



Power.org™ White Paper

Nextgen Multicore/Accelerators SoC Platform defined by ITRS

Version 1.1 – February 2, 2010

Fawzi Behmann
Director of Marketing
Power.org

© Copyright 2010 Power.org. All rights reserved

The Power Architecture and Power.org word marks and the Power and Power.org logos and related marks are trademarks and service marks licensed by Power.org.

Abstract

Internet everywhere, IP everywhere and seamless mobility are key macro trends that will continue drive more complexity in the design of systems to meet the requirements of next generation systems and support new classes of applications and services. This in turn calls for software-based multi-core/Accelerator SoC platform. In addition, as technologies and structures push the limits of Moore's law and productivity, new semiconductor approaches to scaling and new functionality on- and off-chip will be required. This paper provides an overview of recent work done at ITRS – International Technology Roadmap for Semiconductor examining 15-year assessment of semiconductor technology requirements and defines MC/Accelerator SoC Platform.

Driven by such general macro trends such as Internet everywhere, IP everywhere and Seamless Mobility, in its 15-year assessment of semiconductor technology requirements, the International Technology Roadmap for Semiconductors projects that as technologies and structures push the limits of Moore's law and productivity, new semiconductor approaches to scaling and new functionality on- and off-chip will be required. **Figure 1, below** shows macro trends for cellular technology.

The semiconductor technologies that will be required can be broadly categorized into three categories: "Moore" " Geometric Scaling; "More Moore" " Equivalent Scaling and "More than Moore" functional diversification " all of which will have significant impact on the embedded networking space with new Systems on chip architectures that make extensive use of

- 1) Multi-Core (MC), □
- 2) Cache hierarchy, □
- 3) On-chip fabric, □
- 4) On demand Accelerator Engine (AE), and □
- 5) Connectivity,

all engineered to provide a scalable, software-based Multi-Core/Accelerator Engine SoC (MC/AE-SoC) platform that targets a wide range of applications from ultra low-end to high-end that preserve & extend the user experience through new services.

