Cybersecurity and functional safety: the case for embedded analytics

An integrated approach to ISO26262 and ISO21434 compliance

Contents

Introduction

From advanced driver assistance systems (ADAS) to a new generation of robots and medical systems, we are seeing an explosion in the development of cyber-physical systems. Because these systems use advanced software to interact with the physical world, security and safety are paramount concerns. These issues are reflected in many industries by the use of safety and security standards based on a philosophy of risk assessment and reduction.

Functional safety (FuSa) has for many years been a mainstream concept in automotive electronics, as reflected by the ISO26262 standard, which plays a key role in assuring implementors and users that road vehicles are safe. More recently, it has become increasingly clear that cybersecurity for connected vehicles represents an equally important issue: this is captured in a newer, related standard, ISO21434.

Although primarily intended for use in the automotive industry, the methodologies embodied by these standards have a broader application. While this white paper focuses primarily on automotive applications and standards, the same basic principles and philosophies are applicable to other domains: for example medical, aerospace or industrial.

In common with ISO26262, standards such as IEC61508 for industrial control, EN50126 for railways and IEC 60601-3 for medicine are risk management-oriented frameworks. All have ramifications for system-on-chip (SoC) designers working across technologies and industry sectors. They call for the ability to monitor internal behavior to a high degree of granularity not only during development but, increasingly, after deployment.

As it turns out, functional safety systems are effectively a subset of cybersecurity-critical systems. And there is significant overlap between safety and security process elements (Figure 1 shows the links between the two, in the context of SAE J3601). So it is essential to consider the two in tandem.

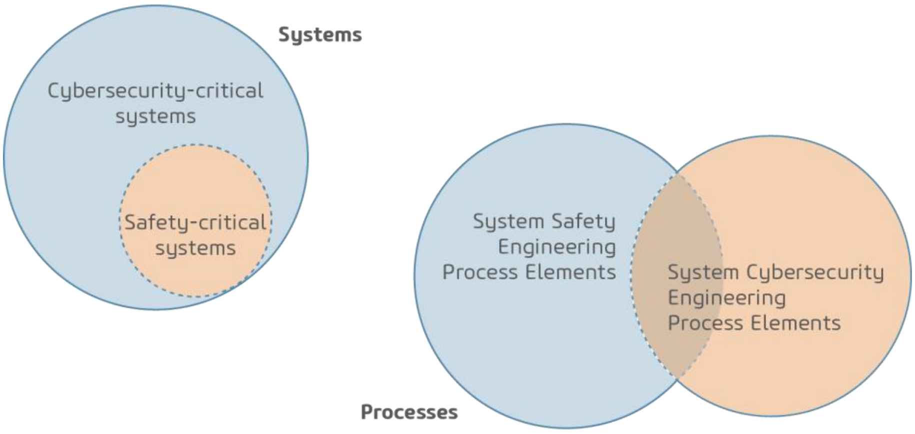
Systems
Cybersecurity-critical systems
Safety-critical systems
System Safety Engineering Process Elements
System Cybersecurity Engineering Process Elements
Processes

Figure 1: The SAE J3061 model of safety and security systems and processes

The cybersecurity challenge represented by connected and autonomous vehicles (CAVs) itself cannot be underestimated. CAVs utilize complex control and communication architectures and are connected to communications infrastructure and the Internet. They include sensor-driven intelligence and autonomous control, generating terabytes of data over hundreds of components. All of this means that detecting threats and providing assurance on security at a system level is a huge challenge.

On top of the technical difficulties, the 'cost of failure' is substantial. Cyber-crimes in the automotive sector have risen significantly recently and, as on-board electronic controls become more complex, it is likely that the sophistication of attacks will increase.

The potential socio-economic impacts of such attacks are tremendous, with the automotive industry alone estimated to lose £24 billion annually by 2023[1]. And as we have already noted, the need for

[1] Upstream Security Global Automotive Cybersecurity Report 2019; https://www.upstream.auto/news/press-release-global-automotive-cybersecurity-report-2019/

systems that are secure and resilient against cybersecurity threats expands to other markets including aerospace, implanted medical devices, Industrial IoT (IIoT), and critical infrastructure.

