

AMD Shows Big Server Plans

Plans Include High-Speed MP Link, 64-Bit Extensions



by Linley Gwennap

With Athlon an established competitor in the performance PC segment, AMD is now eyeing the lucrative server-processor market as an opportunity to move even further upscale. At this month's Microprocessor Forum, AMD Vice President Fred Weber discussed some technologies the company plans to use to break into the penthouse suite. These include a high-bandwidth multiprocessor interconnect and 64-bit extensions planned for the next-generation processor now known as SledgeHammer, formerly called the K8.

Combined with the high-end features already in Athlon, these new technologies will give AMD a strong base from which to launch its server initiative. AMD's challenges in the server market go far beyond technology, however. The company has no experience delivering products in a market that has been dominated by powerful processors such as SPARC, Alpha, IBM's Power, and, more recently, Intel's Xeon line. With Intel's Itanium (Merced) looming on the horizon, AMD will have to work hard to gain design wins in this market.

Athlon Strong for Workstations

Next year, AMD plans to introduce new workstation and server products under the brand Athlon Ultra. It will take little work to position Athlon as a strong competitor to the Pentium III Xeon in x86-based workstations. At 700 MHz, the current Athlon delivers 31.7 SPECint95 (base), which ranks it near the top of all workstation processors, RISC or x86 (see MPR 10/25/99, p. 35). AMD expects this number to rise another 5–10% as it deploys faster system-bus speeds and improved compiler optimizations. Of course, this performance will rise dramatically, along with clock speed, as AMD moves Athlon from 0.25- to 0.18-micron manufacturing. The company is already sampling 0.18-micron parts at 800 MHz and expects Athlon to exceed 1.0 GHz sometime next year.

On the floating-point math used in many scientific and technical applications, Athlon is competitive with Xeon but not with the top RISC processors. At 24.0 SPECfp95 (base), the Athlon-700 lags the best Alpha, PA-RISC, SPARC, MIPS, and Power processors. AMD claims that faster bus speeds and compiler tuning will raise Athlon's FP score by nearly 50%, which would put it ahead of Xeon and in line with at

least some of the RISC chips, but this remains to be seen. Athlon workstations are also likely to cost less than competing RISC-based systems.

These scores show Athlon Ultra could be a solid workstation processor, particularly for applications where extreme floating-point performance is not required. AMD currently lacks a chip set for workstations, however. Both AMD's 750 (Irongate) chip set (see MPR 8/23/99, p.12) and Via's forthcoming Apollo KX133 support only one processor instead of two, a major shortcoming in the workstation market. Both deliver less than a third of the memory bandwidth of Intel's new 840 chip set (see MPR 10/25/99, p. 28).

Alpha Processor (API) is developing a dual-processor north bridge that will support either 21264 or Athlon processors in the Slot B configuration (see MPR 6/21/99, p. 19). The north bridge has two separate processor ports, as the EV6 bus used by both Athlon and the 21264 allows only one processor per port. This requirement drives the pin count of this single-chip north bridge above 1,000. The API chip also supports a 128-bit-wide 100-MHz DDR SDRAM memory system with enough bandwidth to match two 200-MHz EV6 buses. We expect it to include 4× AGP and PCI 66/64 as well. API expects to sample this chip set late this quarter.



MICHAEL MUSTACCHI

AMD VP Fred Weber discloses that the next-generation SledgeHammer will use a 64-bit x86 architecture.

More Work Needed for Servers

Most of today's servers are used for file serving, Web serving, transaction processing, and other integer-only applications. To succeed in these systems, a processor must start with a powerful integer core, a large, fast level-two (L2) cache, and a high-bandwidth bus to main memory. Athlon has all three. The processor's strong SPECint95 score certainly attests to its integer prowess.

The current Athlon modules come with 512K of half-speed external cache, fine for PCs but not for servers. AMD could expand this cache to 8M using commodity SRAM but at no more than half speed, whereas Xeon provides up to 2M of full-speed cache. Weber said AMD will introduce Athlon Ultra processors with 1M and 2M of full-speed cache in 2000. Due to this cache's speed and 16-way associativity, we expect it will be on the processor die. An on-die cache is faster and less expensive to build than external cache, but it will require AMD to deploy a new piece of silicon before entering the server market. Such a product is unlikely to ship before 2H00.

In its fast system (frontside) bus, Athlon has a large potential advantage over Xeon. Using point-to-point connections, this bus currently operates at 200 MHz (100-MHz DDR), twice the speed of the bus in Xeon servers, and Weber disclosed plans to push the bus to 266 MHz next year. The problem, again, is chip sets. Irongate's memory system runs at 100 MHz, supplying only half the potential bandwidth of the current bus. API's forthcoming dual-processor chip set takes better advantage of the fast bus.