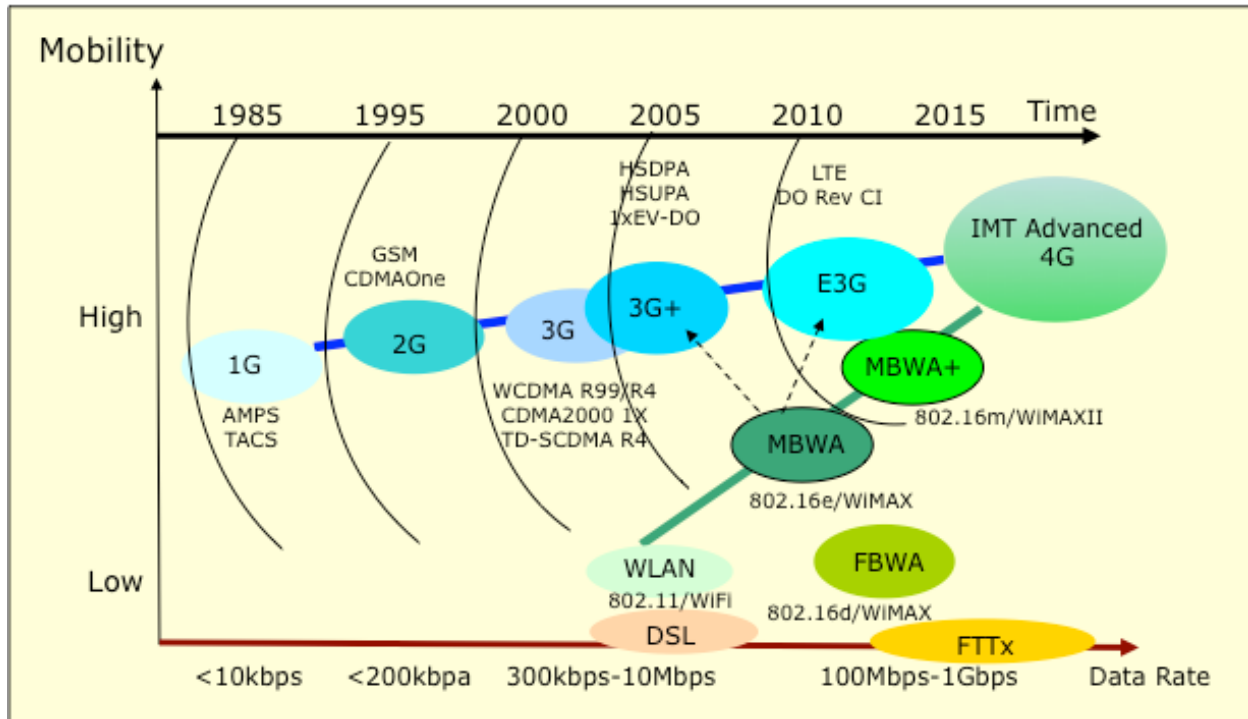


Figure 1. Macro trends for cellular technology.

The three "Moore's" □ While technologies and structures push the limits of Moore's Law and productivity, the ITRS initiated the concept of "More Than Moore," which first appeared in the 2005 ITRS publication, calls for the integration of Functionality that does not scale. It is mostly analog functionality, but also includes passives, high voltage, sensors, actuators and enablement.

During the ITRS summer conference, an overall definition was introduced grouping three aspects of "Moore" concept:

Moore: Geometric Scaling □ **More of Moore:** Equivalent Scaling □ **More Than Moore:** Functional Diversification

While "**Moore's Law**" is mostly focused on geometric scaling in continuing shrinking of horizontal and vertical physical feature sizes of the on-chip logic and memory in order to improve density (cost per function reduction) and performance (speed, Power) and

reliability values to the applications and end customers.

"More of Moore" is about equivalent scaling which occurs in conjunction with, and also enables, continued Geometrical Scaling plus non-geometrical process techniques that affect the electrical performance of the chip. The third element is "More Than Moore" is about functional diversification.

The **"More Than Moore"** refers to the incorporation into devices of functionalities that do not necessarily scale according to "Moore's Law," but provide additional value to the end customer in different ways.

The "More-Than-Moore" approach typically allows for the non-digital functionalities (e.g. RF communication, power control, passive components, sensors, actuators, 3rd party IP/enablers) that to migrate to system board-level/particular package level (SiP) or Chip-Level (SoC) potential solution.

There is increasing tendency to have more functions on a chip which are not scaling according to the same pattern [*as defined in Moore's Law*]. This is functional diversification rather than scaling, but it's part of the same business and same technology.

The combination "Moore's Law" and "More Than Moore" enables the creation of system-on-a-chip and system-in-a-package and, as such, adds value to systems rather than just integrating more of the same functions on a chip.

Functional diversification in SoC design□The ITU-R is currently studying user demand predictions in future systems such as the amount of traffic in the year 2010 onwards in calculating required spectrum bandwidth for the future development of IMT-2000 and IMT-Advanced.

The IMT-2000 (International Mobile Telecommunications) systems are 3rd generation mobile systems, which provide access to a wide range of telecommunication services, supported by the fixed telecommunication networks (e.g. PSTN/ISDN/IP), and to other services which are specific to mobile users. Among the key features of IMT-2000 are:

- 1) Capability for multimedia applications within a wide range of services and terminals□
- 2) High degree of commonality of design worldwide□

- 3) Compatibility of services within IMT-2000 and with the fixed networks
- 4) High quality
- 5) Worldwide roaming capability, and,
- 6) Small terminal suitable for worldwide use

The next 5-15 years will also mark trends towards:

- 1) Scalable networks that deliver high rich multimedia content at broadband speed anywhere and anytime and on any device;
- 2) Markets in which the consumer will play a major role in creating high rich multimedia content;
- 3) Emergence of advanced IP-based applications and services that drive high bandwidth scalable networks;
- 4) Complex multi-processing platforms equipped with multi-core/multi-threading and accelerators that support advanced applications and services;
- 5) Advancement in process technology from 65-45-32, 22 and sub 10nm technology
- 6) Scalable encryption and antivirus everywhere in the network;
- 7) Home networking will be a complex network converging data communications, entertainment;
- 8) Seamless mobility in the home, in the office/vertical market, on the road

In contrast with PC & Server Applications, and due to the fundamental difference between core speeds and memory/IO latencies, today's embedded processor architectures are unable to deliver meaningful performance for the connected computing scenarios outlined earlier.