The emerging ISO21434 standard, which builds on SAE J3061, will provide a framework similar to that of ISO26262 to ensure a consistent, well defined and robust approach to managing cybersecurity risk and building a cybersecurity management system, while also allowing adaptation to a continually changing threat landscape.

All of these approaches and associated standards share a number of core principles. First, they place great emphasis on verification and validation: with especially strict tests and the ability to map each test back to requirements – and conversely, that every requirement be fully tested and verified. Second, after release and shipping, the system should be monitored and should include safety and security mechanisms that can detect issues and trigger a suitable response.

Against this backdrop, developers across multiple industries are recognizing the need for hardware-level monitoring and analysis of system behavior. UltraSoC supports those roles, with embedded analytics that enables both validation and in-life monitoring.

The changing face of engineering for safety and security

The automotive industry defined and adopted the ISO26262 standard, drawing on the deep experience that counterparts in industrial automation had acquired since the 1980s. That experience led to the creation of the International Electrotechnical Commission (IEC) 61508 standard, which marked a radical shift in emphasis in safety engineering that has been embraced across a number of sectors, including medical electronics, railways and the nuclear industry.

Learning from IEC 61508, ISO26262 advocates an approach that is based on risk analysis performed by engineers working as part of the development process. Safety engineers assess the various risks posed by the system, determine which need to be reduced and the level of reduction required and define the tests needed to establish that the risk reduction criteria have been met.

The safety requirements are specified separately from the functional requirements of the system. This makes it possible to compile safety evidence and perform safety planning separately from the core design process. This in turn enables the ability to demonstrate what steps have been taken to ensure safe operation, providing confidence that the risk-reduction measures were appropriate and effective. A vital consideration is whether the system will still be safe if a failure occurs. If risk analysis finds that is not the case, the system needs to be redesigned to allow it to fail in a safe manner.

One of the innovations of IEC 61508 lay in its use of the safety integrity level (SIL), which was adapted by ISO26262 to become ASIL (automotive safety integrity levels).

The emerging ISO21434 cybersecurity standard will take a similar approach, with threat analysis and risk assessment performed by security engineers as part of the development process using a specific methodology for the analysis, assessment and management of risk.

Security requirements are specified separately from the functional requirements of the system, making it possible to compile security evidence and perform security planning.

And finally, cybersecurity assurance levels (CIL) are introduced, which indicate the level of cybersecurity process rigor required to provide “defense in depth”.

To support both the safety and security case for each vehicle, bidirectional traceability is a key requirement – this links functionality in the final system to the functional, safety and security requirements defined in the specification that applies to the target system. Traceability ensures that the results of verification-phase tests are linked to the specifications.

