Beyond Fault Tolerance:

Third Generation SIS Approaches for Optimizing Safety Integrity and Operational Availability

A white paper of RTP Corporation, Pompano Beach, FL

Introduction: new approaches to process safety

When introduced in the 1980s, triple modular redundant (TMR) emergency safety shutdown (ESD) systems established a benchmark for providing high levels of safety integrity while reducing the occurrence of costly nuisance trips. Over the years, these types of systems became known as *safety instrumented systems* (SIS). By any name, traditional TMR safety systems tend to be expensive to purchase, implement, and maintain. Traditional SIS's also add an additional degree of complexity which many of today's downsized process plants cannot afford to handle.

Several new SIS approaches have emerged in recent years. These include new integrated systems that use a common platform for both SIS and distributed control system (DCS) functionality. Some of the new integrated systems depart from the redundant module approach, and instead, use redundant logic solver processors mounted within a common module. When designed and implemented correctly, these systems can help to reduce lifecycle costs while providing the required safety integrity levels (SIL). However, without modular redundancy, integrated systems are not nearly as fault-tolerant as traditional TMR systems and cannot be repaired or upgraded online.

A third option is now available. These are *third-generation*, fault-tolerant SIS's that combine well-proven redundancy approaches with more flexible and modern system architectures. In this manner, these new, third-generation safety systems deliver unmatched safety integrity and operational availability with the reduced lifecycle costs that today's process manufacturer's demand.

Why owner-operators install safety systems

Every owner-operator strives to make the processes performed at his/her facility as safe as possible. When risks exist in spite of that, it may be necessary to install an (SIS). The primary purpose of an SIS is to take the process to a safe state if needed.

Ideally, other than being available to perform the required safety instrumented function (SIF) at the right moment, the SIS should have no impact on the process manufacturing operations. In reality, however, internal faults or errors in the SIS can, and often do, result in spurious safety trips. While safety trips are designed to prevent costly and/or dangerous accidents from occurring, in almost all cases, a safety system trip itself will result in the production of off-spec product, reduced production, or a complete loss of production. In situations where the SIS trip

is in response to an immediate or impending process demand, safety trips are necessary and justifiable. Spurious trips, due to internal faults or errors in the SIS, cannot so easily be justified.

Not only are these spurious "nuisance" trips extremely costly in respect to lost production, but they can in themselves, create a safety hazard. This is because shutdowns and startups – and especially *unplanned* shutdowns and startups – are when most accidents occur in process plants.

Safety availability vs. operational availability

Safety systems typically operate independently of the basic process control system and require a higher degree of integrity, or *safety availability*. Safety availability involves the availability of the SIS to perform the appropriate SIF upon a process demand. Safety availability is measured in terms of average *Probability of Failure upon Demand* (PFDavg).

Safety Integrity Levels

The specific degree of safety availability required for each SIF, the *safety integrity level* (SIL), is determined through a formal process hazard analysis. In refineries and chemical plants, SIL levels typically range in criticality from SIL 1 to SIL 3. Each successive SIL represents an order of magnitude *risk reduction factor*. Specific SIL levels are achieved through a combination of SIS hardware quality and redundancy, internal diagnostics, periodic proof testing, estimated repair times, resistance to common cause failures, and successful field experience. Current IEC and ISA process safety standards focus on the performance characteristics required to achieve specific SIL levels. The manner in which they are achieved is left up to the individual owner-operators. In the US, compliance to these performance-based safety standards is largely optional and heavily dependent on a specific company's operating philosophy. In Europe and some other regions, compliance to process safety standards is mandated by law.

Significantly, the frequency of spurious SIS trips or any other SIS-related issues that can negatively impact operational availability, have little if any impact on the PFD calculations. That's because the existing safety standards are only concerned with the availability of the SIS to perform the appropriate SIF upon a process demand. Spurious trips are of interest only to the degree that they trip in a safe manner. It's safe to say that operational availability is of little interest to the standards bodies that establish process safety guidelines and regulations. On the other hand, owner-operators of process manufacturing plants are very concerned with maintaining high levels of operational availability, since the ability to manufacture product goes right to their bottom lines.

