TECH BRIEFS

SAE INTERNATIONAL

SPECIAL REPORT:
RUGGED COMPUTING

JANUARY 2019

Sponsored by

Atrenne
A Celestica Company

kontron
S&T Group

GMS
COMPUTING ENGINES
GENERAL MICRO SYSTEMS, INC.
TRUSTED AND DEPLOYED SINCE 1979

LAB DEVELOPMENT to RUGGED DEPLOYMENT

- ARINC 404A
- MIL-STD-1275
- MIL-STD-167
- MIL-STD-461
- MIL-STD-704
- MIL-STD-810
- MIL-STD-91403
- RTCA/DO-160

LAB DEVELOPMENT CHASSIS

- 3U and 6U, Air and Conduction Cooled Modules
- Extreme cooling for high-powered modules
- VPX/OpenVPX, VME/VME64x, VXS, Hybrid, CompactPC
- Tower, Desktop, Open Frame, Rackmount, Horizontal Insertion and ATR
- 10 Gbaud Gen-3 signaling
- 25+ Standard VPX backplanes
- 3rd-Party Agnostic
- Optional 110/220 VAC, 28/48 VDC, ATX, 12v and 5V centric power supplies

RUGGED ENCLOSURE CAPABILITIES

- 1/4, 1/2, 3/4, 1, 1.5+ ATR enclosures
- 1U to 14U high Rackmount enclosures
- Small form factor enclosures
- VPX/OpenVPX, VXS, VME64x, VME, CompactPCI and other embedded form factors
- Baseplate, forced-air, heat-exchanger, and liquid-cooled chassis
- 3U/6U air and conduction-cooled modules.
- Backplanes with 10 Gbaud signaling in a range of topologies
- Lightening protection
- Red/black Tempest separation

CONTENTS

FEATURES

APPLICATION BRIEFS

ON THE COVER

A new embedded computing specification called CompactPCI Serial Space was ratified in 2017. Designed for use on defense and space-based projects such as satellites, it takes the proven CompactPCI Serial standard to the next level by addressing requirements like redundancy, radiation hardening, outgassing, and testing/screening. Read the article on page 10.
(*Composite illustration by Ayinde Frederick*)

Beyond VMEbus

A New Concept

Deployability and life-cycle management of large and disparate collections of weapons, combat and C4ISR computing systems is one of the most critical challenges facing the Department of Defense today. Numerous defense programs, each defining separate, and oftentimes unique system architectures, configurations, and compositions, result in a wide range of processing and control systems that create complexity and drive excessive lifecycle costs.

Although serious efforts have been made to standardize on a set of common hardware solutions such as NAVSEA Acoustic Rapid COTS Insertion (A-RCI), Common Processing (CPS), and Common Display Systems (CDS), that limited commonality has not extended across any one ship, let alone an entire fleet. This lack of a truly common modular infrastructure has led to:

- Difficult, asynchronous technology insertion cycles & delayed modernization due to expensive, time-consuming integration and shipboard industrial work;
- DMS/MS-driven logistics that lead to the purchase of inefficient, obsolete, end-of-life products;
- Non-uniform system administration & management creating unnecessary complexity;
- Failure to achieve broad-based economies of scale due to small quantity purchases and a variety of "spares."

These challenges not only make it difficult to upgrade to the latest technology but may also limit the government's ability to repurpose used — but still viable — equipment that may have residual value to other programs with less funding. If the DoD is to get onto, and stay on, the commercial technology cam, current efforts to standardize common hardware platforms must be accelerated. What is needed is a standardized, Modular Open Systems Approach (MOSA) compliant architecture that is at once scalable, extensible, serviceable and available. Once deployed, it should be populated with a relatively small set of mechanically robust, Commercial Off-The-Shelf (COTS), loosely-coupled, hardware modules compatible with mission critical land, sea, and airborne applications—truly COTS products with the end in mind.

The Impact of VMEbus

To understand the importance of standards-based technology in the DoD market, one of the best known examples is VMEbus. In the 1980's and 90's VMEbus overcame a slew of competitors to become arguably the leading embedded bus-board architecture for a wide range of commercial, industrial, military and aerospace applications. Over the years, it became the "party candle" of bus structures—whenever a competitor arose adopting some form of the Eurocard packaging, whether it be Multibus II, CompactPCI, Futurebus, ATCA, etc., the VME community would morph their product just enough to not only keep it alive but also maintain its dominance as the embedded computing architecture of choice. The compact (6U x 160mm deep) form factor, flexible packaging, pin-and-socket connector, and robust commercial/industrial construction enabled it to soldier on as an embedded computing platform for much longer than anyone could have originally predicted—often using the backplane as little more than a physical structure with power and ground distribution.

The success of VME in the DoD market was largely attributed to its:

- inherent ruggedness;
- industry standard architecture;
- composability;
- broad range of suppliers;
- adequate performance for most applications.

However, despite its achievements, VME's days were numbered due to the rise of fast pseudo-serial interconnects such as RapidIO, PCI Express, Ethernet, Infiniband and Myrinet. Additional drawback features also contributed to its decline, such as:

- **Backplane limitations:** signaling features were often underutilized, with the backplane itself being used for power.
- **Upgrade Complexity:** the VME standard allowed custom backplane pin assignments and custom backplanes; board models were rarely interchangeable without backplane or other design changes.
- **Inefficient Sparing:** Lack of fully compatible modules and lack of purchasing coordination between programs led to the need to stock a variety of spares.

Several generations of bus protocol enhancements and higher performance VMEbus interface devices helped hide the inherent liability of VME's relatively low-speed, shared parallel bus structure. In fact, much effort went towards the development of highly integrated VME Single Board Computers (SBCs) and auxiliary local IO bus expansion capabilities, all to remove data traffic from its most fundamental resource — the bus itself.

When blade systems began to take market share and Sun Microsystems introduced the 1U "pizza box," the same capabilities were now available in more cost-effective 1U servers — with commercial motherboard production volumes that were orders of magnitude greater than any VMEbus SBC. The highly integrated systems with packaged ATX-style motherboards, industry standard IO, and Ethernet interconnects were not only cost-effective, but easier to integrate and offered greater flexibility.

Possibly the most obvious example of a successful transition from SBCs to 1U rack-mount servers in a mission critical system is the NAVSEA, Acoustic Rapid COTS Insertion (ARCI) program. By transitioning to 1U rack servers, the ARCI IPT has been able to simultaneously reduce costs in hardware, the technology insertion process, and related shipboard industrial work, while speeding the deployment of higher-performance, contemporary technology unavailable on any suitable bus structure.

Looking across the universe of military embedded systems, one can see the proliferation of 1U "pizza boxes" across virtually every major weapons, combat, and C4ISR system, including:

- The U.S. Navy
- Aegis Weapons & Combat Systems
- AN/SQQ-89 Anti-Submarine Warfare System
- Submarine Warfare Federated Tactical System (SWFTS) & Tech Insertion Hardware (TIH)
- Acoustic — Rapid COTS Insertion (A-RCI)
- Consolidated Afloat and Network Enterprise Services (CANES)
- The U.S. Army
- Distributed Common Ground Station (Army-DCGS)
- Intelligence Fusion System (IFS)
- Warfighter Information Network — Tactical (WIN-T)
- Battle Command Common Services (BCCS)

Only 10″ deep, the Themis HDversa accommodates up to twelve special purpose modules that each share common electric, physical and environmental characteristics. HDversa is suitable for hosting existing bare metal applications while simultaneously providing a platform for incremental transitions to a fully hyper-converged architecture.

Beyond VMEbus

Yesterday's Wisdom is Today's Liability

Clearly, there has been an obvious "sea change" in the way many embedded systems have been architected over the past five to ten years. The only thing certain about computing systems is that change happens, and yesterday's conventional wisdom is today's liability.

Contemporary computing technologies, virtually across the board (Mil-Aero, Commercial Data Centers, Cloud Computing, etc.) are being driven by a seismic shift in how applications are being deployed—driving the way hardware is, or will be, deployed in the very near future. In summary, technology is increasingly leaning to:

1. **Everything-Over-Ethernet**

   Ethernet is ubiquitous and, in many instances, has become the common network, or fabric for virtually all IO in the data center.
2. **Virtualize Everything, Minimize Appliances**
   a. Given that racks and racks of servers were typically underutilized, or utilized in a random fashion, server virtualization has enabled more uniform loading and resiliency of the data center.
   b. Virtualizing storage has resulted in the deployment of fast and highly available NAS and SAN virtual appliances based on standard rack servers and local (direct) attached storage.
   c. Virtualizing the network allows for more efficient use of higher bandwidth "pipes" to the point that it's often less expensive to deploy fewer of the latest, highest performance network or fabric devices.
3. **Highly Composable Special Purpose Modules**
   a. In alignment with MOSA guidelines — modules that may be combined with other modules to provide new types of functionality.
   b. A set of modular building blocks representing a wide range of functional elements (processing, co-processing, storage, graphics, fabric extenders, etc.) that are disaggregated & highly modular with minimal dependencies (i.e., loosely-coupled).

In deference to disciplined standards and availability of broad market commodities, such as those found in web-scale data centers, Themis introduced the ResHD product line, which incorporates five simple, yet powerful concepts:

1. **Micro containers** — A small set of standardized, and composable compute, storage and IO modules. A library of common modules, designed to abstract fast-changing technologies from the platform infrastructure, that may be combined to resolve a much larger array of specific system functions.
2. **Platforms** — A platform is the immutable HW infrastructure that supports multiple generations of micro containers by functioning as a "hotel" structure with integrated interconnect fabric. A means to build N systems from M modules, where N >> M.
3. **Composability** — The ability to combine common, interdependent modules in diverse ways to achieve different functionalities.
4. **Convergence** — The aggregation of two or more functions on a single device by way of virtualization. For example, zero rack volume virtual storage arrays by converging storage and virtual machines on the same servers.
5. **Unified Hyperscale Fabric** — Leaf & Spine interconnect fabric and Equal Cost Multi-Path (ECMP) or Open Shortest Path First (OSPF) protocols to provide predictable performance in terms of bandwidth, blocking ratios, latencies, jitter, and resiliencies.

Currently, the 20″ deep ResHD offers six different types of modules, some of which have integrated key partner technologies: processor, storage, IO, management, two types of switches (one integrates a 100 Gigabit/sec Mellanox switch), and a fabric extender (with Bay Microsystems FX50).

Bringing Datacenter Strategies to Mission Critical Applications

Given that it may be a while before DoD weapons and combat systems embrace virtualization, the seeds can be planted today for near certain future occurrences.

The U.S. Navy, for example, can begin to replace legacy VMEbus systems, while simultaneously deploying an infrastructure able of accommodating future virtualized systems in incremental and manageable steps.

To this end, small form factor platforms like the Themis HDversa (6U x 10″ (254mm) deep x 19″ rack-mount), roughly the volume of a 6U VMEbus card cage, are designed to accommodate small form factor modules. As part of the Themis ResHD family, HDversa provides the ability to configure a number of small, lightweight, special purpose micro containers — each sharing common electrical, physical & environmental characteristics, including:

- SWaP-C;
- Input power and cooling methodology;
- General purpose IO integration;
- Environmental resilience, e.g., kinetics, temperature, humidity;
- Management;
- Life-cycle support, and Technology Insertion (TI).