AMD must also be concerned about more than just today's Xeon. Itanium, due in 2H00, will push bus bandwidth above 2.0 GBytes/s, and Foster (see MPR 10/26/98, p. 16) will reach 3.2 GBytes/s about six months later. Even at 266 MHz, the Athlon bus delivers 2.1 GBytes/s. The Alpha 21264, which uses the same bus as Athlon, is already cranking that bus at 2.7 GBytes/s, and other RISC processors will be topping 2.0 GBytes/s next year. By the time it appears in 2H00, Athlon Ultra will be competitive in bus bandwidth, but it will not have much of an advantage.

I/O Fast As Lightning

Servers require fast I/O as well as fast memory. To solve this problem, Weber disclosed a new high-speed interconnect dubbed Lightning Data Transport (LDT) that will first appear in products next year. Like the Athlon system bus, LDT achieves high transfer speeds using point-to-point signals. A single LDT link has two unidirectional buses that can each be one, two, or four bytes wide, each operating at 1.6 billion transfers per second. This transfer rate is achieved using both edges of an 800-MHz clock. The total raw bandwidth of a bidirectional 32-bit LDT link is 12.8 GBytes/s.

LDT data is enclosed in packets. The headers indicate one of several logical channels, allowing several data streams to use the link at once. The headers also encode special functions, such as PCI bus interrupts, providing support for legacy I/O and system-management functions. Weber did not disclose any details of how these mechanisms will work.

Although LDT's point-to-point design connects only two devices, additional devices can be connected in a chain. This provides flexibility; for example, a system could include one or more PCI bridges simply by chaining them together. Bridges to Gigabit Ethernet or the forthcoming System I/O (see MPR 9/13/99, p. 4) would also be useful.

LDT bears a strong resemblance to the HRC interface being developed by HotRail (see MPR 7/12/99, p. 12) to solve the same problem. HotRail CEO Rick Shriner said his company's initial chip design is too far along to adopt LDT, but he hopes to put LDT in future products. The current LDT specification defines bus protocols but not the physical layer. If AMD agreed to adopt the HRC physical layer, it would be simple to merge the two interfaces.

Both AMD and HotRail hope third parties will develop I/O bridge chips for their high-speed interfaces, and unifying the two is essential to gaining this support. Since HotRail

has been a long-time partner of AMD, we expect the two companies to work together to unify HRC and LDT.

Building Multiprocessor Systems

Servers must also scale to large numbers of processors. Figure 1 shows how AMD expects OEMs to use LDT to combine several north bridges, each with two processors and local DRAM, into a single large system. This design—which resembles the meshes of processors that will be created from Compaq's forthcoming 21364 and IBM's future Power4 (see MPR 10/6/99, p. 11)—uses the LDT links to access memory from remote processor pairs and maintain coherency across the system. Assuming a low latency across each LDT link, the access time to remote memory in the eight-processor configuration shown should be reasonable, avoiding the need for NUMA software optimizations.

A big advantage of this design is that memory bandwidth grows with the number of processors in the system. Every pair of processors gets the full bandwidth of its local memory, so a system with four processors has twice the bandwidth of a dual-processor system, an eight-way system has four times the bandwidth, and so on. I/O bandwidth can also scale if each north bridge includes an LDT link to I/O devices. SMP systems using chip sets such as Xeon's Profusion or Itanium's 460GX must share a single memory subsystem and a single I/O subsystem.

API is developing a chip set to support up to eight Slot B processors that it expects to deploy in 2001. This chip set is likely to embody the design Weber presented, connecting API's dual-processor north bridges using LDT links. If so, the total memory bandwidth in an eight-way system is likely to be at least 12.8 GBytes/s, assuming each north bridge uses a 128-bit DDR SDRAM memory subsystem or a pair of Rambus channels.

This type of MP interconnect has the advantage that it requires no changes to the processor, unlike the 21364 or

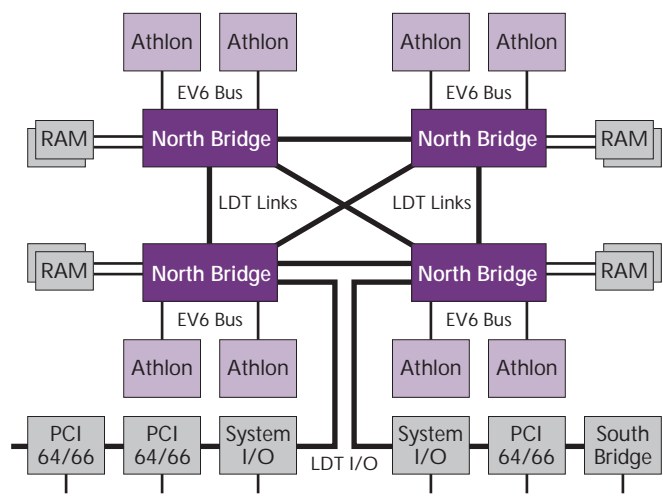


Figure 1. A hypothetical eight-CPU Athlon system could connect four dual-processor north bridges using high-bandwidth LDT links. Additional LDT links connect to daisy-chained I/O bridges.

