

Six-Core AMD Opteron Processor Istanbul

Paul G. Howard, Ph.D.
Chief Scientist, Microway, Inc.

Copyright 2009 by Microway, Inc.

In April 2009 AMD™ announced the first 6-core Opteron processor, codenamed Istanbul¹; delivery began in June 2009, four months ahead of schedule. Istanbul is based on the AMD 64-bit K10 architecture, and is available for 2-, 4-, and 8-socket systems, with clock speeds ranging from 2.0 to 2.8 GHz. Istanbul is the successor to the quad-core Shanghai processor, launched in 2008. It provides up to 30 percent more performance in the same socket and with the same power draw.

Here is a summary of the features of the 6-core Istanbul processor:

- Six cores mean more performance, and also more performance per watt because the power and thermal envelopes are the same as for 4-core processors
- Consistent architecture, scalable to 2, 4, or 8 processor systems
- Energy efficient, incorporates AMD-P power management
- Virtualization support using AMD-V
- Improved bandwidth, through the use of HyperTransport 3.0 and HT Assist.

The current state of AMD Opteron processors

Istanbul is based on the AMD K10 64-bit architecture and manufactured on a 45 nm SOI process². Processors in the K10 line to date are shown in the following table.

Code name	Cores	Sockets	Frequency (GHz)	L3 cache (MB)	Process	Model number range ³	Introduction date
Barcelona	4	4 or 8	1.9-2.5	2	65 nm	8346-8360	September 2007
Barcelona	4	2	1.9-2.5	2	65 nm	2344-2360	September 2007
Budapest	4	1	2.1-2.5	2	65 nm	1352-1356	June 2008
Shanghai	4	4 or 8	2.2-3.1	6	45 nm	8374-8393	November 2008
Shanghai	4	2	2.1-3.1	6	45 nm	2372-2393	November 2008
Suzuka	4	1	2.5-2.9	6	45 nm	1381-1389	June 2009
Istanbul	6	4 or 8	2.1-2.8	6	45 nm	8425-8439	June 2009
Istanbul	6	2	2.0-2.8	6	45 nm	2423-2439	June 2009

Budapest uses the same core as Barcelona, and Suzuka uses the same core as Shanghai. Use of the 45 nm process reduces power consumption compared with the 65 nm process.

1 Like the other processors in the current AMD Opteron line, “Istanbul” is codenamed for a Formula 1 race track.

2 For details of the SOI (silicon-on-insulator) CMOS process, see Appendix A.

3 For details of AMD's processor numbering system, see Appendix B. Model number suffixes are omitted from this table.

The Istanbul product line

The AMD Opteron 2400 and 8400 series product line is shown in the following table.

Model number	Frequency	Sockets	TDP, W	ACP, W
8435	2.6 GHz	4 or 8	115	75
8431	2.4 GHz	4 or 8	115	75
2435	2.6 GHz	2	115	75
2431	2.4 GHz	2	115	75
2427	2.2 GHz	2	115	75
8439 SE	2.8 GHz	4 or 8	137	105
2439 SE	2.8 GHz	2	137	105
8425 HE	2.1 GHz	4 or 8	79	55
2425 HE	2.1 GHz	2	79	55
2423 HE	2.0 GHz	2	79	55

The five standard chips (no SE or HE suffix) were released in June 2009. The two SE chips and the three HE chips were released on July 13, 2009.

AMD provides two power metrics: TDP (Thermal Design Power), the maximum power the cooling system is required to dissipate, and ACP (Average CPU Power). ACP is more relevant for users: it provides a estimate of actual power consumption under real-world operating conditions.

Istanbul processors are built on a 346 mm² die, and contain 904 million transistors. By comparison, Shanghai has 758 million transistors on a 258 mm² die, and Intel's quad-core Nehalem processors (also manufactured on a 45 nm process) have 731 million transistors on a 263 mm² die.

The K10 architecture

The K10 architecture made its debut with the introduction of Barcelona in September 2007⁴. K10 is a 64-bit architecture with native quad-core support. It supports the full AMD64 instruction set, a number of SIMD instructions for both integer and floating point operations (MMX, SSE, SSE2, SSE3, SSE4a, and 3DNow!), ABM⁵ instructions, and the NX bit⁶. The 16 integer registers are each 64 bits wide, and the 16 SSE registers are 128 bits wide. The technology supports 48-bit addresses. Other features are shared L3 cache, 128-bit floating point units, AMD-V Nested Paging virtualization, HyperTransport 3.0, and HT Assist.

K10 cores have 9 execution units: 3 integer ALUs (arithmetic logical units), 3 AGUs (address generation

4 The K10 architecture is officially called "Family 10h Processors", and also sometimes (confusingly) called "Barcelona" in the press.

5 There are two ABM (advanced bit manipulation) instructions, POPCNT and LZCNT, which count the number of 1-bits and the number of leading 0-bits in a word. These instructions are useful in pattern matching.