The modular HDversa embraces Ethernet as the fabric of choice; providing an easy and proven path to virtualization. The standardized compute, storage and IO modules have a firmly defined backplane interconnect, unlike VME's user-customizable IO. Upon release, it will provide twelve module bays that accommodate a range of configurable modules, including:

- Single-wide, compute modules featuring a single Intel Xeon D — 1500 embedded series processor, (4) DIMM sockets, (2) 10GbE, (2) 1GbE, (1) IPMI2;
- Double-wide, storage server modules featuring all of the compute module functions with added (4) removable 2.5" SAS or SATA storage devices;

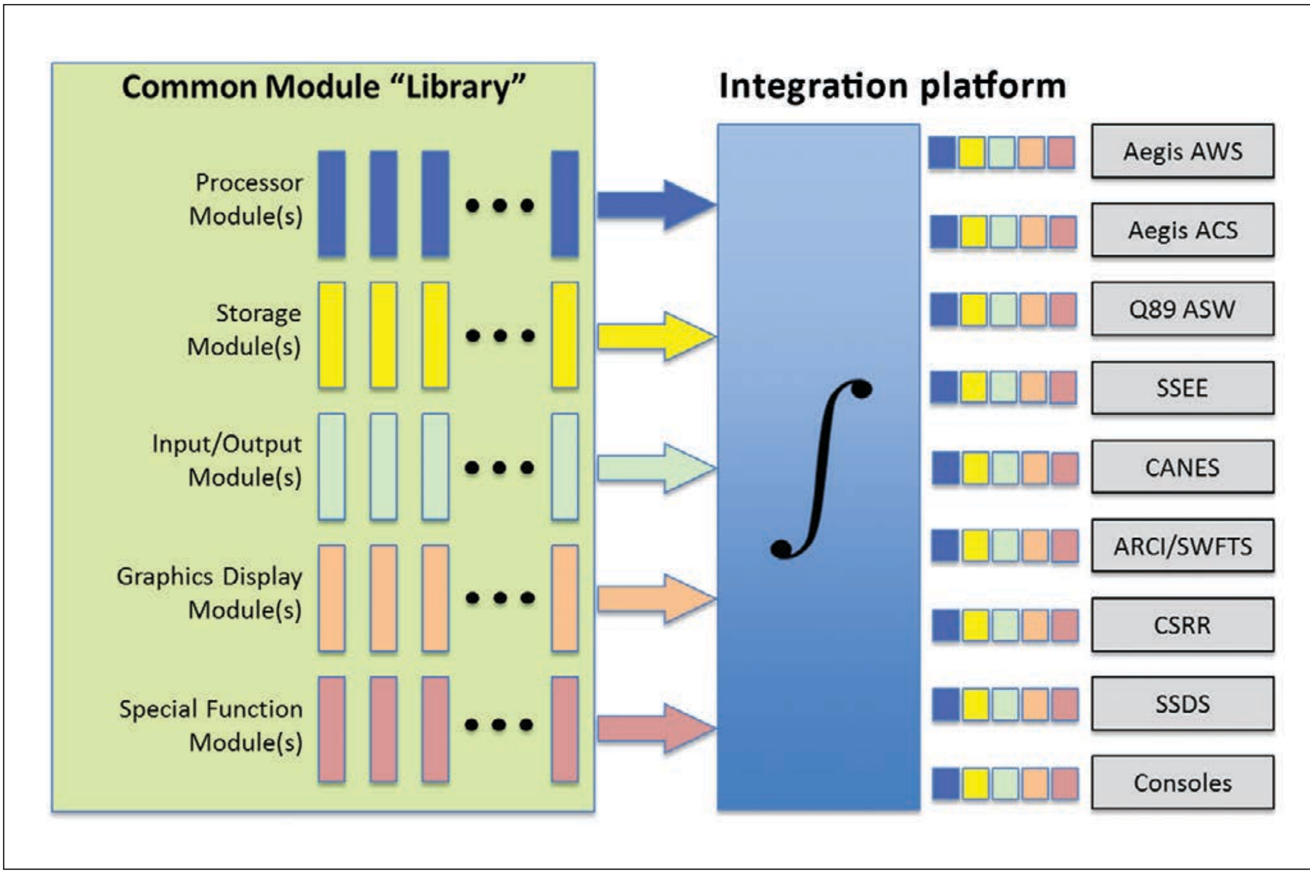

ResHD modules are "micro containers" that package different technology components in separate containers. These disaggregated pieces can then be configured in different ways to deliver a multitude of program specific functionalities. This not only reduces costs associated with keeping spares, but also utilizes less space and opens the door for logistical sharing. Themis has delivered the initial variant of ResHD to a mission-critical Navy weapons system. The figure above illustrates how further utilization of this technology could be employed by the U.S. Navy.

- Double-wide, NVMe storage server modules, identical to (2.) while replacing the SAS or SATA devices with U.2 NVMe storage devices;
- Double-wide, PCIe expansions also featuring all compute module functions, but adding PCIe expansion with either (2) PCIe x8/Gen3 or (1) PCIe x16/Gen3 slot;
- Single-wide, 16-port 10GbE managed switches;
- Single-wide, 16-port 1GbE managed switches and/or resource manager modules.

As with the full-length (20" depth) ResHD, the shallow-depth (10" depth) modules are optimized for all of the same common characteristics, but feature less SWaP and a single processor (8/16c Xeon D-15xx series), better aligning it with target applications best served by this level of performance.

Themis has positioned ResHD and HDversa to bring the latest thinking in datacenter and cloud scalability, extensibility, availability and serviceability within a SWaP–C (cost & cooling) optimized platform to Mil & Aerospace mission critical systems. These systems support existing bare metal applications while simultaneously providing a platform for taking intermediate steps toward server, storage and network virtualization strategies. More importantly, they are designed to ameliorate the previously identified problems associated with a plethora of unconstrained, and ad-hoc system configurations; deployed on a single platform, without regard for any adjacent systems. Whereas the current and future versions of RESHD address the high end, HDversa is targeted for applications where less Space, Weight, and Power (SWaP) are of greater utility than maximizing the power and density of cores, storage or network bandwidth. In effect, it's just good enough…

Voltair once said, "perfect is the enemy of the good." What he meant was, the good should never be held hostage to the pursuit of perfection.

This article was written by Rick Studley, VP Navy Business Development & CTO, Themis Computer (Fairfax, VA). ***For more information, visit http://info.hotims.com/69507-500***. ■

The Bus TOO TOUGH to Die

The venerable MIL-STD-1553B bus has survived remarkably well even as other more advanced solutions gained wide acceptance in the last few years. However, the fact remains that its maximum data rate of 1 Mb/s is orders of magnitude too slow for today's data-intensive systems, so logic dictates that it will soon fade away. That may be a logical assumption, but it's likely to prove wrong, for several reasons.

The most obvious is that MIL-STD-1553B continues to fly on at least 30,000 aircraft, as well as on commercial and military ships, and is widely used in industrial and other applications. The situation is analogous to the International Space Station (where it's also present). The ISS was expected to "last" until 2015 but when it arrived, ISS lifetime projections were extended to 2020 and then to 2024, and the latest consensus is 2028.

Basically, until its weaknesses blatantly outweigh its strengths, the ISS will still be up there, along with MIL-STD-1553B. Down here on Earth, it will take decades before all the platforms MIL-STD-1553B controls are either obsoleted or worn out, and combined with the slow process of defense technology insertion and the high cost of retrofits, MIL-STD-1553B will be here longer than many of the readers of this article.

The standard is so valuable that there are still many sources of every type of component used by MIL-STD-1553B and many IC suppliers have been supporting it for decades. There are even bridges between MIL-STD-1553B and Gigabit Ethernet that allow the existing standard to transfer data to the world's most widely used networking standard.

Many of its key and often unique benefits can be found in its architecture, which makes MIL-STD-1553B reliable and fault-tolerant for connecting processors with real-time sensors and controllers. It's arguable that the most important reason MIL-STD-1553B still retains its stature for mission-critical systems is its command/response protocol that ensures real-time determinism (Figure 1).

Avionics and other systems that operate in real time require determinism to ensure predictable behavior, every time, without fail. That is, a real-time system must behave in a way that can be mathematically predicted, executing functions with no concern that they will be degraded in an unexpected way. For real-time systems in which surprises are intolerable, MIL-STD-1553B is nearly perfect in its ability to predictably perform functions in real-time with microsecond accuracy and very low jitter.

The standard was created to operate in hostile environments that include lightning, wide temperature ranges, high levels of vibration, and the potential for significant interference. The latter is the result of galvanic isolation that its transformers provide to fend off lightning, which has become even more important as many newer aircraft are made from composite materials, reducing and sometimes eliminating the benefit of having a Faraday shield inherently created by an aluminum skin. Finally, MIL-STD-1553 has refined its criteria for validation testing over the years and has not encountered issues with interoperability, even though massive numbers of designers have implemented it in diverse systems.

Beyond these points there is the fact that MIL-STD-1553B, or its protocol, is

CHRIS PARYPA PHOTOGRAPHY/SHUTTERSTOCK.COM used in a variety of other standards, and it's a long list (Table 1). The world of communications bus standards — and MIL-STD-1553B nomenclature — is deep, wide, and often obscure, with countries and defense agencies within them tweaking the bus and renaming it. For example, the upper-layer protocol of MIL-STD-1553B is also used in FC-AE-1553 and High-Speed 1760.

FC-AE-1553 uses the MIL-STD-1553B command and response protocol and supports all its core elements including command and status, sub addresses, mode codes, transfers between remote terminal, error checking, and broadcast. As a result, it allows the reuse of MIL-STD-1553 and MIL-STD-1760 commands and legacy software. In addition, FC-AE-1553 includes extensions and optimizations supporting RDMA to provide direct memory access of remote systems over Fibre Channel.

MIL-STD-1760 is typically used for interfacing weapon stores to an aircraft's control systems, but an enhanced version called High-Speed 1760 (SAE standard AS6653) has a high-speed interface based on Fibre Channel that can deliver data rates up to 1 Gb/s over two 75-ohm coaxial cables. The Fibre Channel upper layer protocols are based on FC-AE-1553, MIL-STD-1553B for command and control messaging, and FC-AV for transferring images, video, and audio files.

The final reason for the standard's longevity it that a lot of time and money has been invested in making MIL-STD-1553B viable in the future. In fact, variants of the standards today are actually delivering data rates of 100 Mb/s — 100 times that of MIL-STD-1553B — and have demonstrated their ability to reach 200 Mb/s.

So, why haven't these variants transformed the standard into something like MIL-STD-1553C? Well, they have, but in a much more limited fashion than might be expected. To better understand this, it helps to trace the long, winding path that this standard has traveled in the last 15 years or so.

Toward a Better Bus

In the 2000s, the Air Force recognized that, as it would cost (at that time) more than $1 million per aircraft to replace MIL-STD-1553, the logical step would be to enhance MIL-STD-1553 to increase data rates, hopefully to 200 Mb/s or even higher while allowing simultaneous transfer on existing MIL-STD-1553B cable. To this end, Data Device Corp. (DDC) and Edgewater Computer Systems expended considerable effort to develop versions of MIL-STD-1553B that would allow it to remain viable. Both were successful, achieving excellent results without significantly modifying the standard's fundamentals. The first of DDC's efforts resulted in what the company called Turbo 1553 that increased the data rate of the bus to 5 Mb/s on standard MIL-STD-1553 terminals over 430 feet with 10 stub connections to three remotes.

The second, called "High Performance 1553" or "Hyper 1553" uses frequency-division multiplexing and other techniques to allow higher speed data to be simultaneously carried along with standard 1-Mb/s data on the same cable. It was envisioned to implement a multi-drop bus and eliminate the need for active hubs or switches. In Hyper 1553, the 1-Mb/s signals are limited to lower frequencies while the higher speed signals occupy higher ones, similar to DSL. DDC determined that enough bandwidth was available to allow the parallel signals to be reliably transferred at higher speeds, depending on the length of the bus and number of stubs.

DDC demonstrated HyPer 1553 in a 2-hr. flight on an Air Force F15-E1 Strike Eagle fighter in 2005, where it was used to transfer imagery between a computer in the forward avionics bay and a bomb mounted on a pylon. The data was transferred without errors at 40 Mb/s over existing cabling along with 1 Mb/s traffic.

Meanwhile, Edgewater was producing similar results with the major benefit of being under contract to DoD to develop Extended 1553 (E1553). Edgewater, along with researchers from the Air Force and Navy, worked on the project for several years, and the technology was flight tested in an Air Force F-16 and Navy F/A-18. The results were very promising, but as the program didn't have the visibility and priority of others, the Air Force cancelled it. It did this even though the goal of a 200 Mb/s data rate was achieved, again without the need for huge changes, while also simultaneously transferring standard data at 1 Mb/s. The company believed it had the potential to reach 500 Mb/s.

