Building Next-Generation Telecom Equipment Based on the AMD64 Processor Family

By Moh ElBaz

AMD Boston Design Center

90 Central Street

Boxborough MA 01719

© 2006 Advanced Micro Devices, Inc. All rights reserved.

AMD, the AMD Arrow logo, AMD Athlon, AMD Opteron, AMD Turion and AMD Sempron and combinations thereof, and AMD-8132 are trademarks of Advanced Micro Devices, Inc. HyperTransport is a licensed trademark of the HyperTransport Technology Consortium. PCIe and PCI-X are trademarks of PCI-SIG. MIPS is a registered trademark of MIPS Technologies, Inc. Other names are for informational purposes only and may be trademarks of their respective owners.

Table of Contents

Introduction

The AMD Opteron™, AMD Turion™ 64 and Mobile AMD Sempron™ processors are now available for high-end embedded designs. The AMD64 superscalar processor family features Direct Connect Architecture with an integrated memory controller, a scalable high-speed HyperTransport™ technology I/O interface and industry-leading x86 processor performance; it is also compatible with the 32-bit x86 instruction set. The AMD Opteron, AMD Turion 64 and Mobile AMD Sempron processors are NEBS/ETSI friendly and are available in a wide range of power envelopes and under an extended longevity program. It is now possible to design dense and high-performance computing and I/O blades for telecom equipment at a price/performance point that is ideal for standardized platforms such as Advanced TCA, MicroTCA, ATCA/AMC, CompactPCI and VME. The elimination of a Northbridge and the flexibility of AMD64 technology's Direct Connect Architecture facilitate designs with a lower chip count, smaller footprint and lower power and cooling requirements at flexible price/performance points.

This paper will describe the changing landscape facing Telecom Equipment Manufacturers (TEMs) and will provide the reader with a detailed overview of the AMD Opteron processor and a brief overview of the Mobile AMD Sempron and AMD Turion 64 processors. Examples are given showing how the AMD Opteron processor's features and benefits can help solve the challenges of building next-generation telecom equipment. The paper will conclude with an overview of the extensive hardware, software and development support for the AMD64 ecosystem.

The Telecom Industry: Changes and Challenges

Over the past several years, the telecom industry has been in the midst of a major transformation. Significant economic, structural and technological changes now require new thinking and approaches to solve the challenges facing the industry. Major shifts in usage patterns and the emergence of packet based systems have resulted in the shrinkage of traditional voice revenue sources. On the other hand, the increased bandwidth and Quality-of-Service (QoS) demands created by new services such Voice over IP (VoIP), media-on-demand and Internet access everywhere represent opportunities for generating substantial revenue growth and for new players introducing disruptive solutions and technologies to enter the market. In order to succeed in this new and highly-competitive environment, Service Providers and Telecom Equipment Manufacturers face the need to deliver products with reduced time-to-market while driving down all associated costs: design, development, delivery, implementation and cost of ownership.

To achieve these goals, there is a growing move in the telecom equipment sector (as well as in many other embedded markets) toward open standards and the use of Commercial-off-the-Shelf (COTS) components for next generation systems. This move is partly the consequence of the recent downturn in the telecom industry which forced equipment vendors to re-think spending practices. It is also driven by the need to more quickly respond to changes in the market and to do so at a low operational cost. As a result, high-end telecom equipment vendors can move away

from specialized products which are based on proprietary chassis and backplanes to systems based on open standards such as the emerging Advanced Telecom Computing Architecture® (a.k.a. AdvancedTCA® or ATCA®). To succeed, the COTS components selected must be reliable, flexible and highly-scalable.

As of 1Q06, a number of NEBS compliant AdvancedTCA platforms have been announced based on the AMD Opteron processor including Sun Microsystems' Netra line and an ATCA reference design from PDSI. Others are expected to follow. These platforms represent a level of modularity, flexibility and performance not seen in prior platforms. (see Figure 5). When AMD64 technology with Direct Connect Architecture is combined with AdvancedTCA, the result can be a value play that is disruptive to the previous status quo and economically compelling with a lower total cost of ownership than platforms based on legacy processors.

Advanced Telecom Computing Architecture®

What is ATCA , AMC and Micro TCA

Advanced Telecom Computing Architecture®, Advanced Mezzanine Card and Micro TCA comprise a series of industry standard specifications created by the PCI Computer Manufacturers Group (PICMG) for the next generation of high-availability carrier grade communication equipment. According to PICMG, AdvancedTCA incorporates the latest high speed interconnect technologies, next generation processors, and improved reliability, manageability and serviceability features (RAS). The ATCA family of standards defines new form factors for full and fractional size system cards, mezzanine cards and a chassis (shelf) optimized for communications.

AMC (Advanced Mezzanine Card) is a well defined standard for a daughter card that fits on the ATCA blade. AMCs, as the name implies, are mezzanine cards that mount on an ACTA platform. Several options (single/dual height, single/dual width) are defined with power budgets ranging from 30 to 60 watts.

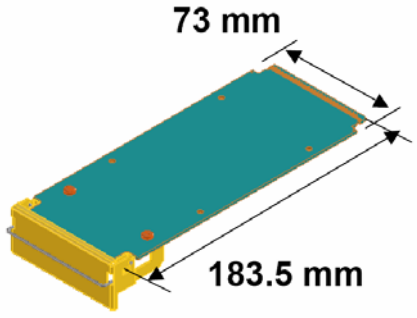

Figure 1 - AMC or MicroTCA Form Factor

MicroTCA is a draft system architecture specification that is gaining traction across multiple market segments (telecom, medical imaging, military, storage, industrial control, etc.). MicroTCA allows AMC modules to be plugged into a stand-alone system shelf. The key features of MicroTCA include hot-swap, low power (<30 watts), small module size, and low cost with advanced management features.