6 Also called "Enhanced Virus Protection". The NX ("No execute") bit marks memory as being code or data, and does not allow execution of data. This feature requires use of the PAE (Physical Address Extension) page table format, adding a little overhead.

units), and 3 floating point execution units that also handle SIMD instructions. K10 has a 12-stage integer pipeline and a 17-stage floating point pipeline. K10 has 32-byte instruction fetch blocks, improved conditional branch prediction, and out-of-order execution algorithms.

The K10 architecture uses Socket F (1207), an organic lidded LGA (land grid array)⁷ interface with 1207 pin contacts⁸. AMD has been using LGA technology for server platforms since 2006. In the LGA interface, there are no pins on the chip, only gold-plated copper pads that contact pins in a socket on the motherboard. This design allows higher pin densities than the older PGA (pin grid array) technology, hence more power contacts, leading to a more stable power supply.

Integrated memory controller

Each Opteron processor has its own memory controller. The memory controller is technically 144 bits wide, but only 128 bits are data bits; the rest are for error correction (double-bit detection, single-bit correction). The memory controller consists of two independent (“unganged”) channels that access 64-bit memory (up to 8 DIMMs, registered DDR2 SDRAM) with prefetch support.

HyperTransport Technology

Communication between Opteron processors and from an Opteron processor to I/O buses takes place over HyperTransport links. In effect, the integrated memory controller and the HyperTransport buses, collectively referred to as “Direct Connect Architecture,” together replace the front side bus. Each Istanbul processor has 3 HyperTransport 3.0 links. In a 4-processor system, two of each processor's HyperTransport links are connected to other processors, and the third is used for I/O. In an 8-processor system, four of the processors are connected only to other processors; the other four use one link each for I/O and two to connect to other processors.

A HyperTransport link is an LVDS (low voltage differential signal) link. It is double-pumped: data is sent on both rising and falling edges of the clock. The HyperTransport clock used in Istanbul processors runs at 2.4 GHz, so the transfer rate is 4.8 GT/s (4.8 billion transfers per second). In Opteron processors, data is sent in 16-bit packets⁹. The peak aggregate bandwidth for one processor is thus 3 links x 2 directions x 4.8 GT/s x 2 bytes/transfer, or 57.6 GB/s.

Cache

All processors using the K10 architecture have a 3-level cache structure, which AMD calls Balanced Smart Cache. Each of the six cores has a 64 KB L1 (Level 1) data cache for each core and a 64 KB L1 instruction cache. The L1 caches are 2-way set associative¹⁰. The data caches support two 64-bit operations per cycle, with a latency of 3 cycles. The instruction caches have advanced branch prediction.

Each core has a 512 KB L2 cache, 16-way set associative, with a latency of 9 cycles. The L2 cache is a “victim cache”: the only data in it is data that was evicted from the L1 cache to make room for other data.

7 In organic LGA technology, some organic material is mixed with copper, providing better thermal expansion characteristics than plastic LGA, the other type of LGA.

8 35 rows x 35 columns, land pitch (i.e., center-to-center distance) 1.10 mm.

9 The spec for HyperTransport 3.0 allows a maximum clock speed of 2.6 GHz and up to 32-bit packets.

10 Set associative cache involves a tradeoff between fully associative cache, which is fast but requires complex hardware logic to search many slots at once, and direct mapped cache, which is simpler to implement but results in more collisions and hence more cache misses. The more “ways” in a set-associative cache, the closer you are to fully associative.

Istanbul's L3 cache is a 6 MB exclusive (i.e., does not mirror L1 and L2 caches) shared cache with sharing prediction. It is 32-way set associative, with a latency of about 45 cycles. It too is a victim cache, containing only data evicted from L2 cache. All the caches are made up of ECC protected memory¹¹ except the L1 instruction caches, which have parity-protected memory.

HT Assist is a technology to speed up inter-core cache accesses. When HT Assist is active, 1 MB of each processor's L3 cache is reserved as an index of the data stored in the L1 and L2 caches of its cores. Cache probes can then query the index instead of being broadcast to all the cores, reducing latency and reducing traffic on the HyperTransport bus. HT Assist makes sense only on systems with four or more processors; it is disabled on 2-processor systems.

Power management

AMD-P technologies built into AMD Opteron processors feature a variety of power optimization technologies:

- AMD Smart Fetch Technology enables inactive cores to write contents of their L1 and L2 caches to the shared L3 cache. Then the inactive cores can enter a halt state.
- Enhanced AMD PowerNow!TM with Independent Dynamic Core Technology allows different clock frequencies for each core depending on application performance requirements.
- AMD CoolCoreTM Technology can cut off power to unused parts of the processor, including L3 cache.
- Dual Dynamic Power ManagementTM provides an independent power supply to the cores and to the memory controller, allowing the cores and memory controller to operate on different voltages.
- AMD PowerCap manager gives the ability to control the maximum power and voltage of each core via a BIOS setting.
- ACPI (Advance Configuration and Power Interface) compliance.
- Two system power states: C0 (operating) and C1 (halt).
- Five processor power states: S0 (working), S1 (caches flushed), S3 (standby), S4 (hibernation), and S5 (soft off).

Chipsets

All chipsets that support quad-core Opteron processors, dating back to August 2006, also support 6-core Opteron processors. Current AMD server platforms use chipsets from NVIDIA and Broadcom. Both NVIDIA and Broadcom will support Socket F platforms through 2010, including systems based on Istanbul processors.