To comply with ISA/IEC safety system standards, internal faults in an SIS must be detected and repaired during a relatively small time window. Unless the faulty module can be repaired or replaced on line, a hastily scheduled production shutdown will be required to allow the fault to be remediated. SIS modifications, upgrades and periodic proof testing must also typically be performed during scheduled outages, creating additional pressures for plant staffs that are already typically working around the clock to complete the given tasks in the shortest possible period of time so that production can resume.

First-generation TMR technology: a huge step in the right direction

Fault-tolerant redundancy approaches developed for the aerospace industry led to the development and introduction of a first generation of triple modular redundant (TMR) PLC-based safety systems in the mid-1980s. Essentially, these were triplicated PLCs using voting schemes that required two of the three (2oo3) logic solvers to agree before the system would initiate a safety trip. These were considered to be fault-tolerant, since the safety system could continue to function (although in a diminished state) even after a faulty logic solver processor was discovered. Hardware options for input/utput (I/O) also provided a degree of fault-tolerance relative to the sensing and actuating elements of the SIF's in the system.

This fault-tolerant TMR architecture with 2oo3 voting logic provides a very high degree of operational availability, since the occurrence of spurious safety trips due to internal SIS errors are greatly reduced.

First-generation, fault-tolerant TMR technology was created before the advent of the international standards that apply today. Originally, they were designed to provide more operational availability while providing the needed safety shutdown functions. As international standards were created, the first generation products were modified to meet them. In some cases, the standard placed requirements on these systems that could not be met, forcing users to implement additional protection outside of the first generation system.

However, the regulations also require that any single failure, even in a triplicated fault-tolerant system, must be repaired within a finite time period. (Mean Time To Repair is one of the parameters that goes into the SIL certification calculations for a given SIS configuration). Thus, unless the faulty component can be repaired online, a previously unplanned shutdown will need to be hurriedly scheduled to be able to perform the repair. This is not generally a problem with the first generation systems, but, as we will see, some second generation systems backslid on this requirement.

Unplanned shutdowns equate to lost production. This can cost the owner-operator many thousands of dollars and really screw up the profit/loss report in any given month. With first-generation TMR systems, design and architectural constraints limited the ability to perform on-line system modifications or upgrades, requiring shutdowns to be scheduled that might otherwise have been avoided.

Second-generation SIS systems

As technology progressed, more SIS systems came on the market. Due to advancing technology, these systems tend to have more extensive diagnostic capabilities. Due to improved diagnostics, many of these products do not offer triple redundancy, the feeling being that the improved diagnostics make TMR architectures obsolete. In many cases, even simple redundancy is not the standard offering. This presents a problem that did not exist with first generation TMR systems. Since the SIS is simplex, it cannot be repaired on line, meaning that according to ISA84, any fault in the SIS will require that a shutdown be scheduled.

On the other hand, new redundant architectures were introduced with 1oo2 and 2oo4 schemes.

With some limitations, most of these systems work very well. The primary fault in these second-generation systems is not in the diagnostics or in the redundancy or lack thereof. The

fault is in the target. These systems were designed to do as well as or to incrementally improve upon the first generation systems. Third-generation systems were designed to provide the user with the safest system available and to provide maximum uptime.

Third-generation fault-tolerant systems bring safety and operational availability to the next level

The latest generation of fault-tolerant, simplex, dual-and triple-modular redundant safety instrumented systems combines all the benefits of first-and second-generation SIS technology with improved diagnostics, safety availability, operational availability, and significantly lower lifecycle costs. This has been accomplished by implementing new, more flexible redundancy approaches, increased diagnostic coverage, better processing and communications performance, plus new online repair, modification, and upgrade capabilities. The goal of these systems is not incremental improvement, but providing the user with the SIS that will provide the very best protection against unsafe situations while providing minimal interference with the operation of the facility.

