Trends and Drivers in Fail-Safe Architectures for Rail Systems

SMART Embedded Computing
www.smartembedded.com
May 2017

The market for embedded computing technologies in rail applications is following a similar trend as has been seen in other embedded market spaces. A layer of the technology value chain becomes 'table stakes'— delivering limited competitive advantage to a point that it makes sense for application providers to reallocate R&D resources to differentiating elements of the end product and buy the base technology from companies who are dedicated to that technology. We are witnessing this transition in the rail market for embedded computers that are certified to safety integrity level four (SIL4), the highest level. These embedded computers offer a certified, commercial off-the-shelf (COTS) generic fail-safe platform allowing rail application developers to focus their R&D resources on differentiating applications.

This white paper examines this trend, its drivers and the technical issues surrounding the emergence of SIL4 COTS computing platforms.

SMART™
Embedded Computing

The market for embedded computing technologies in rail applications is following a similar trend as has been seen in other embedded market spaces. A layer of the technology value chain becomes ‘table stakes’—delivering limited competitive advantage to a point that it makes sense for application providers to reallocate R&D resources to differentiating elements of the end product and buy the base technology from companies who are dedicated to that technology. We are witnessing this transition in the rail market for embedded computers that are certified to safety integrity level four (SIL4), the highest level. These embedded computers offer a certified, commercial off-the-shelf (COTS) generic fail-safe platform allowing rail application developers to focus their R&D resources on differentiating applications.

This trend is driven by a number of emerging trends in the global rail industry.

In the past few years we have witnessed an explosive growth in global investments in public rail transportation, in particular high-speed rail and metro, caused mainly by the effort to reduce a nation’s carbon footprint by replacing inefficient automobile-based transport with efficient mass transportation. This is particularly evident in emerging economies such as China and India, as well as established economies in the Far East, Africa and South America. While less so in Europe and North America, we do witness growth in these markets due to other factors such as pan-European rail standardization as well as modernization of the rail infrastructure to enhance safety.

However, a growing market, while creating an attractive target for COTS products, will not on its own cause an outsourcing trend. Additional safety, technical, and commercial factors come into play.

As train speed increases to 300 kilometers per hour and above, reliance on computers that control the rail infrastructure and the trains themselves increases exponentially. As an example, stopping a train that travels at 300 Km/h will only take two (2) minutes or so, but during those two minutes the train will travel ten (10) kilometers, requiring real-time and continuous monitoring of the rail network to provide early alerts of potentially hazardous events.

High-speed, high-availability, and fail-safe computer-based control equipment must be deployed to guarantee safe operation under all conditions. High-performance and high-availability computing expertise is relatively widespread, however fail-safe computing has been the domain of a few expert companies, located mostly in Europe (Alstom, Bombardier, Siemens, etc.) for SIL4 certified systems, and Japan (Nippon Signal, Hitachi, etc.) for certification to Japanese safety standards and deployed locally. Fail-safe know-how has not been prevalent in other markets that are investing in rail networks, relying on mostly European vendors for acquiring the fail-safe systems (e.g., India, Africa, South Korea) or for forming joint ventures with these same European vendors to develop fail-safe systems for the local market (e.g., China).

The demand for SIL4 certified equipment has been further fueled by safety incidents that have driven governmental bodies to make it mandatory for all new installations to be SIL4 certified, and that non-SIL4 certified equipment in use today must be upgraded to SIL4 certified equipment. For example, the South Korean government mandated that rail equipment be upgraded to SIL4, and the Indonesian rail authorities have recently issued an RFP to upgrade their infrastructure to SIL4 certified equipment.

Another interesting trend in the global rail market is the aspirations of Asian application providers and rail integrators to expand their reach and penetrate overseas markets. Witness Hitachi’s establishment of a design center in London, recent announcements from Chinese vendors of wins in the US and Africa, as well as efforts by South Korean vendors to expand into former Soviet Union countries. Almost without exception, SIL4 certified equipment is a mandatory requirement.

A few major and factors emerge from these trends that are the root cause for the emerging trend to outsource SIL4 certified application platforms:

The lack of SIL4 development expertise by Asian rail application providers and the barrier that poses to aspirations to expand into overseas markets.

The threat to western vendors posed by the entry of Asian vendors into the global rail market and the price erosion that would likely bring (witness the impact Huawei had on the global telecom market).

The prevalent architecture implemented by existing fail-safe computers is no longer capable of handling the required performance, requiring an expensive development effort in ‘table stakes’ base technology.

Lockstep Architectures

Most rail systems today use an architecture called hard lockstep, whereby two processors execute the same instruction at the same time and drive their respective address and data buses in synchronization.

