AMD Opteron in 32 and 64 bit test

Linux Sledgehammer

AMD's 64 bit Opteron processor is making inroads on the world of Linux servers. Linux Magazine and Tom's Hardware looked into the benefits of performing 64 bit processing, and benchmarked 32 bit applications on the Opteron. As a control system we also tested the Opteron against the Intel Xeon processor.

BY MIRKO DÖLLE AND TIMO HÖNIG

The Opteron really does live up to its code name, "Sledgehammer", and has taken the Linux server scene by storm. The first dual-processor boards for less than £1000 [1] became available only a few months after the CPU was released. But the question is, do you really need a 64 bit processor?

The 4 GByte addressable memory space restriction for 32 bit processors is increasingly becoming a bottleneck for servers. Even desktop computers have 128 or 256 MBytes RAM nowadays, and low-end workgroup servers often have upward of 1 GByte.

Current technologies provide very little leeway for enhancing 32 bit processors performance-wise and can no longer be expected to produce major advances in clock speeds, or further miniaturization of chip structures.

The 64 Bit Mission

There is an obvious trend towards 64 bit processors, although Windows is burdened with a legacy of 32 bit applications. Users of commercial software products have no alternative, but to rely on timely ports from their manufacturers.

To soften the impact of migrating software and hardware at the same time, both Intel and AMD have designed their 64 bit processors to support 32 bit programs.

Optimized Compilers and Distributions

We installed SuSE Linux Enterprise 8 on all of our Linux lab systems. This reflects the fact that many organizations insist on a certified distribution for mission-critical applications. Application support is a major factor in this context. And it often rules out recompiling the kernel with an optimized compiler to avoid infringing the certification rules.

Intel is at a disadvantage here, as the 32 bit support for SuSE Linux Enterprise 8 uses the GCC compiler and not Intel's optimized ICC. We used ICC version 7.1 to compile the benchmark program and provide reference figures, but the deviation from the GCC benchmarks proved to be negligible.

AMD avoided the issue of a separate optimized compiler by adding the Opteron optimizations to the mainstream GCC. SuSE enabled the Opteron optimizations to build Enterprise 8 for 64 Bit. Even the NUMA (Non-Uniform Memory Access) kernel that reached our labs in the middle of our test series was certified for the SuSE Enterprise distribution.

It is difficult to say what effect the Opteron optimizations built into the 64 bit kernel had on the test results. We used the (non-Opteron enhanced) SuSE Enterprise Linux 8 for 32 Bit kernel – the one we used on all our Non-Opteron machines in fact – for all of our 32 bit tests.

The Lab Systems

The 32 bit performance of 64 bit CPUs is particularly important at the start of the migration to 64 bit. The Windows camp has very little to offer in the line of 64 bit software at present, and 64 bit Intel versions of commercial Linux programs are also quite rare.

Intel's Itanium I performed quite poorly in our Linux lab tests when it was first launched. The Itanium was easily outpaced by contemporary Pentium and Athlon systems at the time. The reason for this was the emulation that the Itanium used for 32 bit code, which was obviously quite inefficient at the time. But it is hard to say if the Itanium 2 has improved on this, as Intel was unable to provide us with an Itanium 2 test machine on request, and instead supplied 2.8 GHz dual and quad CPU Xeon systems to take up the challenge.

AMD's Opteron can run 32 and 64 bit code natively and in parallel, and claims to be a Pentium 4, when queried by a 32 bit OS. This prompted us to run all our benchmarks both on 32 bit Linux and on the 64 bit kernel.

Both Opteron test systems were provided by AMD. The dual Opteron, that Linux Magazine tested in [3] had pre-series processors running at 1.3 GHz and 2 GBytes RAM. The quad Opteron system had 1.8 GHz series processors clocked at 1.8 GHz and 8 GBytes RAM.

Our other benchmark candidates were a Pentium 4 with a Canterwood core and PC800 RAMs running on an Asus P4C800 motherboard, and an Athlon XP 3000+ with DDR 333 RAMs on an Asus A7N8X motherboard.