AdvancedTCA, MicroTCA and AMCs provide standardized platform architecture for carrier-grade telecommunication equipment, which can scale in the number of ports, bandwidth, power envelope and physical size based on standard chasses. Additionally, ATCA complies with

network operation environmental features such as Network Equipment Building Standard (NEBS), European Telecommunication Standard Institute (ETSI) and meets the 99.999% availability requirement as demanded by Telecom Equipment Manufacturers.

ATCA Overview

ATCA specifications are divided into a base PIGMG 3.0 specification and a series of sub-specifications (PICMG 3.x) for different backplane and switching fabric technologies. The ATCA chassis is designed to fit the latest 600mm ETSI shelf but can also be used with older 19" or 23" shelf systems. The ATCA backplane supports two 16 hot-pluggable front slots with 1.2” spacing, where each slot supports blades of 8U-high and 280mm-deep consuming up to 200 watts each. In addition, rear transition modules and AMCs are defined for end-point flexibility. For example, AMCs can be used as hot swappable, application-specific interface modules supporting Fibre Channel, InfiniBand, Gig/e, or OC-192.

The base specification defines 10/100/1000Base-T Ethernet backplane connectivity across all slots. This interface has a dual-star topology and can support low bandwidth applications up to 1 Gbit/s per slot. This base connectivity can either be used as an inter-module control channel or point-to-point high speed data path; higher-bandwidth applications (>1Gb) can use the fabric interface. This interface is a flexible switching solution supporting redundant switching fabrics, mesh and dual-star topologies. Currently ATCA supports the following technologies as fabric interfaces: Ethernet (3.1), InfiniBand (3.2), StarFabric (3.3), PCI Express (3.4), and RapidIO (3.5).

The ATCA chassis is targeted to the central office. As such, it is NEBS/ETSI compliant and is powered by dual-redundant 48 volt DC power. A unique feature of ATCA is the built-in shelf management which controls power distribution, temperature and cooling. In addition, ATCA supports hot-swap/power-on servicing capabilities for field-replaceable units. This feature is important during failovers and in-service installation of new ATCA blades and AMCs.

A recent addition to the ATCA standard includes the compatible MicroTCA architecture which provides the same benefits as ATCA but in a smaller form factor suitable for remote offices. MicroTCA leverages the ATCA AMC form factor and provides flexibility and economies of scale by re-using or sharing the same form factor either as a mezzanine card on an ATCA or as a full fledged MicroTCA blade. Yet another addition is the Hardware Platform Interface (HPI)-to-AdvancedTCA Mapping Specification. HPI was created by the Service Availability™ Forum (SA Forum) and is a data model and programming interface which describes generic hardware platforms, and in the context of ATCA, exposes the shelf management functionality in a standard, vendor-independent manner.

ATCA in the Market Place

ATCA was developed by PICMG in conjunction with the telecom industry and is rapidly gaining acceptance in other sectors as well. The market for standard chassis-based systems and for ATCA-based systems in particular, is set for significant growth. According to

In-Stat, it could reach near $8 billion in 2009. (For more information on the ATCA market or to obtain a copy of the In-Stat report - go to www.instat.com or email info@instat.com; reference report No. IN0501810NT.)

The AMD Opteron™ Processor and Advanced TCA

The AMD Opteron processor is well suited for the ATCA platform in a number of ways:

- High-performance x86-64 CPU with an impressive 24GB/s of on-chip I/O
- One socket can accommodate single or dual cores with only a BIOS change
- Available in 3 power envelopes that fall within the ATCA specification (see Table #1)
- AMD64 Longevity Program
- Capability of meeting NEBS thermal requirements and ETSI standards
- RoHS compliance
- Support of AMD's Professional Design Support Services team, speeding time to market

