

**2014 NDIA GROUND VEHICLE SYSTEMS ENGINEERING AND TECHNOLOGY SYMPOSIUM
VEHICLE ELECTRONICS AND ARCHITECTURE (VEA) TECHNICAL SESSION
AUGUST 12-14, 2014 – NOVI, MICHIGAN**

**HIGH PERFORMANCE EMBEDDED COMPUTING REFERENCE
ARCHITECTURE FOR C4ISR/EW IN GROUND VEHICLES**

Mr. David Jedynak

Chief Technology Officer, COTS Solutions
Curtiss-Wright Defense Solutions
Austin, TX

ABSTRACT

Presenting a reference architecture for High Performance Embedded Computing for use in Ground Vehicles, based on OpenVPX, up to 40 Gigabit / Second data fabrics (Infiniband and Ethernet), methods of Remote Direct Memory Access, and Open Standard software layers (OFED). How to provide the appropriate chassis and backplanes to accommodate the HPEC modules, Signal I/O, and data fabrics which can then provide sophisticated capabilities, such as software defined radios, active protection systems, electronic warfare, and sensor processing (fusion and analysis). Illustrate paths for technology refresh, showing historical and expected gains in hardware performance across technology refresh cycles and the SWaP-C reduction for a fixed amount of processing capacity over time.

INTRODUCTION

Ground Vehicles are a mix of two vastly different computing requirements. Vetrionics, mission management, and various electro-mechanical control systems leverage highly SWaP-C optimized computing systems, leveraging the benefits of adjacent industry technologies, such as automotive and industrial. On the other hand, increasing demand for complex integration of sensors and systems to provide sophisticated protection, situational awareness, communication, and electronic warfare capabilities requires the use of High Performance Embedded Computing (HPEC) systems. To meet the challenges of incorporating HPEC systems in SWaP-C constrained ground vehicles, an open standard approach can be used, leveraging low-risk parallel development paths with technology refresh of processing hardware to meet the C4ISR/EW capability requirements at the time of fielding.

HPEC REFERENCE ARCHITECTURE FOR GROUND VEHICLES

The HPEC Reference Architecture for Ground Vehicles is intended to serve a general range of high performance C4ISR/EW functions, with strongest focus on the signal processing challenges of Communications and ISR/EW, as well as the computational challenges of data analysis and fusion. This includes such applications like radar, software defined radio, signal / electronic intelligence, jamming, video, and high performance sensor / actuation control

loops. Rather than detailing out the specific requirements and constraints of each application type, the HPEC Reference Architecture for Ground Vehicles can be abstracted with a set of qualitative goals and constraints:

Goals:

- Significant and Scalable Processing Capability
- Significant Data Transfer Capability
- Minimized Data Transfer Latency
- Modular High Performance Signal I/O
- Open-Standard Building Blocks and Interfaces

Constraints:

- Rugged
- Limited Size & Weight
- Limited Input Power
- Limited Cooling Capacity
- Significant Cost Sensitivity

To put it more bluntly:

- “We want a super computer!”
- “No. Here’s a rugged tablet.”

The diagram below shows Curtiss-Wright’s FabricX™ Generic HPEC Reference Architecture, which is a general presentation of the building blocks which go into a HPEC subsystem, and a good start for defining a focused HPEC Reference Architecture for Ground Vehicles. Bear in mind

that the diagram is showing the elements inside a single LRU, not a set of individual LRUs.