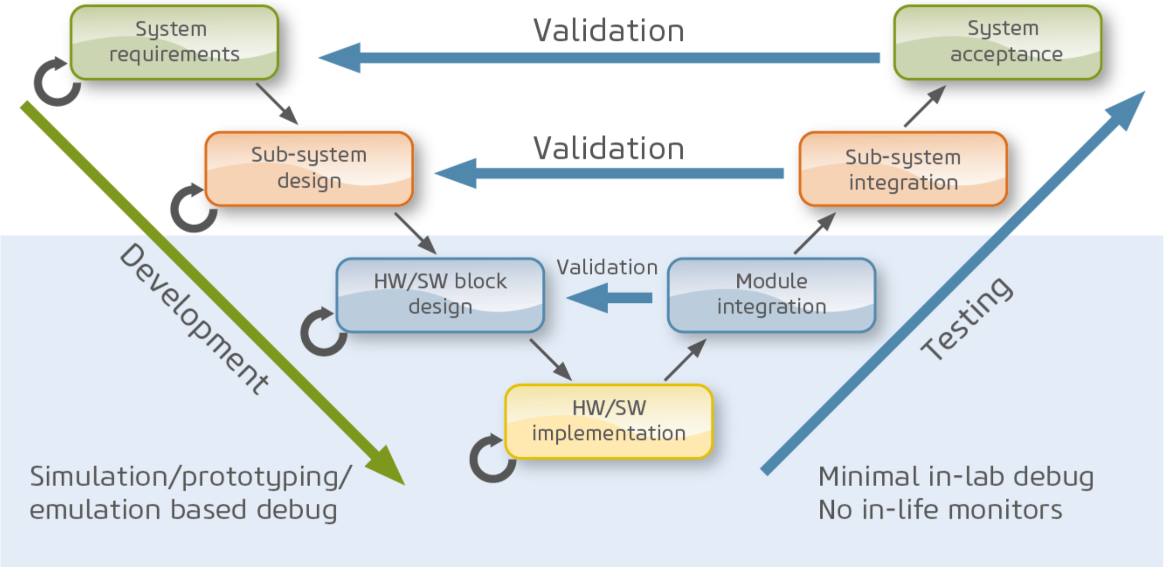

Figure 2 – The V-model development loop

The traceability loop is closed through the use of the V-model development process. The V-model guides the system from initial requirements to system acceptance, taking in subsystem and block design and integration phases on each leg. During the various integration phases, validation tests ensure that the specification has been implemented. With each update, the V-model ensures that changes reflect changes in the specification that apply to that vehicle for its complete set of updates.

Real-time monitoring for safety and security

A clear market need exists to monitor SoCs 'in-life' (in the real operating environment rather than in the lab), and this is further mandated by standards. But once all the modules and subsystems have been integrated into the vehicle, visibility using conventional test harnesses can be highly limited.

Test harnesses that can only make use of processor trace modules and conventional hardware probes have severe limitations. Probes can only be used at the board level. The bulk of hardware signals remain out of reach within the SoCs and 3DIC modules that automotive suppliers are using to reduce system cost and increase functionality. Trace modules are specific to each processor vendor and can only provide information on functions within the processor core and cache complex.

In many cases, accelerators lie outside the reach of a processor-centric trace architecture. Their internal state can only be accessed through software-level instrumentation that needs to be left in the production code if it is to be used to support a safety or security case.

Off-chip software monitoring solutions take hundreds of milliseconds to detect anomalies. This is generally inadequate for safety-critical applications. At 110kph (70mph), a vehicle will travel 3 meters (10ft) in under 100ms, a distance that could be critical to avoid an accident. Software-level solutions also inherently disturb the normal operation of the SoC, are highly visible and more prone to hacking, and are unable to monitor the entire SoC.

What really is required is an on-chip infrastructure that supports the entire SoC, responds in a timely fashion, and can be deployed non-intrusively so as not to perturb or interfere with the normal performance of the system. UltraSoC's monitoring and analytics IP provides such an infrastructure.