Furthermore, unlike earlier generations of safety system technology, the new simplex, dual-and triple-modular redundant SISs can often come configured as SIL certified "right out of the box," with few, if any, restrictions imposed by the certifying body.

The net result is that third generation redundant safety systems can often deliver significantly increased integrity and availability over first and second-generation systems. With safety integrity in excess of 99.9999 percent ("six nines") when configured in triplicated fashion, third-generation SIS can eliminate the one in ten outages attributable to the control system while significantly reducing nuisance trips to provide operational availability in excess of 2000 years.

With third-generation safety systems, any faults are automatically identified by the system without the need for user application programming. Significantly, in TMR configurations, these third-generation safety systems will always continue to operate in a safe manner in the presence of a single fault and, in many cases, in the presence of multiple faults. Unlike earlier PLC-based solutions, third-generation TMR safety systems will always fail in a safe manner, even when faced with multiple faults.

A more flexible (and robust) approach to redundancy

Traditional TMR systems are fairly rigid when it comes to redundancy, since the architecture is usually dictated by the supplier. With third-generation systems, redundancy levels can be fully defined by the user. This includes the ability to freely employ simplex, dual-redundant, or triplicated modules on a SIF-by-SIF basis. I/O redundancy levels can be software-defined on a point-by-point basis. An input or output can be wired to a single input point on a card, to multiple points on the same card, or to different cards in different chassis'. This allows the user to cost-effectively "dial in" the precise level of redundancy and fault-tolerance required at various points across the system. Redundancy is handled by the system and is transparent to system users.

With first-generation TMR systems, triplicated CPUs must all be mounted in the same enclosure. This makes all three susceptible to possible physical damage from the same accident (such as when an inattentive forklift operator smashes into a system enclosure). With third-

generation systems, redundant and/or triplicated CPUs can be mounted in separate chassis or in the same chassis, at the user's discretion.

Redundant communications throughout the system further reduce the chance that a single SIS fault will result in a process interruption.

Increased diagnostic coverage and error-checking

Third-generation TMR SIS's typically offer substantially increased diagnostic coverage relative to first-or second-generation systems. The comprehensive built-in diagnostics automatically pinpoint any faults to the field-replaceable module level. Continuous diagnostics performed with every scan typically test:

- processor module integrity
- communications integrity
- backplane integrity actuator and field wiring
- input/output card field interfaces
- transmitter and field wiring

Forced fault diagnostics ensure that the advanced diagnostics are working properly and the system can remain in operation.

Under today's stringent safety standards, periodic proof testing of the SIS is required to enable latent faults to be identified, isolated, and remediated before they can negatively impact the availability of the safety function. With conventional TMR systems, additional software or other outside resources are often required to be able to perform the proof test and/or analyze the results. In contrast, with third-generation systems, SIS proof test software is provided with the system and SIS proof testing is completed by the customer onsite. In some cases, this can be accomplished simply by turning the power on the logic solvers off and then on again.

Increased on-line diagnostic coverage and frequency in third-generation systems also reduce the frequency with which off-line proof testing must be performed to maintain SIL certification, thus further reducing SIS-related operational downtime. In some cases, logic solver proof test intervals can be as long as ten years, thus providing extensive scheduling flexibility to move proof testing from being a major concern for those responsible for maintaining SIS SIL certification to a minor consideration.

Increased diagnostic coverage also helps improve operational availability by detecting, identifying, and locating SIS faults or errors that could result in spurious "safe" trips before they have a chance to occur. In this manner the increased diagnostic coverage helps avoid the negative impact that nuisance trips can have on operational availability and plant profitability.

To ensure the performance integrity of the SIS, sophisticated error checking routines are now employed to continuously test all system hardware, communications, and calculations. Transmitter and actuator field wiring, I/O field interfaces, processor integrity, backplane integrity, and communications integrity are all tested with *every scan*. Data is exchanged only after both command and address lines are tested. Watch dog hardware and software timers with separate time bases monitor any abnormal program execution.

