

**ITRS – Networking and Microprocessor (MPU) Driver White Paper**  
**Fawzi Behmann/Freescale/fawzi.behmann@freescale.com**

## **Abstract**

As technologies and structures push the limits of Moore's law and productivity, the ITRS 15-year assessment of the semiconductor industry technology requirements and potential solutions indicates new approaches to maintain scaling and introducing functionality on- and off-chip. In this section of ITRS document, will start by highlighting some macro trends such as Internet everywhere, IP everywhere and Seamless Mobility that will have significant impact in defining future requirements and the challenge to address them at Semiconductor Level. This will be followed by focusing on micro trends that influence the networking embedded space and calls for an SOC platform with key drivers such as Multi-Core (MC), cache hierarchy, on-chip fabric, on demand Accelerator Engine (AE), connectivity - engineered as SOC-MC/AE platform to provide a scalable, software-based solution and target wide range of applications from ultra low-end to high-end that preserve & extend the user experience through new services. Next will introduce the new approach defining "Moore" – Geometric Scaling, "More Moore" – Equivalent Scaling and "More than Moore" – Functional diversification and how it applies to SOC-MC/AE Networking platform. The document will conclude with assumptions and performance model for SOC-MC/AE networking platform.

## **1. MACRO Trends**

### **1.1 Internet everywhere - Internet & Web Services**

With convergence between telecommunications and data systems witnessed in recent history, and after the mass commercialization of the Internet and the rapid emergence of the Web during the 1990s, telephony systems began migrating to internetworked client-server architecture and TCP/IP protocols. In this decade the data world has evolved from the client-server architectural approach to the service-oriented architecture (SOA) approach for system organization. In 2007 the telecommunications industry is still in the early adopter of SOA and web services, and accordingly we are seeing many technology innovation and business models being developed as part of the adaptation of SOA and web services and for the coordination of real-time communications services, including telephony, video and multimedia communications.

### **1.2 IP everywhere - Scalable IP-Based networks**

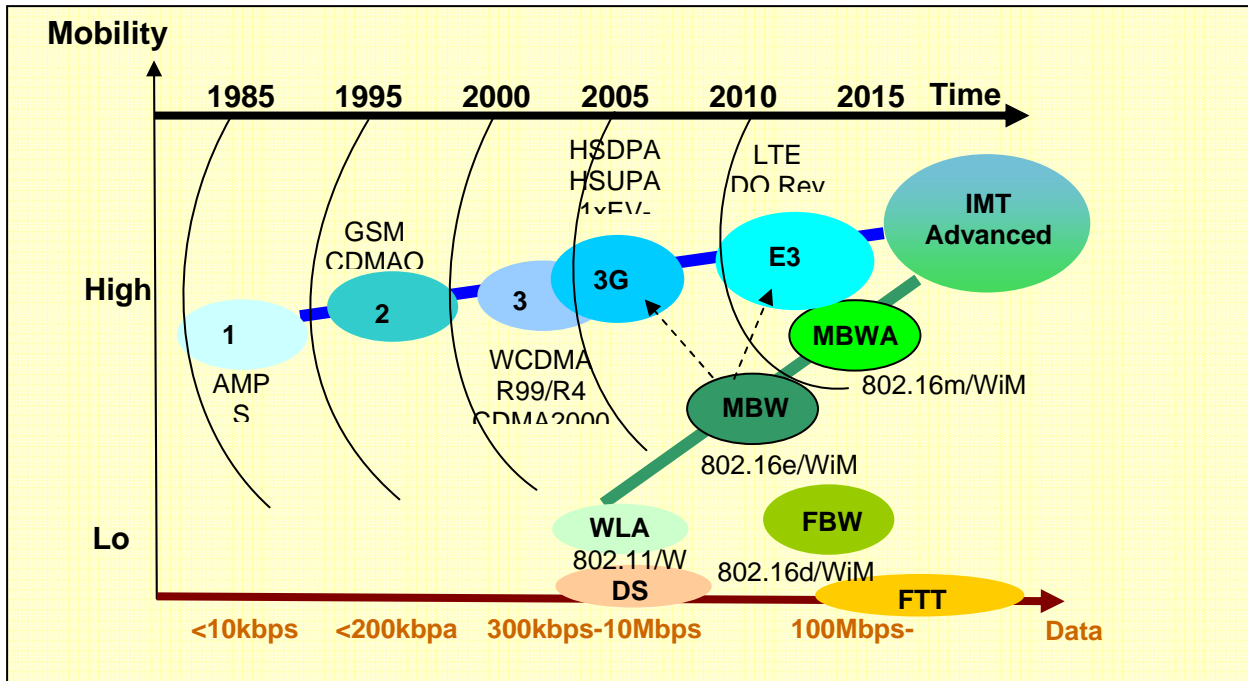
Why IP? IT managers know the headaches of dealing with separate voice, data and video networks that are inflexible, expensive and complex to manage. When disparate networks cannot communicate with each other, control and visibility are limited. Multiple hardware and software vendors with multiple support and maintenance contracts add to the management cost and complexity.