Nearly every commercially available integrated general-purpose processor shipping in volume today is designed using a single-threaded architecture, which is performance and application limited by today's standards. As applications are becoming more and more network-centric, this legacy processor design approach fails to address the throughput requirements of today's converging compute and networking paradigm.

This evolving packet-oriented environment is characterized by high memory access

latencies, which are not effectively managed by conventional processor architectures. This weakness can severely impact processor performance and workload efficiency. When a memory access cannot be serviced immediately and no additional instructions are ready to be executed, conventional processors stall and waste valuable processing cycles.

PE SoC Consumer & MC/AE SoC Networking Architectures

Adding "More of Moore" to the mix (**Table 1, below**) provides a converged/ integrated heterogeneous platform (**Figure 2, below**) that makes possible the creation of a scalable, intelligent, compact value-add ecosystem. This PE SoC Platform implementation based on the scaling achieved by the use of the 3-Moore's is becoming an important paradigm moving forward.

SoC Networking Platform Building Blocks	"Moore"	"More Moore"	"More Than Moore"
	Geometric Scaling	Equivalent Scaling	Functional Diversification
Multi-core (MC)	X		
Cache Hierarchy	X		
On-chip Fabric		X	
Connectivity – NET Interface		X	
Connectivity – Peripheral Interface		X	
On-Demand Accelerators Engines (AE)		X	
Hybrid Simulation			X
Enablement/Ecosystem			X

Table 1. MC/AE SoC Networking Platform and Moore's classification

Early in 2005, ITRS introduced PE SoC Architecture Template, where a PE is a processor customized for a specific function targeting portable and wireless applications such as smart media-enabled phones or digital camera chips, but also high-performance computing and enterprise applications.

To complement this PE SoC architecture, a Multi-Core/Accelerator Engine SoC Architecture template is defined to address the networking embedded space as shown in **Figure 2**

below. As shown, the MC/AE SoC networking platform contains the necessary building blocks to:

- 1) Support Multi-Core (MC) for high processing performance within 30 watt power envelope,
- 2) Support an unprecedented Tri-level Cache Hierarchy, with back-side L2 caches, multiple L3 shared caches and multiple memory controllers
- 3) Support high-speed inter-connectivity
- 4) Introduces a scalable On-chip Fabric for concurrent, non-blocking, hardware-based 100% cache-coherent platform connectivity which scales to more than 32 cores and supports heterogeneous cores.
- 5) Eliminate shared bus contention and supports dramatically higher address issue bandwidth to "feed" multiple cores
- 6) Include an On-demand Acceleration Engine (AE) that offers performance advantages over pure core processing cycles, enables lower power implementations and reduces silicon area / cost
- 7) support a hybrid Simulation Environment combining cycle-accuracy and functional-accuracy that enable ease of software development , performance prediction and optimization
- 8) Network/System Enablement & Ecosystem looking into software partitioning and virtualization leveraging multi-core hardware architecture

