DEFENSE SOLUTIONS

Five Budget-Busting Cost Drivers that Threaten Sensor Processing System Design

cwcdefense.com

CURTISS WRIGHT Controls
Defense Solutions

CURTISS WRIGHT Controls Defense Solutions

Five Budget-Busting Cost Drivers that Threaten Sensor Processing System Design

Recent advances in open architecture standards and the rise in performance and inter-connectivity found in the latest generations of rugged COTS products have made designing very high performance, scalable embedded Sensor Processing systems for defense and aerospace applications much easier, cost-effective and flexible. However, experience has revealed a number of unique hurdles, each a potential threat to reaching their program's goals, that designers of the complex High Performance Embedded Computing (HPEC) systems used for demanding sensor processing applications, need to fully understand. We have identified the five most common "traps" that can increase design risk, drive cost overruns and delay schedules. For each of these traps we have also identified the most effective mitigating solution.

The 5 Cost Drivers that Threaten Sensor Processing System Design:

1. Funding Delays
2. Lack of Software Portability
3. Interoperability Surprises
4. Program Longevity
5. Evolving Requirements

Figure 1: Sensor Processing Systems Provide Back-End Processing for Sensors on Mobile Platforms

CURTISS WRIGHT Controls Defense Solutions

Trap #1: Funding Delays

All System Integrator program managers are dedicated to delivering their designs on-time and on-budget. If developing their hardware in-house, the program manager may maintain a design team to develop the processing modules intended for use in their Sensor Processing System. To bring a CPU, FPGA or other processing module to deployment requires an extensive staff. It may include logic designers, layout designers, project managers, mechanical engineers, systems engineers, system architects, software engineers, DVT engineers, manufacturing engineers, supply chain management, and product marketing among other areas of expertise. Typically, the salaries and overhead associated with an in-house design team is paid through the program budget. Unfortunately, program costs are fixed and don't vary regardless of whether the team is productive or idle. When a design takes longer than expected, the additional calendar time is burdened upon the project. As the schedules proceed, costs accumulate due to fixed overhead and any overruns can threaten a program's viability. Program delays are costly and have many causes, but delays commonly occur due to technical risk or uncertainty in the funding climate and uncertain times can increase schedule variability.

Solution

Increasingly, integrators are moving to a COTS system development model to avoid the fixed overhead associated with in-house engineering. Leading COTS vendors, such as Curtiss-Wright have years of experience designing and manufacturing best-in-class processing modules for the rugged, deployed computing market. One benefit for system integrators is that the volume and competition in the marketplace drives efficiencies into the pricing of COTS vendor products. Even better, when a system integrator uses COTS boards their cost structure shifts from fixed to variable. For example, if their program is delayed, the COTS approach eliminates the need to cover overhead: the purchase of required processing modules is simply deferred until needed. Moving to variable costs with COTS frees up much needed capital, enabling more investment in core competencies to foster competitive differentiation, which ultimately delivers the greatest ROI.