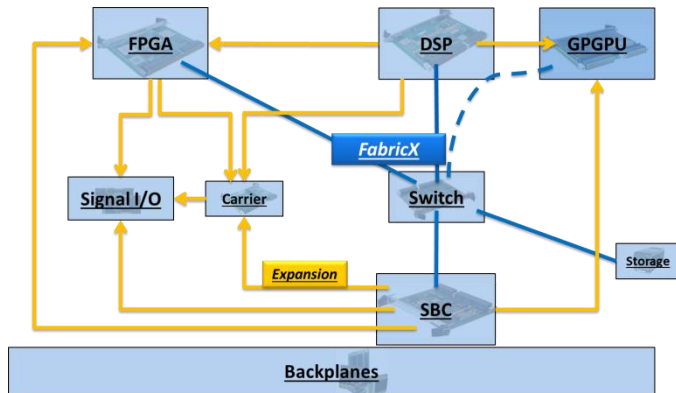


Figure 1: Curtiss-Wright's FabricX™ Generic HPEC Reference Architecture

The architecture shows multiple building blocks interconnected with switched high speed data fabric, various expansion buses, and modules of different types all installed within a backplane. Heterogeneous processing is provided by single board computers (SBC), digital signal processors (DSP), field programmable gate arrays (FPGA), and general purpose computing on graphics processing units (GPGPU). The architecture is intended to be implemented with rugged 6U VPX modules, with such features as air or liquid cooling. There may be multiple types of the various modules, e.g. 8 DSPs, 2 Switches, 2 SBCs, 2 FGAs, and 2 carriers with Signal I/O all in a 16 slot 6U VPX backplane.

This architecture is very well suited to developing very high performance subsystems for mobile / fixed radar, ship-borne radar and targeting, and highly capable signal processing for fast jet and wide body planes. Of course, most of these HPEC subsystems have per unit costs that would easily buy a number of tactical vehicles and even a good portion of combat vehicle, not to mention power and cooling requirements in the kilowatts.

The challenge, then, is to derive a good HPEC Reference Architecture suitable to ground vehicles that fits within the constraints. Making a few assumptions in the original goals and constraints helps to refine the generic architecture, as shown in Table 1.

Table 1: Assumptions for Ground Vehicle HPEC Goals and Constraints

Goal / Constraint	Assumption	Assessment
Significant Processing Capability	As much as possible within the constraints	Thermal will be the limiting factor
Scalable Processing Capability	Use Physical Modularity to scale this up and down	Modularity makes tech refresh for better capability very easy
Significant Data Transfer Capability	Nothing about ground vehicle environment affects this	There is a cost as it pertains to backplanes and interface chips
Minimized Data Transfer Latency	Nothing about ground vehicle environment affects this	There is a cost as it pertains to backplanes and interface chips
Modular High Performance Signal I/O	Nothing about ground vehicle environment affects this	This depends somewhat on the data transfer performance
Open Standard Building Blocks and Interfaces	In line with ground vehicle market	Use VPX
Rugged	Standard combat and tactical temperatures, other environmental requirements	Thermal will be limiting factor.
Limited Size & Weight	Both factors will drive to smaller boxes	6U / 3U mix, but smaller makes thermal harder
Limited Input Power	No more than 10-15 Amps, 28VDC	Thermal will be the real challenge
Limited Cooling Capacity	Conduction Cooled Cards with Natural Convection Cooling for Chassis	300-350 Watts with a cold-plate is an aggressive max for combat vehicles
Significant Cost Sensitivity	Modularity and Commonality are critical to per unit cost	Population / Depopulation depending on vehicle variants

Given these assumptions, the ground vehicle HPEC reference architecture can be derived. For ease of understanding, it's first presented as a reduced and altered version of the FabricX™ diagram, then presented with more details and a final form.

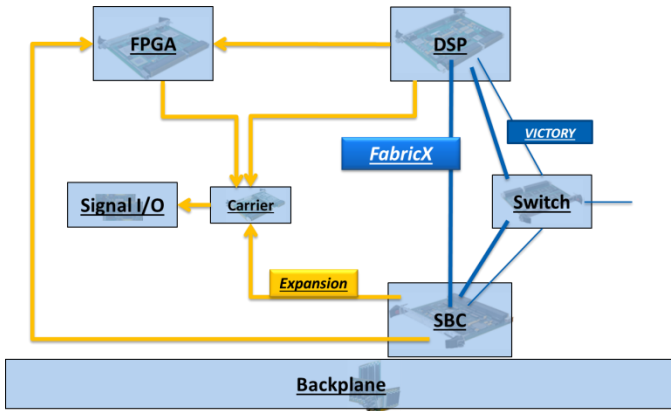


Figure 2: Initial Presentation of Ground Vehicle HPEC Reference Architecture

A couple of major items have been altered, as follows:

- GPGPU has been removed in favor of using more modest GPU capabilities on board single board computers
- External high speed storage has been removed as that is normally for things like unmanned system signal recorders
- The high speed data fabric switch is shown as optional in favor of direct SBC to DSP connections (mesh), but VICTORY (gigabit Ethernet) switching is retained for the OpenVPX Control plane
- Direct high speed data fabric connections to the FPGA have been removed for cost reasons, but expansion fabrics are retained.
- Signal I/O is shown as always hosted via a carrier since that aids both modularity and prevents heat peaks when hosting Signal I/O directly on Single Board Computers

In addition, the intent is that LRUs could be comprised of a mix of 6U and 3U modules. Note that moving to 3U from 6U may reduce size, cost, and weight, but it can present more difficult thermal gradient issues with high power cards. The 6U form-factor is well suited to higher power cards and the heat spreading they require.

