

the INQUIRER

News, reviews, facts and friction

SEARCH
REVIEWS
WEEK TO DATE
EARLIER

Hypertransport is Intel's only hope

Sharing the crowded bus till eternity...

By [Nebojsa Novakovic](#): Sunday 30 October 2005, 18:53



OCTOBER WAS yet another interesting month for Intel with more changes in the roadmaps.

We had the Paxville dual-core Xeon unveiling, no Itanium at all at IDF Taipei but yes it showed Merom, Conroe and Woodcrest and, well, rumours about things to come...

Besides the latest Montecito postponement - which looks a lot like the "mysterious" EV7 Alpha delay just before Alphacide - Whitefield canning was last week's hot snippet. Not so much because Whitefield was to be a single-chip quad-core solution with a nasty 16 MB cache, or because Whitefield is supposedly one of better-off Bangalore suburbs, but because of its CSI interconnect, hyped before as the HyperTransport-killer... or is it?

Iwill

Leading on
memory capacity
per processor

Dual Opteron
Server Board

Max.

64

GB

Memory



Top INQ Stories

1. Cops say hard drives justify 90 day jail without trial
2. IBM claims Intel Paxville beats AMD, sets record
3. Why the Xbox 360 has three cores
4. 7800GTX 512 MB clocked at 550/1800 MHz

SOFTWARE
MEMORY
MICROPROCESSORS
MOBILE/PDA
GRAPHICS
SERVERS
CHANNEL
SET TOP BOXES
CONSOLES
GUTTER WATCH
EMISSION GOALS
ABOUT US
MESSAGE BOARD
LOGIN
TRAINING

RSS

Edited by Mike Magee
Phone
+44 208 248 2800

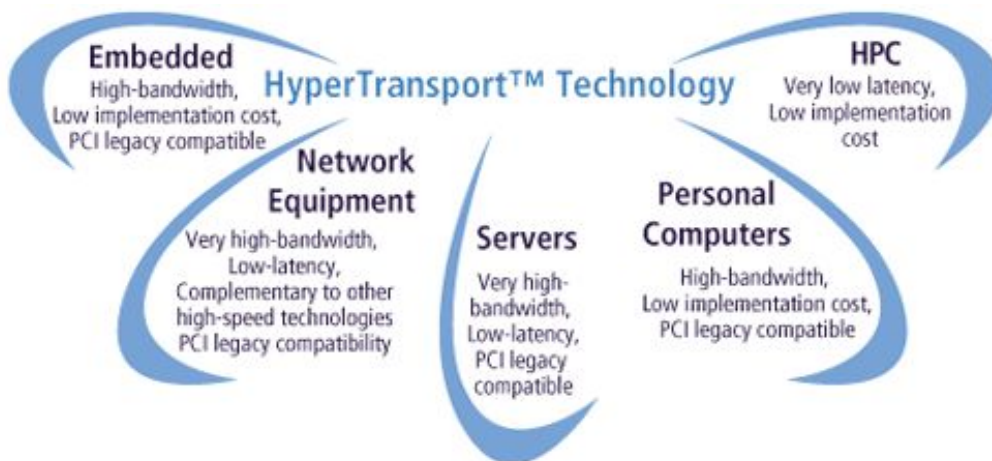
Terms and Conditions
of use.

To advertise in the UK
or Europe email us or
for more info click
here.

To advertise in Asia
email here.

To advertise in North
America email here.

Join the INQbot Mail
List for a weekly guide
to our news stories:
Subscribe/unsubscribe
here.



Whitefield or not, Intel's next X86 core does seem to be a good step forward and, on its own merit, should be superior to the current Athlon64 (K8) core - which of course, may have its own replacement by this time next year, too. In single-CPU config, a fast 1.33 GHz FSB with dual-channel DDR2 memory should be sufficient to feed a dual-core Conroe with 4 MB cache.

For Woodcrest, a dual-bus structure with two 1.33 GHz connections, one per each dual-core chip, should also be fine, coupled with quad-channel DDR2-667 memory (whether it is FB DIMM or not) - Blackford chipset is a good example of such offering. If the core itself is done right, these solutions should bring it a reasonably good competitive position from the system point of view against the upcoming DDR2-enabled Athlon64 and Opteron - not exactly the best, but surely better than now.

What about quad-socket and larger systems? Well, the "EV7 for the masses" or Opteron HyperTransport-based platform, lets each new CPU add its own memory and I/O bandwidth to the pool, with now proven far superior scaling to the current XeonMP platform.

The dual-core PaxvilleMP, despite the 800 MHz FSB push, just aggravates the situation, as now a single FSB has to feed up to eight cores (four sockets x two cores), and that FSB itself has only the same bandwidth as the current low-end Pentium PC FSB.