As IP an important driver for scalable networks, managing moves, additions and changes become easier and consistent extending capabilities to the web and control of features to the end-user. In addition, geographically distributed IP-based networks can connect remote sites/branch offices/corporate/communities with the same look and feel.

### **1.3 Seamless Mobility - Multiple Wireless Access converging to 4G and beyond**

Several mobile access schemes available today are evolving. Time-Division, Frequency-Division, and Code-Division Multiple Access (TDMA, FDMA, and CDMA) are key examples of active access technologies in the 90's when finally CDMA took its leadership in the 3G mobile networks. As a result, the main systems under the larger umbrella of IMT-2000 have all adopted the concept of wideband CDMA or other techniques. A crucial presumption was that mobile communications is for voice communications with data as a secondary or value-added service. Since then many trends have changed; most important is the fact that data traffic is now becoming the main part of mobile communications, in particular high-speed data communications. With such requirements, will CDMA keep its preference?

With the progress in high-speed wireless data systems such as wireless LAN and WiMAX, a new form of multiple access as well as multiplexing has emerged: Orthogonal Frequency Division Multiplexing (OFDM), which is now gradually finding its path into cellular networks and 4G systems. The orthogonally among different users' signals, low computational complexity, and easier intersymbol interference (ISI) reduction, among many other features, make OFDMA a highly reliable multiple access scheme for high speed data communications.



Today, both CDMA and OFDMA have been adopted by various standards, CDMA is used in IS-95, CDMA-2000, Wideband CDMA (W-CDMA), time-division synchronous CDMA (TD-SCDMA), and so on. On the other hand, OFDMA/OFDM is used in IEEE 802.11a/g/n WLAN, HIPERLAN/2, WiMAX, DVB-T, asymmetric digital subscriber line (ADSL), very high rate DSL (VDSL), and others. It has been chosen as the physical layer architecture for 3GPP long-term evolution (LTE). The major advantage of OFDM is its ability to deal with multi-path fading and narrowband interference without using complicated channel equalization.

Going forward, and independent of the powerful access technologies based on CDMA and OFDMA, multiple-input multiple-output (MIMO) antenna systems have been identified as one of the key technologies to support higher data rates through spatial multiplexing and diversity in comparison to single-antenna systems, MIMO techniques can be elegantly combined with OFDMA as well as CDMA.

**Three combined powerful drivers coming together: Internet/Web Services, IP, and Mobility impacting future networking architecture, applications & services and accelerating semiconductor silicon platform architecture.**

The great success of the Internet and Wireless communications has opened a new vista for future all-IP wireless applications, driven by increasing demand for packet data services in worldwide 3G, 3.5G networks. It has been envisaged that the future wireless systems will operate based mainly on burst data services carrying multimedia contents, including voice, data, image and video.

Currently research on beyond 3<sup>rd</sup> and 4<sup>th</sup> generation (B3G/4G) mobile radio systems is in progress worldwide. A future mobile radio system will provide packet-oriented data services carrying multi-media contents. On the other hand, properties like high spectral efficiency, as well as high flexibility and granularity in terms of different data rates are essential. On the other hand, low cost implementation and high-power efficiency are important. The choice of the multiple-access (MA) scheme has a great impact on the achievement features of a future cellular mobile radio system. The candidate MA schemes can be classified as a single-carrier-based and multicarrier-based MA schemes.