Table 2 presents more details about each building block.

Table 2: Ground Vehicle HPEC Reference Architecture Building Block Notes

Building Block	Notes
Backplane	6 Slot 6U central switch backplane as starting point
Switch	6U data and control plane switch, or 3U Control Plane only
SBC	6U with control and data plane, or 3U with control plane only
DSP	6U with control and data plane
FPGA	6U Multi-FPGA module for processing or 3U single FPGA module for I/O and processing
Carrier	6U or 3U for carrying various Signal I/O mezzanines (XMC)
Signal I/O	High Performance Signal I/O mezzanines (XMC) for interfacing to sensors and actuators

With this constrained set of building blocks, the HPEC Reference Architecture for Ground Vehicles is presented in Figure 3, representing the maximum capability.

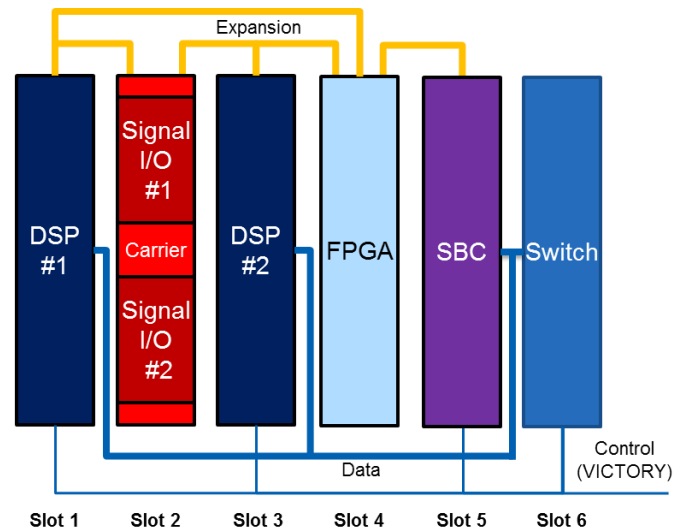


Figure 3: HPEC Reference Architecture for Ground Vehicles

This architecture includes the following items and expected maximum power consumption as shown in Table 3.

Table 3: HPEC Reference Architecture Modules and Maximum Power Consumption

Slot	Size	Module	Max Watts
1	6U	DSP Module #1	160
2	6U	Multi-FPGA Module	150
3	6U	DSP Module #2	160
4	6U	Carrier w/ 2 Signal I/O XMCs	30
5	6U	SBC	80
6	6U	Data / Control Switch	80
Total Maximum Module Watts			660

What’s immediately clear is that this is a very power hungry box, and cooling it via natural convection alone will be exceedingly difficult, requiring the use of such techniques as cold plates to meet the highest ambient temperature environments. Also important to understand is that 660 Watts is only the module power. Typical efficiency of the internal power supply (75-80%) pushes the overall power consumption to about 850 Watts.

It’s important to understand all the other physical parameters, such as size and weight, as well as performance. With the assumption that each module is roughly 2.2 pounds, that’s 13.2 pounds of modules. The rest of the chassis (including internal power supply and backplane) will be around 20-30 pounds, depending on the amount of metal need for thermal control (e.g. fins, cold plate). That means the entire unit will be anywhere from 35-45 pounds. The size will also vary a bit due to thermal control measures, but a rough estimate would be about 13” x 9” x 8” (about 1000 in³).

What’s even more important to understand is that the performance of the box is not at all static. As of the writing of this paper in 2014, assume the following capabilities which were available at the start of the year:

- ~10 Gigabaud signaling on a copper backplane
- PCI-Express Gen 3 (985 MB/s per lane)
- Expansion plane is x4 PCIe Gen 3
- Data plane is Ethernet and Infiniband
- Processors are Intel 4th Generation Core i7 (Haswell) with embedded GPUs
- GPUs are used for GPGPU in DSP modules
- GPU is used for user interfaces in SBC module
- FPGA is Xilinx Virtex 7 (690T)

With these assumptions, the reference architecture performance would be as shown in Table 4.