Not exactly stuff fit for a high-end, bandwidth-hungry server? Just compare, the total memory bandwidth accessible to all CPUs at any one time in a quad-socket Opteron is 25.6 GB/s if using DDR400, while in the quad-socket Paxville it is 6.4 GB/s. On top of it, the inter-processor communication among those Opterons can go across separate 8 GB/s HT channels simultaneously with the memory access, while the Xeon platform has to fit all that, including the big bottleneck of cache coherency traffic, on that same FSB.

No wonder many comparative results are catastrophic, and that doesn't even count the power consumption difference - which is another catastrophe for Paxville in itself.

That's why I was hoping that Intel, with the help of all those great ex-Alpha engineers, would create something better, faster and more scalable than HyperTransport - after all, the proposed EV8 Alpha link system for 2002 would have been far superior to even today's Hypertransport. The CSI (Common System Interconnect) was supposed to be just that, by 2007 - or at least, that

was the order by the "senior management".

After all, with its well-know NIH (Not Invented Here) policy, Intel rarely took something from the market that it didn't develop itself, even when it was both technically superior and a good business decision. For instance, instead of maltreating themselves and the whole industry with the Itanic, they could have simply taken Alpha from DEC, with its fantastic performance, excellent compilers, the best-even Windows NT port and, in 2001, more software titles than Itanic has now, and take the whole high-end market by storm within a year - but no, it had to be something that can kill all the competition just by its very promise, and yet so complicated that no one ever would think of copying it.

Well, they almost succeeded on the first goal (only POWER and Fujitsu SPARC still hang on), and, on the other one, well I am sure no one is rushing to copy Itanic anyway - there is no commercial incentive at all to do it, beside the horrendously complicated architecture only rivalled by X86 itself in its monstrosity.

Talking about DEC, Intel did - probably only after a lot of internal wrangling - take the StrongARM instead, even though it was Not Invented Here - it became the stunningly successful Xscale, with now tens of millions of world's handphones and other appliances using it. Heck, even the arch-enemy Motorola would use it... That gives you the idea how different Intel's own position at high end would have been if they took the Alpha in the same way.

Another give-in by Intel on the NIH front was, of course, accepting AMD's X86-64 platform and extending its own Pentiums / Xeons the same way. It may have been a slap in the face for certain egos within Intel, but it surely didn't hurt the bank balance - on the contrary, 64-bit Intel PC chips sell pretty well despite the performance and power consumption "issues", to put it lightly.

So, rather than wasting another three years - an eternity in this business - trying to beat Hypertransport and quite possibly unsuccessfully, why doesn't Intel simply take it in and create the next wave of its X86 chips around it? After all, HT is not owned by AMD - it is an open consortium with many members. Some Pentium PCs these days also include HT - in the Nvidia SLI chipsets, for instance.

The upcoming Hypertransport 3, with further improvements and maximum bandwidth reaching 45 GB/s per link (45 GigaBYTES, that is), should be a nice fit for Intel to enter the fray - after all, aimed at 2006 release, the HT 3 would fit very well to the follow-ons of Conroe and Woodcrest... and, how hard is it really to put up a memory controller and HT on the CPU? Forget AMD, look at SiByte and, now, PA Semi... lot's of this stuff is almost at the standard cell level.

A CPU chip - Opteron, Tigerton, POWER, Alpha, anything - having four such HT 3 channels on-chip could have up to 180 GB/s total system communication bandwidth in a year or so - on top of its local memory bandwidth. Enough for anything, from multiple ultra high-speed GPUs plugged in parallel, to very, very scalable, glueless SMP systems.

Compare that with the idea of, say in 2007, having a bunch of dual-chip dual-core MCM Tigertons sharing a single bus, whatever that bus is.... looks pathetic, doesn't it? Just imagine the zillion data transfer requests queuing for the "seats" on that bus for eternity.

In summary, Intel should not waste any more time on CSI - it is just too late now - or depend on the NIH syndrome in general. The non-NIH StrongARM performs darn well as Xscale both technically and financially, so could an Intel Alpha (instead of all that Itanium mess - it's actually not too late for Intel to go back, hah). For the X86 mainstream, the non-NIH HyperTransport 3 is the way to go for Intel, in my own opinion. Yes, Intel and AMD X86 CPUs could then even share the chipset and other infrastructure stuff, but so what? Intel can save itself few years of valuable time and jump back instantly into the high-end X86 market with a top-notch product then.

With its enormous firepower, capable design teams, and hopefully, no mingling of wrong people into the engineering decision making process, Intel can (and will) surely do some of the world's best CPU chips, even if using the underlying platforms shared by others - after all, if Conroe would be a better 64-bit X86 CPU than the current Athlon64, using the same X86-64 instruction set borrowed from AMD, why not the same across the whole spectrum? μ

See Also

[Intel to license Hypertransport](#)

[Home](#) [Discuss on our Forum](#) [Flame Author](#)

 [Recommend this article](#)  [Print](#)