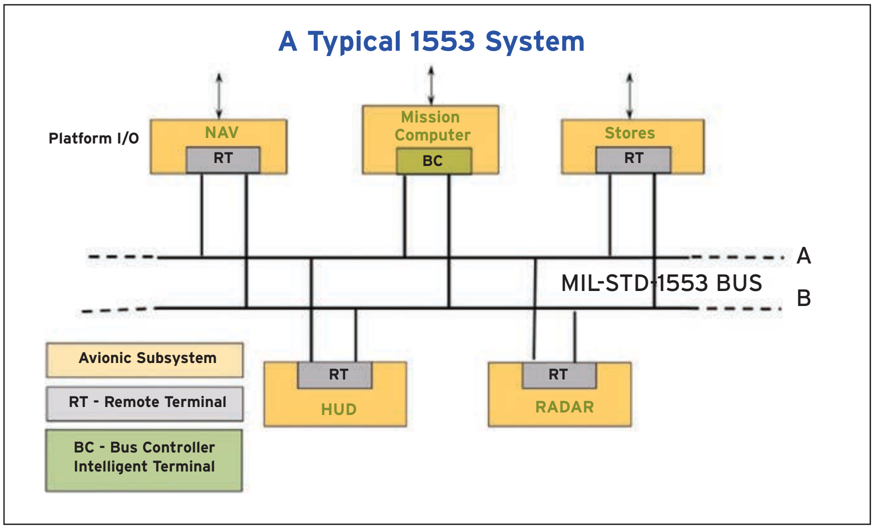

Figure 1. A typical MIL-STD-1553B system including remote terminals and bus controllers serving various portions of an aircraft.

The Bus TOO TOUGH to Die

However, all of this work had not gone unnoticed, and the Assistant Secretary of Defense for Research and Engineering and a consortium from Canada, the U.K., Germany, and others successfully petitioned to complete E1553 development work. It was ultimately tested on fixed-and rotary-wing aircraft, which led the NATO Avionics System Panel (AVSP), chaired by the U.S. Navy, to sponsor standardization efforts within NATO. The ratification process began in 2010.

The result was a NATO Standardization Agreement (STANAG 7221) and in 2015 the "Broadband Real-Time Data Bus Standard" was unanimously ratified. STANAG defines everything from processes and procedures to terms and conditions for design and manufacture of military equipment among NATO member countries. Its goal was to create a NATO-specific set of standards so that all members could work off the same page, so to speak. It was a logical and, perhaps, essential approach considering the alternative.

3DSCULPTOR/SHUTTERSTOCK

Figure 2. MIL-STD-1553B still continues to function high above the Earth aboard the International Space Station.

ARINC 629	Derivation of MIL-1553 used by commercial aircraft.
Miniature Mission Stores Interface (MIMSI)	10 Mb/s MIL-STD 1553 protocol using RS-485 transceivers. Same as EBR-1553
SAE AS5652	10 Mb/s MIL-STD-1553B variant
Enhanced Bit Rate-1553 (EBR-1553)	enhanced version of MIL-STD-1553, which enables 1553 data rates of 10 Mb/s
MIL-STD-1773	Uses optical fiber instead of copper
MIL-STD-1760	Enhanced, 10-Mb/s version of MIL-STD-1553 optimized for data transfer between aircraft and stores
HyPer 1553	100 Mb/s coexisting with legacy MIL-STD-1553
E1553 (Extended 1553)	200-Mb/s version of MIL-STD-1553, developed and trademarked by Edgewater Computer Systems, in NATO called NATO-STANAG 7221
NATO-STANAG 3838 AVS	Originally the same as MIL-1553B; modified later
NATO-STANAG 3910	Derivation of MIL-1553 used in NATO and the UK operating at 20 Mb/s over fiber
NATO-STANAG 7155	Designed to improve maintenance and reduce life-cycle costs of avionics systems using SAE AS15531
NATO-STANAG 7221	NATO standard for E1553
DEF STAN 00-18 Part 2	Modified UK defense version of MIL-STD-1553B. Has six parts, covering different variants

So, E1553 is still alive and well in its standard form as well as delivering 100 Mb/s performance using its infrastructure in avionics and other defense applications it could previously not serve. However, E1553 (i.e., STANAG 7221) has an additional result: limiting the use of a new and improved version of MIL-STD-1553B only to those with permission: NATO countries and their representatives. So, while MIL-STD-1553B has advanced over the years, the odd result is that the hundreds of other applications in which MIL-STD-1553B is used can't benefit from the one delivering the highest performance.

Summary

MIL-STD-1553B has nine lives, and about eight have been used up. However, the sheer breadth of the applications in which this standard is deployed makes it virtually certain that it will continue to be around for many years. Nevertheless, new systems all use something more modern typically based on Ethernet or one of its variants such as Time-Triggered Ethernet (SAE AS6802) or AFDX, as well as Fibre Channel, or IEEE 1394 (FireWire).

Looking back, it's unfortunate that the evolution of MIL-STD-1553B ultimately wound up as a rather closeted standard, highly unlikely to find its way into commercial markets, even though it delivers data rates high enough to serve many applications today and tomorrow.

This article was written by Mark Hearn, Product Manager, MilesTek Corporation (Lewisville, TX). ***For more information, visit http://info.hotims.com/69510-501***. ■

CompactPCI Serial Space

A New Embedded System Specification Takes on Extreme Environments

There are a few open specifications in the embedded industry that have been used in the extreme environments of space over the years. These include VME, CompactPCI, OpenVPX, and MicroTCA. But, prime contractors such as Airbus and others desired a high performance and versatile architecture that was relatively simple and cost-effective. The CompactPCI Serial Space specification was ratified in 2017 to address these requirements, resulting in a compelling architecture for defense and space-based projects.

A Brief History

VITA and PICMG have developed excellent specifications to address the needs of extreme environments. But, CompactPCI and VME are bus-based standards that do not provide the performance and inherent reliability of switched fabrics. OpenVPX and MicroTCA are both excellent, high bandwidth, and low SWaP (Size, Weight Power) standards. But, CompactPCI Serial — the basis for CompactPCI Serial Space (also referred as cPCI Serial Space) — is less complex and has a low-cost design approach. CompactPCI Serial, which was ratified in 2015, is similar to legacy CompactPCI in that it leverages the 3U and 6U Eurocard form factor with 160mm deep cards. But it uses new, rugged high speed Airmax VS connectors. The power is based on 12V (with an option of 5V standby), providing simplicity and cost-effectiveness. Having one main power rail is much simpler to specify the right power supplies and can significantly reduce the costs. The CompactPCI Serial power rails can be either a single plane, or every board could be supplied and controlled individually. Utilizing PCIe (current solutions are up to Gen3) or Gigabit Ethernet (current solutions are up to 10GbE with 40GbE-capable implementations), the aggregate data rates parallel other high-performance standards.

It is a significant benefit to potential users that CompactPCI Serial already has several design provisions for rugged environments. The specification started out in railway applications, having to survive shock vibration, moisture & dust ingress, etc. The transition to CompactPCI Serial Space did not require signification adjustments. See Figure 1 for an example of a conduction-cooled CompactPCI Serial switch card that meets -40° to +85° temperature environments as well as MIL levels of shock/vibration.

Ready for Space

There were important implementations for CompactPCI Serial to meet space-level ruggedization.

Figure 1. CompactPCI Serial has long been ruggedized for extreme environments. This SBC (courtesy of MEN Micro) shows an example of a 3U version in a conduction-cooled clamshell.

These include:

1) Redundancy: A key element for space applications is redundancy. Downtime is extremely costly and hardware maintenance is impossible. Having a Dual Star architecture for both the data traffic and Spacewire was important. Spacewire is a rugged variant of Firewire for space applications. It can be used to connect multiple enclosures in the system, providing a serial point-to-point connection.

2) Control & Monitoring: This is a critical requirement for satellites, positioning, and other space systems. So, each slot of CompactPCI Serial Space allows the use of several dedicated monitoring and control signals.

3) RAD-Hard & Rugged: Naturally, the specification required the rugged implementation of CompactPCI Serial and defines the levels for meeting compliance. Radiation is a problem for all standard CMOS devices, so ensuring the system is RAD-hard is paramount. Significant testing was required.

4) Outgassing: Outgassing is a potential issue for plastic components as they can condense onto optical elements, thermal radiators, or solar cells and obscure them. NASA has a list of low-outgassing materials for use in spacecraft. The issue requires careful selection and testing.

5) Testing & Screening: As mentioned above, testing and screening of components was critical in the specification development process. The PICMG members working on the specification needed to be certain that it could meet shock, vibration, radiation, outgassing, extreme temperatures of -40°C to +85°C, dust and moisture ingress, EMI, etc.

The cPCI Serial Space specification defines a utility connector, which can be controlled and configured via an open management bus. It takes over the hot-plug functionality from CompactPCI Serial. Of course, hot-swap functionality is not required in a satellite in space. But it can actually be an attractive feature on ground-based systems and for test/simulation systems. System monitoring/management is supported via the CAN (I2C bus is also optional for less critical applications).

Applications

With wide use in railway and other rugged applications, ruggedized CompactPCI Serial is common. The wedge lock clamshells for conduction-cooling provide further stability, and the modules meet MIL specs such as MIL-STD 810 and 901D for shock and vibration levels. CompactPCI Serial has already been used in defense applications, including a fully DO-168G qualified rugged natural convection ARINC600 ATR for an airborne network server.

With cPCI Serial Space's additional controlling and monitoring capabilities, there is an even stronger solution for failure detection, isolation, and recovery (FDIR). Another application for CompactPCI Serial in defense applications is a conduction-cooled ATR with a 3U backplane and a supplementary internal fan for additional cooling for a data recorder in a military vehicle. The architecture was also recently chosen on a major naval program as an upgrade from a previous

CompactPCI Serial Space

Figure 2. There are a variety of advanced processors being developed in CompactPCI Serial/cPCI Serial Space, including this SPARC-based Leon4 quad core board (right) and P4080 board (left). SPARC/Leon is a preferred platform computer used by the European Space Agency.

CompactPCI design. There are certainly benefits from leveraging CompactPCI boards for the newer serial solutions.

The types of Mil/Aero applications for CompactPCI Serial and cPCI Serial Space include (but are not limited to):

- Simulators (for satellite environment, military systems)
- Test equipment (for satellites, military systems)
- Satellite platform control (altitude and orbit control)
- Satellite payload/instrument control
- Mass memory (satellite, military data recorders)
- Networking/communications (satellite & ground station, military)
- RADAR & SONAR systems
- Electronic warfare & signal intelligence systems
- C4ISR & situational awareness systems

The space industry faces some great challenges as there is a clear trend for mega constellations, such as OneWeb (900 satellites providing Internet service worldwide). Industrialization has reached the spacecraft manufacturers and therefore, the opportunity is huge for future programs. As these highly rugged systems are developed, the ecosystem will undoubtedly continue to expand. So, Mil/Aero engineers will have a wealth of high-performance cPCI Serial modules to choose from in an architecture that is comparatively low cost and easy-to-use. Figure 2 shows cPCI Serial Space CPU modules (without clamshells). On the right is a Leon4 quad-core processor and on the left is a NXP P4080 QorIQ chipset version that provides the processing power for on-board payload data processing. Both boards have undergone the radiation, thermal, shock/vibration, EMC, and functional tests for space applications.

What is Next?

Most Military embedded computing applications require a proven, reliable, scalable architecture with low SWaP, high-performance options, a simple and easy design that is ruggedizable. With the wealth of ruggedized CompactPCI Serial products, there appears to be increasing interest for the architecture in Mil/Aero projects. For space requirements, additional testing is required for radiation, outgassing, shock/vibration, etc. The cPCI Serial Space architecture has performed these tests, providing a space hardened implementation. There are a significant number of processors, switches, I/O cards, carriers, storage, and graphics modules in the CompactPCI Serial architecture. As more FPGA and digitizer cards are created over time, the industry may see an even larger push in defense applications.

This article was written by Justin Moll, Vice President of Marketing PICMG (Wakefield, MA) and Hans Juergen Herpel, Expert on Advanced Avionics Software, Airbus Defence and Space GmbH (Toulouse, France). ***For more information, visit http://info.hotims.com/69503-502.*** ■

RADIATION TOLERANT "SMART BACKPLANES" FOR SPACECRAFT AVIONICS

In recent years there has been a trend towards the wider use of COTS (Commercial Off The Shelf) equipment in space missions. This trend has been mainly driven by the restrictions in R&D budgets and a growing demand for shorter design cycles. Funding Agencies are encouraging designers of spacecraft systems to identify and overcome the obstacles that previously prevented the use of COTS products for space missions.