It is envisaged that MA scheme will combine the advantages of CDMA and OFDMA into Interleaved Frequency Division Multiple Access (IFDMA). IFDMA combines the advantages of CDMA; namely, low complexity for signal generation, low Peak-to-Average Power Ratio PAPR, and high frequency diversity and advantages of OFDMA; namely, low complexity for user separation and channel equalization. At the same time, it provides sufficient robustness and flexibility in terms of data rates

The need to support high-speed burst type traffic in wireless channels has already posed a great challenge to the current air link technologies based on CDMA

## 2. Next 10-15 Years Needs

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 3<sup>rd</sup> 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:

- Capability for multimedia applications within a wide range of services and terminals
- High degree of commonality of design worldwide
- Compatibility of services within IMT-2000 and with the fixed networks
- High quality
- Worldwide roaming capability
- Small terminal suitable for worldwide use

IMT-2000 and IMT-Advanced have the ability to improve lifestyles, social interactions and productivity. The international trend is to use telecommunications as a means to reduce social and economic differences by enabling entire populations, despite their location and resources, to have full coverage and access to telecommunication services.

### 2.1 Trends will be driven by mass market:

With introduction of cell phone and movement towards mobility, by 2004, consumer semiconductor content has crossed 50% of total semiconductor setting the stage for a direction that target to serve the mass market.

- A new era of advanced innovative applications and services are in the plan taking advantage of collaborative set of innovative technologies and targeting next generation devices to enrich one's experience of high definition content and bringing communications to a new level. For example, a new Mobile Personal Services (MSP) platform is defined making use of collaborative technologies for smart phone, GPS, wireless headset, biometric sensors targeted to offer users feedback on their behaviors and habits. One prototype application, called the Personal Performance Coach, has enormous potential for professionals, particularly in sales. The solution is still in development stages, with a projected release still a few years out.

- The semiconductor industry's global opportunities are shifting to "green everything," an aging population and broadband, ubiquitous connectivity, Freescale Semiconductor CEO Michel Mayer said during a keynote address at the company's technology forum held June 25-28, 2007 in Orlando, Florida. The embedding of intelligence in everyday devices, coupled with wireless connectivity, will be disruptive. "Appliances that stood alone are now being connected," Mayer said. Energy, for example, is a \$200 billion cost to consumers and \$100 billion for industry per year that will grow by 53% by 2010. By that time, emissions will rise to 25 million tons a day. The enormity of the challenge will create "opportunities for innovation and design intelligence," Mayer added.

- Another area that will undergo rapid technological development is in robot intelligence. Robot intelligence will dramatically increase in the coming decades, owing primarily to the progress of high performance and low power chip technologies. The figure below shows a prediction of robot intelligence by Dr. Moravec of Carnegie Mellon University. The robot will provide the biggest challenges for the low power electronics in the future, since the total power consumption is very limited in the range of several watts to several tens of watts. There fore it is required to achieve very high value like  $10^6$  to  $10^7$  in terms of MIPS/Watt.

### 2.2 The next 5-15 years will mark the trends towards

- Scalable networks that delivers high rich multimedia content at broadband speed anywhere and anytime and on any device.
- Consumer will play a major role in creating high rich multimedia content
- Emergence of advanced IP-based applications and services that drives high bandwidth scalable Networks

- Complex multi-processing platforms equipped with multi-core/multi-threading and accelerators that support advanced applications and services
- Advancement in process technology from 65-45-32, 22 and sub 10nm technology
- Scalable encryption and antivirus everywhere in the network
- Home networking will be a complex network converging data communications, entertainment,
- Seamless mobility in the home, in the office/vertical market, on the road

### 3. Semiconductor Approach to meeting forward looking challenges

It wasn't all that long ago that consumers of microprocessors expected Moore's law-driven performance gains of 60% per year as a matter of course. According to industry analysts such gains abruptly ceased in 2002. At that point, a variety of issues conspired to make it extremely difficult to continue scaling clock rates and issue widths as before. Since the sea change of 2002, all the major chip manufacturers have jumped onto the multi-core bandwagon in an attempt to continue growth in functionality, if not in single-thread performance. Unfortunately, this trend completely neglects questions of how such systems will be programmed and raises many issues – such as how to efficiently communicate between processors, how to do so at reasonable power.

#### 3.1 Compute & Embedded Environments:

The bulk of the world's microprocessors appear in embedded applications, not workstation or laptop environments. The following sections will address briefly needs of high-performance computing environment then address in more details the embedded space.