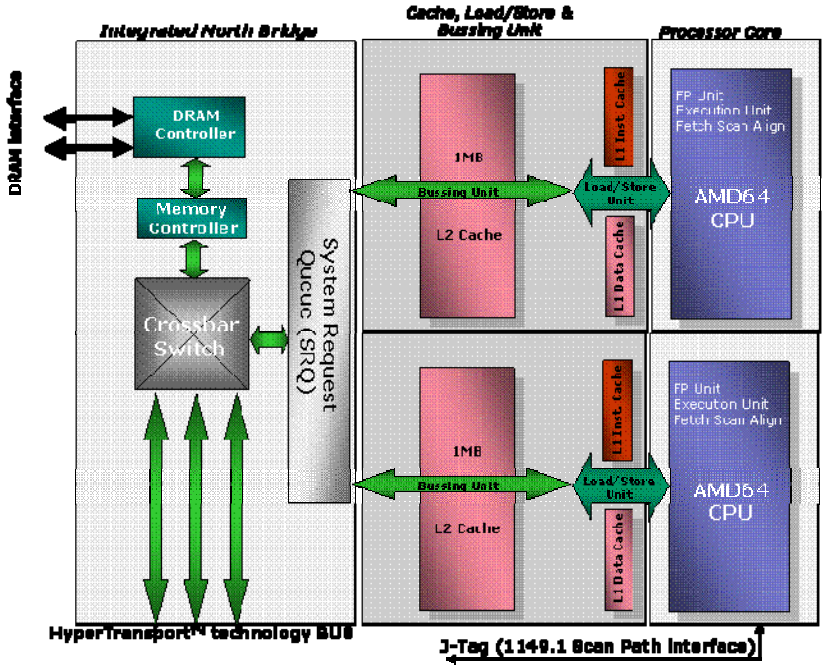

Figure 2 - Data Flow View of a Dual-Core AMD Opteron™ Processor

Figure 1 above depicts the innovative Direct Connect Architecture of the Dual-Core AMD Opteron processor which offers a high-performance and low latency. The strengths of the AMD Opteron processor are its balanced computing power and its data-movement capabilities, putting it in a class of its own and making it an outstanding CPU for telecom applications.

The AMD Opteron™ Processor Core

In the selection of a processor for a next generation ATCA platform, total system performance is much more than the subjective value of how fast the CPU is clocked. Rather, the objective important factors are:

- How much I/O can be streamed without degrading memory performance
- How much work can be done in "n" clock cycles at "x" watts
- How much board real estate is needed to meet the needs of the ATCA application

AMD64 technology is based on an innovative, backward compatible and fully native 64-bit x86 implementation which leverages numerous advancements in CPU architectural development. AMD64 technology supports pin and thermal compatibility between single- and dual-core processors. This flexibility provides TEMs with significant headroom when designing boards for different needs and conserves platform investment, i.e. one platform for:

- One single-core processor
- One dual-core processor
- Two single-core processors
- Two dual-core processors

Many of the elements of AMD64 technology such as branch prediction, shorter pipelines, etc. are based on the extremely successful AMD Athlon™ XP processor. However, the Pick, Fetch, Decode and Execute compute blocks were designed from the bottom-up to be a true x86-64 execution unit. The design was not an afterthought or a quick design turn with "64-bit extensions" rather, a thoughtful execution of a truly native x86-64 architecture. This resulted in a world class processor that evolved from a 130nm Silicon on Insulator (SOI) single core x86-64 processor to a 90nm SOI dual-core processor with leadership in x86 performance-per-watt and throughput.

Why AMD64 Technology for Telecom?

High-speed memory and I/O combined with low memory access latencies are critical to the overall performance of telecom equipment or any computing system. The bandwidth supported by communication devices, often measured in Packet-Per-Second (PPS), is largely determined by processor performance. The AMD64 processor family provides a high throughput and low latency x 86 solutions for next generation telecom platforms. What this means to Telecom Equipment Manufacturers is a platform that is low in complexity (no Northbridge) while at the same time able to scale up and out to meet the ever growing needs of next generation telecom systems.

AMD64 technology is unique because of the care and thought that went into its design and implementation. By integrating the memory controller on the chip, the AMD design team

provided a low latency path between external memory and the CPU. Many applications in the telecom space are data-intensive and have strict real-time requirements. Often, such applications aren't characterized by good memory locality so L2 cache misses will impact the overall performance and throughput, regardless of size. There are two ways to address this problem. One option is to include a very large on-chip cache. However, this approach significantly increases the gate count, which negatively impacts power consumption, die size and cost. Another approach is to design the processor for a broad range of data and I/O intensive applications by dramatically reducing the latency in accessing external I/O and DRAM. AMD chose the latter approach with AMD64 technology and Direct Connect Architecture.

The Integrated Memory Controller

With legacy processor architectures, external memory is addressed though a dedicated external memory controller which in turn, interfaces to the processor via high-speed interface called the Front Side Bus or FSB (the FSB is also known as the Memory bus or System bus). The FSB connects the processor with the main memory and I/O via a device called a Northbridge or Memory Controller Hub (MCH). The MCH contains the DRAM memory controller and also interfaces to the PCI bus and the Southbridge or I/O Hub. In this legacy topology, as CPU speeds increase, one way to reduce the effects of the FSB bottleneck is to eliminate it, by integrating the memory controller into the CPU. Of course, the frequency of the FSB can be increased, but not to the speeds of today's cutting-edge processors, whereas an integrated memory controller operates at the same frequency as the processor. Another issue is the inherent limitation of current FSBs when it comes to connecting PCI devices, especially the new faster versions such as PCIe™. The solution chosen by AMD designers was to link the I/O to a new interface called HyperTransport technology.

In the AMD64 processor family, there is a common internal interface between the integrated memory controller and the CPU. This interface, shown in Figure 2 as the System Request Queue (SRQ), was designed to scale with the AMD64 processor family's increasing system clock. The SRQ interfaces to the HyperTransport technology I/O busses, the internal memory Controller (MCT) and DRAM Controller (DCT), via a five-way, 72-bit wide crossbar switch. This crossbar switch is the heart of the integrated memory controller and I/O interface and it too scales with CPU clock speed. Hence, as the CPU clock rate goes up, the usable bandwidth of the interface scales proportionally. This is not the case with the legacy FSB-based architecture. As an example, a Dual-Core AMD Opteron processor Model 265 has a 16.2GB/s wide interface between the bussing unit and SRQ. The peak HyperTransport technology half-duplex payload bandwidth (per channel) is ~2.8GB/s (at 800MHz DDR), and the peak PC3200 DDR memory bandwidth is 6.4GB/s. Therefore, the peak aggregate half duplex loading of this interface is only 14.8GB/seconds which is well below the maximum bandwidth supported by the SRQ. In this example, there is more than enough system bandwidth to support the peak theoretical memory and I/O loading of the system bus; the bandwidth scales with CPU clock rate. In short, the AMD

Opteron processor has been carefully designed to support low latency, high bandwidth applications with ample bandwidth between the CPU and the I/O interface.

