

Intel Core 2 Extreme QX9650 - Penryn Ticks Ahead

Date: Oct 29, 2007
Type: CPU & Chipset
Manufacturer: Intel
Author: Anand Lal Shimpi
Page 1

We've never seen Intel with such a strong roadmap before, the company is truly firing on all cylinders and executing with amazing precision. Server, mobile, desktop and even new areas like ultra mobility and graphics all have absolutely wonderful roadmaps to look forward to. The biggest complaint we've had about Intel these days is that they kind of botched the X38 launch. Think back to a couple of years ago, what was our chief complaint then? Probably leaving us with power hungry, under performing processors for about 5 years. Today we're looking at a damn good Intel.


Recap: What's a Penryn?

Core, it's the architecture that shook an industry and today Intel is officially doing its first update to it. Prior to Intel's Core architecture, there wasn't much to get excited about when it came to Intel on the desktop.

At the same time, with Penryn Intel is very much the victim of its own success. How do you follow up such a tremendous splash with anything but equal greatness? AMD is close but still has yet to produce a response to Core 2, much less Penryn, and thus Intel's biggest competition today is itself.

In January 2007 Intel first showed off its 45nm High K + metal gate transistors, a dramatic departure from Intel's current 65nm transistors not only in size/switching speed but actual composition. If you remember back to the days of the original P6 processors, with a smaller transistor we saw tremendous improvements in die size, power consumption and performance. These days, such dramatic improvements are much harder to come by given that we're already dealing with such small transistor feature sizes. Gone are the days of the free lunch with each die shrink.

We've gone over [the technical details of Intel's high-K + metal gate enhancements](#), but the end result is that at the same clock speed you can expect dramatic reductions in power. Alternatively, at the same power levels, you can achieve much higher switching rates and thus higher clock speeds.



The core architecture of Penryn remains unchanged from Conroe; with the smaller transistors Intel's able to fit in a few new features and more cache on the chip while still maintaining a smaller die size. Where each dual-core Conroe die measured 143 mm², Penryn is merely 107 mm² despite having 50% more cache (6MB vs. 4MB). Obviously the quad core chips double overall area but you get the point.

Intel also uses a lot more of these new 45nm transistors than before; while a dual-core Conroe was made up of 291 million transistors, the comparable Penryn weighs in at 410 million (582M vs. 820M for quad-core variants). You're getting 40% more transistors and 50% more cache in a 25% smaller package; the latter is obviously most important to Intel as it helps reduce costs and drive profits up. So while it may seem generous, the move is purely self motivated on Intel's part.


The larger cache is a bit different than what we've seen in Conroe. While Conroe's cache is a 4MB 16-way set-associative L2, the 6MB Penryn cache is 24-way set-associative, designed to improve hit rates and keep latency manageable in an already large cache. Intel hasn't revealed whether Penryn's prefetchers have been adjusted to help populate its larger cache any better. As we saw in [our original Penryn preview](#), Penryn's cache performance remains unchanged; latencies in our final stepping are identical to Conroe.

The cache enhancements are by far the biggest consumer of those extra transistors in Penryn, but believe it or not, they aren't responsible for the biggest performance boost. Intel has been fairly steady in adding new instructions to the x86 ISA and Penryn continues the trend with the addition of SSE4. Penryn gets 47 new instructions that make up the first implementation of SSE4; more will come with Nehalem at the end of 2008. We'll talk about SSE4 performance later on in this article, but here are the instructions you get with Penryn:


Sub Group	Instructions	Description	Expected Application Benefits
Packed Dword Multiplies	PMULLD, PMULBQ	New support for four signed or unsigned 32x32 bit multiplications per instruction, as well as signed forms of 32x32+54 multiplication.	Broadly useful for improved automated compiler vectorization of data processing written in high-level languages (like C and Fortran).
Floating Point Dot Product	DPBS, DPPD	Improved performance for AOS (Array of Structs) data processing through support for single and double-precision dot products.	3-D content creation, gaming, and support for languages like CG and HLSL.
Packed Blending	BLENDDQ, BLENDDPS, BLENDQPB, BLENDQPB, BLENDQPB, BLENDQPB	Blending conditionally copies one field in the source onto the same field in the destination. These new instructions improve the performance of blending operations for most field sizes through packing multiple operations in a single instruction.	Broadly useful for automated compiler vectorization of data processing written in high-level languages (like C and Fortran), and applications such as image processing, video processing, multimedia, and gaming.
Packed Integer Min and Max	PMINSB, PMAXB, PMINL, PMAXL, PMINQ, PMAXQ, PMINSD, PMAXSD	Compares packed signed/unsigned byte/word/dword integers in the destination operand and the source operand, and returns the minimum or maximum value.	Broadly useful for automated compiler vectorization of data processing written in high-level languages (like C and Fortran), and applications such as image processing, video processing, multimedia, and gaming.

