



Migrating Mil/Aero Programs to NXP Freescale Multi-core

Author:

John Sterns, Field Applications Engineer

August 2016

Table of Contents

1. Multi-core platforms in modern military applications	2
2. Rehosting a single application to a single core.....	2
3. Mapping multiple applications to multiple cores.....	3
4. Migrating uniprocessor code to parallel multi-core clusters	6
5. It's certifiable.....	7
6. Conclusion	8
7. References.....	9

1. Multi-core platforms in modern military applications

Embedded multi-core platforms and related software support tools are now at a maturity level suitable for use on mil/aero programs. Virtualization technologies have moved from the server farms and clouds to embedded. This means multi-core is ready for adoption and deployment in harsh rugged environments, wherever there is a need for end-of-life mitigation, footprint reduction or improved capabilities from additional processing power.

The rub, of course, is how to get all the goodness of a proven end user application to the new platform complete with the capabilities and critical timing scenarios intact. Introducing parallelism without a plan might change “ready-aim-fire” to “ready-fire-aim”. Fortunately, there are methods and techniques available to mitigate these risks.

Depending on program needs and risk sensitivity, three main porting scenarios are possible:

- Rehosting a single application to a single core
- Mapping multiple applications to multiple cores
- Migrating uniprocessor code to parallel multi-core clusters

2. Rehosting a single application to a single core

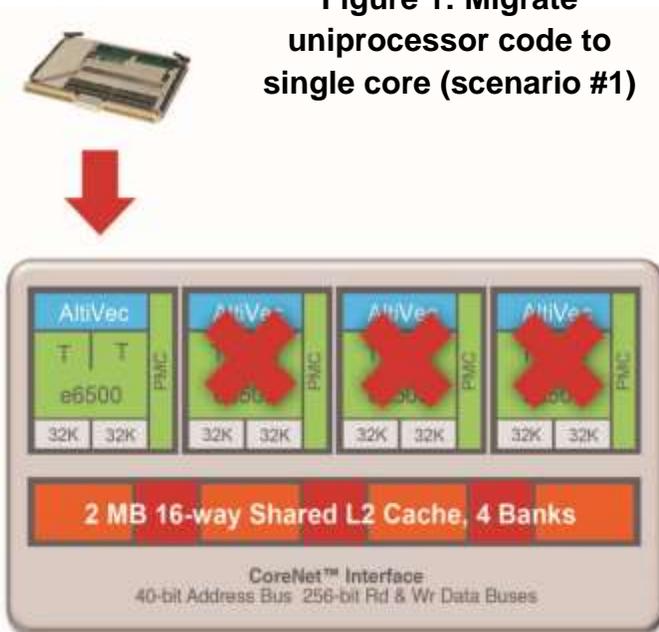
Rehosting a single application to single core is the least efficient use of multi-core capabilities, but also the least risky, because you remove the multi-core parallelism from the equation. In this case, you're essentially repeating the “tried and true” tech refresh path that came with Moore's Law – you get faster execution speed, and mitigate end of life considerations by moving to a processor family that is earlier in its production lifecycle.

For example, the e6500 core featured in the current T-Series family from Freescale provides 16 GFLOPS of AltiVec vector processing speeds, compared to a peak of 10 GFLOPS on the PPC 7448 SIMD unit from ten years ago.

From a software perspective, you need to rebuild the software under a current OS and toolchain platform. One can't expect binary code compatibility with a VxWorks 5.5 or Red Hat 5.2 application. When loading the application image, load it to the single, primary core and leave the other cores disabled. (Figure 1)

Since U-Boot, and its derivatives like AiBoot, operate in single core mode, no explicit code to disable unwanted cores is needed. The application load-from-flash or load-from-network boot methods are done as before.

Figure 1: Migrate uniprocessor code to single core (scenario #1)



If desired, the application boot sequence can be augmented to utilize NXP Freescale’s Trust Architecture 2.0 security features introduced with its multi-core processors. The P- and T-Series of multi-core processors support a secure boot “chain of trust” load method that can be adopted to improve the security and data integrity of systems for foreign deployment or foreign military sales.