The HyperTransport™ Technology Bus

HyperTransport technology is a dual, uni-directional (full-duplex); point-to-point chip and board interconnect technology based on Low Voltage Differential Signaling (LVDS). HyperTransport technology is particularly appropriate as a system and I/O bus, offering low latency, high-bandwidth, and built-in error detection. It offers system designers a very compelling solution for today's high-speed, small form factor, low-powered embedded applications. It was originally developed by AMD but is now an open standard owned and maintained by the HyperTransport Technology Consortium (www.hypertransport.org) which has over 50 member companies. Inherent in the HyperTransport technology specification is its built-in scalability based on a range of frequencies (200Mhz to 1GHz) and support for multiple width lanes (2, 4, 8, 16 bit). Differential signaling results in very low current path to ground, thereby reducing the number of ground and power pins required for the interface. The flexibility inherent in the technology makes it suitable for both high-speed, high-bandwidth applications as well as for low-cost, lower bandwidth uses, all sharing the same base technology. HyperTransport technology is simple to implement and is well understood. It is further strengthened by the vast ecosystem of HyperTransport technology-enabled processors and support chips that are available from a number of sources. Recent additions to the HyperTransport technology specification include a connector definition which allows it to be used between modules (such as between an ATCA carrier card and its AMC). As we will show later, this flexibility allows designers to architect a wide-range of innovative solutions for the ATCA platform.

AMD's HyperTransport technology interface was designed to support both a proprietary coherent processor-to-processor interface as well as an open non-coherent I/O link protocol with the same physical interface. This allows for one set of pins to be used as either a non-coherent I/O interface to industry standard devices (PCI-X bridges, Southbridges, FPGAs and other silicon) or as a coherent high speed, low latency interface connecting multiple AMD Opteron processors. When used as a non-coherent interface for passing data between processors and peripherals, the HyperTransport technology interface maintains full software and operating system compatibility with PCI and CompactPCI, thus enabling investment protection and the ability to benefit from the large PCI ecosystem and knowledge-base. The AMD Opteron processor supports up to three HyperTransport technology links for linear scalability of memory size and I/O connectivity without sacrificing bandwidth.

The AMD Opteron™ Processor With Direct Connect Architecture

The HyperTransport technology interface provides a high bandwidth, low-latency link not only to I/O devices but also to adjacent processors enabling 2 to 8 processors to be connected in a variety of configurations. Such flexibility allows TEMs to design optimized solutions based on their specific compute, data, and I/O requirements while being able to choose from a wide-range of options at various price/performance points.

In figure 3, two multi processing (MP) topologies are compared. On the right is a 2P AMD Opteron processor system interconnected via its HyperTransport technology ports, while on the left is a typical 2P legacy system. As noted earlier, the same HyperTransport technology ports which provide a non-coherent interface to I/O devices can be used to provide a coherent multiprocessor interface to adjacent AMD Opteron processors.

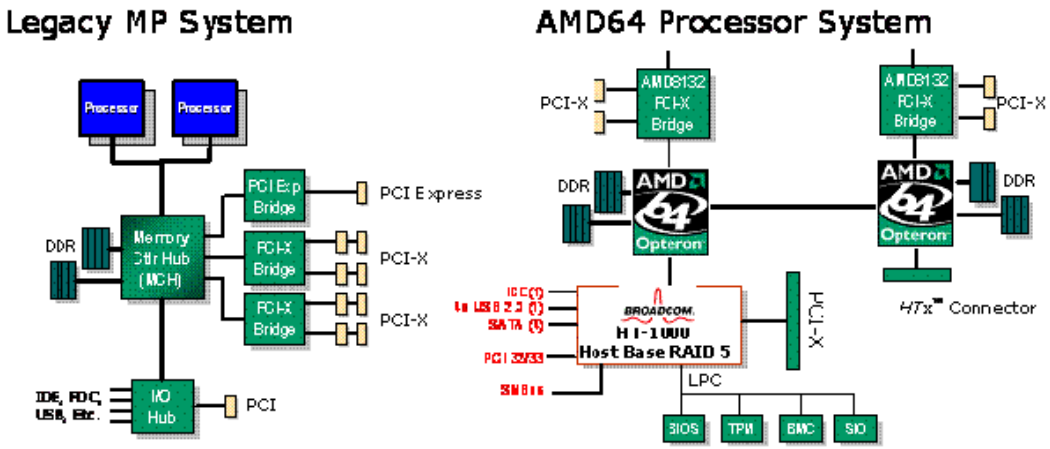

Figure 3 - 2P Legacy vs. AMD64 System

The AMD Opteron processor family is available in a number of models, all of which share the same architecture, come in a lidded 940-pin uPGA package and are available in a number of different configurations as noted in table 1.