Increased performance

As everyone who is familiar with Moore's Law knows, computing technology has progressed at an extremely rapid rate in the years since the first-generation digital SISs were introduced. Microprocessors have become much smaller, more powerful, and more efficient. Tremendous strides have also been made in the area of parallel processing. This enables multiple microprocessors to work in parallel to perform even highly complex calculations at a much higher speed than even much larger and more costly monolithic computers. This is one of the reasons why the latest breakthroughs in supercomputing involve parallel processing.

In contrast to first-and second-generation SIS's, which were largely designed to work in a serial manner, one process at a time, third-generation SIS's can have a hundred or more processors working in parallel to dramatically improve logic solving, communications, and I/O scanning performance. Processing power in now available to solve in excess of 500 control loops every 10 msec with a reaction time of 25 msec. This includes scanning of I/O, logic solving, and alarm handling functions as well as peer-to-peer and other communications functions. (Compare this to a typical system reaction time of 200-500 msec for first-generation SIS).

Unlimited online repair, modifications, and upgrades

On-line repair capability is the ability to swap out faulty modules for good ones without interfering with the manufacturing process. This is an important characteristic for all mission-critical control systems, but especially so for process safety systems. All dual-or triple-modular-redundant safety systems offer some degree of on-line repair capability. However, it's not clear how on-line repairs could be safely performed with single modular configurations, since there would be no backup module available to initiate a SIF if needed during the time that the faulty module was being replaced. Thus, while the system might be able to continue operating when faced with a single fault, a shutdown would need to be hastily scheduled to enable the faulty module to be replaced within the allotted mean time to repair (MTTR).

To be able to perform online module repairs with conventional TMR systems, every other slot must typically be kept free for a hot spare module. This increases the number of chassis' required and takes up a lot of valuable plant real estate. In contrast, with third-generation systems, every chassis slot can be used for an active module and every module can be swapped under power without interrupting operation. Combined with the extensive built-in diagnostics, the ability to swap out faulty modules under power results in the lowest mean time to repair specifications of any class of safety system.

Finally, with first-generation TMR systems, the number of online system upgrades that can be performed is limited by the size of the buffer allocated for this purpose. Even some second-generation SIS systems use change buffers in this manner. Once the buffer is filled, additional upgrades could only be performed during a scheduled shutdown. With third-generation systems, an unlimited number of system upgrades can be performed without requiring a shutdown.

By dramatically reducing the number of nuisance trips through flexible redundancy and reducing system downtime requirements through unlimited online system repairs, modifications, upgrades and reduced proof test requirements, these third-generation safety

systems can help today's process manufacturers to achieve required safety integrity levels while significantly improving availability and delivering appropriate safety integrity levels.

Conclusion

Today's process manufacturers are faced with a confusing array of options when it comes to selecting and implementing safety instrumented system functionality in their plants. These range from safety systems that have no commonality or integration with the plant process control platform, to systems that are virtually indistinguishable from the process control platform. As is often the case, each vendor touts their particular approach as the best approach. In fact, while most vendors can make compelling arguments to support their case, each approach will involve tradeoffs between safety integrity, operational availability, and cost.

For this reason, it's very important that, before selecting a safety system vendor, owner-operators must carefully "look under the covers," to be certain that they understand what tradeoffs might be involved and whether these tradeoffs are compatible with their operating philosophy.

About RTP Corporation

Founded in 1968, RTP Corp. is a developer and manufacturer of high-performance critical control and safety systems. RTP Corporation's products serve applications for both basic process control and safety systems. Markets for RTP's products include Refining, Upstream Oil and Gas, Chemical, Nuclear Power, and Glass Industries. RTP offers a wide range of rugged hardware and a complete suite of software for industrial control solutions that include seamlessly redundant and triplicated systems for mission-critical applications.

For more information on the current range of third-generation TUV-certified simplex, dual-modular, and triple-modular redundant safety systems from RTP Corporation, readers can visit www.rtpcorp.com, send an e-mail to rtpinfo@rtpcorp.com, or call 954-974-5500.