2011 NDIA GROUND VEHICLE SYSTEMS ENGINEERING AND TECHNOLOGY SYMPOSIUM
VEHICLE ELECTRONICS AND ARCHITECTURE (VEA) MINI-SYMPOSIUM
AUGUST 9-11 DEARBORN, MICHIGAN

AFFORDABLE, HIGHLY AVAILABLE, HIGH PERFORMANCE COMPUTING ENVIRONMENT FOR GROUND VEHICLES

Barry Press
Northrop Grumman Information Systems
Carson, CA

Jacob Sealander,
Curtiss-Wright Controls
Embedded Computing
San Diego, CA

ABSTRACT

The U.S. Army has realized the need to change its procurement strategy to demand shorter equipment fielding times, delivering capabilities with clear military value that can be improved in low risk increments. This paper addresses the need for a C4ISR infrastructure based on open standards that can be sourced from multiple vendors and host the full spectrum of applications.

The Distributed Computing Environment (DCE) -- is a computing and network infrastructure scalable from handheld and dismounted form factors through redundant, distributed clusters supporting multiple secure enclaves, and qualified for the most stringent of ground combat environments. The DCE captures the benefits of commercial-off-the-shelf (COTS) capitalization surrounding the x86 architecture, and implements open standards including OpenVPX, POSIX, Gigabit Ethernet, the Object Management Group Data Distribution Service, and the Service Availability Forum.

Packaged COTS subsystems implement the hardware as rugged chassis comprised of industry standard backplanes with fully integrated computing and I/O modules. The result is rapid availability for high technology readiness level equipment that delivers field-proven technology that will last the lifetime of a combat vehicle, and has a clear technology upgrade path when required.

INTRODUCTION

U.S. Army computer hardware acquisition over long procurement cycles and subject to the unique requirements of each distinct C4ISR system has lead to multiple sets of dedicated computing and network hardware in which there is little commonality in systems. Generational differences between systems preclude common standards across system baselines, and ensure there can be no benefits from shared technology, be they portability, availability, or ruggedization.

This state of affairs need no longer be true. The advances in raw compute power brought about by the forward march of Moore's law enable a change in computer and network hardware procurement towards shorter equipment fielding times, delivering capabilities with clear military value that can be improved in low risk increments. This paper describes an infrastructure -- what we call the Distributed Computing Environment (DCE) -- based on open standards that can be sourced from multiple vendors and host the full spectrum of C4ISR applications. The DCE is scalable from handheld and dismounted form factors to redundant, distributed clusters supporting multiple secure enclaves and qualified for the most stringent of ground combat environments.

REQUIREMENTS

The requirements we set for the DCE reflect public, open standards to achieve interoperability with existing and future systems. We defined a single set of environmental specifications for all devices in the infrastructure by aggregating ones from a range of ground vehicles. Those environmental specifications become the core of what makes the DCE flexible, creating a set of thermal and mechanical expectations both vehicle designers and COTS implementers can rely on.

- Black-box specifications characterize parametric values outside the device chassis.

- Translations of those parametric values to the worst-case environment presented to COTS cards within the chassis defines both the requirements on those cards and the functionality required of the chassis itself.

Separate specifications define the remaining characteristics of individual devices, including functionality, size, weight, power, and EMC/EMI.

The devices we characterized implement a complete computing and network infrastructure, including physically small, distributable computers, network switches, and a router/firewall. We chose not to define specific displays or input devices, deferring those choices to individual system implementations. That approach allows system configurations ranging from compact, low-cost implementations through high performance clusters:

- A smart display is a chassis incorporating both a processor and the associated LCD panel, and so the standardized DCE environmental boundaries characterize both the perimeter of the display and the packaging for the internal electronics. The resulting configuration benefits from the environmental models defined for the standard interfaces, and retains the standard I/O interfaces and software stack.

- The same requirements set adapts to characterize medium to high performance clusters, composing them from multiple chassis each enclosing a number of processor cards, storage, and specialized electronics required by the mission configuration.

We had some difficulty characterizing performance requirements because required processor performance for these vehicles may be characterized in terms of the now-obsolete

SPECint 2000 metric, not the current SPEC CINT2006 suite. Using the analysis from http://mrob.com/pub/comp/benchmarks/spec.html (and recognizing the dangers in simple ratio conversion), we converted requirements of 19,000 to 23,000 SPECint 2000 units to the range of approximately 165 to 200 SPEC CINT2006 2006 units. Adding a margin of 50 percent, we came to a final requirement of minimum processor performance in the range of 250 to 300 SPEC CINT2006 2006 units. The SPEC throughput benchmarks — SPECint_rate2006 and SPECfp_rate2006 — are likely better characterizations of the quad-core processors we implemented in our systems, but do not have equivalents in the SPECint 2000 suite.