		These new instructions simplify data insertion and extraction between GPRs (or memory) and XMM registers	Useful for improved automated compiler vectorization of data processing, image and video processing, 3-D content creation, multimedia, and gaming
Floating Point Round	ROUNDPS, ROUNDSS, ROUNDPD, ROUNDDQ	Efficiently rounds the scalar and packed single- and double-precision operands to integers, with enhanced support for Fortran, Java and C99 language requirements	Image processing, graphics, video processing, 2-D/3-D applications, multimedia, and gaming
Register Insertion/Extraction	INSERTPS, INSERTQ, PEXTRB, PEXTRW, PEXTRQ, PEXTRD, PEXTRQ, PEXTRQ	These new instructions simplify data insertion and extraction between GPRs (or memory) and XMM registers	Broadly useful for automated compiler vectorization of data processing written in high-level languages (like C and Fortran), and applications such as image processing, video processing, multimedia, and gaming
Packed Format Conversion	PMOVSQBQ, PMOVSQDQ, PMOVSQDQ, PMOVSQDQ, PMOVSQDQ, PMOVSQDQ, PMOVSQDQ, PMOVSQDQ, PMOVSQDQ, PMOVSQDQ	Converts from a packed integer (from XMM register or memory) to a zero- or sign-extended integer with wider type	Broadly useful for automated compiler vectorization of data processing written in high-level languages (like C and Fortran), and applications such as image processing, video processing, multimedia, and gaming
Packed Test and Set	PTSET	Faster branching from SIMD decisions to support conditionally vectorized code	Useful for improved automated compiler vectorization of data processing, image and video processing, 3-D content creation, multimedia, and gaming
Packed Compare for Equal	PCMPEQB, PCMPEQD	Performs SIMD compare for equality of the packed QWORDS in the destination and the source operand	Broadly useful for automated compiler vectorization of data processing written in high-level languages (like C and Fortran), and applications such as image processing, video processing, multimedia, and gaming
Pack QWORD to Unpacked WORD	PACKUSDW	Converts packed signed QWORDS into packed unsigned WORDS using unsigned saturation to handle overflow condition. This new instruction completes the set of other instructions in this type	Broadly useful for automated compiler vectorization of data processing written in high-level languages (like C and Fortran), and applications such as image processing, video processing, multimedia, and gaming

Penryn also implements a new divider that impacts both integer and floating point divides using a radix-16 algorithm. The algorithm computes more bits of the result of a divide each pass (four bits per iteration vs. two bits in Conroe), decreasing divide latency.



The faster divider is a very specific enhancement that should manifest itself as a performance boost in 3D and imaging applications.




Penryn's Super Shuffle Engine should also improve SSE2, SSE3 and SSE4 applications that use a lot of shuffle operations. Cache performance is also improved slightly for misaligned stores, which should improve performance, once again, in 3D and imaging applications. Finally, there are some power enhancements made to Penryn, but these are mobile-specific and thus don't apply to any of the desktop variants.


Page 2

What do we have here today? Yorkfield


The problem with Intel's codenames these days is that we've got processor family codenames and then actual chip codenames. Penryn refers to the entire family of 45nm Core architecture products that have been/are being announced, but the actual chips themselves all have their own codenames. For example, the 45nm Penryn based quad-core Xeon processor is codenamed Harpertown. Penryn on the desktop carries two names: Yorkfield and Wolfdale.

Yorkfield is quad-core Penryn for the desktop, Wolfdale is simply dual-core. Yorkfield isn't actually a different die, because a Yorkfield chip is just made up of two Wolfdale die on the same package (just like current quad-core Kentsfield Intel CPUs). This won't change until Nehalem.