Model Number	Dual Core	Revision	Frequency	Max Thermal Power	Max Tcase[1] °C	Cache Size (L2)	Package
x52		**Ex 90nm**	**2.6 GHz**	**95W**	**49-67°C**	**1M**	**940-pin oPGA**
x65	**X**	**Ex 90nm**	**1.8 GHz**	**95W**	**49-67°C**	**2*1M[2]**	**940-pin oPGA**
x48 HE		Ex 90nm	2.2 GHz	55W	78°C[3]	1M	940-pin oPGA
x65 HE	X	Ex 90nm	1.8 GHz	55W	83°C[3]	2*1M[2]	940-pin oPGA
x44 EE		Ex 90nm	1.8 GHz	30W	78°C[3]	1M	940-pin oPGA

Note(s):
1. Refer to the AMD Opteron™ Processor Power and Thermal Data Sheet, PID# 30417 for the latest product information.
2. Cache size reflected is "per processor" core.
3. Thermally enhanced processor

Table 1 - Embedded AMD Opteron™ Processor Product Family

As noted in figure 2, all the AMD Opteron processor models feature up to three HyperTransport technology links which can be configured as coherent or non-coherent interfaces, depending upon the model number. Coherent links are used for inter-processor communication while non-coherent links are used to connect to external devices such as PCI bridges (assigned during the enumeration phase of power-up reset). The three series differ only in the number of coherent links supported. All the links on the AMD Opteron 100 Series are non-coherent and can therefore be only used in a single-processor system. The AMD Opteron 200 Series has one coherent link and two non-coherent links, making it suitable for a dual-processor system. Finally, with the AMD Opteron 800 Series, all three HyperTransport technology links are coherent. This allows up to 8P systems to be created in different configurations each supporting a different balance of compute and I/O requirements. The HyperTransport technology links make it possible to design 8P topologies where CPUs are never more than two hops away from external memory, helping ensure minimal latencies. Of course, with AMD64 technology, each of the 8 possible processors implemented can be either single- or dual-core.

This diversity of processor models provides a new level of scalability currently unmatched by competing products. The fact that all AMD Opteron processor models under the AMD64 Longevity Program share the same technology and architecture, are thermal and pin-compatible, and are available as either single- or dual-core, allows system designers to use one platform across multiple possible configurations.

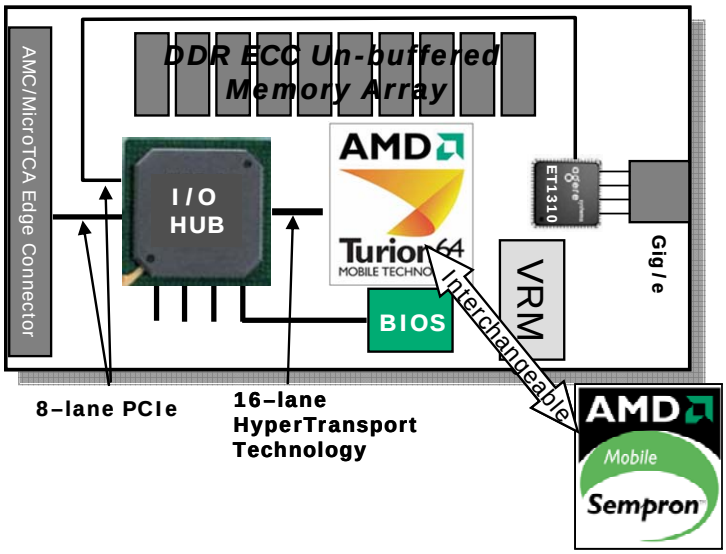

Figure 4 - AMC Based on AMD Turion™ 64 or Mobile AMD Sempron™ Processor

AMD Turion™ 64 and Mobile AMD Sempron™ Processors

Both the high-performance and value offering in AMD's mobile processor lines have the same rich feature-sets as the AMD Opteron processor (with exceptions as noted in the following table.) These processors were selected for the AMD64 Longevity Program - a high-end embedded roadmap - to facilitate lower power value platforms such as AMCs and MicroTCA.

Feature Set	AMD Opteron™ Processor	AMD Turion™ 64 Processor	Mobile AMD Sempron™ Processor
Architecture	AMD64 technology with Direct Connect Architecture	AMD64 technology with Direct Connect Architecture	Based on Direct Connect Architecture with 32-bit capability
L2 Cache	1MB/core	1MB	128KB
L1 Cache	128KB Data and Instruction Cache	128KB Data and Instruction cache	128KB Data and Instruction cache
HyperTransport Technology Links	Up to Three (3) 1GHz	One (1) 800MHz	One (1) 800MHz
Memory Interface	144-bit ECC Registered DDR	72-bit ECC un-Buffered	72-bit ECC un-Buffered
Peak Power	30, 55, 95 watts	25 watts @ 2GHz	25 watts @ 2GHz
Peak power at boot	30, 55, 95 watts	7.9 watts @ 800MHz	8.9 watts @ 800MHz
Package	940-pin Lidded	754-pin Un-lidded	754-pin Un-lidded

As shown in figure 4, due to the fact that there is no bulky traditional Northbridge to design around, both processors fit well into an AMC or MicroTCA form-factors. The CPUs come out of reset at very low power and are interchangeable, with only a BIOS change. Even at a peak power of 25 watts, they fall well within the power envelope of AMC and MicroTCA at 35 watts.

Summary

To conclude, the AMD Opteron processor provides a cost-effective high performance x86 solution. In an MP system, memory bandwidth and size scale with the number of processors while at the same time offering a significant reduction in access latencies. It provides a flexible, open architecture based on industry standard interfaces, which allows low-cost, small, and less complex designs while at the same time allowing for large headroom for growth. The AMD Opteron processor family is fully compatible with the x86 instruction set, provides industry-leading performance and comes in a wide range of power combinations. Based on low-power and cool operating technology, and with all the standard dynamic power saving features of AMD PowerNow!™ technology, the AMD Opteron processor provides an ideal platform for demanding, next generation telecom applications.

