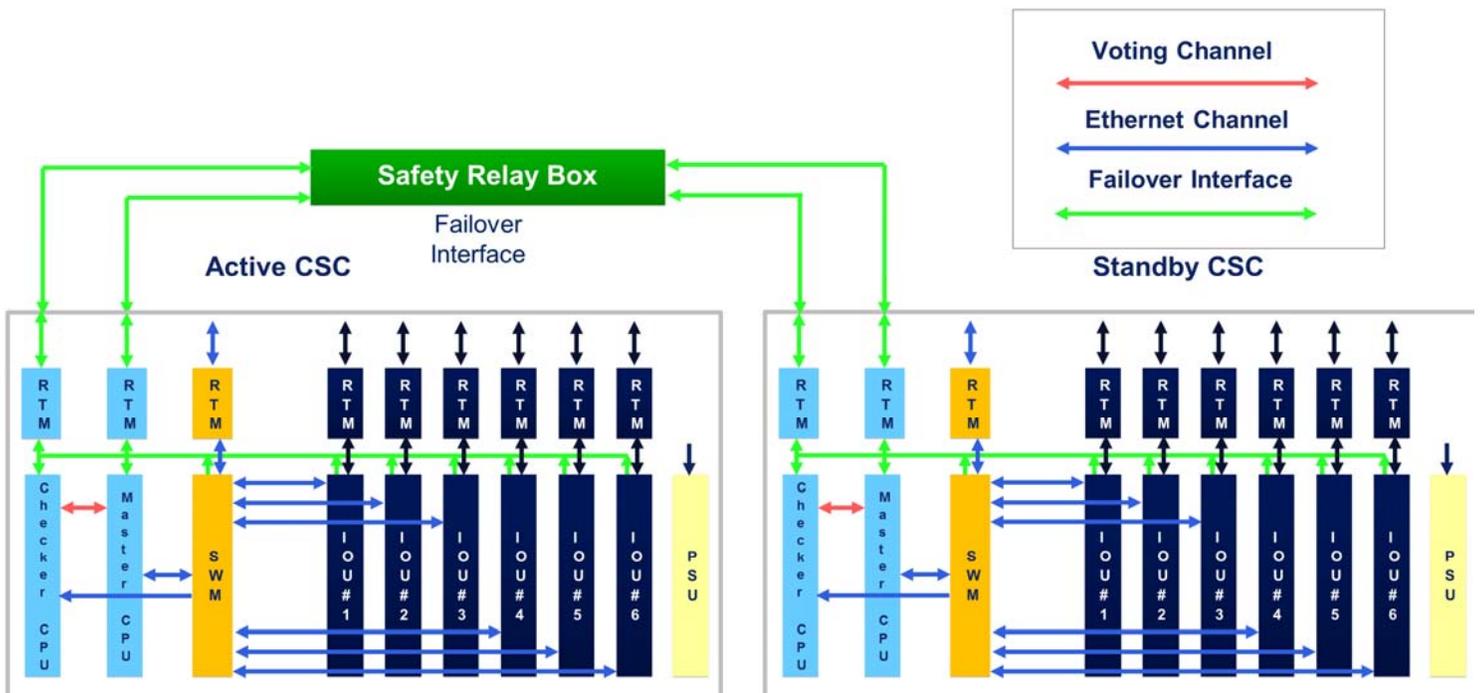
Similar to high availability systems, fault tolerance offers the ability to continue operating when a failure occurs. A fault-tolerant system is designed from the ground up for reliability by building multiples of all critical components, such as CPUs, memories, disks and power supplies into the same computer. In the event one component fails, another takes over without skipping a beat.

Different System Properties

One might assume that safety is increased by increasing the reliability of the individual system components. If components do not fail, then accidents will not occur.

However, as explained above, safety and reliability are different system properties. One does not imply nor require the other — a system can be reliable and unsafe, or safe and unreliable. In some cases, these two system properties are conflicting, i.e., making the system safer may decrease reliability and enhancing reliability may decrease safety.

