

Embedded Computing Enhanced by OpenVPX Multi-Processing Systems

Pete Jha, Curtiss-Wright Controls Embedded Computing (www.curtisswright.com)



The ANSI/VITA 65-2010 OpenVPX System Specification was officially recognized as an industry-wide standard on June 15th, 2010 (a date it auspiciously shares with the signing of Magna Carta in 1215), after its ratification by the American National Standards Institute (ANSI). While the Magna Carta became the basis of Western Democracy and our Judicial System, OpenVPX provides order to the rugged deployed embedded community with a framework that enables disparate computational entities to co-exist within a larger system while defining compliance to adjudicate disagreements.

OpenVPX was built on the legacy of several decades of open standards progress, beginning with the VMEbus, followed by VPX (VITA 46) and advanced most recently by the OpenVPX Technical Working Group. The OpenVPX system architecture defines a Multi-Planar, highly-rich switched fabric, low latency and high bandwidth system interconnect solution for heterogeneous systems composed of multi-board, multi-vendor and multi-chipset modules. The focal point of OpenVPX was to address system interoperability and codify a set of compliance rules.

Historical Perspective

Prior to 1980, the embedded industry was in state of disarray with many vendor-specific proprietary solutions which created interoperability issues within the community. The VMEbus Standard combined the Eurocard mechanical form factor with electrical specifications from VERSAbus. VMEbus provided a ruggedized and interoperable solution with support for real time control. Several enhancements, including VMEbus IEEE-1014, VME64, VME64x and VME320 improved the addressable data and address space of the computational node with speed improvements culminating with VME320 at 320 MByte/sec.

Following the advent of High Speed serial fabrics, the bandwidth supported by the VME Parallel Bus proved insufficient for the needs of higher-performance embedded systems. VPX addressed these emerging performance requirements with a highly scalable and flexible COTS standard that provides a multitude of high-speed pins, high speed board-to-board communications and higher power delivery. VPX vastly increased performance from VME320 to 6.25 Gbit/s, replacing VME's pin-connectors with new blade-style MultiGig RT2 connectors, while retaining popular 3U/6U

Eurocard form factor. VPX also introduced support for two-level maintenance, highly desired by military customers, through the complementary VITA 48 (VPX-REDI) mechanical specification.

Because the VPX specification was very flexible, the embedded industry once again faced the challenge of numerous, disparate vendor-specific implementations which threatened to impede interoperability. OpenVPX was specifically developed to address these interoperability challenges at the Systems Level. Figure 1, describes the Multi-Planar Systems Architecture, introduced by OpenVPX: expansion, data, control, management and utility (not shown) planes.