The memory size and network bandwidth requirements were much simpler to define, largely because the existing baseline technology has capability far in excess of that needed anywhere in the range of configurations we surveyed. We ended up stating a requirement of 16GB RAM per processor and 1Gbps for the standard network port.

Similarly, we used the availability of standards for virtualization and I/O ports to set the requirements we'd target, requiring hardware support for virtualization, along with SATA, USB, and RS-232/422 ports.

We avoided setting other requirements in too much detail. Recognizing, for example, that specifying graphics support at the hardware level would lock in specific vendors, we required only that a range of support be available ranging from integrated chipset graphics with DVI outputs through an XMC site to host a multicore GPU.

ARCHITECTURE

We defined our infrastructure in layers, based on a few guiding principles. Key among them is the assumption that the vehicle operates as an integrated network, possibly separated into two security enclaves. The computing environment expects the network environment to provide high performance, reliable connectivity, and expects as many devices as possible to be connected through that network.

That assumption, combined with the significant improvements in computing power possible in very small packages, makes it possible to distribute network, computing, storage, communications, and other functional blocks throughout the vehicle, and makes practical both physical redundancy and enclave separation by hardware.

We layered on top of that basic architecture the assumption that we would have a reasonable fraction of the computing resources to dedicate to improving overall system availability, taking both hardware and software failures into account. Spare processor resources make it possible to accommodate the overhead from virtual machine hypervisors, and from middleware supporting high availability. Systems implementing the architecture have high enough inherent reliability and availability that we did not need to set explicit specifications.

Network Architecture

Our baseline choice for the vehicle network is 1 Gbps copper Ethernet. Although we expect fiber optic Ethernet to become cost effective in the future, we found that so many devices still expect a copper interface, and the cost of copper-to-fiber media converters is so relatively high, that fiber is still impractical without some other factor such as EMI/EMC mitigating the added cost. We did find, however, that a high port count switch incorporating a dual-homed 10 Gbps fiber optic backbone and 1 Gbps general service ports made for a compact, relatively inexpensive, highly available architecture capable of carrying both C2 and video traffic. Dual homing is also readily implemented between the computers and the switches, ensuring availability of the network connection.

Computing Architecture

Choosing the x86 architecture and Linux operating system is essentially a given for ground vehicles, and is required by specification in some procurements. In our baseline configuration we added to those open standards a requirement for hardware virtualization support and 16GB of memory per processor, helping to ensure a bare-metal hypervisor could run virtual machines with the least possible overhead. Handheld and other very small form factor packages not required to implement the full C4ISR software suite could be relieved of the memory requirement if necessary, although the process isolation made possible by virtual machine technology is still highly desirable. We layered high availability middleware on top of Linux in all but the smallest configurations to permit processes to be restarted after failures, and to enable process migration from a failed processor onto one still in operation. This computing architecture, when combined with local wireless networks, supports scalability from Android or other Linux-based handhelds to standard single board computers in various package formats through clusters of multiprocessor systems able to support real-time image and video processing.

STANDARDS

Although we built the specifications for the DCE on a large number of supporting documents, including both commercial and military standards, only a few choices were necessary to shape the resulting infrastructure in a way that strongly supports the goals of open COTS-based systems, multiple source availability, network ready, and high TRL supporting short fielding timelines. The key hardware choices were these:

- X86 instruction set including hardware support for virtualization
- OpenVPX in a 3U form factor
- Networking via 1Gbps and 10Gbps Ethernet
- Device interfaces via Ethernet (preferred) or either USB or RS-232 (deprecated)

Requirements for shock and vibration, thermal limits, sand and dust, humidity, and the like were readily derived from the ground environment and known vehicle characteristics. More interesting was that the choice of the 3U OpenVPX form factor coincided well with those requirements, meeting them all with no exceptions to the standard.

We imposed only a small number of software standards, focusing on those necessary to create an open, highly available, networked environment. In addition to Linux, we required the Object Management Group Data Distribution Service [1, 2] and both the Hardware Platform Interface (HPI) and the Application Interface Specification (AIS) developed by the Service Availability Forum [3].