Table 4: Reference Architecture Performance for 2014

Module	Giga-FLOPS	Data Plane	Expansion Plane (x4)
DSP Module #1	1400	40 Gbps	31.5 Gbps
Multi-FPGA (x3) Module	2900	N/A	31.5 Gbps
DSP Module #2	1400	40 Gbps	31.5 Gbps
Carrier w/ 2 Signal I/O XMCs	N/A	N/A	31.5 Gbps
SBC	350	40 Gbps	31.5 Gbps
Data & Control Switch	N/A	40 Gbps & 1 Gbps	N/A
Total	6000		

That’s 6 Tera-FLOPS with a data fabric running at 40 Gigabits / second. This represents a significant amount of computational ability for the 660 Watts of modules (850W total). Comparing the performance to the various SWaP parameters provides an important set of metrics for understanding how performance changes with time.

Table 5: Performance / SWaP Metrics for 2014 Reference Architecture

Metric	Value
GFLOPS / Watt	6000 / 850 = 7
GFLOPS / Pound	6000 / 45 = 133
GFLOPS / in ³	6000 / 1000 = 6

PERFORMANCE IMPROVEMENTS OVER TIME

The current reference design is capable of the performance as shown in the previous section, roughly 6 Tera-FLOPS. What’s important is to understand how this will change over time given Moore’s Law (doubling in power performance roughly every 18-24 months).

For reference, the top supercomputer in the world in 2000 was the ASCI Red, designed to simulate nuclear detonations for Sandia National Lab. Its specifications and metrics are shown below in Table 6 alongside those of the HPEC Reference Architecture for Ground Vehicles as presented for 2014.

Table 6: Historical Performance Comparison

Metric	ASCI Red	2014	Factor Improved
Performance	1 TFLOPS	6 TFLOPS	6
Size	~27 Million in ³	~1000 in ³	27,000
Power	850 kWatts	850 Watts	1,000
Data Plane	8Gbps	40Gbps	5
GFLOPS / Watt	1.17 x 10 ⁻³	7	6,000
GFLOPS / in ³	3.7 x 10 ⁻⁵	6	160,000

What’s noteworthy is that the performance improvements shown above are across roughly 14 years, which is within the normal lifecycle of ground vehicle production programs, and only twice that of the a typical 7 year material acquisition (TD / EMD / LRIP).

What’s even more noteworthy is the performance gain expected by changing just the assumption on the processor to the expected mid-2015 baseline (Intel Broadwell). Updated values are in bold in Table 7.

Table 7: Expected Reference Architecture Performance for mid-2015

Module	Max Watts	Giga- FLOPS
DSP Module #1	130	2400
Multi-FPGA (x3) Module	150	2900
DSP Module #2	130	2400
Carrier w/ 2 Signal I/O XMCs	30	N/A
SBC	65	350
Data & Control Switch	80	N/A
Total	605 (780 total)	8000

A comparison of the performance metrics of 2014 and mid-2015 are shown in Table 8.

Table 8: 2014 HPEC and Mid-2015 HPEC Performance Metric Comparison

Metric	2014	Mid-2015	Factor Improved
Performance	6 TFLOPS	8 TFLOPS	1.33
Size	~1000 in ³	~1000 in ³	No change
Power	850 Watts	780	1.09
GFLOPS / Watt	7	10.25	1.5
GFLOPS / in ³	6	8	1.33

Although improved, the overall thermal problem is still unresolved at this point in time. Assume the goal is to get the overall power consumption down to 300 to 350 Watts before fielding in order to meet the highest ambient temperature requirements, but the HPEC system requirement is the current baseline (6 TFLOPS). If the mid-2015 system depopulates one DSP, then the system’s performance will drop to 5.6 TFLOPS with 475 Watts of module power (616 Watts). This is going in the right direction, but it does change the nature of the system a bit (one DSP instead of two). It’s advantageous to continue to extrapolate the improvements forward with assumptions that the performance will continue to improve, even if the overall power stays the same. This is very conservative, because it only focuses on the improvements in the GPU portion of the DSPs and ignores additional power efficiency gains. Table

9 extrapolates the performance improvements through mid-2018, using the performance improvement factor of 1.33.