When it comes to space vehicle engineering, the tolerance of onboard electronics to radiation effects can be one of the most challenging aspects of the system design. The risk of failure for avionics equipment on-board spacecraft due to radiation exposure is determined by the vehicle's orbit trajectory and flight duration, during which the vehicle is exposed to trapped radiation as well as solar and cosmic radiation sources.

The overall impact at equipment level is determined by a complex interaction of shielding, circuit design, device technology and particle energy spectra and the tolerance to radiation effects is one of the key criteria for selecting the onboard equipment and subsystems. A novel approach to radiation mitigation at the board level, as opposed to component level, enables the use of existing high-performance COTS electronics systems in a space radiation environment, thereby lowering the cost involved in the design, certification, manufacture and deployment of such a system. This new approach, called "Radiation Tolerant Space COTS," is based on a new type of backplane called the "Smart Backplane." The Smart Backplane enables integrators to use proven, affordable COTS modules in a space radiation environment. In addition to faster time to deployment, the numerous benefits of this game-changing design strategy include significantly lowered costs for deployment, design, certification, and manufacture of space avionics.

Space Radiation

Radiation events can result in temporary or permanent disturbances to the function of a device, a phenomenon known as a Single Event Effects (SEE). Over the years, as the density of ICs has increased, the size of elementary semiconductor structures has shrunk to the level where a spurious current spike produced by a single radiation particle can result in SEEs capable of disrupting the circuit's operation. Two types of SEE are most relevant to protecting spacecraft avionics:

- **Single Event Upset (SEU)** - occurs when a radiation-induced current causes a memory structure to change its state. This results in a temporary error in device output or its operation and is commonly referred to as "soft error." In the case of an SEU, the device is not damaged and will function properly in the future, but the data processed by the device can be corrupted.
- **Single Event Latch-Up (SEL)** - occurs when a radiation-induced current activates a parasitic structure (e.g. transistor), which forms an undesired low-impedance path in the semiconductor structure. It disrupts proper functioning of the circuit, and if not corrected, can possibly even lead to its destruction due to overcurrent. The circuit typically remains latched up until it is powered off and afterwards it may continue to function properly.

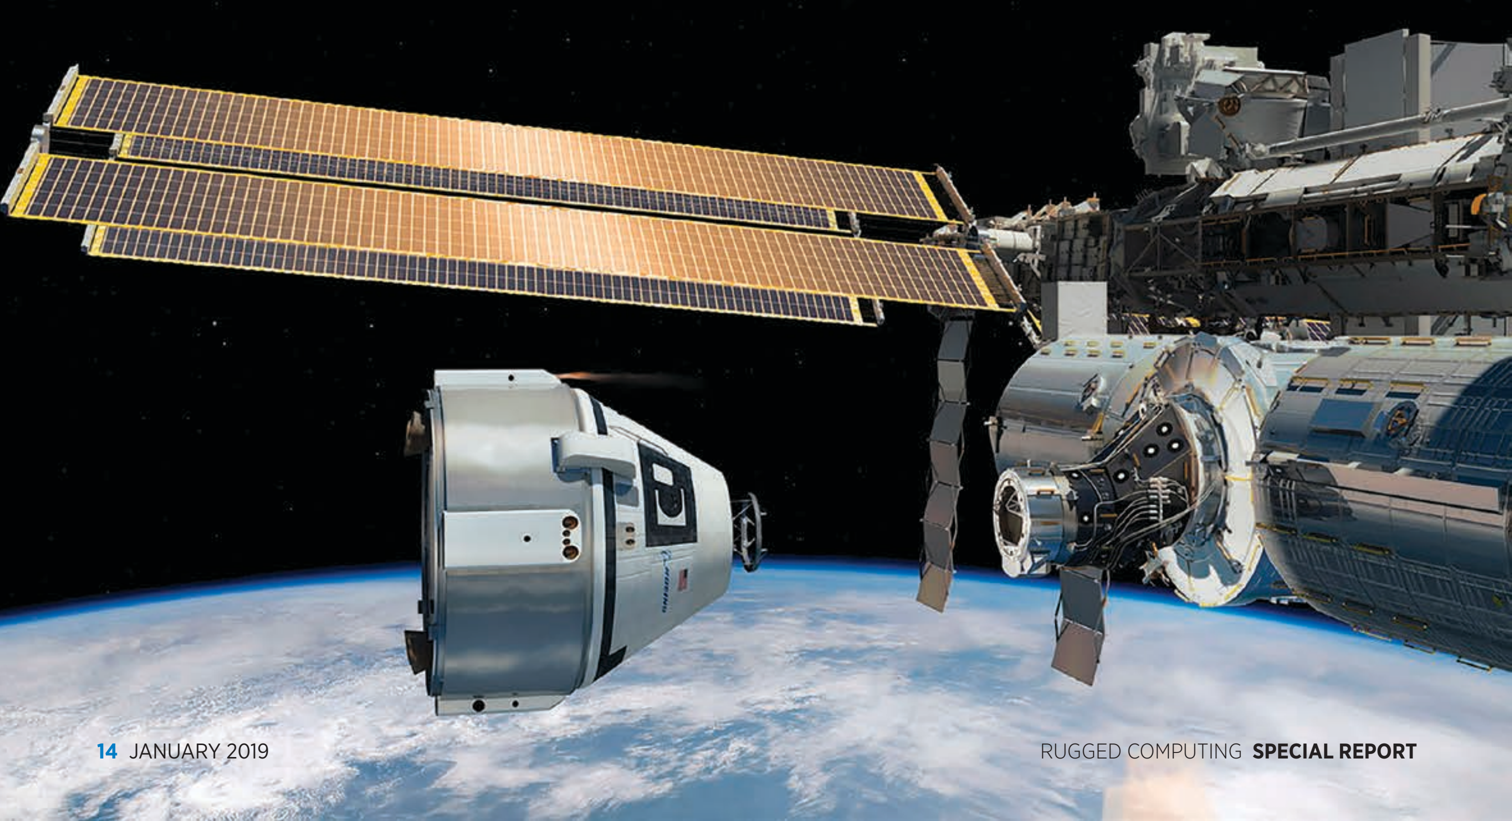

The Curtiss-Wright radiation tolerant Smart Backplane data acquisition unit in flight configuration.

The Space COTS Approach

While the lowest-risk method for preventing radiation damage is the use of hardened components throughout the system design, this approach is very costly and time-consuming. The Radiation Tolerant Space COTS approach provides an alternative method for mitigating against SEUs and SELs that ensures the reliability and mission assurance requirements of the system while respecting the program's budget and schedule requirement constraints.

In order to mitigate against SEUs and SELs, Curtiss-Wright has designed the 'Smart Backplane' chassis, a rugged 12-user slot chassis for data acquisition in a radiation-intensive environment. Based on the Acra KAM-500, a modular rugged data acquisition unit (DAU) and recording system designed for use in flight test market, the Smart Backplane enables the use of COTS data acquisition plug-in modules while at the same time preventing against the harmful effects of ionizing radiation. The system's unique design enables the use of standard plug-in COTS modules in a space environment. This eliminates the need for modules to be built using radiation-hardened components. This minimizes the cost of the overall system and enables the space system designer to leverage over 100 plug-in modules already designed for data acquisition in aircraft flight testing.

The DAU operates as a collection of synchronized FPGA based state machines. Once per acquisition cycle, the RAM is refreshed. Therefore, any SEUs that occur in RAM are overwritten within one acquisition cycle time. An acquisition cycle time can be anywhere from 100 microseconds up to 2 seconds in length.

The plug-in COTS modules used in the system are manufactured with commercial components but are protected by the radiation hardened Smart Backplane. The Smart Backplane can detect an SEL event on one of the modules in real-time and correct for that event before any damage can be done, thereby ensuring normal data acquisition is resumed without component damage and with minimal data loss. The system then recovers from the SEE, and normal operation of the entire data handling subsystem is uncompromised. The breakthrough concept behind the system design is that instead of seeking to prevent radiation events, the Smart Backplane instead quickly detects and then corrects for these events in real-time when they happen, with no damage to the equipment and most importantly, it ensures that the mission assurance, safety and reliability requirements are met.

The radiation tolerant Smart Backplane KAM-500 data acquisition system has already been selected for deployment in mission critical spacecraft avionics systems for manned and un-manned re-entry vehicles as well as launcher upper stages and is being considered for the instrumentation system on future planetary re-entry vehicles and as the basis for low cost COTS-based small satellite avionics systems.

In addition to supporting space applications aboard the International Space Station and the European Space Agency's IXV (Intermediate Experimental Vehicle) launch vehicle, Smart Backplane was selected by The Boeing Company to supply rugged data handling avionics for use in the Crew Space Transportation (CST)-100 spacecraft. The Boeing CST-100 spacecraft will provide transportation for up to seven passengers or a mix of crew and cargo to low-Earth orbit destinations such as the International Space Station (ISS) and the Bigelow planned station.

More recently, the Smart Backplane has been selected by Rocket Lab for use on the Electron launch vehicle. The DAU will acquire data from various analog and digital sensors onboard the Electron, Rocket Lab's dedicated vehicle for launching small satellites and other payloads to Low Earth Orbit. The two-stage Electron vehicle is a dedicated launch service for small satellites to Low Earth Orbit. Its innovative design uses advanced carbon composites for a strong and lightweight flight structure. The Electron is the first oxygen/hydrocarbon engine to use 3D printing for all primary components. Rocket Lab recently successfully launched the first test flight of Electron from its orbital launch site on New Zealand's Mahia Peninsula.

As space vehicle system designers look to lower the cost of space flight through the use of COTS electronics equipment, those solutions must be provably able to survive in demanding space radiation environments. Fully radiation hardened designs may yield the most protection, but they are also the most expensive option. The Smart Backplane chassis allows the use of high-performance COTS data acquisition user-modules in radiation-intensive space applications, lowering the cost of such a system by up to 75% while still meeting with the mission reliability requirements and minimizing the loss of telemetry data due to radiation events to less than 2% for a typical low Earth orbit (LEO) application.

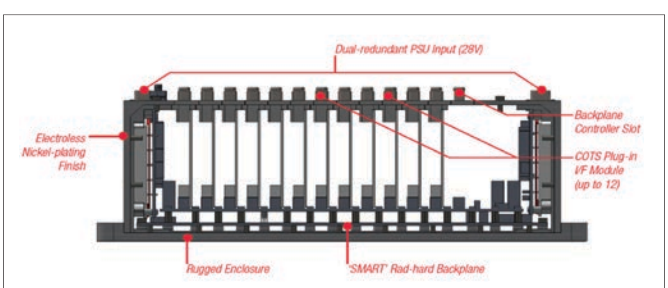

Cross-section showing Curtiss-Wright COTS modules in the radiation tolerant Smart Backplane.

This article was written by Daniel Gleeson, Space Business Development Manager, Curtiss-Wright (Dublin, Ireland). ***For more information, visit http://info.hotims.com/65855-500.*** ■

Bringing RF into the Embedded World: *It's Time*

Embedded systems have been almost entirely digital throughout their long history, while RF and microwave technologies were separate subsystems with no effective interface between the two. For many reasons, this "RF/digital divide" should finally be connected.

Mercury Systems, which makes embedded systems (i.e., board-level digital and RF subsystems), proposed that manufacturers of embedded and microwave subassemblies participate in its initiative called OpenRFM™. The company's goal is to make it possible to integrate RF and microwave technology into current "digital-only" embedded form factors for the first time. If you're not in the embedded systems business, you're probably wondering, "Why is this just happening now?" If you're into the designing of RF and microwave subsystems, you might ask, "Why bother?"