Military users can benefit from new network security features added for the Internet of Things (IoT). Even in single core mode, there is still plenty of possibilities for improved features and value to be added for the end user.

Hardware compatibility with existing system enclosures can be attractive since it limits upgrade costs. COTS vendors, as before, are supporting the tech refresh of single board computers (SBCs) with pin compatible replacements.

Aitech, for example, offers the C111 T4 Series VME single board computer as a pin compatible replacement for its

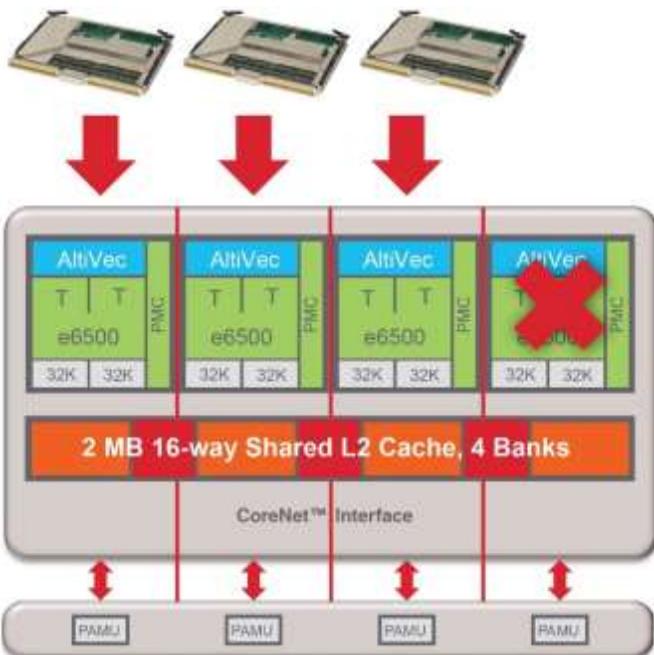
C102 PPC dual node 7448. A single T4080 replaces the two distinct CPUs and several supporting chips sets.

Table 1–Typical T-Series SBC Power Dissipation

SBC Model	Cores and Hardware Threads	Typical Power Dissipation
C102 7448 @ 1 GH	2 (two discrete CPU chips)	42 watts
C111 T4080 @ 1.5 GHz	4 cores/8 threads	35 watts
C111 T4160 @ 1.6 GHz	8 cores/16 threads	42 watts
C111 T4260 @ 1.5 GHz	12 cores/24 threads	50 watts

Even nicer, though, is accomplishing performance improvement in the same thermal envelope. The C102’s dissipation, with two discrete CPU chips for two cores, is 42 watts for a typical application. A C111 T4160 model at 1.5 GHz with eight cores will typically dissipate 42 watts, as well. (Table 1)

3. Mapping multiple applications to multiple cores



More efficient use of multi-core processing resources can be achieved by mapping multiple single core applications 1:1 with multiple cores. Each core becomes a “virtual single board computer” with associated memory and I/O. (Figure 2)

The MMU is set up to provide the address space protections that prevent other cores and applications from inadvertent or malevolent corruption of application code and data.

In addition, the NXP Freescale P- and T-Series processor families have PAMUs (Peripheral Access MMUs) to prevent I/O devices from “invading” the virtual card’s address space with remote DMA from the network or backplane fabric.

Figure 2: Mapping one or more applications to multiple single cores (scenario #2)

In legacy applications it's typical for discrete SBCs to communicate via a parallel bus backplane such as VME or CompactPCI, or via Ethernet connections. (Figure 3)

When migrated to discrete P- or T-Series cores, the interconnect fabric between the cores—with additional software tools—provides the same communication mechanisms.