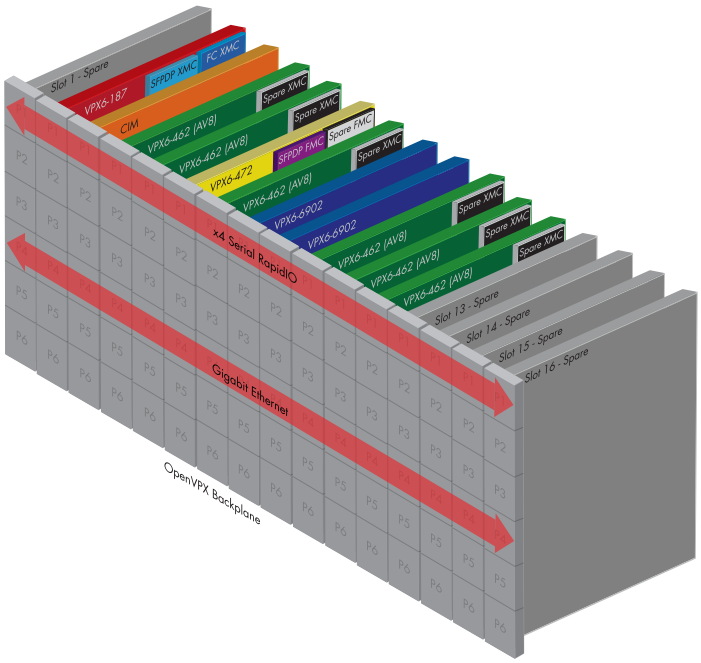

Figure 2: The OpenVPX Platform Enables a Mix of Heterogeneous Processing Elements, Interconnects and IO

CURTISS WRIGHT Controls Defense Solutions

Trap #2: Lack of Software Portability

Over the years, the computing world has witnessed the emergence and passing of numerous architectural platforms. Each successive generation of innovation has introduced new performance features that impact system software, sometimes in small and sometimes in significant ways. For example, in the mid-1990's, a leading CPU vendor introduced a SIMD architecture that for a time became a dominant processing element in defense and aerospace embedded computing. In response, a large installed base moved to this numerical processing approach. When the vendor removed the popular SIMD engine from their device architecture, a competing processor vendor emerged with their own SIMD engine and an alternative API library. Similarly, the recent emergence of multi-core processing required many legacy applications to be re-architected around the new multi-core SMP model.

In a recent example, the popularity of OpenCL™ and CUDA™ devices has driven significant amounts of software recoding and porting. We've seen it occur repeatedly: As each new technology emerges to claim performance leadership, it arrives with its own programming model, its own architecture, and its own APIs. In an ideal world application code would simply recompile on each new generation of hardware. While that promise might be obtainable in some static situations, Sensor Processing Systems are performance hungry embedded applications, always seeking the next great performance boost. As a result, they regularly, face the monumental and expensive task of re-architecturing and porting code. As a rule of thumb, it is useful to consider that the cost of one engineering month can far exceed the cost of a single hardware module.

Solution

One proven method for mitigating the ceasless cycle of architecture adoption and re-coding is through use of a middleware platform that will remain constant over succeeding generations of processing technology. Examples of stable middleware include Gedae®, MPI, Continuum™ IPC, and DDS. Applications that are architected around these software components have a much greater chance of minimizing the potentially disruptive change brought about by technology insertion. Use of a stable middleware platform translates into increased flexibility for adopting new technologies, which ultimately delivers greater performance and lower cost.