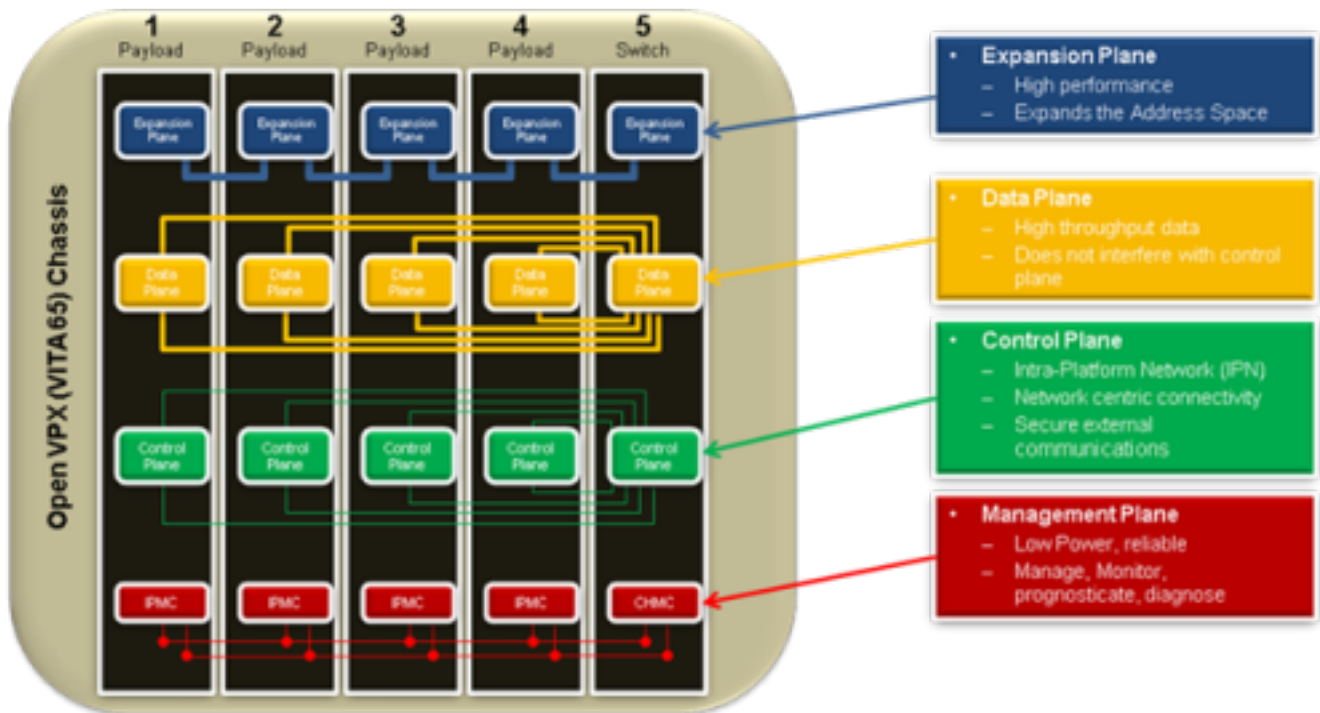


Figure 1: OpenVPX Multi-Planar Systems Architecture

Multi-Processing Systems Trends

Multi-Processing Systems are typically composed of Single Board Computers (SBCs), Digital Signal Processing (DSP) boards, Field Programmable Gate Arrays (FPGAs), Memory Boards, Switches, etc. A direct relationship exists between the ability to aggregate these computational entities within a chassis, and the number and types of complex applications that can be addressed.

These systems are employed in the military, aerospace and heavy industrial sectors such as radar, multi-INT (e.g. radar data exploitation and information dissemination), image processing, avionics, homeland security, telecom and transportation. Multi-Processing Systems are driven by the amount, size and frequency of data required to address the needs of the application.

The need for greater computational performance is driven by the higher frequency at which data is being collected, the need for better imagery and the proliferation of sensors deployed on a given platform. All of these trends increase the need for faster data rates. As a result, there is an ongoing, burgeoning demand for higher levels of intelligence, faster I/O and external networking. These drivers are

monotonically increasing as the industry continues to evolve. One noticeable effect of these computational drivers is the requirement to cool such systems due to the increased amount of generated heat.

Figure 2 describes a VMEbus-based Multi-Processing System. This system is composed of the fundamental computational entities, described above. This Multi-Processing system uses a StarFabric Network to interconnect entities to create a larger system. A StarFabric bridge and switch, located on a PMC mezzanine card, is utilized to communicate on the network and provides a 64-bit 133 MHz PCI-X Interface with up to 880MB/sec of I/O Bandwidth. PCI traffic is translated to a serial frame format for transmission across the network. The bandwidth is unsatisfactory for the modern computational needs of the embedded community.

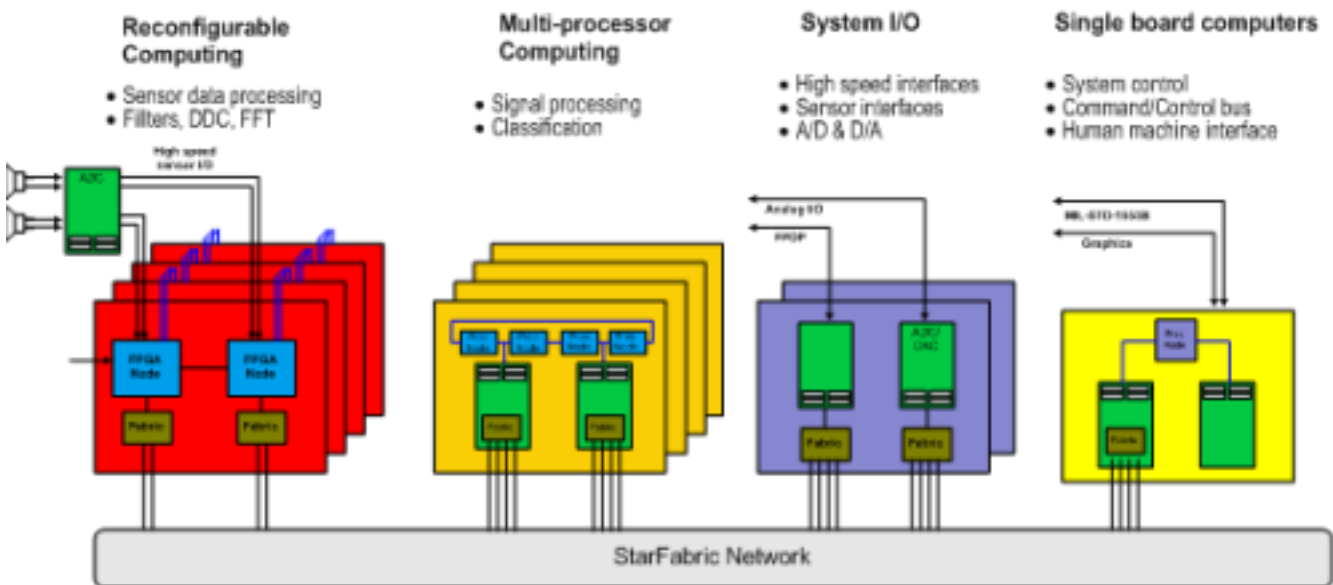


Figure 2: VME-based Multi-Processing System

Figure 3 describes an OpenVPX Multi-Processing System which uses the Multi-Planar Systems Architecture. It uses Serial RapidIO as the data plane, PCI Express as the expansion plane, and Gigabit Ethernet as the control plane. In comparison, the VME based Multiprocessing system had two serious limits - the amount of data that could be sustained by the VME signaling pins and the amount of power that could be dissipated per board slot.