AMD Opteron™ Processor Example Applications

The AMD Opteron processor and HyperTransport technology low latency interface enable the logical extension of the processor core to include FPGAs and ASICs, which make it possible to develop ATCA systems without the need of an intermediate PCI bus. Two sample configurations are shown below. They demonstrate how the inherent flexibility of AMD64 technology with Direct Connect Architecture enables the design of low-latency, low-cost and highly differentiated ATCA blades and AMCs as compared with legacy solutions.

The first configuration, shown in figure 5, illustrates a modular ATCA system comprised of a single motherboard with two AMD Opteron processors, attached memory, their associated Voltage Regulation Modules (VRMs), three HyperTransport technology connectors and two eight-lane PCIe™ AMCs. All the I/O devices, including the Southbridge, are placed on mezzanine cards. This approach facilitates a design methodology that is both modular and flexible in several ways.

1. It is possible to populate the motherboard with two single- or dual-core processors; both have the same 55 watt power envelope, which is well within the ATCA 200 watt power budget. It is also possible to populate the ATCA board with only one processor without sacrificing I/O connections or compromising system I/O performance since the Southbridge and upper mezzanine card are still accessible by the one remaining processor. If the lower processor is not populated, access to the lower HyperTransport technology mezzanine card (holding the FPGAs) can easily be obtained by inserting a bridging socket. A bridging socket is a simple PCB of the same height and width as an AMD Opteron processor with a specific number of pins and copper PCB runs carefully placed to allow the HyperTransport technology interface to pass through to the lower mezzanine card. The bridging socket is available from Interconnect Systems, Inc. (www.isipkg.com).
2. By having all I/O on mezzanine cards, it is possible to have one ATCA Stock Keeping Unit (SKU) for multiple compute functions. Adding new I/O or offload compute functions only involves the design of another mezzanine card. This modularity of design applies to the

Southbridge mezzanine card as well. For example, it is possible to replace the nVIDIA Southbridge mezzanine card with a next generation Southbridge, (the BIOS is also on the same card).

3. Finally, this modular approach makes it possible to rapidly deploy custom platforms designed to address telecom, storage, security and/or imaging applications with nothing more than the addition of a mezzanine card, be it a PCIe AMC or a HyperTransport technology mezzanine card.

Also shown in Figure 5 is a proposed FPGA module. The first FPGA is statically programmed at reset to implement the HyperTransport technology protocol stack interface and a low latency high-speed LVDS interface to the second FPGA, which is dynamically programmed by the host processor. A 16-lane HyperTransport technology interface on the FPGA can support a 600MHz DDR connection (a 17 Gbit/s half duplex full payload interface). The latency to the FPGA is in the order of 400ns.