The chip that Intel is launching today is the first Yorkfield: the Core 2 Extreme QX9650. The quad-core QX9650 runs at 3.0GHz with a 1333MHz FSB, much like its predecessor the QX6850. Like all Yorkfield CPUs, the QX9650 is made up of two independent dual-core die on a single package, each one with a shared 6MB L2 cache for a total of 12MB of on-chip L2 cache.



QX9650 (left), QX6850 (right)

The QX9650 will work in a number of presently shipping motherboards; we tested ours in an ASUS P35 board, but you'll have to check with your board vendor - or check [our Penryn Compatibility article](#) - to make sure that there's BIOS/board support for it. The chip is still physically an LGA-775 processor so it'll fit in any LGA-775 socket; it's just up to the motherboard guys to implement hardware and BIOS level support for the processor.

Pricing hasn't been announced yet but we expect the QX9650 to come in at the \$999 mark, the same as previous Core 2 Extreme parts.

Test Configuration

CPU:	Intel Core 2 Extreme QX9650 (3.00GHz/1333MHz) Intel Core 2 Extreme QX6850 (3.00GHz/1333MHz)
Motherboard:	Gigabyte GA-X38-DQ6 (Intel X38)
Chipset:	Intel X38
Chipset Drivers:	Intel 8.1.1.1010 (Intel)
Hard Disk:	Seagate 7200.9 300GB SATA
Memory:	Corsair XMS2 DDR2-800 4-4-4-12 (1GB x 2)
Video Card:	NVIDIA GeForce 8800 GTX
Video Drivers:	NVIDIA ForceWare 163.75
Desktop Resolution:	1600 x 1200
OS:	Windows Vista Ultimate 32-bit

Recap: When's a Penryn?

In our first Penryn preview we laid out the launch schedule for Intel's new chips; now we can finally bring you an update with more specifics including model numbers and clock speeds. Let's first look at a table that will round out the rest of 2007:

CPU	Clock Speed	FSB	L2 Cache	Availability	Pricing
Intel Core 2 Extreme QX9650	3.00GHz	1333	6MBx2	Nov 12	\$999
Intel Core 2 Extreme	3.00GHz	1333	4MBx2	Now	\$999

QX6850					
Intel Core 2 Extreme QX6800	2.93GHz	1066	4MBx2	Now	\$999
Intel Core 2 Quad Q6700	2.66GHz	1066	4MBx2	Now	\$530
Intel Core 2 Quad Q6600	2.40GHz	1066	4MBx2	Now	\$266
Intel Core 2 Duo E6850	3.00GHz	1333	4MB	Now	\$266
Intel Core 2 Duo E6750	2.66GHz	1333	4MB	Now	\$183
Intel Core 2 Duo E6550	2.33GHz	1333	4MB	Now	\$163
Intel Core 2 Duo E6540	2.33GHz	1333	4MB	Now	\$163
Intel Core 2 Duo E4600	2.40GHz	800	2MB	Q4	\$133
Intel Core 2 Duo E4500	2.20GHz	800	2MB	Now	\$133
Intel Core 2 Duo E4400	2.00GHz	800	2MB	Now	\$113
Intel Pentium E2180	2.00GHz	800	1MB	Now	\$84
Intel Pentium E2160	1.80GHz	800	1MB	Now	\$84
Intel Pentium E2140	1.60GHz	800	1MB	Now	\$74

The big new introduction here is the Core 2 Extreme QX9650, the very first Yorkfield and the first Penryn we'll see on the desktop. The QX9650 will officially launch on November 12 and although Intel hasn't revealed pricing, we're expecting it to be at \$999 thus replacing the QX6850. Given that it costs Intel less money to make than a QX6850, we'd expect Intel would want to sell more of the QX9650 anyways, and pricing it more than \$999 just isn't going to help that cause.