The answer to the first question is that digital and microwave designers have always operated in separate domains and ignored each other. The answer to the second is that integration of these two disparate technologies simply makes sense regardless of what "camp" you reside in, and for the following reasons:

- It would allow these two technologies (digital and microwave) to be integrated in a standard form factor followed by the embedded systems industry.
- It would be a major step toward realizing two major goals of the U.S. Department of Defense (DoD): to dramatically increase function integration, and to allow radar, electronic warfare, and other systems to be used in multiple platforms without major redesign.
- All embedded systems manufacturers and prime contractors could follow a single design roadmap while retaining the flexibility to differentiate their products from others by using their own OpenRFM-compatible products.
- Systems could be constructed that are smaller, lighter, consume less power, are less expensive, and shorten the time to market for both embedded systems and end products.

Figure 1. This is an OpenRFM module riding atop an OpenVPX board. Traditionally, the functions performed in the OpenRFM module would have been implemented in a larger separate housing.

Although Mercury targeted OpenRFM to sectors of the defense industry in which it participates, there is no reason why it could not be adopted in other embedded markets such as rugged industrial, scientific, and medical systems, and telecommunications, for example. As some of these applications make use of wireless communications in some form, OpenRFM should be useful for manufacturers of commercial systems that incorporate it.

To affect this integration on a broad scale, Mercury will need to have it appended to the current family of standards used in defense embedded systems — OpenVPX, whose champion is VITA (formerly the VMEbus International Trade Association). There is at least a reasonably good chance this will occur as it fits neatly into the association's roadmap, which has already expanded to include software-defined radio and space qualification.

OpenRFM is a modular, open architecture built around OpenVPX that combines hardware, firmware, and software, and allows high channel density, advanced interconnect technology, and employs a "building block" approach. The company has already produced OpenRFM products in 3U and 6U OpenVPX form factors, as shown in Figures 1 and 2. The real estate available for OpenRFM in the OpenVPX environment is sufficient to allow a broad array of integrated microwave subsystems (today referred to as Integrated Microwave Assemblies or IMAs) to be accommodated.

One of the most common is the downconverter, which effectively forms the front end of every system in which a signal must be captured over the air and lowered in frequency to one that can be handled by an analog-to-digital converter.

The downconverter's opposite is the upconverter, which does just the reverse: taking a signal at a lower frequency and increasing it to one required by a specific application. Both of these IMAs are required elements of every type of military and commercial system, from satellite communications terminals to microwave point-to-point links, and almost every type of communications system.

These are complex products that include many active and passive microwave components, and range from relatively small to quite large, depending on their requirements. As these subsystems typically don't reside on a board-level product but rather are separate functional blocks, OpenRFM has the potential to make the whole system smaller. Other typical IMAs include complete digital receivers that, when combined with traditional signal processing, general-purpose processing, signal distribution, and other functional blocks, form a larger subsystem.

Figure 2. The 6U VME carrier (bottom) and three OpenRFM modules (top): a quad downconverter (left), wideband tuner (center), and direct-digital synthesizer (right).

Why Is This Just Happening Now?

As we are all painfully aware, the defense industry has been driven primarily by technology rather than cost. This is because new radar, electronic warfare, and other systems required by the DoD are extremely complex and push the envelope of achievable performance. It's also because the Pentagon has to combat both current and perceived next-generation capabilities of its adversaries and because these systems remain in the field for decades with upgrades over the years. Both of these requirements don't come cheap.

However, the days of this "just do it" mentality are rapidly coming to an end as the DoD scrambles to squeeze as much as possible from its annual budgets. It's obviously not the first time Pentagon spending has been under the microscope, but it may be the first time that, technologically speaking, there are ways purveyors of digital and microwave technology can broadly help make it happen. OpenRFM is well suited to become one of these enablers.

By standardizing the electromechanical, software, control plane, and thermal interfaces used by IMAs, OpenRFM could almost completely change the way defense systems are built, providing a standards-based roadmap that would allow systems to be used on more than one platform, which is the opposite of why they are designed and built today. Even though a radar built by manufacturer A for one aircraft may perform very similar functions as one built by manufacturer B on another aircraft, they are invariably almost totally different, incorporating proprietary designs that can't easily be "ported over" from one platform to the next.

There Are Precedents... Sort Of

There are market sectors in which digital and microwave technologies comfortably coexist in a quasi-standard form factor — smartphones and other wireless devices being a classic example. In all these cases, integration was driven by need: There's simply no room in these small, typically battery-powered devices for disparate technologies singing different tunes.

Bluetooth products and Zigbee owe their very existence to wireless communications, as do Wi-Fi access points, some medical devices, and other products in which size or some other constraining factor dictates that digital/microwave integration is essential. Their small size and specific digital and microwave requirements allow this integration to be accomplished at the wafer level using a System on a Chip (SoC) approach that can integrate a truly astonishing amount of functionality in a single device. A good example is Silicon Labs' Si468x CMOS digital receiver ICs that are essentially low-power, multi-band, digital broadcast receivers that support AM, longwave, shortwave, FM, AM and FM HD radio, digital audio broadcast, and digital multimedia broadcast.

Among other things, they include an OFDM channel demodulator, audio DSP processing, on-chip source decoding, I2S digital audio output, a stereo audio DAC, VCO, PLL, frequency synthesizer, AGC SPI, and I2C control interfaces. Typical applications include multimedia handsets, media players, GPS devices, tablets, all types of fixed and portable radios, boom boxes, and entertainment systems.

One other SoC that packs broad functionality in a single device is Broadcom's BCM20736 WICED SMART Bluetooth SoC that allows OEMs to develop applications on its ARM Cortex-M3 processor. It includes an RF and embedded Bluetooth stack, support for wireless charging, software for profiles, the stack, APIs, and application software development kits, and has two serial peripheral interfaces (SPIs). The device, designed to be powered by a single coin cell at 1.2 VDC, is pin-compatible with Broadcom's Bluetooth Smart SoCs and allows secure updates to be implemented over the air.

Challenges

It's important to note that other defense electronics manufacturers have, out of necessity, developed their own ways to integrate RF and microwave technology onto board-level products. The difference is that, unlike Mercury, none has gone so far as to put a standards proposal on the table.

The first major challenge is getting the member companies within VITA to give approval for incorporating OpenRFM in OpenVPX after the usual standards development process. Interestingly enough, OpenVPX was also proposed by Mercury in 2009 as a way to broaden the acceptance of the VPX standard, and a working group of 28 companies participated in its development. The challenge will be to get many of these

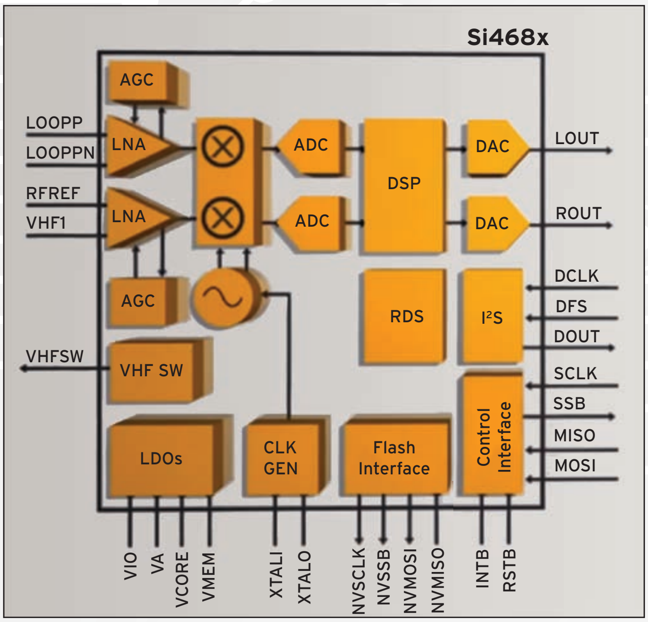

Figure 3. Silicon Labs Si468x Digital Radio Receivers are an all-in-one solution.

same companies — almost none of which have interest or technical capability in RF and microwave technology — to buy into the concept. If they do, they would effectively be signing up to adhere to the standard if and when they do choose to add this technology to their product portfolio.

If OpenRFM takes root after adoption by VITA, the next and unquestionably most formidable challenge will be enticing the microwave industry to participate. This will be orders of magnitude more difficult than with OpenVPX for two reasons. First, VPX was already a VITA standard and required introduction of no new or drastically different technology. The challenge was getting Mercury's competitors on board. Second, the microwave industry, at least those parts of it serving aerospace and defense, has no standards body like VITA and is unaccustomed to dealing with any standards other than those dictated by DoD or other government agencies. It also has no relationship with VITA or the embedded systems business. Consequently, OpenRFM is likely to be a tough sell to these companies, although it may not initially matter whether this industry participates or not as it takes its orders from DoD and prime contractors; however, if DoD mandated its use, IMA manufacturers would have to conform to it.

Limitations and Opportunities

OpenRFM is inherently best suited for applications in which RF power output is low, as higher power requires larger components, resulting in subsystems too large to be integrated within OpenVPX. This means it will be used in receive rather than transmit sections of systems and in the low-power driver stages of RF power amplifiers, low-noise amplifiers, and other small-signal subsystems.

It would, however, be appealing for use in subsystems operating at very high frequencies (the millimeter-wave region), where components are much smaller than at lower frequencies, and power levels are lower as well. While it may not be possible to incorporate higher-power RF and microwave subsystems within OpenRFM, it may still be possible to make them compatible with it, which would still allow systems from various microwave manufacturers to build to the standard.

What's Next

OpenRFM was first exposed to the embedded community at a VITA meeting called Embedded Tech Trends. For the microwave industry, however, the next step is for Mercury to present its case for OpenRFM at IEEE's International Microwave Symposium. Sponsored by IEEE, the industry's primary international symposium and exhibition brings together almost the entire industry throughout the world. At the 2018 event, the company introduced its first 500-to-18 GHz digital microwave tuner targeted at electronic warfare applications. The Ensemble RFM-1RS18 consists of three OpenRFM modules in a single-width 6U VXS-format package. Within this space is a converter with four IF outputs and a frequency synthesizer.

It's important to note that Mercury wins regardless of whether or not OpenRFM is accepted within the embedded systems and microwave communities or is made part of OpenVPX, as it can still use this formula for its own products. Nevertheless, the impetus for integrating RF and microwave IMAs within board-level products that, with few exceptions, are totally digital, was to meet the looming challenges facing both industries as DoD is serious about changing the design paradigm. It makes this case at every available opportunity, in programs conducted by the Defense Advanced Research Projects Agency (DARPA) and at R&D laboratories in the Army, Air Force, and especially the Navy. It may take some time before the defense industry broadly accepts this challenge, but in the long term, it has little choice.

Another factor behind the development of OpenRFM was the most basic: to integrate digital and microwave technologies that have existed in their impenetrable stovepipes since World War II. After all, design engineers are only human and typically resistant to the disruptive challenges of change. Without some force driving these two technologies together in the embedded space, who knows when — or if — it would ever occur, and OpenRFM may just be such a catalyst.

This article was written by Barry Manz for Mouser Electronics, Mansfield, TX. Reprinted with permission from Mouser Electronics. ***For more information, visit http://info.hotims.com/69511-542.*** ■

New Technologies Tackle UAV Challenges

PHOTO: U.S. ARMY

Unmanned vehicles for air, land, and water represent an increasingly-important capability for virtually all military services worldwide. Because each new generation of device technology promises capabilities and higher levels of overall performance for unmanned vehicles of all types, military customers expect these benefits in the latest UAV offerings. To meet these expectations, systems engineers must deal with increasingly difficult problems of thermal management, power budgeting, EMI emissions and susceptibility, weight, size, cabling, connectors, antenna management, and command/control functions. Several new technologies and innovative packaging concepts now provide better solutions to these issues.

Thermal Management Solutions

Since unmanned vehicles have no requirements for human occupants, SWaP penalties for life-support systems and crew or pilot quarters are eliminated. The downside for system designers is that UAVs are therefore expected to operate over temperatures ranging from -50°C to +75°C. Often unappreciated is the fact that these limits can be more difficult to meet when the vehicle is on the ground and powered down for sustained periods of time between missions. Primary design concerns here are non-operational effects of mechanical stress and packaging integrity. These are usually well documented as storage temperature specifications, which can help predict survivability on mid-winter airstrips in Alaska and mid-summer runways in the mid-East.