#### 3.2 High-Performance computing environment Needs:

With the Internet's rapid growth, the convergence of voice, data, video and associated services has redefined the traditional computation model for platform design. Today, business enterprises need high-bandwidth, high-connectivity, low-cost, highly secure, power-efficient, and reliable systems for front- and back-end machines. These systems are typically deployed in data centers as pedestal, rack, and blade servers. Examples include file servers, firewall and secure servers, storage, e-mail servers, and Internet servers. Similarly, the high-performance computing (HPC) sector and scientific community require teraflops of computing power for computation-intensive applications such as flight simulators, computational fluid dynamics, finite-element analysis and modeling, oil and gas exploration, medical imaging, and computer-aided manufacturing. In addition, desktop and workstation computing have ushered in an era of trendy applications in music and video on demand, multi-media authoring and content creation, animation, and gaming, all of which require new levels of power, performance, and efficiency.

#### 3.3 Embedded Space Challenge/Needs:

In contrast with PC & Server Applications, and due to the fundamental difference between core speeds and memory/IO latencies, today's processor architectures are unable to deliver meaningful performance for the connected computing solutions. Why? These processors are usually waiting for data. This delay also known as a "stall" can be as much as 90% of the time. Up until now, today's processor designs attempt to overcome these barriers by using deeper pipelines, superscalar (multi-issue) operations and larger caches. However, those efforts are yielding diminishing returns because the data processing occurs at the macro or packet level not at the instruction level.

#### 3.4 The need for Multi-core:

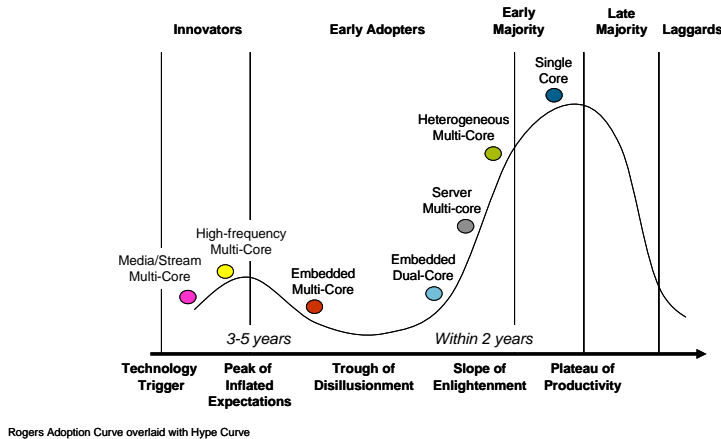
Increasing performance under 30 watts is a major challenge for embedded applications since core performance under 30 Watt can't increase in frequency without increase of power. Other techniques such as multi-core, multithreading, hardware & software accelerator are used to increase performance within power constraints.

The continuing drive to IP convergence and service-oriented networks will create monumental performance demands on infrastructure and access networks. Tomorrow's networking needs can no longer be met by increasing the operating frequencies on single-core architectures.

The diagram below highlights multi-core hype vs adoption, Multi-core processing is establishing itself as the solution of choice in a number of electronic industries, yet the Technology itself is still in its infancy.

- The **majority** of embedded systems are based on single-core, migrating to **dual-core**.
- Early adopters of current multi-core architectures are imbalanced, suffer high-latency and **cannot make up the deficit** with “more cores”.
- **All cores are not created equal!**

.A related challenge is that the majority of the networking installed base is still operating on a mixture of single-core processors, ASICs and DSPs. Developers have the enormous task of effectively migrating millions of tested, proven and fielded lines of code to multi-core architectures before the full benefits of any multi-core system-on-chip (SoC) can be realized.



Thermal management challenges are overwhelming the performance improvements achievable by increasing CPU frequency. The answer, however, is not simply adding cores to a die either. As many current implementations show, more does not necessarily mean better. There may be contention for bus bandwidth and memories, scalability problems, and perhaps even worse, unused processing cycles due to lack of programming visibility.

To meet the needs of next-generation systems and applications while continuing to provide improved performance, new processor designs must incorporate innovative architectural concepts.

Some multi-core platform uses short pipe-line for efficient cores and consequently stalls frequently. Other multi-core platform provides key capabilities to reduce latency and thus length of stalls. Features include: Out Of Order Execution, better branch prediction, wide bandwidth and multiple concurrent accesses leads to low-contention on the bus, low-latency inline L3 shared cache(s), and stashing enables major accelerators and network blocks to write data into the L3 instead of external DDR.