A system that is safe will have a very different design from a system that must only be reliable. Imagine a brake control system on a train. A highly reliable braking system will incorporate redundancies to operate without failure for very long periods of time, but it can still fail, and potentially do so when lives are at stake. A safe braking system is designed on a “de-energize to trip” principal where pneumatic pressure is required to prevent the brakes from operating. If the pressure system fails, the brakes are applied by design, stopping the train. More modern electrical systems are designed with the same fail-safe principle: the system must be electrically energized to prevent the application of the brakes.



Artesyn’s ControlSafe Platform (CSP) consists of two redundant CSCs, each of which delivers fail-safe operations. They are linked by a Safety Relay Box (SRB) that monitors the health of the two CSCs and designates one of them as ‘active’ and the other as ‘standby’. The user application running on the active CSC has full control of all I/O ports, while the same user application running on the standby CSC can monitor input ports but has no ability to drive any output port (unless specifically configured for output in Standby mode). When the active CSC fails, it signals its state to the SRB, which in turn causes the standby CSC to become active. The unhealthy CSC is taken out of operation and, once it has been repaired by service personnel, can be brought back into service. Monitoring the health state of the two CSCs and controlling fail-over operation between them enables a fail-safe and fault-tolerant computing system.

Guaranteeing Levels of Behavior

In a safe system, safety relevant faults must be detected in order for the system to transition to a pre-defined state that is safe. In train interlocking systems, the fail-safe state renders all outputs inactive.

This is known as ‘guaranteed behavior’ — if a system fails, we guarantee to a quantifiable level of probability that the system will fail-safe in the event of a safety relevant failure. In some cases, guaranteeing that a failure in a circuit will not cause an unsafe situation requires that two or more redundant circuits be deployed. While such internal redundancy increases safety it does not increase the system’s reliability or ensure availability.

In Artesyn’s ControlSafe™ Platform (CSP), high availability behavior is achieved by deploying a second redundant ControlSafe Computer (CSC) which operates in standby mode. A fully fault-tolerant Safety Relay Box (SRB) which is connected to both CSCs guarantees that only a single healthy CSC is operating in active mode. If a CSC fails, it sets its outputs to a safe state and declares itself unhealthy.

The SRB will then assure the remaining healthy CSC becomes active such that the system continues to deliver its prescribed service.

Standardizing Safety Levels

There are two major sets of standards related to railway applications, managed by the International Electrotechnical Commission (IEC) and the European Committee for Electrotechnical Standardization (CENELEC).

IEC 61508 is a generic standard for the functional safety of electrical, electronic or programmable electronic safety-related systems. One of the first statements within the IEC 61508 specification is that, “a major objective of this standard is to facilitate the development of application sector international standards by the technical committees responsible for the application sector”.

For railway applications, CENELEC has produced a number of standards that address the functional safety of railway applications. To the extent that electrical, electronic or programmable electronic systems are involved, the CENELEC standards can be regarded as the “application sector standards” referred to in IEC 61508.

This is the case for the family of standards:

- EN 50126 The specification and demonstration of Reliability, Availability, Maintainability and Safety (RAMS)
- EN 50128 Software for railway control and protection systems
- EN 50129 Safety related electronic systems for signaling

Safety Integrity Levels (SIL)

The Safety Integrity Level (SIL) concept is a way of categorizing safety functions into four discrete levels: SIL1 to SIL4. The SIL determination follows a complex, although systematic, process outlined in the relevant standard.

These SIL levels affect both the lifecycle processes used to create a safe system, and the methodology for calculating the probability that a dangerous fault will be undetected. The standards include techniques for the probabilistic analysis that determines the likelihood of a dangerous failure, and set increasingly challenging requirements on this probability.

Safety Integrity Level	Safety	Probability of Failures on Demand	Risk Reduction Factor
SIL 4	> 99.99%	0.001% to 0.01%	100,000 to 10,000
SIL 3	99.9% to 99.99%	0.01% to 0.1%	10,000 to 1,000
SIL 2	99% to 99.9%	0.1% to 1%	1000 to 100
SIL 1	90% to 99%	1% to 10%	100 to 10

SIL are order of magnitude bands of average Probability of Failure on Demand (PFD) as defined in IEC 61508. PFD measures the amount of risk reduction performed by a Safety Instrumented Function (SIF). Corresponding to PFD, each of the four SIL levels can also be denominated as Safety Availability and Risk Reduction Factor (RRF). Safety Availability is the complement of PFD (i.e., 1-PFD), and the RRF is the inverse of PFD (i.e., 1/PFD). EN 50129 has a similar parameter called Tolerable Hazard Rate (THR).