Opteron Inside

The Opteron does not use a Northbridge as the memory controller is located on the CPU. This means that each Opteron has its own memory area. Memory access is handled by the crossbar (XBAR, Figure 1), which not only handles the data streams to the memory controller, but also to the CPU core (via the system request queue), and the three hypertransport ports. In other words, the crossbar is a kind of central corridor on the Opteron chip.

Non-Local Memory

The crossbar also handles the data exchanges between the other Opteron CPUs on a multiprocessor system. Each processor has a maximum of four local memory modules (two DDR 333 channels) and is linked to its immediate neighbor via a hypertransport channel. In other words, the hypertransport channels on a quad system build a ring. If the required data are not stored locally, but in remote memory, the local crossbar communicates with the remote processor and asks it to transfer the data. This request does not impact the remote processor, however, as the transfer is offloaded onto the crossbar.

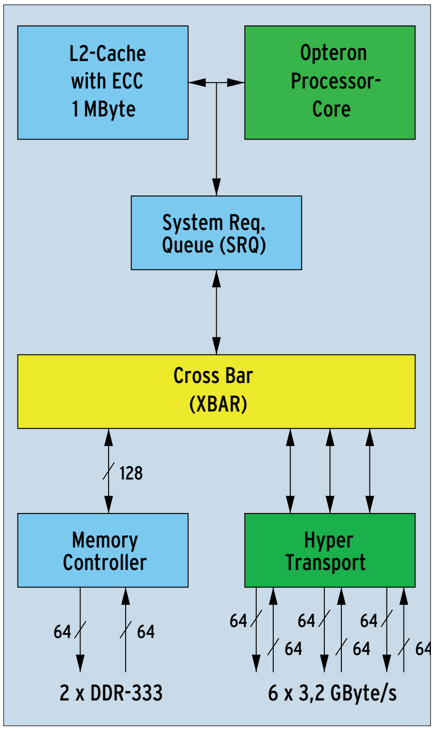

Figure 1: The crossbar is the nexus at the heart of the Opteron. It handles the data streams between the memory controller, the three hypertransport channels, and the CPU core (via the system request queue)

The hypertransport channels provide 3.2 GByte/s per channel and direction. A Pentium 4 with a 533 MHz frontside bus achieves a transfer speed of about 4 GByte/s, but only in one direction, whereas a hypertransport channel can achieve an amazing 6.5 GByte/s during full duplex operations.

In contrast to this, Intel uses a memory bus shared by all the processors on a multiprocessor Xeon system. The Northbridge, which takes care of memory management, can only serve one CPU at any given time.

Good Memory Link

The differences between AMD's non-uniform memory access, or NUMA for short, and Intel's memory bus immediately become apparent when you look at a multiprocessor system. In single or dual-processor mode (see Figure 2), the benchmark results are fairly level; as was to be expected, the Pentium 4 beat the other test candidates.

Quad CPU systems are a different issue, however. Our benchmark used four *stream_d* processes launched simultaneously with a problem size of 20 million array entries and 100 iterations, which is equivalent to memory usage of 457 MBytes per process.

As you can easily see in Figure 3, the four stream processes have access to a far greater memory bandwidth on the Opteron when compared to the equivalent processes on the Xeon system. The memory bandwidth a given application can use will depend on the parallelization capabilities of the individual processes, and the data distribution in memory – in a worst case scenario, the Opteron will need to retrieve this data from the remotest CPU.

D-Bench Hammer

All our benchmarks returned better results in the Opteron's 64 bit mode than its 32 bit mode. Having said that, the

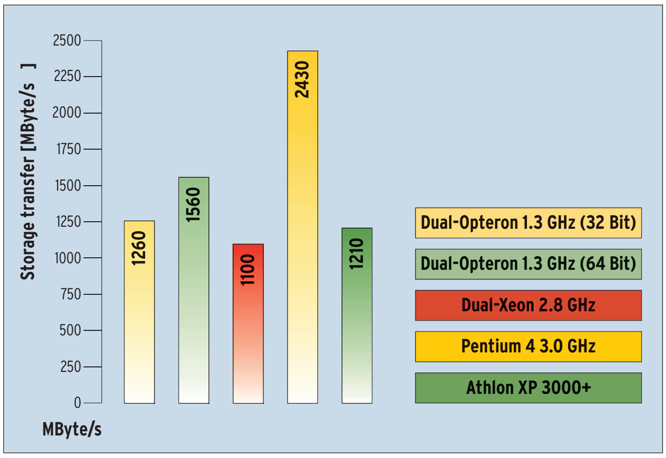

Figure 2: The memory benchmark champion is Intel's Pentium 4 with its Canterwood core and PC800 memory. The Opteron came in second