There are numerous operational temperature issues. If the unit is not provided with a standby heating system, starting up from -50°C often requires a warm-up delay and the use of heaters or partially powering up some equipment to heat the more temperature-sensitive equipment. Notorious for problems at low temperatures are crystal oscillators, batteries, capacitors, and some semiconductor devices like A/D converters.

At high start-up temperatures, similar care must be taken to cool down the equipment before applying full power. Otherwise, components like processors and FPGAs can sustain permanent damage, completely disabling critical subsystems within the vehicle. Cooling strategies must transfer heat to a heat exchanger coupled to the outside surface of the vehicle or to outside air. Methods of moving the heat include heat pipes, carbon nanotubes, circulating air or liquid, refrigerators, thermionic coolers, and some promising new nano-technology cooling engines.

Once the UAV is operational, these same structures can continue to regulate internal temperatures, often consuming relatively little energy because of self-heating of the payload systems and the normally cold skin temperature of the vehicle once it reaches operational altitude.

Often a simple, independent vehicle thermal management processor capable of operating across the entire temperature range controls the heating and cooling systems and initiates operational power to the UAV systems when ready.

Open Packaging Standards

Even though each UAV targets a specific class of applications and missions, systems designers can reap significant benefits by exploiting the latest open standards for the many internal subsystems. An outstanding example is the VITA OpenVPX standard, now also adopted by ANSI. It defines numerous mechanical and electrical profiles for circuit boards, backplanes, chassis, connectors, as well as cooling and power distribution methods, all capable of withstanding severe military environmental conditions.

Particularly appropriate for UAVs are the numerous cooling methods for OpenVPX defined in the VITA 48 standard, which includes conduction, liquid flow-through, air flow-through, air flow-by, and variants. Designers can select the most appropriate cooling technique for a given UAV by surveying vendors for availability of VPX solutions sharing a common VITA 48 cooling method. This can greatly simplify the overall thermal design of the vehicle.

Another important benefit of open standards is improved life cycle support, especially for military programs looking for multi-year acquisition and installation phases, followed by ten or more years of operational life, that can be fully supported with maintenance and repairs. Obsolescence of critical components like memories, processors, or FPGAs is all too common and quite difficult to predict. In some cases, redesign of modules or subsystems is the only solution, and compliance with a well-defined standard helps ensure success.

Upgrades become far easier when an older module can be replaced with a new one that exploits the latest technology and delivers new performance levels, but yet still complies with the OpenVPX infrastructure to minimize system integration efforts.

Making Good Connections

A peek inside a military UAV reveals a staggering array of wires, cables, harnesses, and connectors, accounting for a significant share of vehicle weight, and having a major impact on both operational costs and mission endurance.

Because UAVs are loaded with sensors, antennas, processors, cameras, telemetry systems, radios, radars, navigation systems, jammers, power supplies and cooling systems, the necessary interconnects are often highly specialized to match the link.

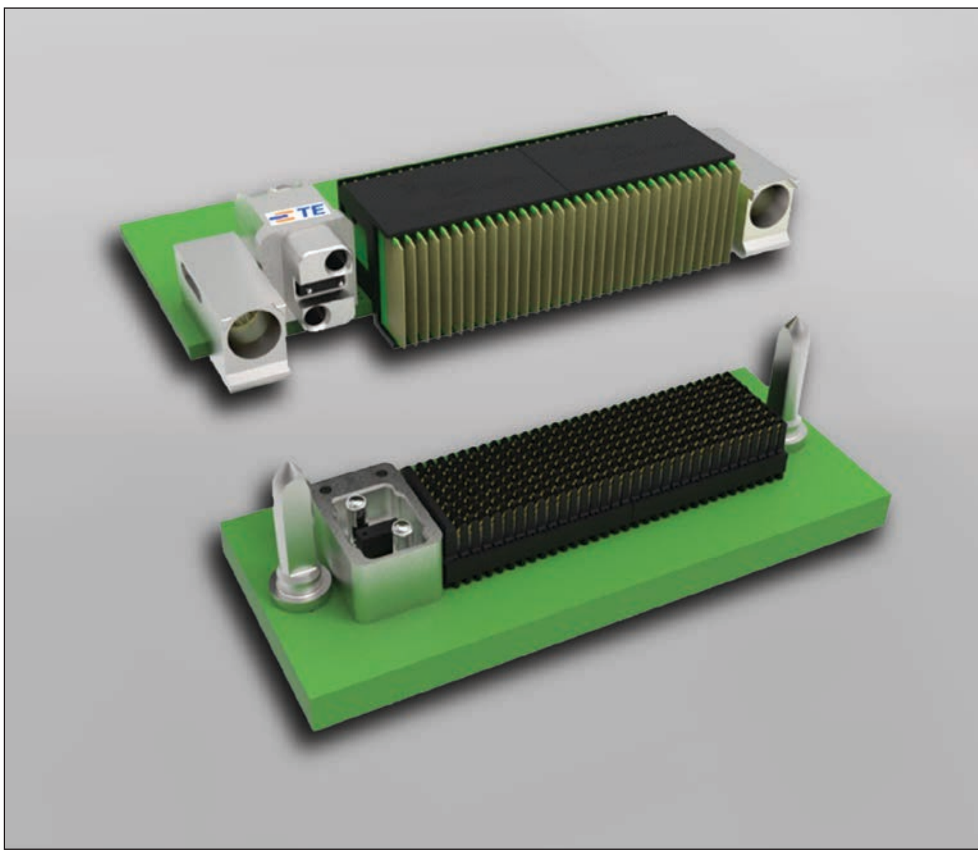

Figure 1. VITA 66 optical connector blocks for a 3U VPX module (top) and backplane (bottom) are installed in place of electrical RT connectors (black). Each holds an MT ferrule with 24 fiber cables.

TE CONNECTIVITY LTD.

Some new power systems distribute higher voltages using smaller diameter wires to minimize weight. Many new classes of POL (point-of-load) switching regulators drop the nominal 24 or 48 VDC distribution bus to lower local voltages required for each subsystem, while maintaining high efficiency across a wide range of supply voltages and load currents. More complex devices can maintain regulation across high pulse current loads to support radar and countermeasure equipment.

New Technologies Tackle UAV Challenges

Parallel digital lines for high-speed data connections are being replaced by gigabit serial links at virtually every level of embedded systems. The benefits are fewer wires, smaller space, and higher rates. Within board-level products, gigabit serial links join processors, FPGAs, data converters, network interfaces, and storage interfaces. Within a chassis, these same serial links stretch across the backplane for connecting boards and for bringing I/O signals to bulkhead connectors. UAV subsystems are now exploiting these serial links to replace fat, parallel data cables to save space and weight. Two of the most popular gigabit protocols in UAVs are PCIe and Ethernet at speeds of 1, 10, and 40 GB/sec.

Switching to Light

A major trend for high-capacity data links is replacing copper cables with optical cables that offer many compelling benefits.

Optical cables are completely free from EMI (electromagnetic interference) emissions. This not only eliminates unwanted radiation that can induce signals in other cables and sensitive equipment, it also prevents unauthorized interception of sensitive information by cable "sniffers." The same trait makes optical cables immune to interference and contamination from generators, antennas, and other noise sources.

Optical cables also offer less weight; smaller diameter; lower cost per foot; immunity to water, salt, and corrosion; as well as greater tensile strength. Connectors do require care in cleaning, installation, and handling, but once installed, are quite reliable.

All these benefits of optical cables are vitally important to UAVs, where many different subsystems, sensors, and antennas are crowded together in very close proximity.

Over the course of several years, the VITA 66 group has spawned several variations of optical interfaces between VPX modules and backplanes. The most recent ones use industry-standard MT ferrules, each typically containing 12 or 24 optical fibers. Metal housings on the modules and backplane replace the electrical RT connectors, and provide precision engagement and spring-loading of the mating MT ferrules for reliably aligning the polished ends of each fiber when the module is inserted. (Figure 1)

Several component vendors are offering new compact, power-efficient optical/copper interface devices directly compatible with the gigabit serial ports on FPGAs and processors. Their products are competing for design wins for VITA 66 systems, which helps advance performance levels and lower costs.

Taming RF Signals

RF signals to and from antennas traditionally require bulky coaxial cables not only to minimize signal loss, but also to protect against interference when passing near and between power generators, switching power regulators, and transmit antennas.

New major U.S. DoD initiatives like SOSA (Sensor Open System Architecture), MORA (Modular Open RF Architecture), and CMOSS (C4ISR/EW Modular Open Suite of Standards) share a common goal of digitizing RF and IF signals as close to the antenna as possible, and then delivering digital streams via

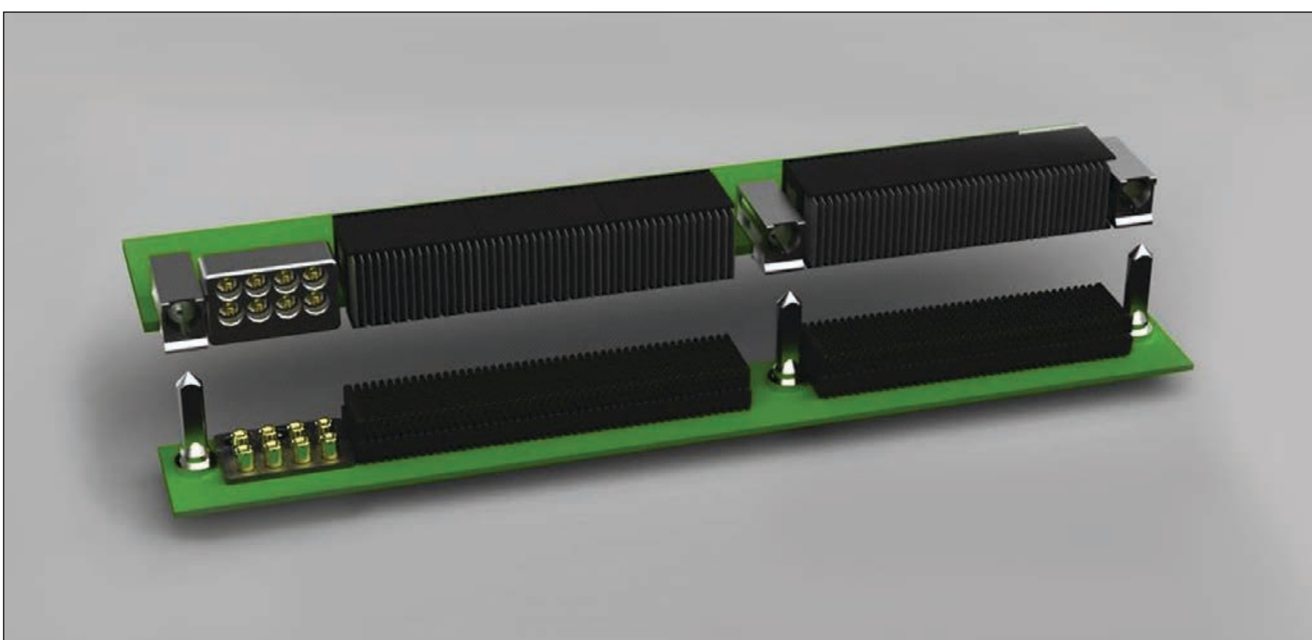

TE CONNECTIVITY LTD.

Figure 2. VITA 67 RF connector blocks for a 6U VPX module (top) and backplane (bottom) are installed in place of electrical RT connectors (black), and support eight coaxial connectors.