Wind River Systems

Multi-Core Interprocessor Communication (MIPC) framework provides for direct analogs between legacy wired interconnects and on-chip core fabrics. It includes:

- MND simulated Ethernet-over-core fabric interface
- MFS file systems-over-core fabric interface
- MSD pseudo-serial ports
- Socket-based message channels
- Virtual bus-over-core fabric interface
- VxMP shared memory objects

Single Board Computers

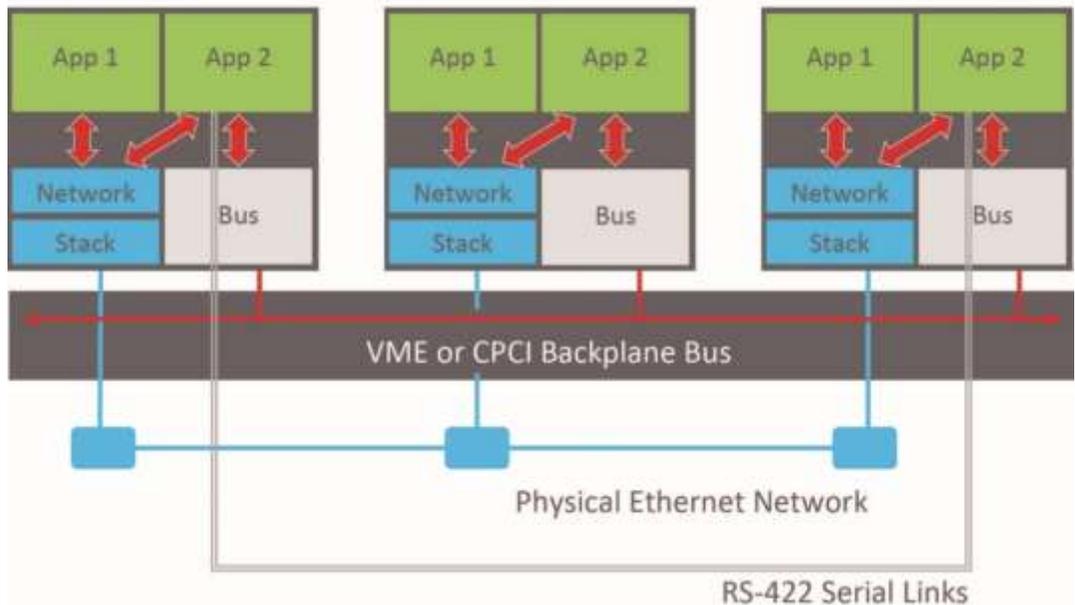
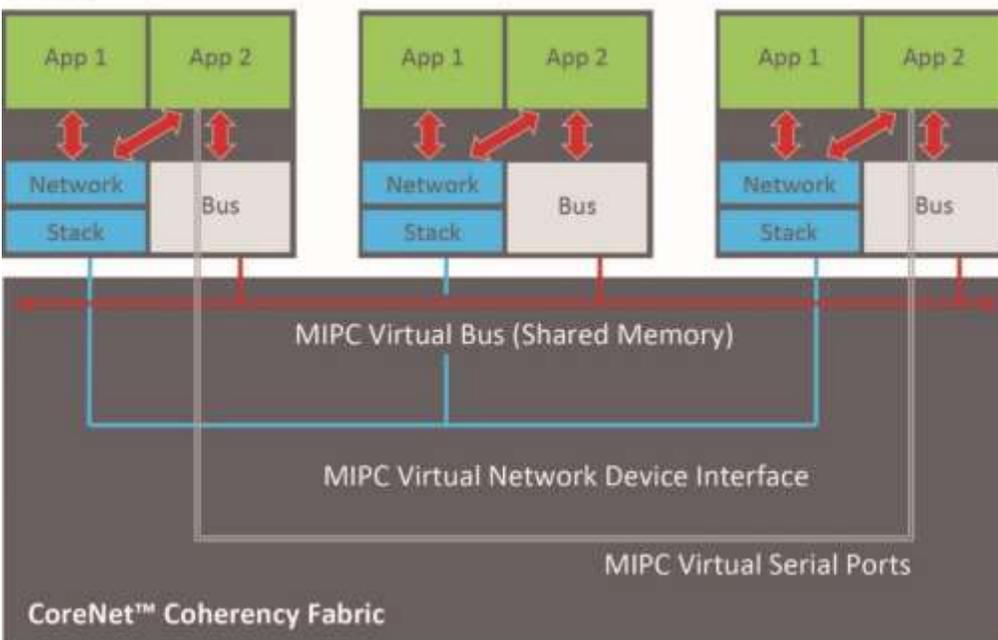


Figure 3: Common Interconnect Paradigms for SBCs

Cores



To complete the virtual SBC (and eliminate the corresponding physical cards), the boot and initialization sequence must, of course, be modified to load and enable the additional cores. (Figure 4)

As before, the bootloader loads software into SDRAM for the one core that is enabled on power up.