The more interesting table however is what happens starting next year, because that's where we get some of the more mainstream Penryn chips in the market:

CPU	Clock Speed	FSB	L2 Cache	Availability	Replaces?
Bloomfield	TBD	N/A	TBD	Q4 '08	TBD
Intel Core 2 Extreme QX9770	3.20GHz	1600	6MBx2	Q1 '08	QX9650
Intel Core 2 Extreme QX9650	3.00GHz	1333	6MBx2	Nov 12	\$999
Intel Core 2 Quad Q9550	2.83GHz	1333	6MBx2	Q1 '08	Q6700
Intel Core 2 Quad Q9450	2.66GHz	1333	6MBx2	Q1 '08	Q6600
Intel Core 2 Quad Q9300	2.50GHz	1333	3MBx2	Q1 '08	Q6600
Intel Core 2 Duo E8500	3.16GHz	1333	6MB	Q1 '08	E6850
Intel Core 2 Duo E8400	3.00GHz	1333	6MB	Q1 '08	E6750
Intel Core 2 Duo E8300	2.83GHz	1333	6MB	Q2 '08	E8200
Intel Core 2 Duo E8200	2.66GHz	1333	6MB	Q1 '08	E6550
Wolfdale 3M	TBD	1066	3MB	Q2 '08	E4700
Intel Core 2 Duo E4700	2.6GHz	800	2MB	Q1 '08	\$266

Look at those availability dates: Penryn is coming for your children in Q1 '08; only the E8300 and Wolfdale 3M parts won't be out until Q2. Clocks move up a little, the QX9770 goes to 3.20GHz thanks to a 1600MHz FSB (which will be enabled by Intel's upcoming X48 chipset - yep another one), and the E8500 brings mainstream chips up to 3.16GHz.

Now look at the "Replaces?" column to get an idea for where these things will be priced. Intel's own roadmap shows the Q9550 slotting in next to the current Q6700, which means that we may be able to find it priced at around \$530.

More interesting is that Intel seems to have segmented the affordable quad-core market a bit, by replacing the ever-popular Q6600 with two chips: a Q9450 and Q9300. While functionally identical, the Q9300 is built off of two Wolfdale 3M cores (meaning it only has 6MB total L2 cache) while the Q9450 is built off of two Wolfdale 6M cores giving it 12MB of total L2 cache. Obviously the 9300 will be cheaper to make, so we'd expect to see that at or below the \$266 price point of the Q6600. The Q9450 would slot in right above the 9300 in pricing.


We would hope that Intel will price the Q9300 below the \$266 price point (*can we have a sub-\$200 quad-core, please?*), because it will actually have less cache than the current Q6600 making it cheaper for Intel to make, but possibly reducing performance over current chips. Granted it will have all the Penryn enhancements which, as you will soon see, do improve performance but we generally like having our cake and eating it too.

The dual-core market also gets interesting, with the E8000 line replacing the current E6000 series. If Intel's pricing structure remains the same then it looks like at today's prices you'll end up with an extra 166MHz, 50% more cache, SSE4 and some other tweaks for the same money. We also have to mention how well the model numbers work out with the Core 2 products; everything is in nice increments of 100, just like when Conroe first launched. Ah those were the days...


The other important item to note on the roadmap going forward is that top line in the table - yep, the one that says Bloomfield. Bloomfield is none other than Nehalem, the 45nm successor to Penryn. It's a brand new architecture complete with an on-die memory controller, SMT (Symmetric Multi-Threading - 2 threads per core) and 8MB of shared cache (probably L3 shared among all four cores). While it's still a year away, it's very nice to see it on an Intel roadmap this far in advance of its launch.

Everything You Need to Know: Yorkfield vs. Kentsfield

The big question is: how does Yorkfield stack up to Kentsfield, the current 65nm quad-core part from Intel? Thankfully, the QX9650 runs at the same clock speed/FSB as the Kentsfield-based QX6850 making our job a little easier. We put these two bad boys up against one another and waited for the smoke to clear.




As we saw in our initial Penryn preview, the general application suite tests just don't show any real improvements of Penryn over Conroe. SYSMark 2007's overall performance went up by 2.6% but that's far from significant. If you look at the individual tests however, you'll see that the 3D benchmarks went up by 4.5% and the productivity suite went up by a similar 4.6% margin. The improvements in the 3D suite are expected, given the Radix-16 and Super Shuffle Engine enhancements to the Penryn core. The productivity suite most likely benefits from the Radix-16 divider as well as the larger cache. Overall the performance boost just isn't that significant, but given that Penryn is expected to arrive without a price premium, any performance improvement is better than nothing.



We see the same situation under PC WorldBench 6: the overall performance boost of Penryn over Conroe is basically nil. Digging deeper we see performance increases on the order of 4 - 7% if we look at the individual 3D rendering or media encoding tests.