The publish/subscribe software architecture is well adapted for C4SIR systems, making it possible to not only distribute data efficiently, but also to readily decouple data producers from consumers and therefore simplify adding new applications to the system. It is readily added to systems via OpenSplice DDS. The Service Availability Forum middleware is similarly available as open source using the OpenSAF code base.

DESIGN

Packaged COTS (PCOTS) represents the next logical step in Commercial-Off-The-Shelf (COTS) solutions for defense systems. PCOTS provides the benefits inherent to COTS products (e.g. affordability, versatility, latest technology), while meeting the unique requirements needed for defense systems (e.g. reliability, configuration control, environmental protection, life cycle support). PCOTS based subsystems provide a means to satisfy strict program affordability goals and accelerated development timelines while meeting the technical requirements with highly reliable application ready solutions.

Provided as a family of general purpose computing systems, PCOTS systems follow three fundamental strategies to solve the system integrator's most challenging tasks:

- System Integration and Interoperability: Through years of experience in developing systems for various Defense Systems, a set of standardized input and output (IO) interfaces are provided that allow for ease of integration of PCOTS based subsystems into an overall system design. By adhering to these standard interfaces, a system integrator's risks of interoperability issues is reduced since these standardized interfaces have been proven in the field, maintained by industry standards organizations and are available from multiple suppliers.

- Module Integration and Interoperability: PCOTS systems are comprised of open standards based COTS components. The modules we used represent the standard in rugged deployed COTS modules available from multiple suppliers. PCOTS systems adhere to industry standards within the subsystem to ensure both physical and logical interoperability between the electronic modules and the system enclosure. This approach not only reduces subsystem integration risk, but enables common modules to be used across multiple subsystems. This not only provides for reduced logistics support (i.e. fewer part numbers), using modern open standards such as OpenVPX, ensures interoperability with future technology thereby avoiding costly obsolescence issues.

- Standardized Environmental Model: PCOTS systems are developed to withstand the most common environmental conditions required for Defense Systems. By isolating the COTS modules from the external environment and providing a well define environmental interface to the COTS Modules, interchangeability between various COTS products is achieved. As new desired functionality becomes available, the system integrator can easily include that functionality as long as the COTS module selected meets the standardized interfaces provide by the PCOTS system. Providing a scalable solution to meet different application requirements, COTS modules are installed in rugged

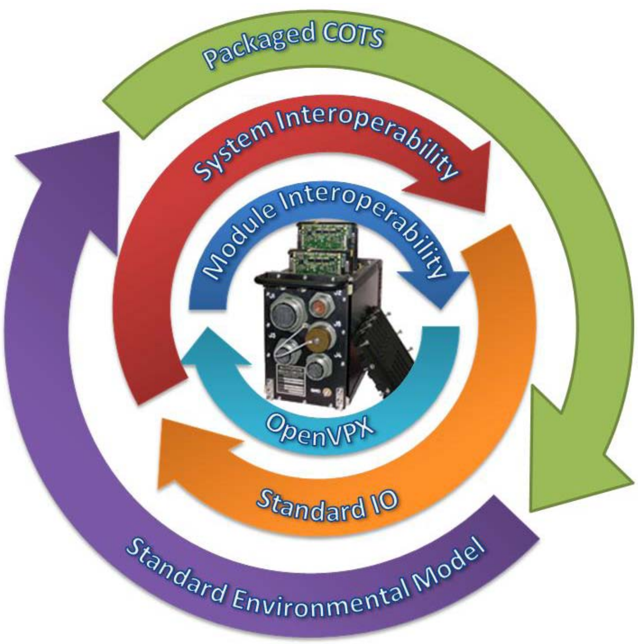

enclosures, which range in size from one to five payload slots. The PCOTS system selection is a function of performance requirements; the number of processing nodes, routing/switching capability, and interface functionality that a system requires.

Designed to Industry Standards

The PCOTS subsystems are designed to adhere to industry standard form factors, such as OpenVPX, allowing for simple integration of any standards based COTS modules. This Packaged COTS (PCOTS) approach of using fully developed and qualified components benefits customers by reducing development costs and deployment time, and reducing risk in general. By implementing open systems and standard form factors, PCOTS systems are designed with a component commonality approach that allows chassis to accommodate COTS payload modules such as Single Board Computers (SBC), secure routers and power supply modules. These COTS components are manufactured in high volumes which decreases unit costs due to economies of scale. Standard COTS hardware is also readily available which facilitates modular sparing and repairing for logistics management.