Table 9: Extrapolated HPEC Reference Architecture Performance Improvements

Metric	Baseline 2014	Mid-2015	End-2016	Mid-2018
Performance (TFLOPS)	6	8	10.6	14.1
Power (Watts)	850	780	780	780
GFLOPS / Watt	7	10.25	13.6	18.1
Performance (TFLOPS) @ 300-350 Watts	2.1 – 2.45	3.1 – 3.59	4.1 – 4.76	5.43 – 6.34

By Mid-2018, the system can run at 50% performance (via under-clocking & throttling) and achieve the desired thermal power of 300-350 Watts. This is important, because it is within a typical procurement cycle of TD and EMD. An HPEC subsystem design starting in a TD phase with this reference architecture in 2014 will be able to run at half-power by the time of an EMD phase environmental qualification program in 2018.

This is a conservative approach in that it changes nothing in the underlying architecture. There are still same number and type of cards, just upgrades to the two DSPs to latest generation processors. The topologies and I/O are unchanged. This approach can significantly de-risk the development, integration, and evolution of the subsystem, as fundamentally, the block diagram does not change.

USING THE REFERENCE ARCHITECTURE

The reference architecture is not just a collection of modules in a chassis. There are a number of features and capabilities to leverage in developing applications.

One of the most important aspects of an HPEC system is the data fabric, which allows high bandwidth communication between various modules. An essential portion of that is the concept of Remote Direct Memory Access, or typically referred to as RDMA. A major issue with high speed data fabrics is the process by which data is moved on and off the fabric by the various end nodes. Without DMA, the processor is involved (consumed, really) in moving data to and from the fabric interface and memory. With DMA, the processor is not involved. This is shown in Figure 4.

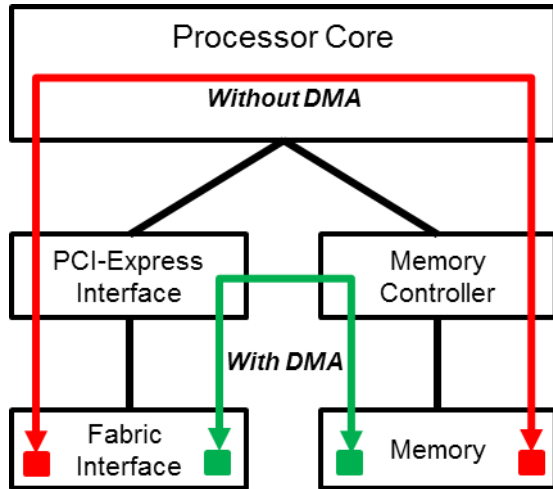


Figure 4: Data Path With and Without DMA

Remote DMA simply means that a remote node can set up and initiate a DMA through the fabric interface into another node's memory without that node's processor even being involved. This is highly beneficial for data which moves through multiple stages of an HPEC design, as a computing node can finish a task with data and transfer it to a partner for the next step without having to interrupt the partner's own work to coordinate the transfer. In addition, this means extremely low latency transfers, as there is very little overhead to transfer data from one node to the next. The Infiniband protocol supports RDMA natively, whereas Ethernet requires the addition of RDMA over Converged Ethernet (RoCE) software driver and fabric interface modifications to support this. Although they are similar in performance, Infiniband generally outperforms Ethernet with RoCE in both throughput and latency.

In addition to the RDMA features of the data fabrics, the OpenFabrics Alliance (OFA) has created a software stack, the Open Fabrics Enterprise Distribution (OFED), to utilize and leverage RDMA for HPEC applications, regardless of the particular underlying protocols, as shown in Figure 5.

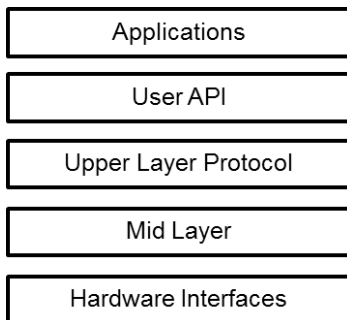


Figure 5: OFED Layers

Developers can build HPEC applications for ISR/EW applications using the OFED stacks, dramatically easing the development and integration risks.