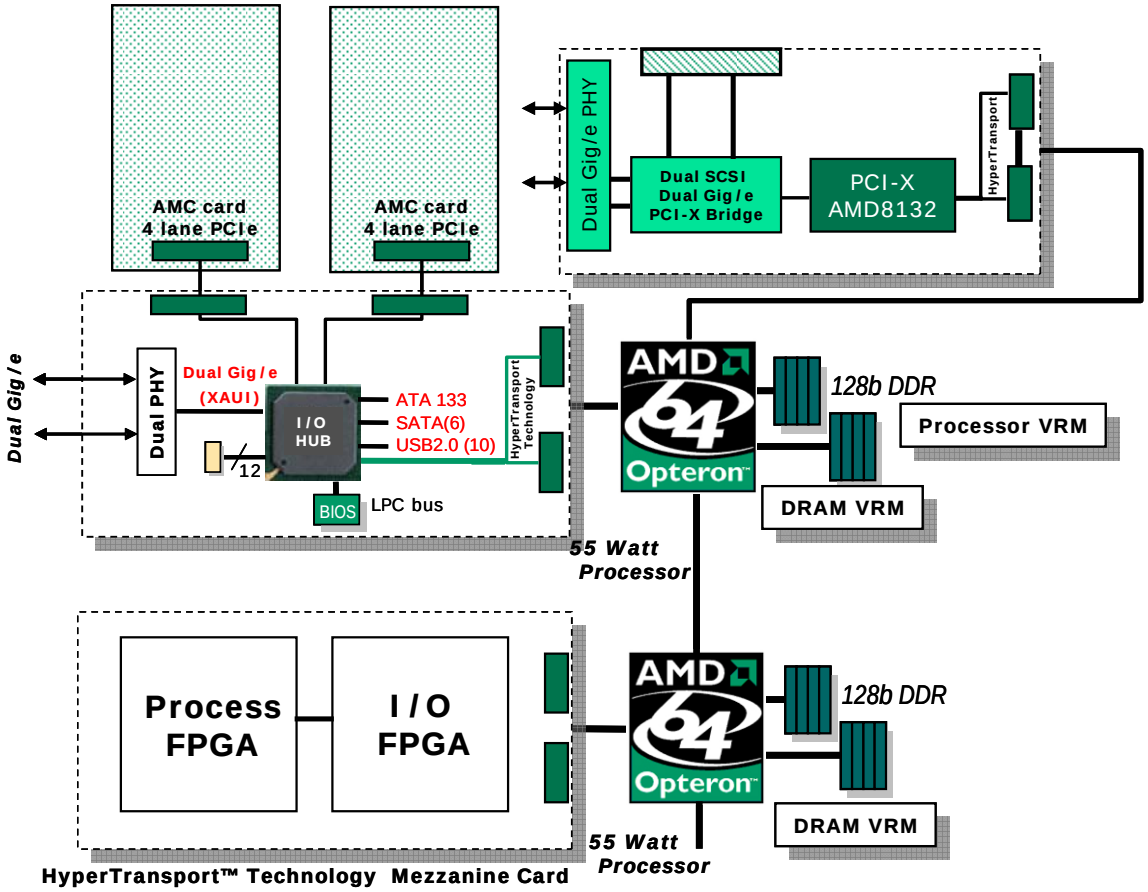

Figure 5: A Modular AMD Opteron™ Processor-based ATCA Blade

Of unique interest to the telecom industry is the recent announcement by Pinnacle Data Systems (PDSi) and AMD on the availability of an AMD Opteron processor-based, ATCA form factor reference board. The TS260 board (shown in figure 6) is NEBS compliant, supports 1 or 2 single- or Dual-Core AMD Opteron processors, two hard drives, and up to 16 GB of memory. It can be used with multiple operating systems. Two four-lane PCIe controllers from Broadcom (ServerWorks HT1000/2000) interface to the two AMC on-board connectors. In terms of available power, this high-performance platform still leaves a lot of headroom. For example, the two high-performance processors only require 55 watts each. If one adds the power requirements for the DRAM (two banks of 128MB DIMMs require about 10 watts each), the total is still only 130 watts. The power budget for a full-size ATCA board is 200 watts. About 70 watts remain available for miscellaneous logic and for future growth. An interesting variation of the PDSi board for applications that require less computing power and more I/O can easily be obtained by removing one of the processors (and its attached memory and VRMs). For more information on the AMD/PDSi ATCA blade, please visit:
http://www.pinnacle.com/products2/advancedtca/bladerdk/

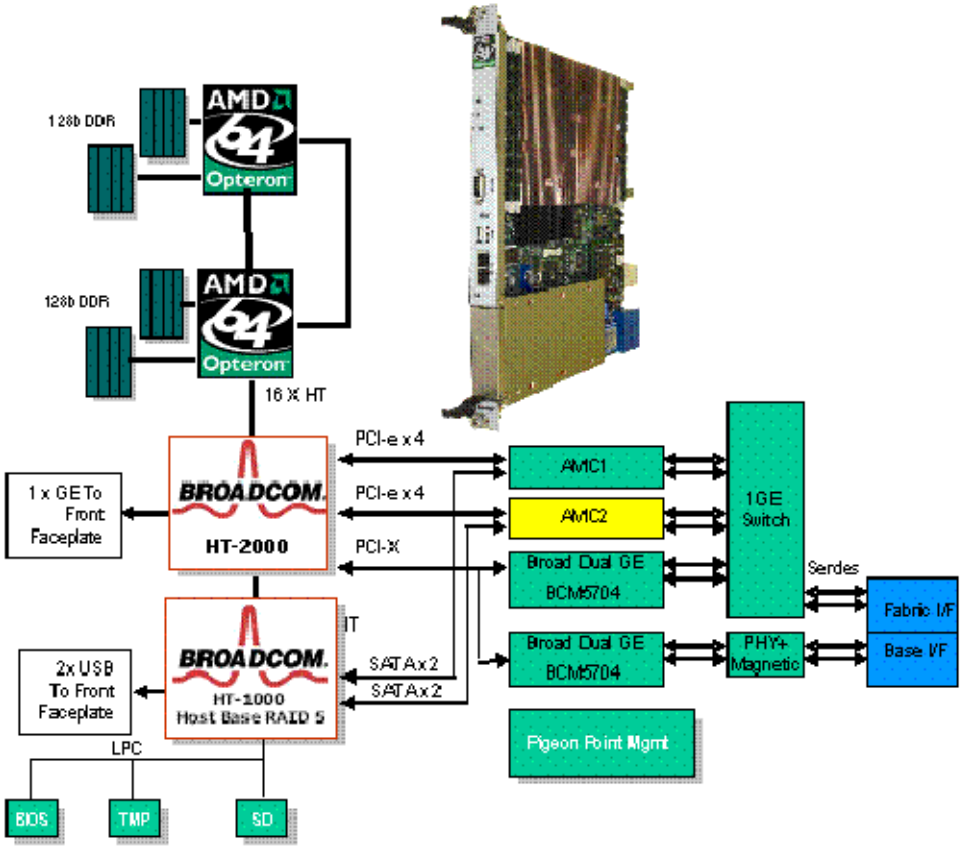

Figure 6: AMD/PDSi 2P ATCA Blade Using Broadcom Chipset

The AMD64 Ecosystem

The AMD64 ecosystem is extensive and growing. AMD is engaged with over 300 providers, including ISVs and IHVs to offer system designers with the tools required to build a successful product. In addition, designs based on AMD64 technology can immediately take advantage of the wealth of hardware and software options already available in the established x86 ecosystem. The embedded AMD Opteron processor line is available today for high-end telecom designs and is derived from standard server and workstation product lines. This provides investment protection for TEMs who will be able to benefit from future developments in the AMD product line, maintaining upward compatibility while leveraging existing infrastructure.