Power4, and the north bridge can be derived from the basic two-CPU version. The disadvantage is that the CPU-to-DRAM latency is likely to be higher than in processors, such as the 21364 and Power4, that connect directly to main memory. Moving much beyond eight processors will require adding a routing protocol to each north bridge, as it would quickly become impractical to fully interconnect them all as they are connected in Figure 1.

HotRail is developing an alternative chip set for eight-way Athlon systems. This chip set implements a single-chip multiport switching fabric that connects up to eight Athlon processors with memory and I/O. HotRail implements a more traditional shared-memory subsystem with a bandwidth of up to 12.8 GBytes/s. This chip set, like API's, is about a year from production; HotRail has not announced a target availability date.

Ultimately, AMD plans to put multiple processors on a single chip to improve MP scalability. This approach is similar to that of IBM's Power4 and should provide a big boost in server performance. AMD has not disclosed whether the first SledgeHammer chip will include one or two cores. We expect the company to start in 2001 with a single-core version for the high-end PC market, with a dual-core model appearing in 2002. In any case, the plan for multiprocessor chips underscores AMD's commitment to reaching the high end of the server market.

SledgeHammer Adds x86-64 Extensions

Weber, the chief architect of SledgeHammer, disclosed that this next-generation processor will extend the x86 architecture to 64 bits when that chip appears sometime in 2001, solving a major weakness of x86 in big servers. A small but growing number of applications see a sizable performance benefit from 64-bit addressing, which is implemented in all of the major RISC architectures. In part because of x86's 32-bit limit, Intel's Xeon has been most successful in servers with no more than four processors.

Intel's solution is to move its customers to IA-64, which implements a new 64-bit architecture. This approach requires an enormous investment in developing new IA-64 operating systems and applications; although Itanium and other IA-64 processors will execute x86 code, we estimate performance in this mode will be roughly half of full native-mode performance. Intel is probably the only processor vendor that has a chance to accomplish such a major software conversion, and it is spending hundreds of millions of dollars to drive it forward.

AMD sees simplicity in extending the x86 architecture while maintaining compatibility. Applications must still be recompiled to take advantage of the new mode, which AMD calls x86-64, but 32-bit applications can run at full speed on the same x86 core. This allows customers to buy SledgeHammer systems and run a full suite of 32-bit code to start, then later upgrade to a 64-bit operating system. In contrast, IA-64 customers must make the OS transition from the start.

Weber said these extensions will be straightforward and analogous to the earlier x86 transition from 16 bits to 32 bits. For example, what was originally the 16-bit AX register and is now the 32-bit EAX register will become the 64-bit LEAX register. The processor will allow both 32-bit and 64-bit code segments, with the latter supporting flat 64-bit addressing with no segment base or limit registers. Most instructions will continue to operate on 32 bits of data, with new instructions for manipulating 64-bit pointers. This design should minimize both code and data expansion.

More Extensions Open Bottlenecks

While x86-64 should solve the 64-bit addressing problem, more difficult is matching the myriad of other performance-oriented features in IA-64, such as its predication and speculative execution. SledgeHammer will include what Weber called "minor additions" to the base instruction set to handle speculative loads and other "specialized operations." These extensions could reduce the gap between x86-64 and IA-64 in performance, but the x86 instruction set would require massive reconstruction to match IA-64's much larger register file and full predication, for example.

The x86 floating-point architecture, with its small stack-addressed register file and no multiply-add instruction, is clearly a hopeless case. SledgeHammer will include a completely new floating-point instruction set called TFP, for technical floating point. TFP will be a RISC-like instruction set with a large flat register file (though not as large as IA-64's) and three-operand instructions. For compatibility, SledgeHammer will continue to support the old x86 FP instructions as well. Weber believes SledgeHammer will achieve leading-edge SPECfp results on applications that use the new TFP instruction set.

This strategy shows what Intel could have done were it not focused on creating a new instruction set. While x86-64 processors may not reach the same performance level as IA-64 processors, they might come close, particularly on applications that don't have enough instruction-level parallelism (ILP) for IA-64 to exploit. Applications that don't recompile for the "minor additions" in x86-64, however, are less likely to keep pace with IA-64's performance.

Head-to-Head With IA-64

AMD believes SledgeHammer will deliver pure x86 performance that approaches the native IA-64 performance of future Intel chips. If this is true, only applications that need 64-bit addressing or TFP will need to be recompiled, reducing the burden on ISVs and end users compared with Intel's IA-64 strategy.