## 4 Semiconductor Networking Driver and Modeling (embedded Space)

### 4.1 Architecture Directions for Embedded Space

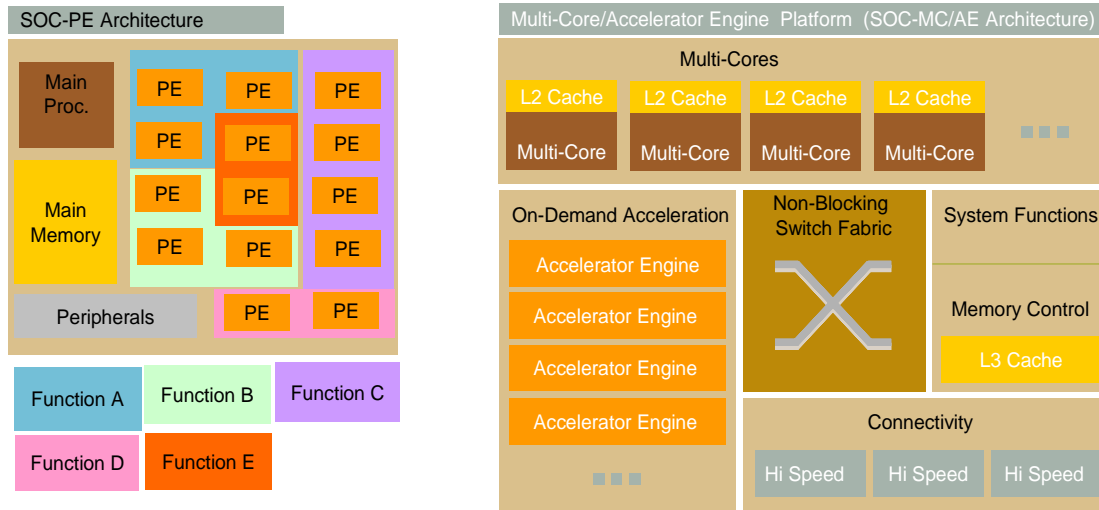
In contrast with PC & Server Applications, and due to the fundamental difference between core speeds and memory/IO latencies, today’s processor architectures are unable to deliver meaningful performance for the connected computing solutions. Why? These processors are usually waiting for data. This

#### 4.1 Traditional Processor Design

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.

## 4.2 Emerging Architecture/Platform Approach - SOC-PE Consumer Architecture & SOC-MC/AE Networking Architecture

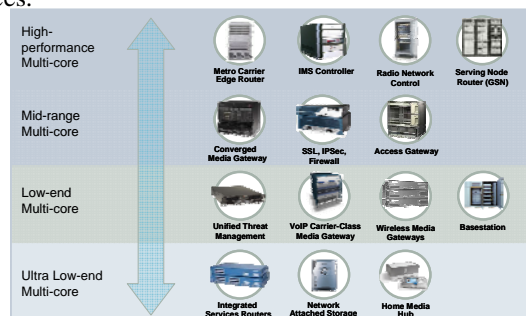
Early in 2005, ITRS introduced SOC-PE Architecture Template, where a PE is a processor customized for a specific function targeting portable and wireless applications such as smart media-enabled phones or digital camera chips but also high-performance computing and enterprise applications. To complement this SOC-PE architecture, a Multi-Core/Accelerator Engine SoC Architecture template is defined to address the networking embedded space.



The MC/AE SOC networking platform contains the necessary building blocks to:

- Supporting Multi-Core (MC) for high processing performance within 30 watt power envelope,
- Supporting an unprecedented Tri-level Cache Hierarchy
  - MCs with back-side L2 caches
  - Multiple L3 shared caches
  - Multiple memory controllers
- Supporting high-speed inter-connectivity
- Introducing a scalable On-chip Fabric for concurrent, non-blocking, hardware-based 100% cache-coherent platform connectivity
  - Eliminates shared bus contention and supports dramatically higher address issue bandwidth to “feed” multiple cores
  - Scales to support more than 32 cores
  - Can support heterogeneous cores
- Including On-demand Acceleration Engine (AE)
  - Offers performance advantages over pure core processing cycles, enables lower power implementations and reduces silicon area / cost
- Hybrid Simulation Environment combining cycle-accuracy and functional-accuracy that enable ease of software development , performance prediction and optimization,
- Network/System Enablement & Ecosystem looking into software partitioning and virtualization leveraging multi-core hardware architecture

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



## 5. SOC-MC/AE Networking Platform & Modeling

The Multi-Core/Accelerator Engine Platform represents a balance approach to Multi-Core/Accelerator-Engine system-on-chip (SoC) design. It introduces an advanced on-chip connectivity fabric that is able to support multi-core and accelerator engines cores. This section will align the approach to that of “Moore”, “More Moore” and “More the Moore” concept introduced in ITRS 2005 publication and has become a working item from application perspective into 2007 working groups. Next, will map the SOC-MC/AE Networking platform into the “Moore” new approach followed by detail description the building blocks and conclude with platform performance model support by assumption.