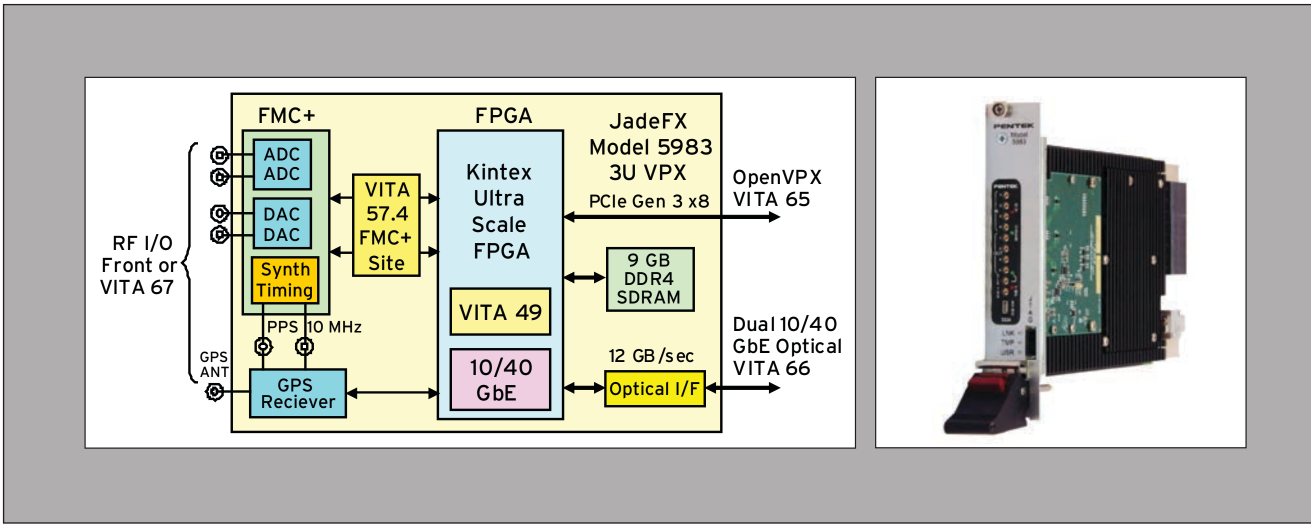

Figure 3. Pentek Model 5983 3U OpenVPX Kintex UltraScale FMC+ Carrier supports numerous VITA standards for new DoD initiatives for software radio embedded systems.

gigabit serial network links. The same applies to transmit signals, which are sent as digital streams to transmitters, where they are converted to analog, power amplified, and delivered to the local antenna. And, tracking the migration to optical, these standards call for optical interfaces for the high speed links.

This new architecture brings many benefits. RF circuitry, amplifiers, data converters, FPGAs, and network interfaces can all be incorporated within compact subsystems directly behind the antennas. These "radio heads" are connected via network switches to the appropriate equipment, neatly eliminating coaxial cables and RF switches, and traditionally hard-wired sources and destinations.

Not surprisingly, these new DoD initiatives incorporate OpenVPX as the hardware platform, capturing the many benefits listed earlier. To facilitate designs of radio heads, VITA 67 standards define several generations of RF backplane interfaces to simplify system integration. Mating metal housings on the modules and the backplane contain multiple coaxial connectors that are installed in place of some of the VPX electrical RT connectors. This eliminates the need for front panel coaxial cables, which greatly complicate service and maintenance operations.

Some of the earlier VITA 67 standards supported four or eight connections, as shown in Figure 2. But the new smaller Nano RF connector variants now under discussion provide up to 26 coaxial connections, which help support phased-array antennas that require one RF signal for each element. These electronically-steered arrays are particularly well suited for UAV radar and EW systems, where mechanical steering is often impractical.

PENTEK, INC.

Observing Protocol

With so many diverse subsystems within the typical UAV, orchestrating them to conduct a coordinated mission is a daunting task. One major initiative is the new VITA 49.2 Radio Transport Protocol approved within the last year. It defines standardized methods for delivering control, status, and payload data for digital software radio subsystems so that information streams to and from different types of radios, radars, EW, and SIGINT systems share common formats.

This allows signals acquired by one radio head to be delivered across switched optical networks to one or more consumers, so signals can be shared for different purposes. VITA 49 adds one more degree of consistency, not only between vendors, but also for upgrades and maintenance.

Figure 3 shows a recently announced 3U VPX software radio module targeting the new DoD architectures like SOSA and MORA. Front end data converters operating at 3 GS/sec capture and generate wideband analog RF signals, matching requirements for emerging radar, EW and communications bandwidths required for advanced UAV subsystems. It also incorporates 10 or 40 GbE interfaces to transfer VITA 49 digital IF/RF signals across VITA 66 optical backplane links to an OpenVPX signal processing system.

Looking Forward

The expanding role of UAVs in military operations, coupled with the rapid evolution of new device technology ensures a vigorous pursuit of extending UAV capabilities and performance levels to meet new requirements. It is clear that open standards are becoming increasingly important, not only to help integrate these new technologies, but also to help secure customer adoption of next generation vehicles.

This article was written by Rodger Hosking, Vice President, Pentek, Inc. (Upper Saddle River, NJ). ***For more information, visit http://info.hotims.com/69506-500.*** ■

Protecting Critical Data on Unmanned Underwater Platforms

Emerging mission requirements from global defense forces are driving new programs and applications for Unmanned Underwater Vehicle (UUV) platforms. Like their airborne counterparts, UUVs are ideal for deploying Intelligence, Surveillance and Reconnaissance (ISR) mission payloads. To speed the development of these autonomous vehicles, system designers are turning to small form factor (SFF) Commercial off the Shelf (COTS) technologies previously proven in Unmanned Aerial System (UAS) deployments. These low-power SFF subsystems, including miniature network switches and mission computers, are ideal for use in UUVs for which any additional weight or power consumption can have significant detrimental effects on mission distance and duration. By selecting proven rugged COTS solutions, already tested and qualified to the extreme demands of MIL-STD-810G, MIL-STD-461, MIL-STD-704 and/or RTCA/DO-160G standards for environmental, power, and EMI compliance, UUV system developers have found that they can greatly accelerate their program integration and reduce overall risk.

The Data Transport System 3-Slot (DTS3) is a rugged Network Attached Storage (NAS) file server for use in Unmanned Aerial Vehicles (UAV), Unmanned Underwater Vehicles (UUV), and Intelligence Surveillance Reconnaissance (ISR) aircraft. Easily integrated into network centric systems, the DTS3 is a turnkey, rugged network file server that houses three Removable Memory Cartridges (RMC) that provide quick off-load of data.

CURTISS-WRIGHT DEFENSE SOLUTIONS

Another important concern for UUV platforms is how to protect the critical data that they capture during an ISR mission. To protect sensitive data-at-rest (DAR), size, weight and power (SWaP) optimized COTS solutions, such as Network Attached Storage (NAS) devices that support data encryption, can mitigate the risk of mission data falling into an adversary's hands. Even better, supporting the NAS with Netbooting (NetBoot) techniques further reduces SWaP by eliminating the need for multiple storage devices and increases data security.

In one recent example, a platform developer defined a common reference system architecture for a new family of Larger Diameter UUVs (LDUUV) using small form factor COTS mission processors, network switches, and NAS line replacement units (LRU). To support the various control, monitoring, and network functions of the UUV platform, the developer specified robust technical requirements for the mission computer and network switch LRUs. The processor systems required low-power multi-core Intel CPU architectures supported with a large number of Ethernet, serial, and digital I/O interfaces, together with a VxWorks real-time operating system (RTOS). The managed Ethernet switches, used to network the computers with onboard sensors and storage devices, required advanced Quality of Service (QoS) traffic prioritization and IEEE-1588 Precision Timing Protocol (PTP) support to enable time stamping with nanosecond accuracy.

For the platform's mission computer and network switch requirements, which needed to be able to meet the program's rigorous technical, cost, and schedule requirements, the UUV developer selected multiple Curtiss-Wright SFF COTS-based systems. The LDUUV's mission computer processing is provided by two Parvus DuraCOR 311 units, one of the smallest rugged mission processors on the market. Network switch functionality is provided by a miniature "pocket-sized" Parvus DuraNET 20-11 8-port Gigabit Ethernet switch, which weighs a mere half a pound (0.23 kg). The fully managed 10/100/1000Base-T switch provides carrier-grade network management together with IEEE-1588v2 precision timing capabilities. Both the mission computer and network switch were pre-qualified to a very comprehensive range of MIL-STD-810, DO- 160, MIL-STD-704, and MIL-STD-461 tests for extreme environmental and EMI conditions.

To protect the onboard mission data, the platform developer required a NAS device that could encrypt DAR to national standards using Government Off the Shelf (GOTS) Type 1 devices or

The Parvus DuraCOR 311 is an ultra-small form factor (USFF) rugged embedded computer/controller based on a low-power, quad-core Intel Atom 3845 (Bay Trail-I) processor equipped with a rugged flash disk and PCIe-Mini Card I/O expansion slots. Featuring a fanless IP67 design with MIL-performance connectors and extended temperature operation (-40 to +71°C), this miniature multi-core rugged Commercial Off the Shelf (COTS) processor is an ideal x86 mission computing solution for size, weight, power and cost (SWaP-C) sensitive applications.

The Parvus DuraNET 20-11 is an ultra-small form factor (SFF) rugged Commercial Off the Shelf (COTS) 8-port Gigabit Ethernet (GbE) switch optimized for extremely demanding size, weight and power (SWaP) constrained vehicle and aircraft platforms exposed to harsh environmental and noisy electrical conditions (e.g. high altitude, extreme shock and vibration, extended temperatures, humidity, dust and water exposure, noisy EMI, and/or dirty power).

Protecting Critical Data on Unmanned Underwater Platforms

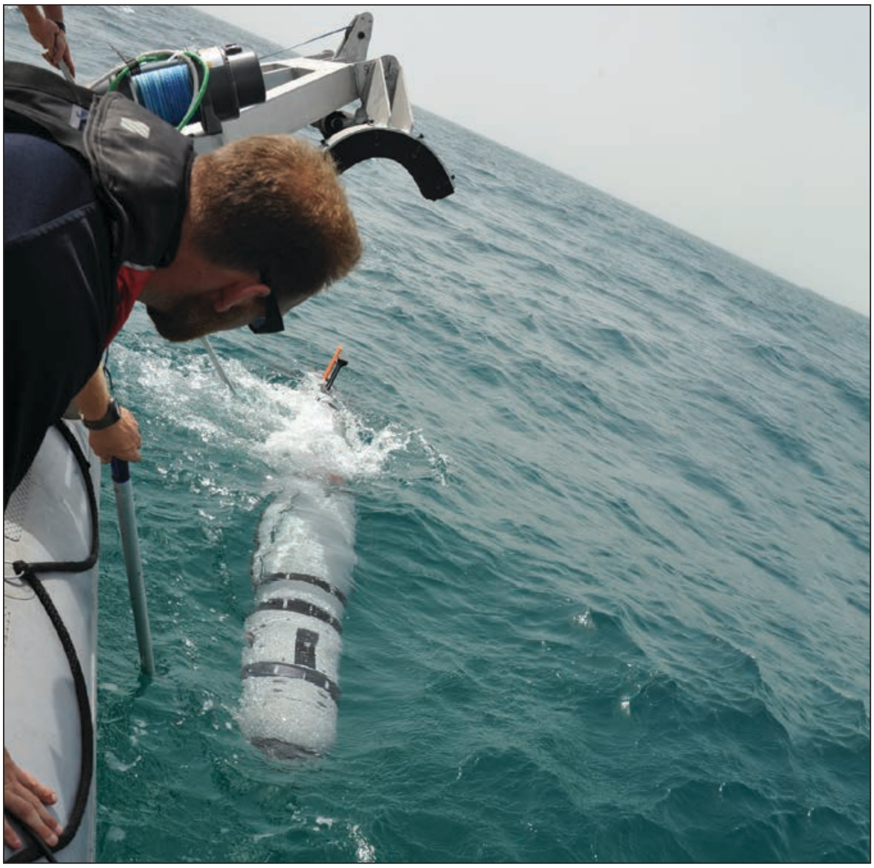

Protecting the critical data a UUV has collected is a major concern should the platform fall into an adversary's hands.

commercial encryption methods able to meet NSA guidelines. Hosting the Operating System (OS) software for the LDUUV's embedded computers on a NAS server, instead of on local media, would enable the elimination of multiple Direct-Attached Storage (DAS) devices, which delivers the benefits of significantly reducing SWaP and simplifying maintenance and future software upgrades.

The NAS device also needed to support NetBoot of network clients. The combination of encryption and NetBoot would ensure that the runtime software used to boot all of the platform's embedded computers was secure. Without encryption, if the UUV was captured, the deployed software on each module or system could be susceptible to intrusion, potentially enabling it to be reverse-engineered. By using a single NAS server that encrypts all of its data, the likelihood of malicious access is eliminated or greatly reduced. The use of NetBoot can limit the potential points of intrusion to the single point of an encrypted server protected with higher levels of security.