That software is responsible for enabling the other cores, configuring the MMUs, and loading their "guest OS" applications.

Figure 4: Wind River Tools Map Common IPC Paradigms to the Core Fabrics

Two management strategies are readily available for this – a hypervisor running at a privileged hardware level, and a manager core running a Linux Kernel Virtual Machine (KVM) or RTOS, such as VxWorks.

The hypervisor strategy places a software layer between the hardware and OS to allocate and control global hardware resources utilized by each virtual card. The hypervisor sets up the partitions, MMUs and PAMUs for each virtual card, loads the guest OS, and “kick starts” application execution on the assigned core. Data path portals map direct I/O between virtual card partitions and the physical device hardware.

NXP Freescale P- and T-Series processors have hardware support for hypervisors. The Supervisor-User privilege levels have been expanded to include a third privilege level-Hypervisor. Manipulation of the virtual card's MMUs and PAMUs can be restricted to this privilege level.

NXP Freescale Embedded Hypervisor

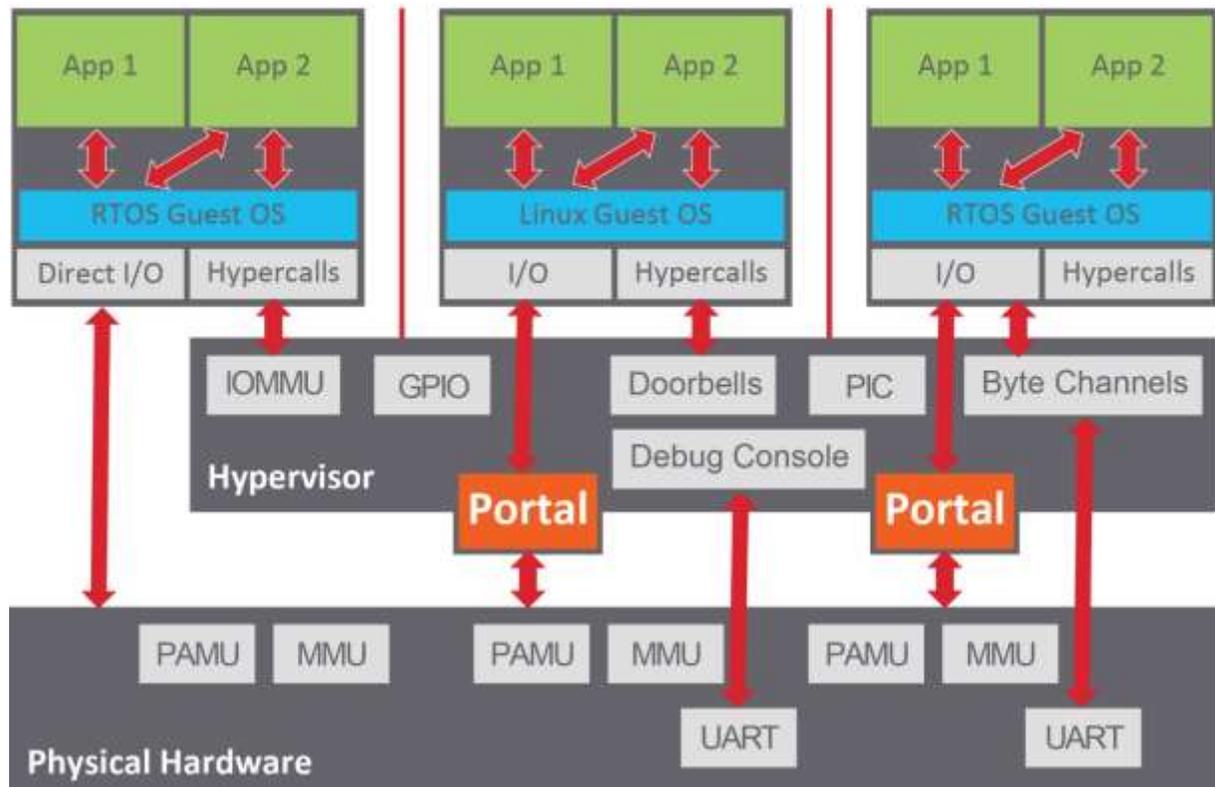


Figure 5: Hypervisor software can eliminate routing interrupt latency using PAMUs

This ensures hardware enforcement of the software virtual partitions. Specific interrupts can be routed to a core using PAMU settings, eliminating the latency of routing all interrupts through the hypervisor software. (Figure 5) NXP, Green Hills and Wind River VxWorks 653 3.0 hypervisors are available for hosting Power Architecture applications.

The Linux KVM manager strategy utilizes the OS as “core manager” in place of the hypervisor. Virtual machines are set up to host guest OS and applications. A kernel module provides the facilities for controlling cores and setting up virtual cards in isolated partitions. QEMU allows emulation of other architectures.

This open source approach to virtualization was originally developed for x86 and has been in the Linux mainline since 2.6.20. NXP Freescale has made it available for Power Architecture with their QorIQ Linux Software Development Kit (SDK). Because all partitions are ultimately under the control of the Linux OS and its KVM module, this approach is really only suitable for non-realtime and soft realtime applications.