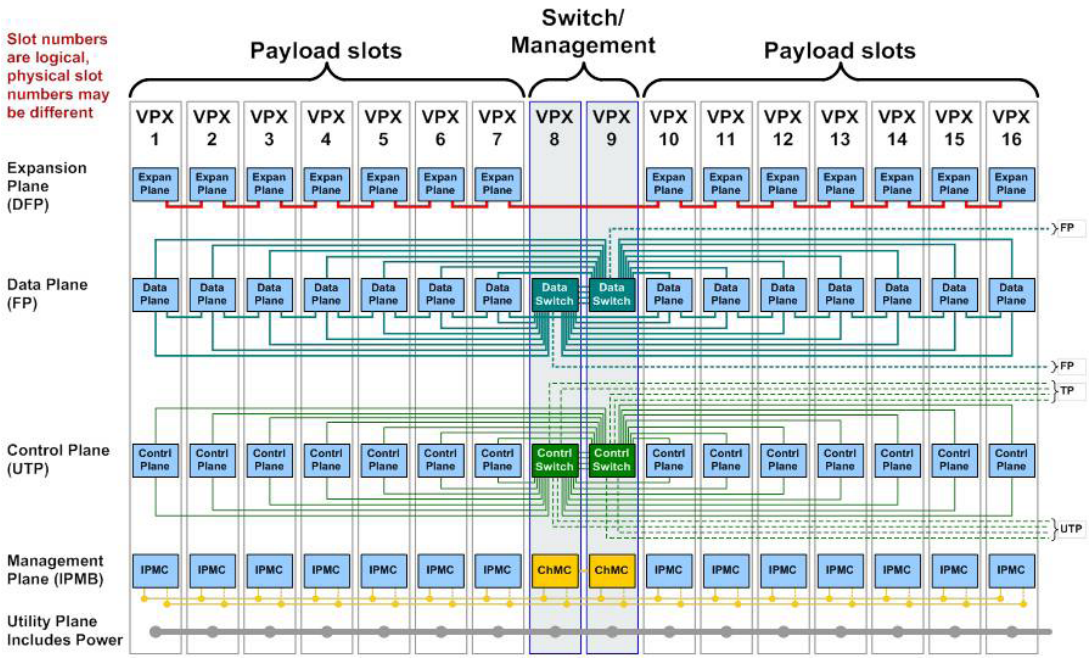

Figure 3: The 16-slot OpenVPX Profile Fits in a Standard 19" Rack and Contains Provisions for both Payload and Switch Modules

CURTISS WRIGHT Controls Defense Solutions

Trap #3: Interoperability Surprises

Sensor processing systems typically comprise a mix of heterogeneous functional elements such as OpenVPX™-based CPU, FPGA, and GPGPU cards. These modules use a variety of interconnects such as Ethernet, Serial RapidIO®, InfiniBand™, and PCI Express®. In addition, these modules require a range of I/O devices, including displays, peripherals, and storage, as well as a variety of software types such as operating systems, stacks, and middleware. The good news is that, compared to what existed previously, the OpenVPX system architecture has successfully brought a beneficial level of order to this array of open architecture modular solutions. OpenVPX defines standardized backplanes and protocols to facilitate interconnection and the needed profiles from which designers can select to plan their system. While it constitutes a true sea change in terms of easing system design, system designers who think that OpenVPX eliminates all of their system interoperability worries may be in for a rude awakening.

For example, while two different types of boards from different vendors may support a particular interconnect standard, each of those standards will typically define numerous protocols. RapidIO, for instance, supports message passing, logical I/O and data streaming protocols. In the Ethernet domain, its numerous protocols include TCP, UDP, sockets, and RDMA, among many others. And protocols themselves may support multiple sub-types, as within RDMA we find iWarp and ROCE. Once the challenge of matching all interconnect protocols is addressed there still remains the issue of system software. One potential critical hurdle is whether or not the applications use the same protocol in the same way.

Put into economic terms, the "interoperability surprise trap" can be daunting. If we calculate an engineer's time as costing a company $100-200 per hour ($4000-$8000 per week), it isn't hard to see how software development work, such as coding, test suite development, regression testing, validation, and remediation can end up costing significantly more than the recurring cost of the underlying updated hardware.

Solution

One effective approach for mitigating system interoperability risk is to work with suppliers that apply a system-level approach to their module designs. Curtiss-Wright modules, including Single Board Computers, DSP, FPGA, GPGPU, I/O and Fabric interconnects, are architected to work together using a common set of interconnects, protocols and software components. Our philosophy is to maintain a common software platform over continuing generations of module design. While this approach can't insulate a project from every possible integration challenge (for example, the challenge of migrating between CUDA and OpenCL), it does enable the system integrator to choose which elements of the architecture to keep stable and where to introduce change. This system design approach lets the integrator decide where to focus their limited technology insertion investment while keeping the rest of the system as stable as possible.

Image courtesy of Department of Defense

CURTISS WRIGHT Controls Defense Solutions

Trap #4: Program Longevity

Stability and predictability are essential for deployed COTS-based mission-critical military applications such as those served by Sensor Processing systems. These systems tend to have an in-service life far longer than the typical commercial production period of some of the system's key semiconductor components. To ensure that a system design can be supported for the full duration of the long lifecycles demanded by these programs, system designers need to know their lifecycle management plan options, which should include both the options offered by their COTS vendor partners as well as services they can internally develop themselves.