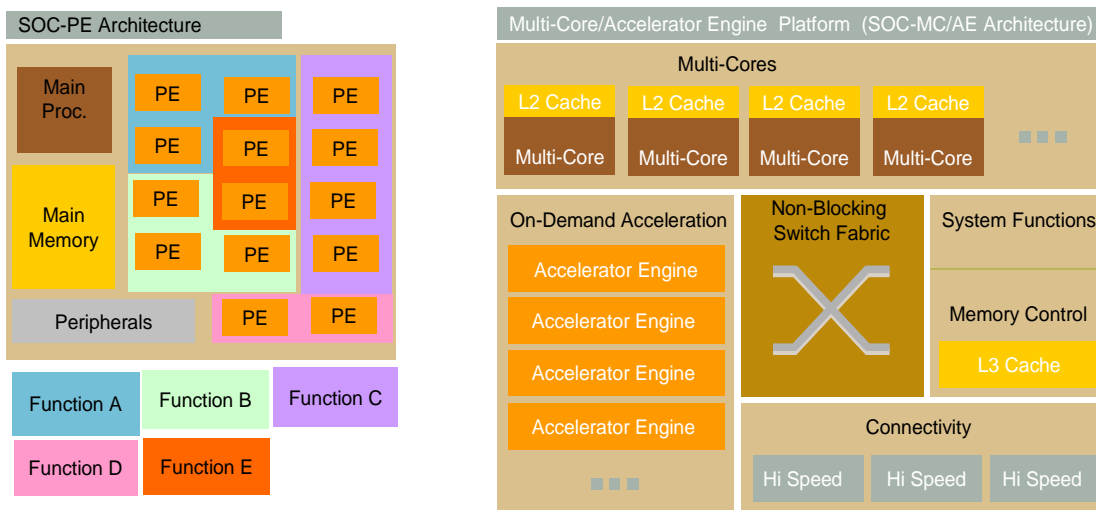


Figure 2. A Multi-Core/Accelerator Engine SoC Architecture template defined to address embedded networking apps

The MC/AE SoC network platform contains the necessary building blocks to provide a scalable, software-based solution and addresses a wide range of applications from ultra low-end to high-end that preserve & extend the user experience through new services.

Multi-Core (MC). The Multi-Core's frequency in a wide multi-core product will be targeted over one GHz. This platform targets the highest instruction-per-cycle (IPC) and highest frequency for a given watt per area.

The MC's are also designed to offload repetitive and computing intensive operations to high-performance acceleration blocks, increasing the number of processing cycles for higher throughput or new services and applications.

Each MC core in the platform will have its own L2 backside cache. Backside cache is connected to the CPU through a direct channel, enabling extremely high application performance.

It allows the cache to match the full speed of the CPU, resulting in latency improvements well over 50 percent of "shared bus/shared-cache" architectures. L2 backside cache also enables tuning the contents of the cache between instruction and data, according to different application needs, easing partitioning and improving performance by drastically reducing CPU stalls.

In addition, the L2 backside cache reduces traffic on the on-chip fabric and main memory, which reduces latencies and improves bandwidth for other users of the fabric and system memory.

Multithreading and multiprocessing are closely related. Indeed, one could argue that the difference is one of degree: Whereas multiprocessors share only memory and/or connectivity, multithreaded processors share those, but also share instruction fetch and issue logic, and potentially other processor resources.

In a single multithreaded processor, the various threads compete for issue slots and other resources, which limit parallelism. Some "multithreaded" programming & architectural

models assume that new threads are assigned to distinct processors, to execute fully in parallel.

Cache Hierarchy. Recognizing the limitations of existing processors that rely on a shared cache model, a new approach calls for incorporating a three-tiered cache hierarchy into the MC Networking Platform. L1 cache is retained on the core.

As previously mentioned, L2 cache is attached to the cores as a backside implementation that can significantly improve performance. Each core has own back-side L2 cache that provides:

- 1) an aggregate bandwidth that could never be sustained by a single shared cache
- 2) Results in latency improvements vs. front-side (shared) cache
- 3) Back-side cache enables tuning of policies by core(s) according to different work sets for easier implementation of performance, isolation, priority, and QoS
- 4) A private cache is more self-contained (vs. a single shared cache) and can serve as a natural unit for resource management (e.g., powering off to save energy)

However, there are some tasks for which a shared cache is desirable, such as inter-processor communication and operating on shared data structures. For those instances, we are also providing a multi-megabyte L3 cache. This high-bandwidth, shared cache maximizes hit-rates while providing fast memory access for input/output (I/O) and accelerator blocks.