The AMD64 product family includes supporting chips such as a PCIe bridge (AMD-8132). In addition, a large selection of supporting chip-sets and single board computers is available from AMD partners such as Broadcom, nVIDIA, SiS, QLogic (PathScale), Xilinx, NetLogic Microsystems and others. These include Southbridges, PCIe/PCI-X bridges, HyperTransport technology connectors, InfiniBand bridges, classification processors and chipsets supporting the various flavors of the ATCA fabric interface. Equally valuable for board designers is AMD's association with leading BIOS vendors such as Phoenix, American Megatrends, and the open Linux BIOS, and carrier-grade OS and real-time kernels from Wind River, FSMLabs, Montavista and LynuxWorks. In addition, a large library of development tools and drivers optimized for the AMD Opteron processor is available directly from AMD and other vendors.

AMD recognizes the unique requirements of TEMs in high-end embedded markets. To accommodate their needs, AMD has established an extensive support infrastructure which includes the availability of reference boards for various target applications, Reference Design Kits, and tailorable Professional Design Support Services, providing advice and guidelines on boards and chip selection, thermal and layout issues, DRAM trade-offs and more.

AMD offers its customers a powerful simulation tool. SimNow™ is a fast and configurable x86-32/64 dynamically-translating instruction-level platform simulator. SimNow allows users to connect complex software models to emulate a fully-functional platform environment. It is also used by AMD engineering to emulate all AMD64 processors. SimNow runs on all AMD64 based platforms under commercially available x86-64 operating systems at a rate of 1/10 actual speed.

AMD has publicly released its Open Platform Management Architecture (OPMA) specification and is now working on the embedded version of OPMA (eOPMA) together with the VITA organization. OPMA defines a common hardware interface between the server platform and its management subsystem. OPMA allows AMD’s customers to build accessible and affordable management solutions.

The AMD64 Longevity Program is specifically valuable for embedded system OEMs. It offers a clear and stable roadmap as well a commitment to the availability of select AMD64 processors, including the AMD Opteron processor. This means OEMs can design with

confidence, secure in the knowledge that the designs they create today can remain viable for years to come.

The AMD Opteron™ Processor Value Proposition

The AMD Opteron processor has already been proven as a viable, high-performing option for the enterprise market. The processor's value proposition is derived from its unique architecture, performance and technical benefits as well as AMD's leadership position in the x86-64 market and a clear vision and commitment to AMD64 technology. AMD was the first to market with an x86-64 architecture (AMD64) and was again first with a native x86 dual-core product. AMD continues its push for performance, price/performance and performance-per-watt market leadership. Its past achievements and execution history is the best validation of its future road-map and commitment to the success of its customers' vision. AMD64 technology allows easy upgrades from single- to dual-core, open management and scalability. For example, usually only a BIOS change is required to move from a single-core AMD Opteron processor to a dual-core, or to switch to a higher frequency model since they share the same pin-out and thermal envelopes. Such flexibility coupled with AMD's ability to deliver volume NEBS-friendly products for up to seven years, helps Telecom Equipment Manufacturers have investment protection in their fast-changing and dynamic industry.

The choice of AMD Opteron processors with Direct Connect Architecture provides TEMs nearly unlimited flexibility in designing solutions for diverse needs, while offering fast core-to-core and core-to-memory, processor to processor, and board to board communication with little or no bottlenecks. The low-power AMD64 processors offer exceptional performance per-watt. Coupled with flexible power management schemes, these processors are particularly suitable for the typical rack-dense environment common in the telecom central office.

Several of the AMD Opteron processors available for embedded applications are NEBS/ETSI friendly and can operate in stringent environment conditions. Reliability, Availability and Serviceability (RAS) capabilities include: ECC Registered memory, Chip-Kill, advanced fault tolerant architecture, excellent MTBF and extensive hardware support for error monitoring and recovery. The built-in RAS features of the CPU make it possible to design products with the 24/7 availability required in the telecom space.

Finally the AMD Opteron processor benefits from strong ecosystem support. Leading chip, board, and systems manufacturers as well as software vendors support AMD64 technology. This helps ensure supply-chain stability and demonstrates AMD's commitment to the embedded markets in general and the telecom industry in particular.

Conclusion – The AMD Opteron™ Processor for High-End Telecom Systems

Over the past 10 years, the telecom industry has experienced major economic, structural and technological changes. The ATCA specification of chassis and boards was developed to facilitate the faster introduction of next-generation telecom products which need to have higher-performance, more functionality and must be less costly to build, implement and manage. With its open and scalable architecture and rich set of unique features, the AMD Opteron processor is an ideal solution for next generation high performance products. The AMD Opteron processor family features low-power, high performance processors with a scalable architecture, integrated memory controller(s), and the open HyperTransport technology. The AMD Opteron processor is a 64-bit processor which is fully x86 compatible, allowing TEMs to benefit from existing know-how and from the vast ecosystem of the x86 technology while concentrating their efforts on differentiating their offerings. AMD recognizes the unique requirements of TEMs in the telecom industry and offers extensive professional design support services which reduces R&D costs and accelerates time to market.

For more information on the embedded AMD64 Embedded products and programs, please visit www.amd.com/amd64embedded.

About AMD

AMD (NYSE:AMD) designs and produces innovative microprocessors and low-power processor solutions for the computer, communications and consumer electronics industries. AMD is dedicated to delivering standards-based, customer-focused solutions for technology users, ranging from enterprises and governments to individual consumers. For more information, visit www.amd.com.