When operating in hard lockstep, the processors’ clocks are synchronized and, before allowing a transaction to drive external equipment, all data and address bits driven by the two processors are compared. If the bits are exactly the same, then the address and data information are allowed to change the state of external equipment. If they do not compare, then a failure is declared and the system is brought to a safe state and is prevented from driving external equipment.

Since, in hard lockstep, comparison is performed at the address and data bits of the processors, a primary and mandatory requirement is that the two processors must execute the same instruction, at the same time, to the same external resources (memory, cache, I/O, etc.). To do so, the processors themselves must be deterministic. We call the boundary created by the comparators the deterministic boundary (Figure 1).

Unfortunately, hard lockstep cannot be implemented using modern processors. The first problem is that modern processors do not guarantee deterministic behavior.

Multi-threading creates multiple paths for the execution of the program. Responses to soft errors in memory and I/O will cause divergent execution paths and timing. For example, errors that are caused by cosmic rays and change a bit in the register are not synchronized and not deterministic. This is more prevalent in current technologies because of the geometries of the transistors, which are so small that cosmic rays can flip bits. Also, other CPU features such as power management and cache operations introduce non-determinism.

The second problem is that it's practically impossible to synchronize the data pairs of two different modern CPUs. The use of on-chip devices to multiply clocks prevents synchronized operation, multiple memory channels and serial peripheral interfaces also make it impossible and it's not practical to synchronize buses operating in excess of 1GHz.

Another problem with a hard lockstep system is that it is fundamentally a closed system. Everything is tuned to work together, and it has to be all synchronized such that it's very difficult to upgrade technologies without affecting the total system. So the bottom line is that hard lockstep is just not possible any more with advanced processors.

SMART Embedded Computing has developed an alternative approach we call data lockstep architecture, whereby a deterministic boundary is created at the output stage of the processor board to the system data fabric that connects the processors to external devices. Before the processor boards are allowed to change the state of external equipment by driving packets on the data fabric, their packets are compared to ensure that they are the same. If they are the same, then the transaction is forwarded to external equipment; if the packets do not compare, then a failure is declared, and the system fails safe; i.e., it is prevented from changing the state of external equipment.

As shown in Figure 2, the deterministic boundary is not at the processor itself but rather at the edge of the processor and before packets are placed on the data fabric.

The benefit of data lockstep is that it makes it possible to use modern processors and deliver the performance required by modern rail applications.

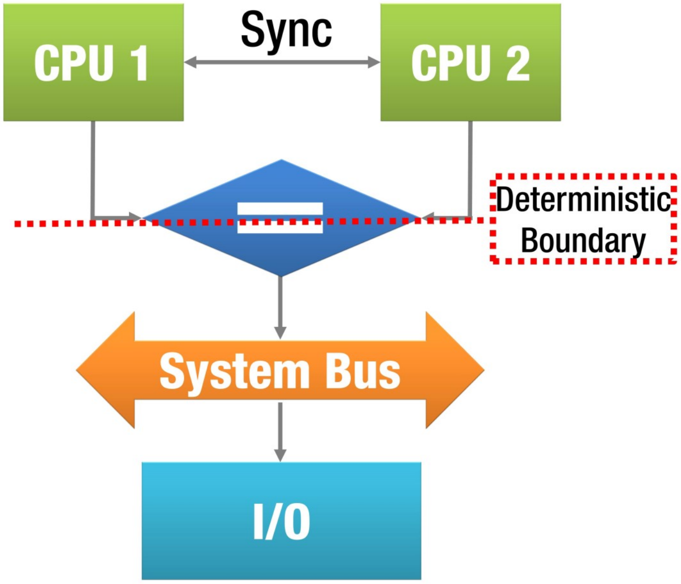

Figure 1. When operating in hard lockstep, the processors' clocks are synchronized and, before allowing a transaction to drive external equipment, all data and address bits driven by the two processors are compared.

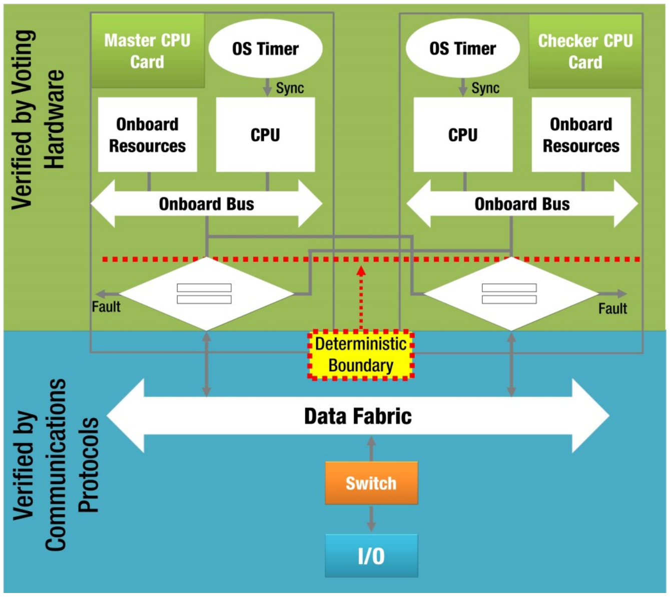

Figure 2. The deterministic boundary is not at the processor itself but rather at the edge of the processor and before packets are placed on the data fabric.