On-Chip Fabric. The on-chip fabric works in concert with the caching hierarchy to enable cache-coherent and concurrent accesses. The innovative backside cache implementation combined with the fabric is designed to enable data replication, modified intervention and full hardware coherence tracking.

The MC Networking Platform will employ highly scalable and modular on-chip fabric, the result of multi-year research and development, which enables cache-coherent, concurrent, low-latency connectivity among cores.

Unlike a shared bus as interconnecting medium among cores, memory and peripherals, the on-chip fabric helps to reduce the bus arbitration and contention issues that other multi-core architectures face as more traffic is introduced into the system. It behaves like a mesh, allowing concurrent traffic to enter and exit the system from any point within the fabric rather than through a single point.

Inherently scalable, the fabric is designed to sustain multiple, fully-coherent transactions every cycle and easily expand to accommodate more cores. On-Chip fabric also supports the option for heterogeneous clustering, allowing full portfolio of MCs, which spans a wide range of power and performance design points, to be mixed and matched in a product with full coherency among the cores.

Connectivity. The MC Networking Platform integrates an extensive set of networking and I/O resources to support its high-throughput architecture. This section focuses on these resources, which provides system designers a wide range of choices for scalable, high-performance systems.

MC/AE SoC Networking Platform Interfaces and building blocks □ The MC/AE SoC Networking Platform supports multiple interfaces including RGMII, XGMIII, and SPI-4.2 Interface controller. Additional high speed interfaces include: PCI-X interface and serial RIO interfaces.

Peripherals Interface. Peripheral devices and ROMs are connected to the MC Networking Platform through the various ports of the Peripherals Interface. The ports are created with different combinations of a 32-bit Peripherals I/O Bus and the programmable General-Purpose Input/Output (GPIO) signals.

The MC Networking Platform has essential standard busses such as standard I2C bus ports where each consists of two bidirectional bus lines; the Serial Data (SD) line and the Serial Clock (SCLK) line.

On-Demand Accelerator Engine (AE). On-demand acceleration provides accelerator Engines (AEs) technologies to take MC networking architecture to a new level of

performance and flexibility. An asynchronous, shared-resource architecture enables lower-latency and multi-task handling without the overhead of thread switching.

On-demand application acceleration offers performance advantages over pure core processing cycles, enables lower power implementations and reduces silicon area thus reducing cost. On-demand, high-performance Accelerator Engine's (AE's) technologies include:

- 1) Pattern matching for deep packet inspection and full content processing□
- 2) Decompression/Compression to unpack data for inspection and pack it for delivery□
- 3) Crypto security for confidentiality, integrity and authentication□
- 4) Table lookups for packet parsing and flow classification□
- 5) Data path resource management to efficiently allocate on-chip resources□
- 6) Packet distribution and queue management

Hybrid Simulation Environment□ The MC/AE SoC networking platform will require full system simulation model, a hybrid that combines cycle-accurate modeling technology with functional modeling technology that enables ease of software development, performance prediction and optimization of customer applications for the MC networking platform.

Using the hybrid simulation environment, which allows easy switching between functional and cycle accurate models, developers will be able to migrate and partition operating systems, middleware and applications onto the virtualized MC networking platform for development, debugging and benchmarking - even prior to silicon availability.

The environment also enables safe and easy experimentation with partitioning, parallelizing and optimizing systems and applications. Software developers can perform "what if" scenarios and tune the performance for specific situations without real-world hardware constraints. The hybrid simulator provides a programmer's view of the hardware, and features the following elements:

- 1) A fast, functional model for the MC networking platform□

- 2) A detailed cycle-accurate model of the MC networking platform□
- 3) A comprehensive package with infrastructure and tools for software development, code partitioning and debugging, profiling and visualization□
- 4) Visibility into system state both architectural and micro architectural including caches and registers pipelines. □
- 5) Run-time control of execution software including break pointing, stepping and reverse execution□
- 6) Ability to boot multiple operating systems

A major advantage of a hybrid simulator is its ability to dynamically switch back and forth from a high speed functional mode to a more detailed cycle-accurate mode.