Future platforms, architecture, and processors

Platforms. In the second half of 2009, AMD expects to release the new Fiorano¹² server platform for the Istanbul and Shanghai processors. Fiorano will be a Socket F platform with the AMD SR5690/SP5100 chipset¹³, and will feature support for HyperTransport 3, PCI Express Gen 2, and IOMMU virtualization

11 ECC (error correcting coding) is a sophisticated (and costly in terms of bits) form of parity checking that allows correction of some memory errors. Ordinary parity codes can detect single bit errors but not correct them.

12 AMD server platforms (except Kroner) are named after Ferrari facilities and race tracks.

13 Based on mutual business decisions between AMD and NVIDIA and Broadcom, NVIDIA and Broadcom will not be producing new server chipsets in 2010.

for memory accesses and interrupts performed by devices. It will support up to 48 cores. The low power version of Fiorano, codenamed Kroner, will be a twin server design, with two 2-socket motherboards in one chassis. Kroner will feature AMD's APM (Advanced Power Management) technology, and will not support standard or SE parts.

In 2010 there will be a platform split. The high-end Maranello platform for 2- and 4-socket systems will feature a new G34 socket and the AMD SR5690/SP5100 chipset. It will support 4 channels of DDR3 SDRAM, with up to 12 DIMMs per socket. The low-end San Marino platform for 1- and 2-socket systems will have a new C32 socket, the AMD SR5670/SP5100 chipset, 2 channels of DDR3 memory, and up to 4 DIMMs per socket.

Architecture. The new Bulldozer core architecture for processors with 4 or more cores is planned for 2011. It will include the SSE5 instruction set, and will be targeted at the 10 to 100 watt power range.

Processors. In 2010 AMD plans on bringing out a new processor on the Maranello platform, codenamed Magny-Cours. Magny-Cours is an 8- and 12-core processor design based on an MCM (multi-chip module) design, comprising two dies in one package. It will support DDR3 memory. In addition, there will be a new 4- and 6-core processor, codenamed Lisbon, for the San Marino platform.

In 2011 AMD plans to ship the Interlagos 12- and 16-core processor, based on the Bulldozer core and manufactured on 32nm process technology, for the Maranello platform. The 6- and 8-core Valencia processor, also on a 32nm process, is planned for shipment in 2011 on the San Marino platform.

Conclusion

The new 6-core AMD Opteron processor Istanbul is an evolutionary step in the Opteron server line. It provides a significant performance improvement, up to 30 percent, over the 4-core Opteron (Shanghai), with the same power consumption. Istanbul is fully compatible with other Opteron models. Especially for heavily threaded applications that can make use of lots of cores, upgrading to Istanbul makes sense.

Appendix A: SOI CMOS process technology

CMOS processes are based either on bulk silicon or silicon-on-insulator (SOI). In bulk silicon CMOS processes, devices are built directly on silicon wafers. In SOI technology, a thin insulating layer of silicon dioxide is buried in the silicon substrate; devices are built in a thin silicon layer on top of the buried insulator. SOI provides better electrical and thermal isolation of the devices; this prevents latchup and reduces parasitic capacitance effects, thus improving power consumption. SOI processes have fewer processing steps, and are more tolerant of defects, so yields can be higher. Device density can be higher than in bulk CMOS.

Why doesn't everyone use SOI? Cost. SOI wafers cost about 4 times bulk silicon wafers. This translates to roughly a 10 to 15 percent increase in the total manufacturing cost, although for 32 nm or 22 nm processes the difference will be less after factoring in the higher device density and increased yield. AMD has been using SOI technology for CPUs since 2001, but they still use bulk CMOS for GPUs and chipsets.

Appendix B: AMD's processor numbering system

Current AMD Opteron processors have 4-digit model numbers, sometimes followed by a power-efficiency rating suffix. Suppose the model number is of the form "SGFF".

The first digit “S” represents the number of sockets in systems that the processor is designed for: S=1 means single-socket systems, S=2 means dual-socket systems, and S=8 means 4- or 8-socket systems.

The second digit “G” is the processor generation: G=2 means dual-core DDR2, G=3 means quad-core DDR2, and G=4 means six-core.

The last two digits “FF” represent the feature set of the processor in question, particularly the CMOS process technology and clock frequency. The digits don't have a direct relationship with the clock speed, but for processors in the same generation with the same process technology, higher numbers mean higher frequencies.

If there is no suffix, the processor has the standard power consumption characteristics. A chip with the “SE” suffix (special edition) is performance optimized, and typically has a higher TDP. An “HE” suffix (high efficiency) means that the chip has a lower TDP than the standard edition, and an “EE” suffix (energy efficiency) means the chip has the lowest TDP in the family.

Disclaimer

The information presented in this white paper is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors.

Trademarks

AMD, AMD CoolCore, AMD Opteron, AMD PowerNow!, AMD-V, and combinations thereof are trademarks of Advanced Micro Devices, Inc. HyperTransport is a trademark of the HyperTransport Consortium. Other names are for informational purposes only and may be trademarks of their respective owners.