We believe, however, that the gulf between today's x86 instruction set and the more-modern IA-64 design is too wide for AMD to bridge with an efficient processor implementation. If SledgeHammer requires a boost from tweaking the instruction set, performance-sensitive applications will need to be recompiled for maximum performance. Once

ISVs decide to recompile their x86 code, Intel's market power dictates that IA-64, not x86-64, will be the primary target.

AMD's advantage is that uncompiled x86 code should perform better on SledgeHammer than in IA-64's compatibility mode. This will help ISVs and end users that don't want to touch their code. For simple applications, however, this difference in performance doesn't matter. The question is how many performance-sensitive applications will be able to resist the IA-64 juggernaut.

Where 64-bit addressing or high-performance floating point is critical, vendors will need to recompile their applications to perform well on x86-64. AMD must convince these vendors to make the effort to support the new extensions. AMD's first foray in this regard was 3DNow, which has been adopted by Microsoft in DirectX and by many vendors of 3D games. 3DNow prospered in part because it reached the market before Intel's competing SSE extensions, but x86-64 will be a year behind IA-64. Adding to AMD's challenge, vendors of scientific and enterprise applications tend to be more conservative than makers of killer 3D games.

Before looking to application vendors, AMD must win at least some operating-system support. Linux is an obvious choice, as the company can do its own port. Linux is also popular among the free-spirited users that are most likely to adopt AMD processors. Microsoft has supported AMD extensions in the past but has not yet committed to adding x86-64 support to Windows 2000. Burdened by a plethora of platforms, other Unix vendors are unlikely to pick up x86-64.

Many Challenges in Server Market

AMD's biggest challenges in the server market have little to do with technology. First, the company must demonstrate that it can deliver a server-class processor. In addition to strong performance at a moderate cost, this processor must be reliable in large configurations. AMD has no experience in multiprocessor validation and will need help from its initial customers to handle this very complex problem. Intel itself has had problems with its Xeon line, discovering latent bugs after it has shipped various products. Server customers become furious when their expensive system goes down, taking their entire operation with it.

Design flaws aside, server processors must detect transient system errors and correct them if possible or accurately report the affected process if correction is impossible. Athlon fares well in this regard, with parity on all internal data structures and ECC on both the frontside and backside cache buses. It matches most of Itanium's reliability features, except for that chip's "data poisoning" ability (see MPR 10/6/99, p. 1).

At a time when nearly all major server vendors have their hands full with IA-64, AMD must convince them to design new motherboards to support Athlon, which is not plug compatible with any of Intel's processors. Compaq would be a good target customer, as Athlon is compatible with Compaq's 21264 motherboards, but no other major server vendor uses the EV6 bus. Since the design cycle of a

For More Information

AMD has not disclosed the price or availability of its future workstation and server processors. We expect the first Athlon Ultra products to appear in 2H00. AMD expects the first SledgeHammer products to ship in 2001. For more information on AMD's Athlon processor, access the Web at www.amd.com/athlon.

typical high-end server is two to three years, AMD's efforts will take some time to bear fruit.

Furthermore, server customers are a notoriously conservative lot, as choosing the wrong system can have huge business ramifications. Corporate buyers are still not buying AMD-based PCs; it will take even more work to convince them to buy AMD-based servers. AMD has made some headway recently in the small-business segment, and it is only a matter of time before Athlon shows up in corporate PCs. It will take even longer before AMD gains a presence in the corporate server market, but with the right products, the barriers should eventually fall.

Start Small, Work Up

Initially, Athlon servers are more likely to be successful among the small businesses and Internet service providers that have already adopted AMD-based PCs. Two- and four-slot Xeon systems dominate this segment, but AMD could get a foothold with a competitive processor and an off-the-shelf motherboard. Athlon is also likely to see some use in workstations, where reliability is less critical and price is more of a factor. The margins in these segments are not as good as in the premium server segment, but they are still better than in the PC segment.

To succeed here, AMD must deliver a strong product under the Athlon Ultra brand and position it against Intel's Xeon line, offering a price/performance advantage. AMD must also have chip sets and board designs available that allow turnkey deployment of workstations and servers, as OEMs will initially avoid making a significant design investment in a new and unproved platform. Intel will, at least initially, cut prices and otherwise encourage OEMs to stay in the fold, so AMD must be prepared for a long and difficult campaign.

Ultimately, AMD should be able to carve out a niche in the high-end market. But the company's plans to extend all the way to the top of the server market are ambitious and will require significant investment that the company, which just announced another quarterly loss exceeding \$100 million, can ill afford. AMD hopes that Athlon revenues from the PC segment will turn its financial picture around and fund development of future server products. If the company can pull off this trick, it will reap the benefits of being a full-service processor vendor. 