This allows software developers to quickly boot an operating system and execute code at critical points and then switch to the more detailed cycle accurate mode to analyze specific areas of interest - no more waiting days for results.

As a development platform for multi-core systems, the hybrid simulation environment is designed to enable an extensive amount of flexibility and experimentation in a non-invasive environment - no instrumentation is needed in the operating system or application.

Software developers are able to decrease bring-up time for the target system all while improving the overall quality of their code.

The MC/AE Enablement Ecosystem (EE) □MC/AE Networking platforms will require software engineers to spend significantly more time thinking about software architecture. Exploiting the performance potential of MC processors means embracing parallel processing, which can be a challenge given the long and successful history of single core systems that are largely self synchronizing.

Networking applications offer coarse grained parallelism in the form of packet processing, and the interactions between a networking data path and the control plane are sufficiently decoupled to create an additional level of parallelism.

While this immediate parallelism is easy to envision, things get interesting when the

performance requirements of a data path flow exceed a single CPU's capabilities, or when a single core can't provide sufficient control plane responsiveness. Load balancing and mixed asymmetric/symmetric multi-processing environments on the same device are challenges that MC Networking Platform is designed to address.

While software architects are thinking about distribution of tasks, the processing densities offered by MC Networking Platform will cause hardware architects to think about consolidation and re-partitioning of functions that have been distributed across discrete CPUs or modules.

These decisions will interact strongly with the introduction of new services and capabilities in the system. For both software and hardware architectures, there is a need for a great deal of flexibility in a multi-core processor and for good mechanisms to help facilitate experimentation with future architectures.

The MC/AE SoC networking platform implements cores, each with their private L2 cache, also known as backside cache. In addition, the platform is equipped with on-demand accelerator engine that can be application specific.

While the Multi-core Platform is designed with aggressive performance targets, ease of use has also figured prominently in the platform definition. One of the significant obstacles in multi-core implementations today is programming efficiency and debugging. The two most likely scenarios (shown in Figure 3 below) are::

Scenario 1: Number of cores are normalized to 1-core in 2007 and System performance normalized to 1-core in 2007.

In this scenario, system performance at 45nm delivers 3.6x the performance at 65nm that required 3.7 cores against 1 core at 65nm. Similarly, at 32nm, system performance is 13.5x performance with 7.5 cores compared to 1 core at 65nm. The graph shows that performance is linear.

Scenario 2: Number of cores are normalized to 4-core in 2007 and System performance normalized to 4-core in 2007.

In this scenario, system performance at 45nm delivers 14.7x the performance in at 65nm that required 10.9 cores against 4 cores at 65nm. Similarly, at 32nm, system performance is 54x performance with 30 cores compared to 4 core at 65nm. The graph in **Figure 3 below** shows that performance is linear.




Figure 3. Two likely scenarios for the evolution of multicore SoCs as process technologies move from 65 down to 32 nanometer and below geometries.

The MC/AE SoC Platform Value Proposition □ Tomorrow's networking needs can no longer be met by increasing the operating frequencies on single-core architectures. Adding cores (MCs) will improve performance (Geometric Scaling).

But thermal management challenges, in the embedded space, are overwhelming the performance improvements achievable by increasing CPU frequency. Hence the need to look at the challenge from the SoC Platform perspective.

There may be contention for bus bandwidth and memories, scalability problems, and perhaps even worse, unused processing cycles due to lack of programming visibility.

Adding Accelerator Engines (AEs) will continue to add incremental improvement to performance (equivalent scaling) in the context of MC/AE SoC networking platform. But leveraging the hardware require greater investment in software enablement and simulation environment (Functional Diversification).