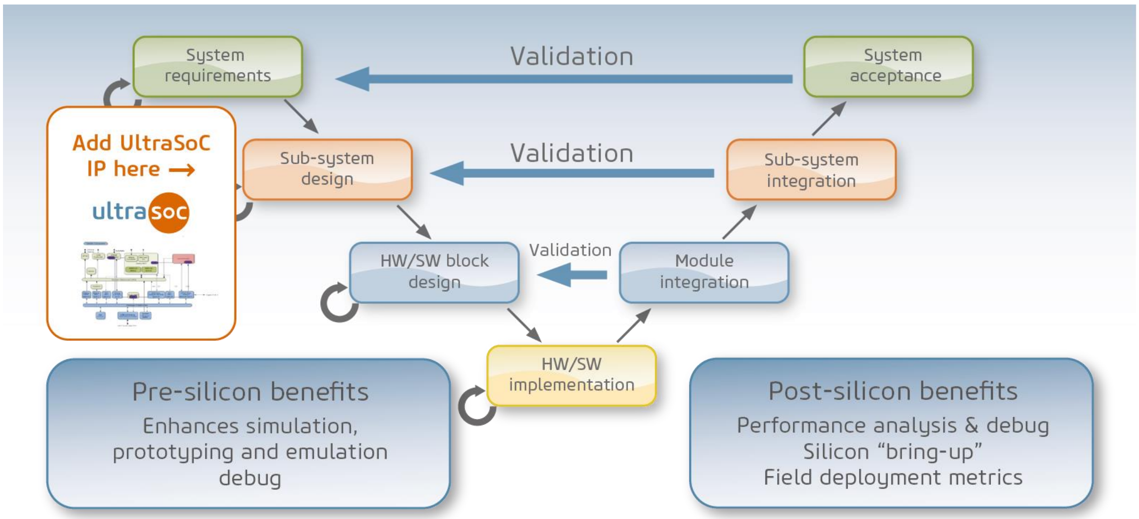

Figure 3 – The UltraSoC architecture supports the V-model

An UltraSoC monitoring infrastructure is built from small, 'smart' hardware blocks incorporated into the SoC to monitor its operation. They operate across the entire SoC, reporting rich real-time information captured via subsystem components, including processors, accelerators, buses and even custom logic. The system is vendor-neutral, providing an efficient way to integrate IP from different vendors and incorporate them into one coherent framework.

The UltraSoC architecture fully supports standards such as SAE J3061 that link security to safe operation and call for the monitoring and control of security though the entire product lifecycle from development and into usage in the field. Systems need to be able to monitor incidents and attempts to penetrate the system and report them.

Use-cases for safer and more secure systems

As an infrastructure for real-time SoC-focused monitoring, the UltraSoC architecture provides many options for implementing risk-reduction strategies for safety-critical and security-critical systems. UltraSoC monitors can check the long-term behavior of functional blocks, effectively becoming part of the security and safety mechanisms that protect them.

> "Semico believes the emergence of the UltraSoC SIP, unique in the market today, is greatly beneficial to the semiconductor industry, SoC designers and architects, and end users alike due to the improvement in response times and the ability to find and correct problems more efficiently than was previously the case. A new day is dawning in the industry which promises to bring a renewed sense of confidence in the electronics we use on a daily basis because of the ability to troubleshoot problems in a more timely and comprehensive basis."
>
> *Semico Research, Quantifying the Value of UltraSoC On-Chip Debug Capability for Time-to-Market, Revenue And Profitability*

Key characteristics to consider for in-life safety and security use-cases are "non-intrusiveness" (the ability to monitor, detect and raise an alarm **without disturbing the operation of the SoC**) and "timeliness" (**reacting most rapidly to a situation**). Because UltraSoC monitors operate in hardware, alerts can be raised in micro-seconds.

One example of the UltraSoC IP's use in safety and security is in supporting lock-step mode for groups of processors. A bus monitor can inspect the transactions performed by each of the protected processor cores to see whether they are consistent. If a mismatch is detected, the monitor can flag an error that is then picked up by recovery functionality in the SoC (Figure 4). This mismatch may be due to a failure within the system, or may be due to a malicious intervention, so detecting it quickly is imperative.

The Lockstep Manager can be extended to include processor trace and even register status monitoring, allowing behavior comparisons to be performed at any level of granularity.

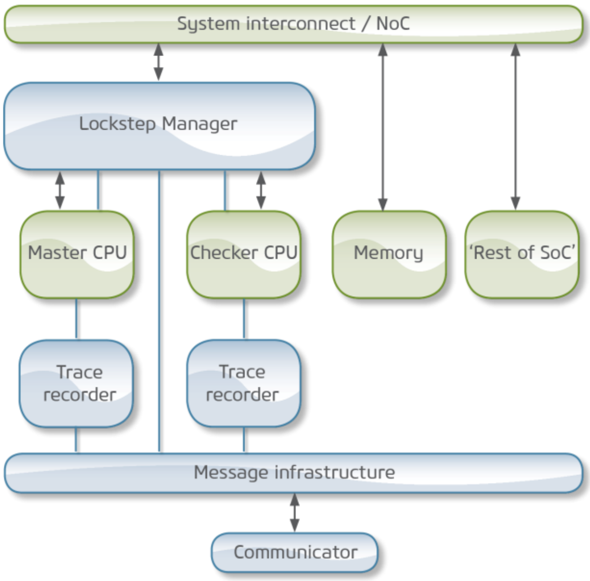

Figure 4 – The UltraSoC Lockstep Manager can monitor any CPU or custom block non-intrusively and at wire speed. It provides flexible and powerful options if a mismatch is detected, and is extendable to other error tests, redundancy or voting schemes