2oo2 or 2oo3

There are two methodologies for voting in a fail-safe system. They are called two-out-of-two (2oo2) and two-out-of-three (2oo3).

As shown in Figure 3, in 2oo2 voting, two computer elements compare the results of their computation and, if they compare, the transaction is driven to external equipment. If they don't compare, a fail-safe state is entered.

SMART EC's ControlSafe™ Platform implements a dual 2oo2 architecture to deliver high availability. In case the first ControlSafe Computer fails, the second redundant one takes over and continues running the application.

In 2oo3 voting, (Figure 4) three computing elements execute the application, and if the three don't agree then the system determines which one is at fault, disables it, and continues running with two. If the two disagree, then the system enters its fail-safe state and is prevented from changing the state of external equipment.

While both of these voting methods deliver the required safety and availability, the 2oo3 method is more complex to implement than the 2oo2 method. In the 2oo2 method, in case of a mismatch, the failed CSC enters its fail-safe state and the second CSC is enabled to run the application. No failure analysis, or fault isolation, hot-swap or re-integration is required.

On the other hand, in case there is a mismatch in a 2oo3 voting, failure analysis, fault isolation, switching to 2oo2 voting mode, module hot-swap, module reintegration, and re-enabling 2oo3 voting are all required. This is complex, and complexity leads to design errors.

For this reason, SMART EC's ControlSafe Platform chose the 2oo2 voting method. A simple design is a safe design.

ControlSafe Architecture Highlights

SMART EC's ControlSafe™ Platform employs data lockstep synchronization and 2oo2 voting. The system runs Wind River's VxWorks 653 operating system, which has been deployed in many fail-safe avionics-certified applications, including extensions to assure the task level synchronizations needed to implement data lockstep. All voting is implemented by hardware using proprietary FPGAs, making it transparent to application software, and easing porting of existing applications.

The architecture is flexible and expandable. All intra

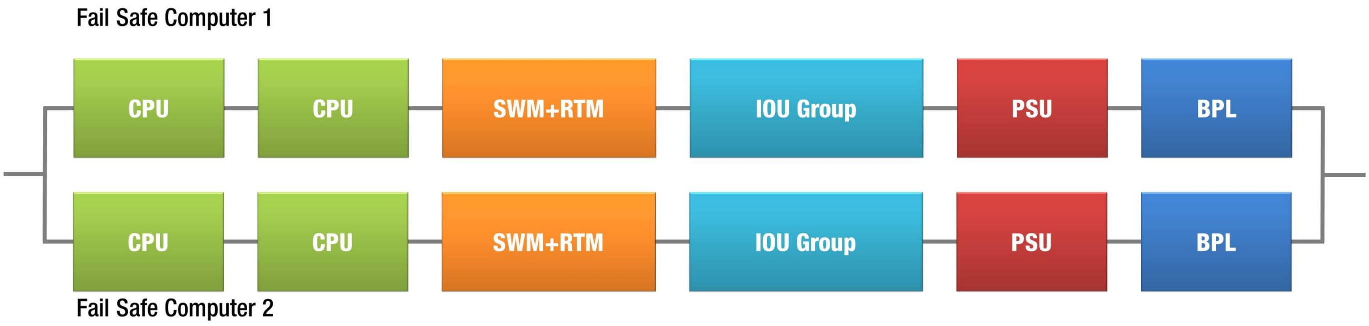

Figure 3. The implementation of a dual 2oo2 architecture to deliver high availability.