Standards Focused on Processes

The functional safety standards focus on the overall development processes. These processes cover the entire lifecycle from requirements, development and testing, operations and maintenance, and training. The standards focus on processes because a rigorous and formal process is seen as the best way to prevent systematic failures.

Systematic failures occur because of design or manufacturing process errors, and these failures can occur in both channels of a redundant system. Rigorous processes reduce the possibility of human error during the lifecycle to reduce the probability of their occurrence and increase the probability of detection.

Certification Accelerates Time-to-Market

Achieving these standards can be very costly, and the effort required to get a product certified at the highest SIL levels is significant. To be certified, extensive documented evidence that comes from all phases of the lifecycle must be produced and reviewed, a significant amount of statistical analysis must be performed to assure that the system will behave to the required SIL level.

All of which must be audited by independent internal auditors before the whole package is delivered to a 3rd party certification agency such as TÜV-SÜD.

Such processes, documented evidence, statistical analysis, and test and validation methodologies are significantly more than is normally required to develop a non-safety product. It has been demonstrated that the cost and time to develop a safety product can be more than double that of a normal product.

Consequently, a significant investment is required to get a safety product developed and certified. Offering

an officially certified safety platform including supporting evidence (the safety case) to application developers and system integrators simplifies the certification of the end product and makes it possible to significantly accelerate time-to-market of end solutions incorporating Artesyn's ControlSafe Platform and application software.

Conclusion

The ControlSafe™ Fault-Tolerant Safety platform is being released to rail system integrators and rail application providers. The ControlSafe platform is modular and designed to accommodate additional I/O interfaces that will be required throughout the product life cycle.

By offering a system that meets all of the functional safety, reliability and availability requirements and deploying the required processes and analysis throughout the development phases, Artesyn's ControlSafe™ platform can deliver a safe application environment and a high availability platform, giving rail system integrators and rail application providers a real competitive advantage.

Leveraging our deep understanding and experience in developing safe and reliable embedded computing platforms, and by offering a platform that is certified to SIL4 standards, Artesyn Embedded Technologies is positioned as a leading supplier of a commercial off-the-shelf (COTS) fault-tolerant safety platforms to application developers and system integrators.



Artesyn's ControlSafe™ Platform



About Artesyn Embedded Technologies

Artesyn Embedded Technologies is a global leader in the design and manufacture of highly reliable embedded computing solutions for a wide range of industries including communications, military, aerospace and industrial automation.

Building on the acquired heritage of industry leaders such as Motorola Computer Group and Force Computers, Artesyn is a recognized leading provider of advanced network computing solutions ranging from application-ready platforms, single board computers, enclosures, blades and modules to enabling software and professional services.

For more than 40 years, customers have trusted Artesyn to help them accelerate time-to-market, reduce risk and shift development efforts to the deployment of new, value-add features and services that build market share.

Artesyn has over 20,000 employees worldwide across ten engineering centers of excellence, four world-class manufacturing facilities, and global sales and support offices.

Contact

Americas +1-888-412-7832

EMEA +49 89 9608 2552

APAJ 400-88-88-183 (China)
+86-29-8874-1896 (outside China)

www.artesyn.com

© Copyright 2015 Artesyn Embedded Technologies, Inc. All rights reserved.

Artesyn Embedded Technologies, Artesyn and the Artesyn Embedded Technologies logo are trademarks and service marks of Artesyn Embedded Technologies, Inc. Intel, Xeon, Core and the Intel logo are trademarks of Intel Corporation in the U.S. and/or other countries. All other names and logos referred to are trade names, trademarks, or registered trademarks of their respective owners. © 2015 Artesyn Embedded Technologies, Inc. All rights reserved. For full legal terms and conditions, please visit www.artesyn.com/legal.

Contact Address:
Artesyn Embedded Technologies
Marketing Communications & Public Affairs
2900 S. Diablo Way, Suite 190
Tempe, Arizona 85282

Artesyn—Maximizing Safety Without Compromising Reliability White Paper R1D1