To fulfill the LDUUV's data storage needs, the system developer selected a Data Transport System 3-slot (DTS3) COTS-based rugged NAS file server previously field-proven in mobile vehicles, field ground stations, and aircraft. The NAS features three high-density removable memory cartridges (RMC) that enable data to be quickly off-loaded after a mission. The RMC supports SSD memory and features a 100,000-insertion cycle connector that includes a SATA interface. It enables Ethernet-based mission storage from the mission computer's clients and other devices. Also, because the NAS system supports PXE Booting, a form of NetBOOT, it enables the platform's x86 network clients to boot directly from the NAS instead of needing to boot from each individual LRU.

With PXE Boot, multiple network clients can be centrally managed and updated from a single location. For data security, the NAS appliance uses two separate hardware and software encryption layers, meeting the program's encryption requirements. An AES-256 bit FIPS-certified ASIC encryptor on-board the NAS provides the first layer in the form of hardware full disk encryption, while a FIPS certified AES-256 bit algorithm provides software full disk encryption for the second layer. Because this two-layer encryption approach follows the NSA's guidelines set forth in their DAR Capability Package it was able to meet the LDUUV system designer's DAR security requirements.

These SFF COTS systems delivered the CPU performance needed to support the LDUUV's vehicle control and data processing, along with the fast networking and specialized I/O interfaces needed to support its current ISR mission. By using a modular, open architecture COTS-based design, the system developer "future proofed" the LDUUV's architecture, easing the need for any later integration of expanded mission payloads as more sensors are inevitably integrated onto the platform. Cost-effective processing, networking, and storage LRUs are ideal for unmanned maritime platform use, ensuring that SWaP is reduced as much as possible, missions are optimized and critical data is protected. Use of these types of small, compact subsystems will help enable future deployment of new capabilities by naval forces for a multitude of potential UUV missions.

This article was written by Mike Southworth, Product Manager, Small Form Factor Systems, Curtiss-Wright Defense Solutions (Ashburn, VA). ***For more information, visit http://info.hotims.com/69506-502.***

APPLICATION BRIEFS

Ruggedized Computer Rack Cases

Verotec Inc
Manchester, NH
603-821-9921
www.verotec.us

Verotec has been selected by Lockheed Martin to provide custom 19" 6U 600 mm deep rack cases to house part of the ISTAR system in the AJAX next generation armored fighting vehicles developed for the UK Army.

A General Dynamics UK project, the AJAX program provides British troops with an improved level of protection and is at the core of the British Army's deployable Intelligence, Surveillance, Target Acquisition and Reconnaissance (ISTAR) capability. It enables the soldier to be at the point of collection of accurate all-weather commander information within a network-enabled digitized platform. Lockheed Martin UK has been awarded a $1 billion contract to deliver 245 technology leading turrets that contribute significantly to the ISTAR, survivability and lethal combat capabilities provided by the AJAX Reconnaissance variant.

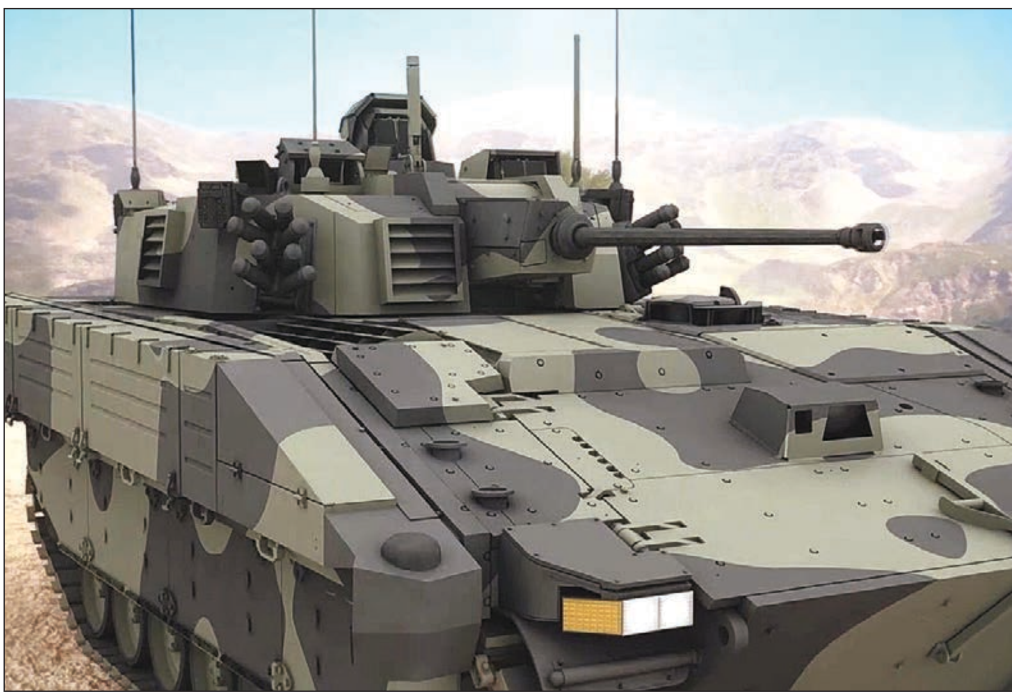

Verotec is supplying nine different versions of extended 600 mm deep steel rack cases, fully earth bonded and fitted with steel engraved front panels, custom chassis trays, and heavy-duty bar handles. The ruggedized rack cases are designed to withstand the high levels of shock and vibration that will be found in the tracked vehicles.

Originally known as the SCOUT Specialist Vehicle (SV) program, the AJAX program includes six variants: AJAX, ARES, APOLLO, ATHENA, ATLAS and ARGUS. Each AJAX variant will be a highly-agile, tracked, medium-weight armored fighting vehicle, providing British troops with state-of-the-art best-in-class protection. AJAX vehicles are developed upon a highly-adaptable and capable Common Base Platform, maximizing commonality in mobility, electronic architecture and survivability. A sophisticated, neatly packaged Electronic Architecture makes it the first fully-digitized land platform that is able to seamlessly integrate both current and future open system ISTAR and communication products.

Each AJAX platform variant has extensive capabilities, including acoustic detectors, a laser warning system, a local situational awareness system, an electronic countermeasure system, a route marking system, an advanced electronic architecture and a high-performance power pack. The panoramic Primary Sight provides advanced all-weather imaging technology capability, which allows the AJAX variant to find, engage and target at far greater ranges than the current UK Ministry of Defense core legacy platforms. Enhanced and modular survivability technologies ensure it will survive both current and future threats. Lethality is provided by the 40mm cannon integrated into a revolutionary, user-defined, fightable turret. Where the operation dictates, a fully stabilized Remote Weapons Station can be fitted to the turret instead of the Primary Sight.

The AJAX family of vehicles also has growth inherently built in. With an upper design limit of 42 tons of driveline capability, scalable and open electronic architecture and a modular armor system, it has the potential to combat future threats and incorporate new technology throughout the lifespan of the platform. As a result, AJAX provides the kind of growth capability needed to face the uncertain challenges of Future Force 2020 and beyond. AJAX will replace the less capable CVR(T), providing broad utility throughout the balanced Army 2020 force across all operations.

***For more information, visit http://info.hotims.com/65856-573*.**

APPLICATION BRIEFS

Rugged Server and Display System

General Micro Systems
Rancho Cucamonga, CA
(800) 307-4863
www.gms4sbc.com

General Micro Systems Inc. (GMS) recently announced that the U.S. Army will exclusively deploy powerful rugged server and display systems from GMS to run the multifunction video display (MVD) software within Type II medium mine protection vehicles (MMPV). The GMS system encoder, the system is a complete full motion video and control system with storage.

A program of record for the U.S. Army's Product Manager Mine Resistant Ambush Protected Vehicle Systems (PdM MRAP VS) and co-developed directly with the Army's Night Vision and Electronic Sensors Directorate (NVESD), the GMS rugged server and display system gives the MRAP mine-clearing vehicles a distributed platform with smart displays that run the Army's

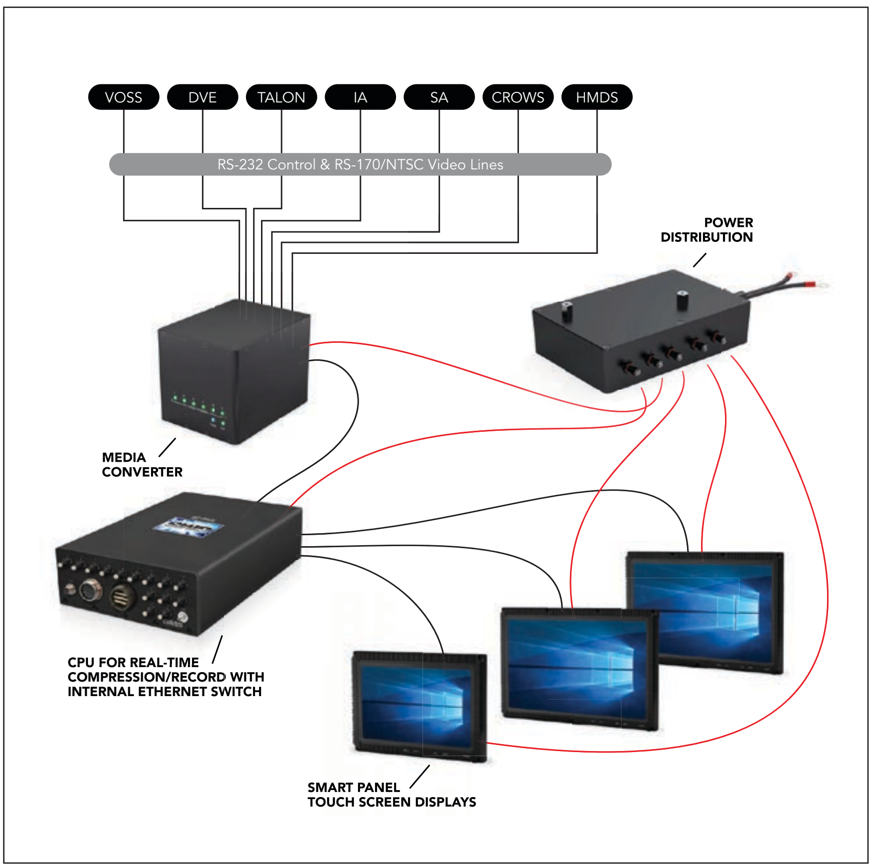

comprises four components — two chassis and two displays. It also includes an enterprise-class, ultra-rugged, secure server with an intelligent 12-port 1/10 Gigabit Ethernet switch, a router, mass-media storage, CITV/DVR, video-over-IP, and two ultra-thin, rugged smart-panel PCs. When coupled with a video portable MVD software. The system's hardware and software enable the seamless distribution of full motion video and control in real time with low latency from all sensor systems mounted on the MMPV Type II trucks to each crew station.

Compared to the traditional system used on the mine-clearing MRAP, the U.S. Army needed a combined hardware/software system that could provide full situational awareness at all times while also improving crew efficiencies. According to the contract synopsis from the Army Contracting Command - Warren (ACC-Warren), the MVD system "integrates full-motion video from all sources at all vehicle crew stations. The MVD system efficiently distributes images and sensor control to all crew stations within a vehicle, resulting in a single touchscreen display for each crew station capable of viewing and controlling all vehicle enablers, and creating a seamless common interface across all enablers."

With the MVD system, each networked crew station operates independently such that one crew member can control one sensor system while another crew member simultaneously controls or views another. The ultra-low latency system enables warfighters to drive "head down" in the vehicle, using only cameras and sensors without inducing motion sickness. The MVD system is integrated with the truck's radio so that it too can be controlled from any crew station or set up during pre-mission checks. Because the system is designed to enable soldiers to navigate without direct sight, it can be used in other programs as well.

The key to the system's anticipated success is a combination of low-latency networked video and data processing coupled with NVSED-created modular software. The software presents a standardized view of sensor feedback that is common across all workstation consoles. This gives operators immediate familiarity with different sensors and enables cross-training and cross-operation should the need arise. Moreover, new sensors and counter-IED payload processing can be added while the user interface remains consistent.

For more information, visit http://info.hotims.com/69505-552**.** ■