Network Ready

PCOTS systems are also designed to support defense systems network ready requirements. By installing a rugged network switch module within the subsystem, the system integrator now has access to an ample number of 1Gbps and 10Gbps Ethernet ports with dual homing in a small, inexpensive form factor. This approach results in a completely integrated switched-fabric based system architecture providing for a unified data, control and system management solution.

Ready to deploy, these rugged network ready PCOTS subsystems allow system integrators to reduce development costs and Time-To-Integration (TTI) by deploying the network switch within their system, connecting other subsystems to the appropriate Ethernet ports, and focusing on the optimal partitioning and segmentation of their applications within the resulting networked computing environment.

System I/O and Connectivity

To provide for ease of integration, all required industry standard I/O for mission computing, including Gigabit Ethernet, Serial Ports, SATA, Graphics, MIL-STD-1553, ARINC 429, and more are available within PCOTS systems. The I/O from the COTS modules installed within the system is routed through the system backplane to the front panel and made available externally through standard military connectors. System connectivity is provided through standardized connector arrangements or can be custom-designed to meet specific program requirements such as interface pairing or specific connector arrangement requirements. Within the system, the integrated modules communicate with each other over standards based backplane interfaces such as PCI Express or Gigabit Ethernet.

Advanced Thermal Management

Critical to providing highly reliable COTS based subsystems is the ability to manage the thermal environment of the COTS components included in the system. This is the most critical aspect of the Packaged COTS paradigm. Care must be taken to ensure the COTS modules are kept within their specified operating temperature ranges. To accomplish this, PCOTS Rugged chassis are designed to minimize the thermal path from the module to the cooling source, whether it be a vehicle cold plate or into the atmosphere.

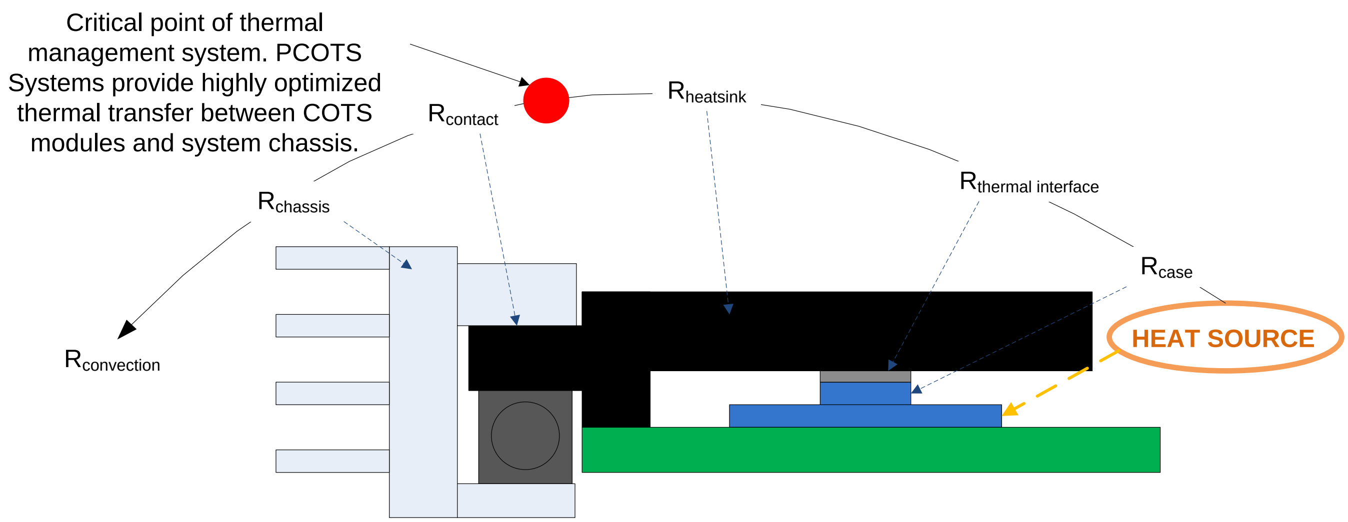

Modules installed within these environmentally sealed chassis are arranged in a manner that is most efficient, based on the cooling approach dictated by the program requirements, whether it be natural convection, baseplate, etc.

Cooling is extremely important on ground mobile platforms. PCOTS systems meet these demanding requirements with innovative techniques that are much more efficient than standard approaches used on legacy chassis. For example, PCOTS systems can cool 180W at 55°C without the need of a fan or cold plate. Further, all of these chassis are environmentally qualified and designed for operation in harsh environments typical of ground mobile applications.