The UltraSoC IP can go further by supporting features that reduce the risk caused by aging effects such as electromigration and bias temperature instability (BTI). As automotive manufacturers start to use nanometer processes to support critical functionality, their SoCs are increasingly prone to usage and temperature related aging effects that can compromise functionality after a period of time. In combination with process, voltage and temperature monitors, the UltraSoC monitors can provide essential information on conditions within the SoC in real-time. These measurements can be used by the operating system or hypervisor to balance the load across parallel cores more evenly and so avoid the development of aging "hotspots" – ensuring longevity and continued safe operation. UltraSoC is working with IP suppliers such as Moortec to provide solutions that leverage advanced PVT monitors to improve SoC performance and reliability.

In addition, monitors on sensor interfaces can help support consistency checks that indicate whether peripherals are beginning to fail and so warn the software algorithms that employ their data. A specific example of this is the ability to detect, in hardware, that some pixels on a camera are "stuck" (Figure 5).

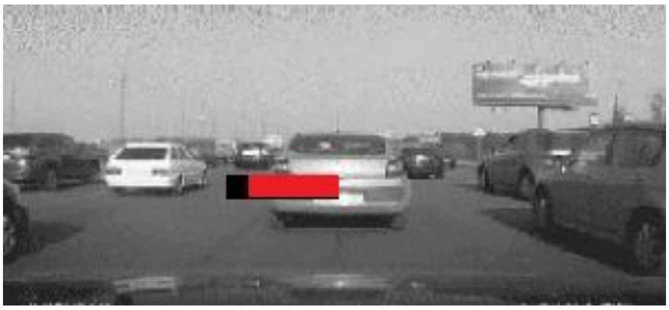

Figure 5 – On-chip analytics is a particularly effective way of spotting anomalous states in sensors, such as stuck pixels within a camera system. This can be achieved non-intrusively and at system speed, providing early warning of fault conditions before they manifest as operational failures

The security features in UltraSoC's IP allow monitoring of incidents even in the event of a successful intrusion, providing a forensic ("black box") recording that can be used to help build countermeasures against future similar attacks, and allowing reporting by non-compromised subsystems within the vehicle.

Such use-cases represent a small selection of the possible safety and security functions that can be supported by the UltraSoC monitoring and analytics solution. The combination of on-chip status sensors, non-intrusive hardware-level instrumentation and efficiently implemented analytical "smarts" provides many other options for supporting safety-critical systems and ensuring that such systems are secure by design.

Collateral support for ASIL (and CIL) readiness

An on-chip monitoring and analytics infrastructure is not expected to achieve an ASIL rating in isolation; but it provides key support for achieving the target ASIL rating for each system in which it is inserted. To support this process, UltraSoC is working to a program of ASIL readiness: providing features and functions appropriate to each ASIL rating.

Together with the functions and features provided by the semiconductor IP itself, UltraSoC can provide customers with safety collateral to support the certification process and the creation of the appropriate safety case. As the CIL definitions mature within ISO21434, we expect to be well positioned to be rapidly compliant.

Summary

The combination of market demands for reliability, flexibility and updateability in product design together with the need for safe and secure operation are leading to major changes in both the nature of, and the development process for, cyber-physical systems. The automotive industry has been at the forefront of these changes, but they affect most businesses where technology is deployed.

In-life monitoring of system behavior – in addition to improved in-lab validation before products are launched – is becoming an increasingly essential element in system design.

Hardware-based on-chip monitoring and analytics capabilities such as those provided by UltraSoC can address these challenges. Not only do they provide an effective monitoring framework which will help improve reliability and ensure safety and security, they also have significant advantages over traditional solutions – most notably in their ability to 'understand' the behavior of the entire system, to do so non-intrusively, and to do so at wire speed.