Thus, MC/AE SoC Networking Platform is not only designed to provide superior performance and energy efficiency, but also to help make the transition to multi-core processors as quick and as painless as possible with an industry leading enablement ecosystem.

Thus, Multi-Core (MC), Accelerator-Engine (ME) and Simulation/ Enablement/ Ecosystem (SEE) are three ingredients that will change the landscape of networking and will deliver a scalable & sustainable performance to meet next generation advanced application and services.

References□

1. ITU-R M.1645 "Future Development of IMT-2000 and IMT-Advanced", WG Spectrum, Document 8F revised draft, July 2007□
2. IEEE Communications, "Web Services in Telecommunications", "Orchestration in Web Services and Real-Time Communications", July, 2007 PP. 26-27, 44-50□
3. IEEE Wireless Communications, "New Generation Heterogeneous Mobile Networks", April 2007, PP 2-3□
4. IEEE Wireless Communications, "The Multiple Access Scheme for Wireless Communication", June 2007, PP2-3□
5. IEEE Wireless Communications, "Next Generation CDMA vs OFDMA for 4G Wireless Applications", June 2007, PP 6-7□
6. IEEE Wireless Communications, "IFDMA: A Scheme Combining the advantages of OFDMA and CDMA", June 2007, PP 9-17□
7. Communications News " Enterprise Network Solutions "Are you ready for converged IP?", July 2007 PP40-41□
8. Semiconductor International, "Semicon West 2007", June 2007, PP20□
9. Mobile Enterprise " Connecting Enterprise Solutions to Business Strategy, "Bettering Behavior, Mobile Tools", July, 2007, PP8, 19-25□
10. EE Times, "Freescale CEO: IC growth drivers shifting", July 2, 2007, PP8□
11. IEEE Micro, "Hot Chips 18", March-April, 2007, PP 7-9, "The AMD Opteron Northbridge Architecture", PP 10-21, "The Blackford Northbridge Chipset for the Intel 5000", PP 22-33, "ARM996HS: The First Licensable, Clockless 32-bit Processor core", PP. 58-68□
12. Power Architecture " Cell BE, "Cell Microprocessor", Wikipedia□
13. IEEE Computer Society, "Synergistic Processing in Cell's Multi-core Architecture", 2006, PP10-24□
14. ACM, "Evolution of Low Power Electronics and its Future Applications", ACM, 2003, PP2-5□

15. IEEE Comp Society, "CMOS Scaling for sub-90 nm to sub-10 nm", 2004, PP1-6□
16. IEEE Journal of Solid State, "CMOS Technology " Year 2010 ad Beyond", 1999, PP 357, 366□
17. IEEE " Proceeding of 8th IPFA 2001, "Direction of Silicon Technology from Past to Future", 2001, PP 1-35□
18. ITRS 2005 Publication " Introduction of "More than Moore" concept□
19. ITRS 2007 Summer Working Group Workshop/Public Conference " Work in progress on "more than Moore" □
20. Semiconductor International, July 18th ITRS Summer Conference " Panel Focus on "More Than Moore", by Peter Siger, Editor-in-chief□
21. ITRS 2007 System Drivers Publications, Networking Driver " SoC Multicore/Accelerators Platform , Pages 3-5

To learn more about designing embedded systems designs using multicore technology, check out a number of classes and presentations on this topic at the [Embedded Systems Conference Silicon Valley 2009.](#)

This paper was recently published by Embedded.com. You can view it at <http://www.embedded.com/design/multicore/216200003?pgno=1>

Author:



***Fawzi Behmann** is currently the Director of Marketing at Power.org. the open community driving collaborative innovation around Power Architecture® technology. He is working with member companies advancing the roadmap of the Power architecture and ecosystem. Prior, Fawzi has been the Chair of the Marketing Committee of [Power.org](#), Fawzi has been the Director of Strategic Marketing for the Networking Systems Division within **Freescale's Networking and Multimedia Group**. Also, he has chaired the networking System Drivers Working Group at ITRS (International Technology Roadmap for Semiconductors and developed the work on MC/AE SoC platform.*