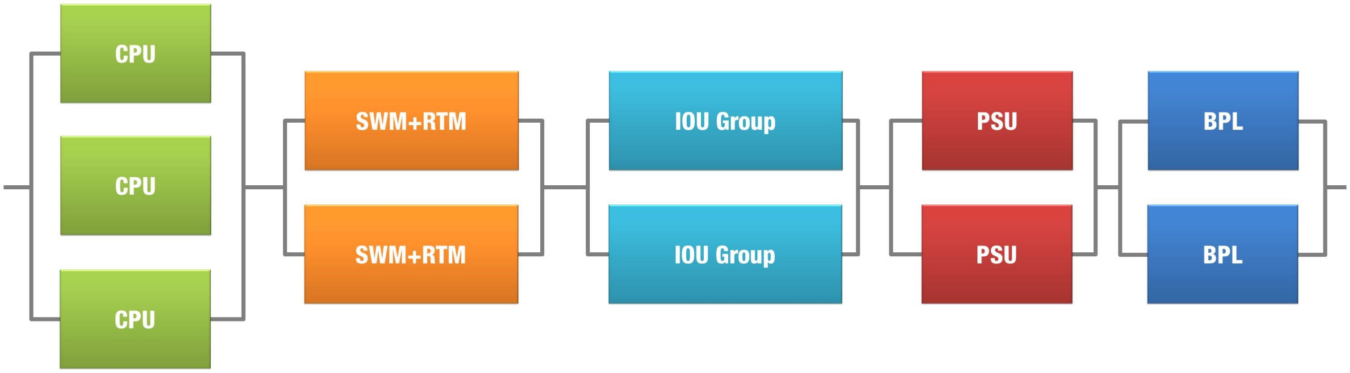

Figure 4. In 2oo3 voting, three computing elements execute the application, and if the three don't agree then the system determines which one is at fault, disables it, and continues running with two.

Dual Redundant 2oo2 System	Feature	Single Redundant 2oo3 System
2oo2 Only	Voting Logic	Must switch form 2oo3 to 2oo2 and back to 2oo3
Box Fail-over	HA Model	Module Fail-over
Not Required	On-line Fault Isolation	Required
Not Required	Hot Swap	Required
Not Required	On-line Module Re-integration	Required and performed by User Application
Failing CSC is diagnosed off-line; low risk of total outage	Human Factors	Risk of total outage due to hot-swapping wrong module
Simplex	Backplane	Redundant

Figure 5. Voting method comparison.

system communications are over the data fabric and are based on Ethernet. All I/O modules are connected via Ethernet such that expanding the system from local to remote or expansions in the I/O environment is straightforward and scalable (Figure 6).

Conclusion

In conclusion, the ControlSafe Platform from SMART Embedded Computing is a cost-effective, modular and scalable system that is based on open industry standards. The system is future-proof and provides protection for the customer’s investment because the architecture enables upgrades to both the CPUs and the I/O modules independently of each other.

Certified to SIL4 by TÜV SÜD, one of the most trusted certification bodies worldwide, SMART EC's ControlSafe Platform brings to customers all the benefits of outsourcing table-stake technology – accelerated time to market, significant savings in R&D and certification costs, and the ability to focus their effort and their R&D on differentiations from their competitors.

For more information on the SMART™ Embedded Computing ControlSafe™ platform, please visit

https://www.smartembedded.com/products/category/controlsafe

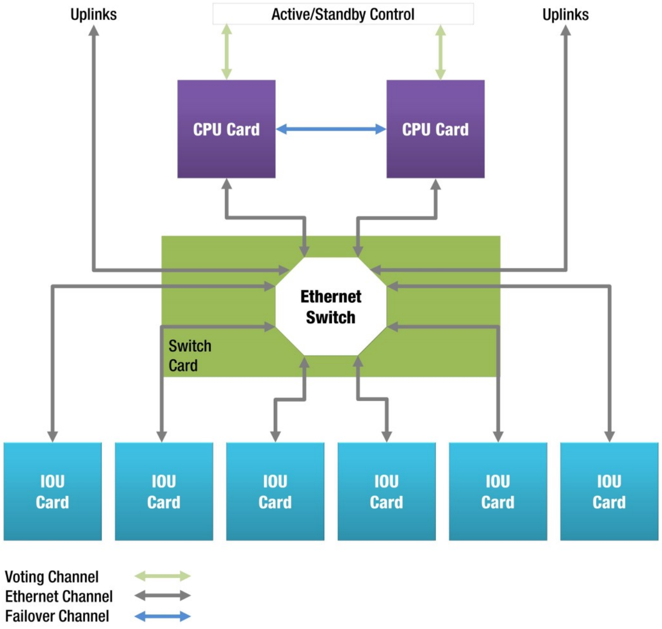

Figure 6. All I/O modules are connected via Ethernet such that expanding the system from local to remote or expansions in the I/O environment is straightforward and scalable.

Figure 7. The SMART EC ControlSafe Platform uses an OS, VxWorks 653 from Wind River, with a track record in numerous fail-safe applications, including avionics.

About SMART Embedded Computing

SMART Embedded Computing (SMART EC) is part of the SMART Global Holdings, Inc family of companies.

We are a global leader in the design and manufacture of highly reliable embedded computing solutions for a broad range of defense, industrial IoT (IIoT), edge computing, and communications customers.

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

For more than 40 years, customers have trusted us 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.

Our engineering and technical expertise is backed by world-class manufacturing, global sales offices and advanced worldwide logistics capabilities that can significantly reduce time-to-market and help customers gain a clear competitive edge.

Contact

+1 602-438-5720

info@smartembedded.com

www.smartembedded.com