VPX addresses both of these problems. One inherent issue with Multi-Processing Systems is data movement between processing nodes and subsystems. Switched fabrics directly address this problem by parallelizing the data movement through the Multi-Planar Architecture, thus minimizing both the processing stalls and the interruption of the input data streams. By reducing processor stalls and increasing incoming data streams, the cumulative computational performance of the system is significantly improved.

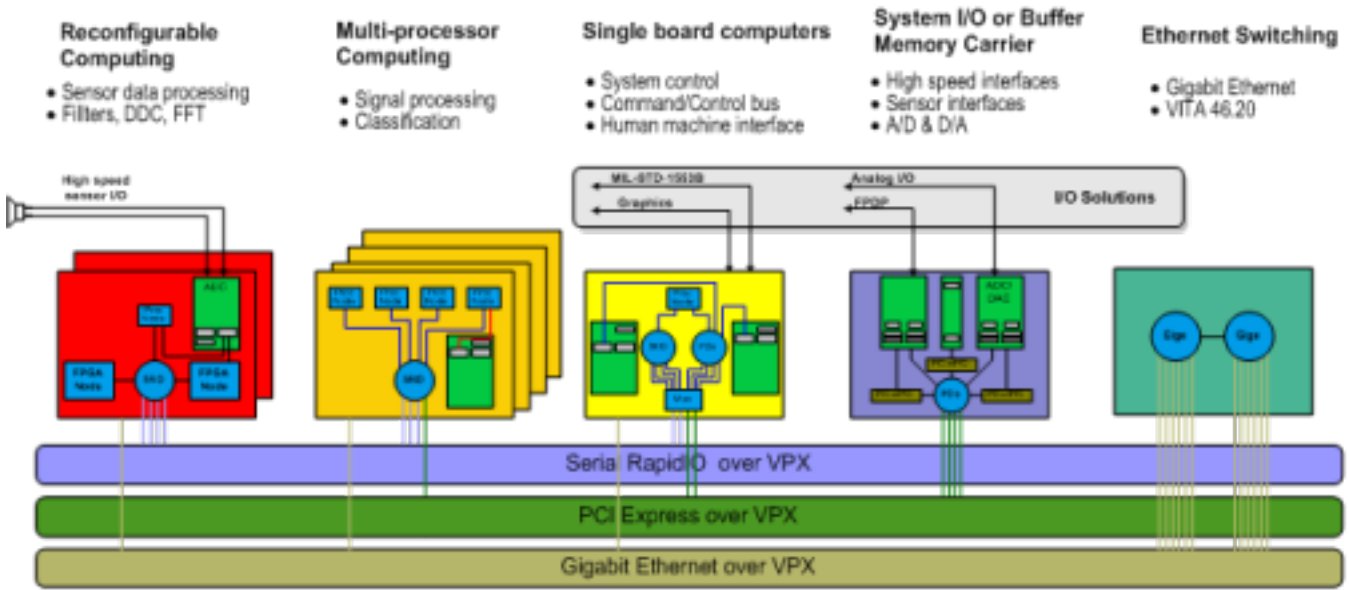


Figure 3: OpenVPX-based Multi-Processing System

Future trends in Multi-Processing Systems will continue to be dominated by the need to compute ever increasing quantities of data, faster interconnects for board-to-board communication, and the incorporation of these systems in an ever more wireless world. The future of computing in harsh environments is expected to include Optical Interconnects, higher speed backplanes to incorporate a 40 Gigabit Ethernet control plane, and the incorporation of special types of user I/O, such as RF.

These advances will help ensure that the Embedded Industry remains ahead of the curve in its ability to turn data into actionable intelligence. Because OpenVPX was established as a “living” specification, readily able to embrace new technologies, it is envisioned that these next-generation performance advances will be folded into the OpenVPX Specification as the technology is adopted by the VITA Standards Organization (VSO).

Conclusion

There have been several pivotal technology trends in the Embedded Industry, beginning in 1980 with VME and leading up to the present with the recent ratification of the ANSI/VITA 65 - 2010 OpenVPX System Specification. OpenVPX was envisioned as a framework to enable the aggregation of various disparate computational entities into larger systems with compliance metrics to adjudicate disagreements within the embedded community. Hindsight is always 20/20 and the future is always obscure. For those reasons, the OpenVPX Specification was envisioned as a living document that will readily conform to the embedded community’s future Multi-Processing System requirements.

Source URL (retrieved on 02/01/2015 - 8:52am):

<http://www.ecnmag.com/articles/2010/08/embedded-computing-enhanced-openvp-multi-processing-systems>