TECHNOLOGY READINESS AND OBSOLESCENCE MITIGATION

There is a balance between technical maturity and product obsolescence that must be achieved for program success. Incorporation of technology that is not mature adds technical risk to the program potentially increasing development time and cost. Incorporation of technology that is too mature adds obsolescence risk to the program potentially increasing life cycle costs. PCOTS systems mitigate both of these issues by providing the latest technology from proven COTS standard modules, while providing an open system architecture based on industry standards. System functionality upgrades by means of technology insertion is another design feature of The Modular Open System Approach (MOSA) employed for PCOTS. Products are designed to a technology roadmap that follows the technology improvements of processors and components while maintaining the same form factor, software and pin outs (when possible). While some integration work may be required, this process supports future technology refresh cycles while causing minimum current-day impact on system integrators. Upgraded modules from the same product roadmap can be installed in the systems without causing any changes for system integrators – same chassis, same pin out, and same application software. Embedded computing modules should be specified to follow an obsolescence management program that guarantees product availability from five to seven plus years. Therefore, PCOTS based systems provide both technical currency and longevity of supply supporting defense program life cycle requirements.

QUALIFICATION

Environmental Qualification

System level qualification testing is essential when dealing with harsh operating environments. All components within the system, as well as the system in its entirety, must prove to be reliable in harsh operating conditions before deployment. For PCOTS systems the qualification testing risk is reduced because all the COTS modules undergo environmental testing individually before they are integrated into PCOTS systems. Additionally, following full functional testing, newly manufactured COTS modules undergo an Environmental Stress Screen (ESS) running functional test software. Hot and cold starts and variations of supply voltage are applied to weed out any early component failures or manufacturing defects.

Once the PCOTS system is fully integrated, a Test and Evaluation Program is conducted to verify that the equipment meets the technical and environmental requirements. Design verification, environmental and electrical performance, and EMI/EMC testing are performed in accordance with the conditions of the fielded application. MIL-STD-810 is the starting point for these tests, covering temperature extremes, rain, humidity, salt fog, dust, shock, vibration, and more. EMI and EMC testing uses MIL-STD-461.

To ensure program specific functionality is verified, a functional test procedure is created based on program requirements that describe how the equipment is being exercised, what inputs are being provided to the unit, and what the expected responses are to the various inputs. These procedures identify how each of the significant modules within the equipment is tested. While conducting the environmental exposure testing, test engineers monitor Continuous Built-in-Test (CBIT) data from the integrated modules for internal circuitry health, continuous read/write functionality of disk drives and more.

Acceptance Testing

An Acceptance Test Procedure (ATP) is also created in accordance with the program's functional requirements. This testing covers all the interfaces necessary for integration and applications. Software stress tests load the system interfaces to accurately simulate fielded applications. All the system interfaces are tested both individually and simultaneously to ensure robust operation.

SUMMARY

The Distributed Computing Environment (DCE) introduced in this paper is an affordable, highly available, high performance computing environment for ground vehicles. DCE realizes a common computing architecture supporting multiple C4ISR system requirements by adopting industry accepted open standards and capitalizing on the technological advances in processing and networking. The open systems architecture central to DCE replaces dedicated and bolt-on systems.

System integrators can greatly reduce both technical and obsolescence risk by implementing the DCE with scalable Packaged COTS systems. The PCOTS paradigm of standardizing both internal module characteristics and external system interfaces lets system integrators rapidly develop, integrate, and field system solutions based on existing, mature COTS technology with a well defined path for technology insertion. Standardized environmental qualification methods for PCOTS systems ensure products operate in the harshest conditions.

Program managers and system integrators will benefit most from our approach by adopting DCE and PCOTS solutions early in the development cycle. Significant program cost and schedule savings are a direct result of tracing functional and interface requirements to the open, modular DCE standards, because the existing catalog of tested PCOTS components greatly reduces the non-recurring engineering necessary versus the one-off, specification-based development approach.

REFERENCES

[1] A. Corsaro, "The DDS Tutorial Part I", http://forums.opensplice.org/index.php?/topic/39-the-dds-tutorial-part-i/, 27 January 2010.

[2] A. Corsaro, "The DDS Tutorial Part II", http://forums.opensplice.org/index.php?/topic/38-the-dds-tutorial-part-ii/, 11 February 2010.

[3] Service Availability™ Forum, "Service Availability Interface", SAI-Overview-B.05.03, February 19, 2010,