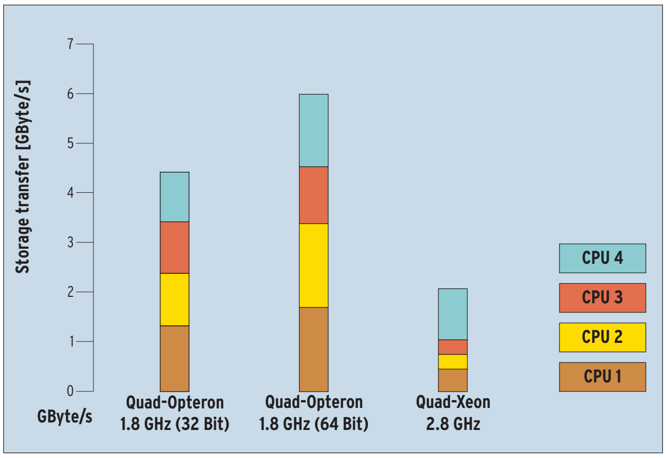

Figure 3: The Quad CPU Opteron with its NUMA memory architecture really shines when handling four parallel memory benchmarks

processor still beat its competitors with one hand tied behind its back. The multiprocessor Opteron negotiated the D-Bench benchmark particularly well (see Figure 4). This emulates SMB clients accessing a Samba share. The transfer rate was around 1.3 GByte/s, dropping to 1.2 GByte/s with 100 simulated clients, and remaining constant at that level until reaching the maximum load with 256 clients.

The Quad Opteron achieved far lower transfer rates per client in 32 bit mode, but still won hands down against its competitors. The D-Bench benchmark caused the Quad Xeon no end of trouble, with the transfer rate per client plummeting to less than 10 MByte/s for 100 simultaneous clients.

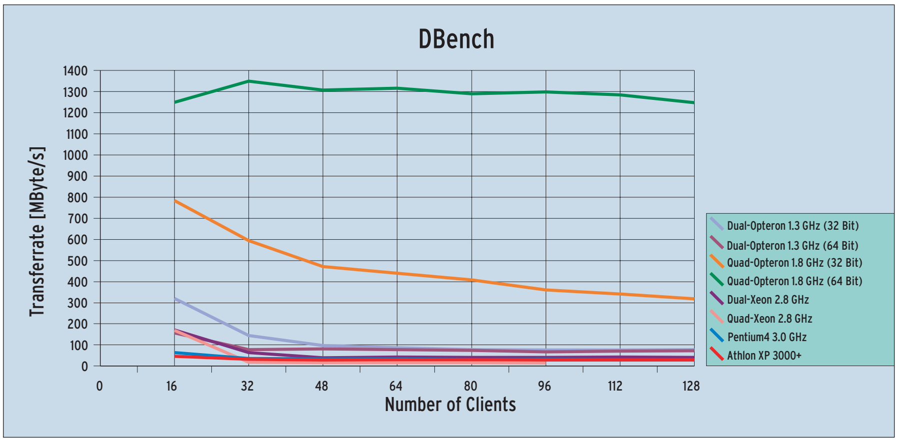

Figure 4: AMD's Quad Opteron really overachieved when performing the D-Bench, and serve up 1.3 GByte/s to a 100 clients. In contrast the Quad Xeon dropped drastically to less than 10 MByte/s when faced with this load, whereas the other systems still achieved rates of around 30 to 40 MByte/s

Good T-Bench Results

The results returned by the T-Bench benchmark, which simulates Samba network and socket I/O, were closer but still quite clear (see Figure 5). The Quad Xeon achieved roughly the same transfer rate as the Dual Opteron in 32 bit mode, and both left the Dual Xeon slightly behind. The T-Bench also clearly indicates that the Opteron is far quicker in 64 bit than in 32 bit mode.