The VxWorks RTOS approach is similar in that a primary core (or core cluster) is responsible for loading and starting the additional cores and applications. The Wind River `wrload` tool provides this capability. Green Hills Integrity-178 tuMP has its core loading tools as well.

Once the main manager application is up and running, additional cores are enabled, loaded with software and brought on-line as asymmetric multi-processing (AMP) elements. These elements then communicate through the RTOS multi-core inter-processor communications facilities.

4. Migrating uniprocessor code to parallel multi-core clusters

The processing resources available to legacy applications can be beefed up by making more than one core available to run the application. The power of NXP Freescale’s multi-core and hardware parallelism is demonstrated by Green Hills Software obtaining the best EEBMC CoreMark® benchmark results ever on the T4240 processor [GHS 2013].

For SIMD vector processing with the AltiVec execution unit, 172 GFLOPS is possible with the 12-core T4240, more than 7x the performance of prior generations of AltiVec [EETIMES2013]. This is largely because each physical e6500 core has its own AltiVec SIMD unit.

In order to realize this performance benefit, concurrency and SMP migration issues must be addressed, since the pseudo-concurrency of multi-tasking software is being replaced with true hardware concurrency.

With the possible exception of Ada this requires the ability to manipulate and modify the source code.

Operating system calls and task/process synchronization mechanisms will have to be reviewed and updated with SMP equivalents.

Critical code threads can still be assigned a single core affinity as needed to maintain a required process sequence and to keep frequently accessed data in the most local cache possible – the L1 cache.

To set expectations for speed improvements, it is worth reviewing the problem of your application space. The returns for software parallelism have their limits, as demonstrated by Amdahl. Adding a few cores may actually have a bigger return on investment than adding 12 cores. Analyzing key sequential threads of code will inform you of the limits for your particular application. (Figure 6)