It's also worth noting that there is some support for using RDMA to transfer into the memory of another device on the host's PCI-Express tree, rather than into memory. This would be useful for things like transferring data in an out an FPGA which is attached via PCI-Express to the host which is on the RDMA capable fabric.

EXAMPLE DATA FLOWS

The HPEC Reference Architecture can be used a number of different ways. The following example data flows provide insight into the flexibility of the architecture. The switch and carrier are omitted for clarity.




Figure 6: Input to Output Signal Processing Flow

The data flow of Figure 6 is well suited to sensor and actuator type systems, in which an input signal is analyzed in order to create an appropriate response output. This type of system is useful for things like jamming and active protection. The central box shows that DSPs and SBC are all connected as a cluster on the data fabric, and each is connected to one of the FPGAs on the multi FPGA module. The inputs and outputs are attached to the DSPs. Not shown is a management and user interface provided by the single board computer.

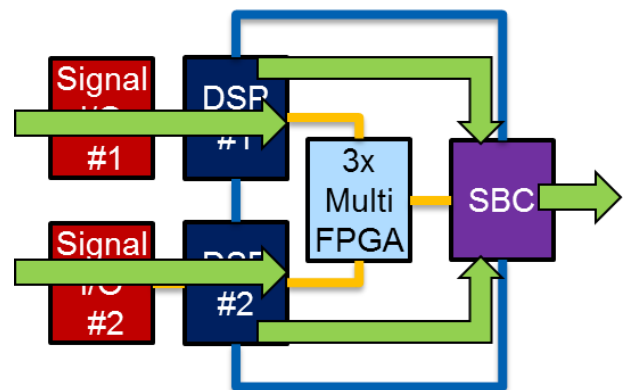


Figure 7: Combining Multiple Inputs

The data flow of Figure 7 is well suited to the fusion of multiple sensors, or the coordinated and coherent input of multiple sensor subunits. This type of data flow is well suited to phased array, large area image processing, or various signals intelligence. The DSPs are ingesting signals in parallel, using the FPGA resources as co-processors, and moving data in parallel to the SBC for final fusion and presentation to other systems or a user.

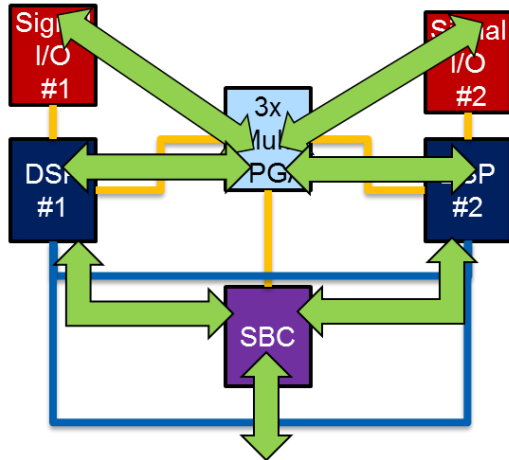


Figure 8: Separated I/O Signal Chains

The data flow of Figure 8 is well suited to implementing multiple separate I/O signal chains, where the I/O on each is completely independent of the other. This is useful for creating something things like multiple software defined radios that eventually feed data to a central management and user interface application, or a bridging / routing application.

This could also be used as a cross-domain transfer system between two separate waveforms, as the data could be placed into the 3rd FPGA of the multi FPGA module via the data fabric from each of the DSPs. That central FPGA could have a cross-domain rule set, which is managed only by the SBC.

Since the architecture has a central switch, it means that security boundaries can be created via fabric segmentation (VLAN). The last example highlights that capability.

Note that in all cases, the topology remains the same, which means the HPEC system could serve multiple purposes with only software changes if the I/O can be multiplexed properly. If enough I/O is available, then it's possible that the HPEC system could provide multiple functions at once if there's enough processing capability available.

CONCLUSION

A highly capable and scalable HPEC Reference Architecture is possible for ground vehicles. The continued advancement of processing capability means that applications can be developed using open standard modules and open standard development frameworks with the full expectation that performance metrics will improve. This allows the developers to focus on capability with planned technology refresh cycles to bring capability in line with vehicle constraints. The architecture is flexible, allowing for the development and deployment of HPEC systems which drive toward commonality for combat and tactical ground vehicles.