Key lifecycle management services essential for supporting the duration of Sensor Processing system platforms, include an ongoing review of product configuration changes and component obsolescence, a quarterly bill of materials (BOM) health check, and a longevity of repair plan. A closer examination of the volume production phase, the post-production phase, and component storage and handling is also vital when deciding on which road to obsolescence management best meets a company's lifecycle management needs.

Image courtesy of DefenseImagery.mil

Solution

A comprehensive lifecycle management strategy is the key to safeguarding programs and mitigating the challenges associated with COTS technology deployed in long-term mission-critical systems. In addition to reducing risk, lifecycle management services cut costs by ensuring timely purchase and banking of End-Of-Life (EOL) components and greatly reducing the logistical burden. Without these services, to avoid program disruption, the system designer must maintain ongoing visibility at a piece/part level with multiple OEMs. All too frequently, these OEMs have no process for providing proprietary data about component lifecycles. Today, leading COTS suppliers offer lifecycle management services that enable customers to access many of these services via a dedicated website that provides 24/7 delivery of potentially critical lifecycle management information such as product health reports and baseline configuration data packages. Even better, Web-based lifecycle management services enable system designers to easily and quickly access, approve, or reject engineering change proposals via the Internet.

There is no question, winning the fight against obsolescence, while obtaining the full performance and economic benefits of using COTS electronics, requires a comprehensive lifecycle management plan. A successful strategy should include an ongoing review of product configuration changes and component obsolescence, a quarterly BOM health check, and a longevity-of-repair plan. The result, especially if the lifecycle management plan is put in place at the beginning of the program development cycle, will be application stability and predictability that system integrators desire. Early adoption of these services is less costly in the long run: it enables costs to be more effectively amortized over the program's overall budget. Proactive lifecycle management, especially in today's budget environment, enables deployed systems to stay in service far longer than they could otherwise, an increasingly attractive option compared to the high cost of system redesign.

CURTISS WRIGHT Controls Defense Solutions

Trap #5: Evolving Requirements

For system integrators, the opportunity to optimize a new Sensor Processing design to meet the needs of one specific program is a temptation that is hard to resist. Unfortunately, that inclination is add odds with the cost conscious environment in which we work. Today, program managers should make every effort possible to think ahead and avoid getting locked into a single technical solution that is unable to evolve and adapt for other later uses. Furthermore, while Sensor Processing systems are typically deployed for many years, the threats for which they were originally deployed are unlikely to remain the same 10 or 20 years later.

Solution

The best approach, now and for the future, is to ensure that deployed systems are designed to support upgrades for new uses, eliminating the need for the program office to start an entire new system design from scratch. System integrators that use open architecture rugged COTS modules, such as those produced by Curtiss-Wright, are tapping into a rich and expansive ecosystem of processing elements, interconnects, I/O and peripherals that is continuously being upgraded from generation to generation. COTS boards that are designed for general purpose computing tend to be more flexible and more upgradable than boards designed to meet the limited and unique requirements of one single function or one single program. Even better, the open architecture approach followed by vendors like Curtiss-Wright enables system integrators to much more easily add features and expand their existing system (by adding additional processing modules, for example) as their program requirements change. And since the ancient days of the Greek philosopher, Heraclitus, the one thing that is certain to be constant is change.

Contact Information

To find your appropriate sales representative:

Website: www.cwcdefense.com/sales

Email: defensesales@curtisswright.com

Technical Support

For technical support:

Website: www.cwcdefense.com/support

Email: support@curtisswright.com

The information in this document is subject to change without notice and should not be construed as a commitment by Curtiss-Wright Controls Defense Solutions. While reasonable precautions have been taken, Curtiss-Wright assumes no responsibility for any errors that may appear in this document. All products shown or mentioned are trademarks or registered trademarks of their respective owners.

*Other names and brands may be claimed as the property of others.

CURTISS WRIGHT Controls
Defense Solutions

defensesales@curtisswright.com
cwcdefense.com