### 5.1 Introduction of 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 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, 3<sup>rd</sup> party IP/enablements) 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. Adding “More of Moore” to the mix, provides a converged/integrated heterogeneous platform that enable the creation of a scalable, intelligent, compact value-add ecosystem. This is Platform concept of 3-Moore’s is becoming an important paradigm moving forward. Thus, the application of this Moore’s platform concept is becoming a critical activity in many of ITRS working groups in 2007.

### 5.2 SOC-MC/AE Networking Platform and Moore’s classification

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

## 5.3 SOC-MC/AE Networking Platform Modeling

### 5.3.1 Process Technology Roadmap

The process technology roadmap is defined by ITRS positioning 65nm for 2007, 45nm for 2010, 32nm for 2013, 22nm for 2016, 16nm for 2019 and sub-10nm for 2022. Thus, the next process technology node is 45nm.

Advanced semiconductor research and design enablement shows 45 nm for both with proven capabilities in high-performance SOI wafer processing. 45 nm SOI is a rapidly maturing technology, expected to reach full process certification in 2008. The 45 nm SOI technologies include a range of high-performance transistor offerings and static random access memory (SRAM) bit cells that provide an excellent balance of performance and low power. Some key enablers of this technology are 193 nm immersion lithography for scaling and reducing sources of device variation, porous low-k ( $k=2.4$ ) dielectric for minimized back-end wiring delay, and advanced strain techniques for enhanced transistor performance. Achieving the highest level of process technology performance in 45 nm SOI is a key building block for multi-core success within a usable power envelope.

### 5.3.2 Geometric Scaling (GS)

#### A. Multi-Core (MC)

The Multi-Core's frequency in a wide multi-core product will be targeted over one GHz. This platform target 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.

#### B. 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. Level 1 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:

- Back-side caches provide an aggregate bandwidth that could never be sustained by a single shared cache.
- Results in latency improvements vs. front-side (shared) cache.
- Back-side cache enables tuning of policies by core(s) according to different worksets for easier implementation of performance, isolation, priority, and QoS
- 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 Level 3 (L3) cache. This high-bandwidth, shared cache maximizes hit-rates while providing fast memory access for input/output (I/O) and accelerator blocks.

### C. 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.

### D. 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.

#### D.1 Networking Interfaces

The SOC-MC/AE Networking Platform supports multiple interfaces including RGMII, XGMII, and SPI-4.2 Interface controller. Additional high speed interfaces include: PCI-X interface and serial RIO interfaces.

#### D.2 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.

## 5.3.3 Equivalent Scaling

### A. 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:

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

## 5.3.4 Functional Diversification

### A. Hybrid Simulation Environment (S)

A full system simulation model for the MC networking platform is required. The simulation is a hybrid 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 can 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:

- A fast, functional model for the MC networking platform
- A detailed cycle-accurate model of the MC networking platform
- A comprehensive package with infrastructure and tools for software development, code partitioning and debugging, profiling and visualization
- Visibility into system state both architectural and micro architectural including caches and registers pipelines.
- Run-time control of execution software including break pointing, stepping and reverse execution
- 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.

### **B. Enablement/Ecosystem (EE)**

MC/AE Networking platforms 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 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.

## **5.4 System Model Assumptions**

To model multi-core by itself will only provide a theoretical number. The objective is to model the SOC platform which would be a complex task to do. Since multi-core is a major key driver for performance, it is essential to start with two key variables and a set of assumptions.

The two key variables: Process technology roadmap and number of cores.  
Core frequency is estimated to contribute 5% improvement over time.

The set of assumptions include:

Target Market Segment: Mid-range segment of the embedded space. In the context of several key market segments identified in section 4.2 addressing the embedded space, multi-core requirement start appearing in mid range.