The waveform visible in the results shown in Figure 5 is caused by deviations in the distribution of processes between CPUs. This kind of fluctuation cannot occur on single CPU systems like the Pentium 4 and the Athlon.

SCSI Issues

The GDT-8523RZ SCSI RAID controller by ICP Vortex, which we intended to use for disk benchmarking caused quite a stir, causing reproducible kernel panic during data transfer on the 64 bit kernel. However, the controller and the kernel module, *gdth.o*, both worked fine on the 32 bit kernel.

As our reference disk subsystem we used an Axus 16012 IDE SCSI RAID [2]. This system used 16 200 GByte Maxtor hard disks. The RAID system was optimized for data throughput for our benchmark. To do so, we first created a RAID 0 array with all 16 disks, and then distributed a 20 GByte slice across all 16 disks to use only about 5 GBytes on each disk.

The disk size of 20 GBytes was predetermined by the fact that the Bonnie++ system reads and writes about twice the RAM size. The Intel and AMD quad systems had 8 GBytes of RAM apiece. We then created a single partition on the SCSI drive, and formatted it with Ext2.

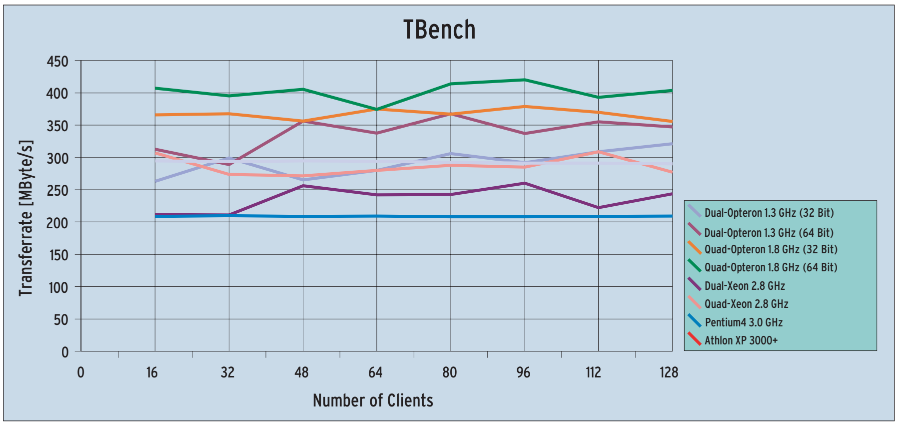

Figure 5: Again the Opteron outshone its competitors in the T-Bench. The Quad Xeon was nearly up to the pace, clocking in at 300 MByte/s. The waveform of the benchmarks only occurs on multiprocessor systems and is due to process distribution

The Xeon was a close second in this benchmark: 65 MByte/s write and 79 MByte/s read access are good values, but the Opteron achieved 80 MByte/s write and 73 MByte/s read throughput in 64 bit mode. In 32 bit mode, the Opteron had to hand over the blue ribbon to the Xeon, achieving a "mere" 67 MByte/s (write) and 59 MByte/s (read).

Conclusion

With its Opteron CPU, AMD has clearly laid down the foundation for migrating to 64 bits. It is impossible to harness the true capabilities of the CPU using 32 bit binaries, although it does achieve two thirds to three quarters of its full 64 bit power, and compares favorably with the field of 32 bit competitors.

The Opteron's high-performance memory access makes it highly recommendable for memory intensive applications, although having said that, it does lose some ground to Intel's Xeon on file I/O.

AMD intends to bring 64 bit power to the desktop with the Athlon 64 and Athlon 64 FX, both of which are slimline versions of the Opteron – and of course, AMD's ultimate aim is to send Intel's desktop processors off to the happy hunting grounds. ■

INFO

[1] Aria Technology Ltd: *http://www.aria.co.uk*

[2] Axus Microsystems Inc: *http://www.axus.com.tw/raid.htm*

[3] Mirko Dölle: "Sledgehammer", Linux Magazine, Issue 32, July 2003, page 44