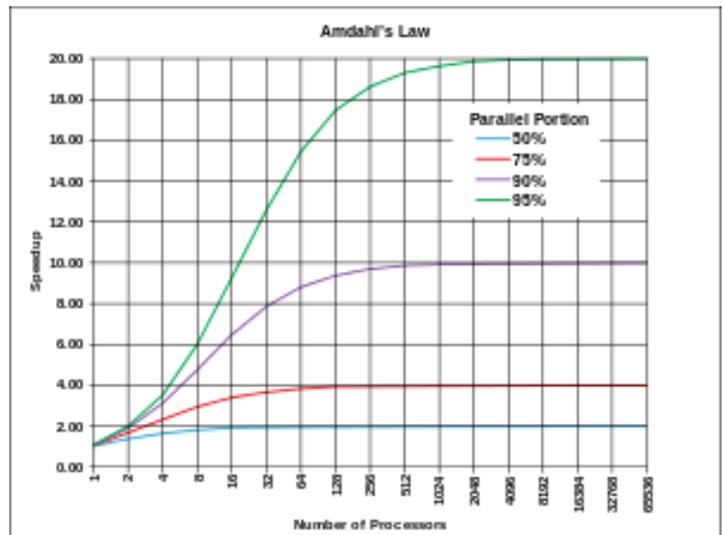


Figure 6: Example of Amdahl's Law Core Limits for Different Degrees of Parallelism
credit: Wikipedia

If, for example, it takes 1.5 seconds for antenna control, sweep and return, don't assume an improvement to 0.5 seconds with additional cores in the loop. Note that this analysis also helps in the assignment of “core affinity” to critical timing threads.

Having estimated your “bang for the buck” return, proceed with the application migration to Symmetric Multi-Processing (SMP). The general concerns of this process are well identified – programmers on the x86 platforms have been doing this for years under Windows and Linux OSs [GRAH2009]. For a real-time OS, such as VxWorks, the vendor typically has well documented guidance for migrating to hardware parallelism. (Table 2)

Elements of AMP (asymmetric) and SMP (symmetric) processing can be combined to provide both application consolidation and the performance benefits of additional SMP cores.

The Aitech AT4 VxWorks BSP, for example, will by default allocate four cores to an SMP “core cluster”. By leaving behind the 1:1 mapping of applications to cores and converting code to SMP, a 1:4 application to core mapping is possible, providing both *more* and *faster* cores to get the application’s job done.

Uniprocessor (pseudo-parallelism in software)	Multi-core (hardware parallelism)
taskLock()	spinLockTask()
taskSpawn, rtpSpawn	Make any implicit synchronization with spawned children explicit
Legacy PCI bus API	VxBus API
KernellIdle, vxIntStackBase, windPwrOffCpuState	No equivalent – remove

The number of core clusters depends on the T-Series processor selection – 1 for the T4080 (4 physical cores), 2 for the T4160, and 3 clusters for the T4240. Core clusters communicate in AMP mode; i.e. each cluster operates independently using the Wind River MIPC facilities to synchronize and communicate as needed.

There is no escaping, of course, the need for thorough integration testing with end equipment to verify critical process timing and reliability. As always, critical code sections should be migrated early and tested often as more code gets migrated to SMP.

Remember that initial timing and sequences that look good may be upset later by the addition of more tasks and processes. “Trust but verify” is, as always, the mantra when porting systems. Tool support for spotting and resolving issues is available. NXP Freescale and RTOS vendors, like Wind River Systems and Green Hills, have augmented their toolsuites for multi-core platforms, allowing multiple GDB CodeWarrior and Workbench core connections from a single integrated development environment (IDE).

5. It’s certifiable

DO-178C does not address multi-core, so additional planning is required to ensure acceptance. But make no mistake, RTOS companies and their clients are confident that, with good analysis and preparation, applications hosted on multi-core platforms can be safety certifiable.

Green Hills Software's INTEGRITY-178B Time-Variant Unified Multi Processing (tuMP) Multicore Operating System has been selected for Level A certification and FACE-compliant applications [GHS2013]. A basic tuMP BSP with minimal interfaces has been ported to the Aitech C912 and is available for prototyping.

DAL A certification artifacts are being developed by Wind River Systems for the QorIQ T2080 with 4 physical e6500 cores and no virtual cores (dual thread operation is disabled) [JONES2015]. Aitech will be providing a VxWorks 653 3.0 BSP without artifacts for its T4080 variants of the VME C111 and VPX C112 products. VxWorks 7 Safety Profile offers another path using a Real Time Process (RTP) approach instead of a hypervisor.