Workload: Mid-range switching/routing workload

Cache Hierarchy: System memory will continue to scale

On-Chip Fabric: On-chip fabric will support system scaling

Connectivity: Inter-connect will continue to scale

Cost/Area: Assume constant-cost/constant-area die

On-Demand Accelerator: Factored in and contribute an average of 5% toward system performance

Simulation Environment: Hybrid simulation (Function & Cycle)

### 5.5 SOC-MC/AE Networking Platform – Performance Model

The platform implements cores, each with their private Level 2 (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 our platform definition. One of the significant obstacles in multi-core implementations today is programming efficiency and debugging.

Two scenarios are shown below:

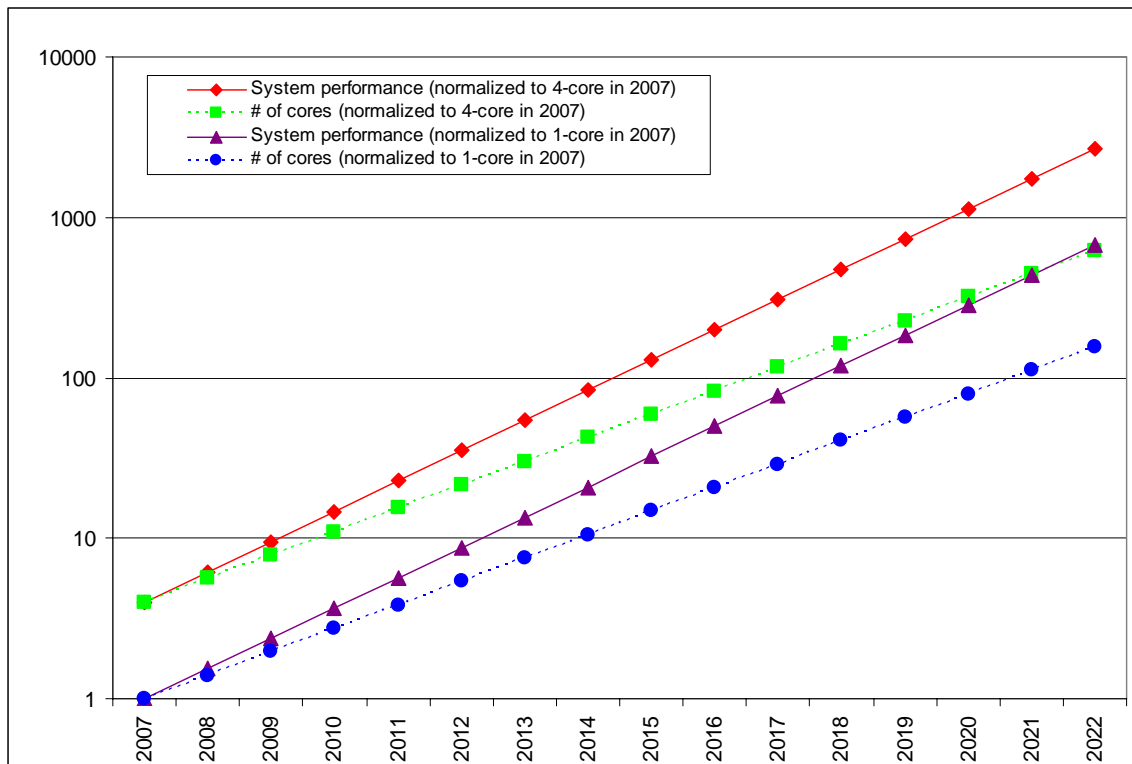
**Scenario 1:** Number of cores are normalized to 1-core in 2007 and System performance normalized to 1-core in 2007. In this scenario, system performance at 45nm delivers 3.6x the performance at 65nm that required 3.7 cores against 1 core at 65nm.

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 shows that performance is linear.

### SOC-MC/AE Network Platform Performance



### 6. SOC-MC/AE Networking 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 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 SOC-MC/AE networking platform. But leveraging the hardware require greater investment in software enablement and simulation environment (Functional Diversification). Thus, SOC-MC/AE 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.

Investment in process technology roadmap and software enablement that will enable:

Geometric scaling:

- Scalable on-chip fabric
- Scalable MCs
- Three-level cache hierarchy
- High speed connectivity

Equivalent scaling:

- On-Demand AEs

Functional Diversification:

- Hybrid Simulation environment
- Enablement/Ecosystem

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

## 7. ITRS – Networking 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 8<sup>th</sup> 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 18<sup>th</sup> ITRS Summer Conference – Panel Focus on “More Than Moore”, by Peter Siger, Editor-in-chief