RTOS vendors like Wind River Systems and Green Hills Software offer multi-core certification services. Non-artifact BSPs and platforms facilitate the development of DO-178B/C compliant applications, while deferring certification costs to the last pre-production tech refresh cycle.

6. Conclusion

Migrating uniprocessor code to parallel multi-core clusters may seem intimidating, but the tools and techniques for doing so are well established and mature. There is a fair amount of control over the degree of parallelism to adopt.

Disabling cores removes the parallelism and allows for the single core migration to current hardware under familiar terms. Hypervisors and OS "core managers" are available and proven for hosting multiple discrete applications on multiple cores in asymmetric processing mode.

Symmetric processing may require more source code modifications, but the performance pay-off can be substantial, either in the execution speeds of existing applications or with the addition of new features and applications. Software and hardware vendors alike have platforms and tools available for making the transition to multi-core both manageable and successful.



Embedded Computing *without* Compromise

7. References

- [GRAH2009] Bill Graham, QNX Software Systems, "Migrating legacy applications to multicore: Not as scary as it sounds", Military Embedded Systems, March 20, 2009
- [EETIMES2013] Julien Happich, "Altivec engine advances code performances with 172 GFLOPs of vector processing", EETimes, June 17, 2013
- [FREE2008] Freescale Semiconductor White Paper, "Freescale's Embedded Hypervisor for QorIQ P4 Series Communications Platform", October 2008
- [FREE2010] Jacques Landry, "The Freescale Embedded Hypervisor", http://cache.freescale.com/files/training/doc/dwf/AMF_ENT_T0616.pdf?fsrch=1 (November 2010)
- [GHS2013] Green Hills Software, "Freescale QorIQ T4240 communications processor leverages Green Hills optimizing compiler to top its own CoreMark® Record", http://www.ghs.com/news/pdfs/Freescale_T4240_CoreMark_050613.pdf (May 6, 2013)
- [GHS2013] Green Hills Software Press Release, "Green Hills Software INTEGRITY-178B tuMP Multicore Operating System Selected by Esterline CMC Electronics for DO-178B Level A Certification", http://www.ghs.com/news/20130402_tump_cmc.html (April 2, 2013)
- [GHS2013] Green Hills Software Press Release, "Green Hills Software INTEGRITY-178B tuMP Multicore Operating System Selected by Northrop Grumman for the U.S. Marine Corps H-1 Mission Computer", <http://finance.yahoo.com/news/green-hills-software-integrity-178b-183113325.html; ylt=A86.JyFB9UxVGhwADggnnlIQ; ylu=X3oDMTByNzhwY2hkBHNIYwNzcgRwb3MDMgRjb2xvA2dxMQR2dGlkAw--> (February 13, 2013)
- [HUYNH2010] Khoa Huynh and Stefan Hajnoczi, "KVM/QEMU Storage Stack Performance Discussion" (PDF) Linux Plumbers Conference, 2010
- [JONES2015] Kevin Jonestrask, "Avoid the Pitfalls in Designing a Multi-Core Safety-Critical Avionics System", Wind River Systems and OpenSystems Media webinar, May 7, 2015
- [WIKI] Wikipedia, Amdahl's Law, http://en.wikipedia.org/wiki/Amdahl%27s_law
- [WRS2011] Wind River Systems Product Note, "Wind River Hypervisor", June 2011
- [WRS2011] Wind River Systems, VxWorks Kernel Programmer's Guide 6.9, 2011

CoreMark and EEBMC are registered trademarks of the Embedded Microprocessor Benchmark Consortium.
Freescale, QorIQ, Altivec and CodeWarrior are registered trademarks of Freescale Semiconductor, Inc.
Green Hills, INTEGRITY and tuMP are registered trademarks of Green Hills Software.
Linux is the registered trademark of Linus Torvalds.
NXP is a registered trademark of NXP, B.V.
QEMU is a registered trademark of Fabrice Bellard.
VxWorks and Wind River are registered trademarks of Wind River Systems, Inc., an Intel company.