Next up we've got our encoding tests, and given what we've already seen, we should expect some nice gains here:



And nice gains we do get. The clock-for-clock improvement under DivX 6.6 is a nice and round 10%, although curiously enough we get a meager 3% from our Windows Media Encoder test. Our QuickTime H.264 test shows a more average 4% performance increase, similar to our WME test.


Professional 3D rendering and image manipulation are also areas where we expect to see reasonably high gains from Penryn. Our 3dsmax, Cinebench and Lightwave tests remain unchanged from previous articles. The CS3 benchmark is the same Retouch Artists test we've used in previous CPU reviews. In order to

avoid having a graph scale that ranged from zero to the thousands we simply included percentage improvements here rather than actual numbers (e.g. 3dsmax scores are single digits while Cinebench scores are 4 digits):



As expected, we're seeing reasonable gains here. 3dsmax shows a 5% increase, Cinebench 8%, Lightwave 6% and Photoshop CS3 has a healthy 10% performance boost in store for us. Without a doubt encoding and 3D/image manipulation are the real strengths of Intel's Penryn architecture.

The final Penryn test is all about 3D gaming and thankfully we have a number of new titles to play with. All of our benchmarks were run at 1024 x 768 to avoid being GPU bound and with high quality settings (the exceptions being Crysis and World in Conflict which both used Medium quality defaults).



* Denotes performance in minutes, lower is better

Penryn's gaming performance is really all over the place. Titles like Oblivion and Bioshock see absolutely no performance increase, while Crysis gives us a mild 3.7% performance boost. Half-Life 2: Episode Two and Unreal Tournament 3 enjoy a 5.1% and 6.4% increase respectively, while World in Conflict and Quake Wars are both at around 9%.

If anything the gamut of gaming benchmarks sums up Penryn's performance improvement over Conroe quite well: it really varies from nothing to something.

Page 5


Diving Deeper: SSE4 Performance

One of Penryn's real strengths is in its support for SSE4, which has the ability to really provide a tremendous performance advantage for some time to come. Unfortunately, as is usually the case with new instructions, it's going to take a while for applications to actually utilize them. Such is the case with SSE4 as the only benchmarks we're able to bring you come directly from Intel, but thankfully they are of real world usage models. Both tests we've actually showed you in the past, during Intel's own sanctioned Penryn previews, and both involve some sort of encoding.

The most important test is a DivX encode using VirtualDub 1.7.6 and DivX 6.7. SSE4 comes in if you choose to enable a new full search algorithm for motion estimation, which is accelerated by two SSE4 instructions: MPSADBW and PHMINPOSUW. The idea is that motion estimation (figuring out what will happen in subsequent frames of video) requires a lot of computation of sums of absolute differences, as well as finding the minimum values of the results of those computations. The SSE2 instruction PSADBW can compute two sums of differences from a pair of 16B unsigned integers; the SSE4 instruction MPSADBW can do eight.

According to Intel's own research on motion estimation with SSE4, the same search algorithm can take 71 cycles per 16x16 pixel block using the SSE2 SAD (sum of abs differences) instruction, compared to only 26 cycles using the SSE4 version. The latency reduction results in an obvious performance increase.






We used VirtualDub 1.7.6 and DivX 6.7 with SSE4 Full Search enabled to measure the impact of this motion estimation optimization. Note that the motion estimation that's taking place here is more accurate than the default DivX setting, so both SSE4 and SSE2 versions of the algorithm result in slower performance (but better quality) than with it disabled.

	SSE2 Search	SSE4 Search
Intel Core 2 Extreme QX9650 (3.0GHz)	21.9 seconds	15.1 seconds
Intel Core 2 Extreme QX6850 (3.0GHz)	35.2 seconds	N/A

On our QX9650, the full search with SSE4 enabled runs about 45% faster than with SSE2 only - impressive! Note also that the Penryn QX9650 offers better SSE2 performance in this test as well, coming in about 61% faster than the QX6850. The total performance increase from QX6850 SSE2 to QX9650 SSE4 in this test is an incredible 133%. Obviously, this is not going to be the norm in many other applications, but there's definitely some potential for meaningful optimizations in certain applications.

It's important to note that the PHMINPOSUW instructions doesn't appear to be in AMD's proposed SSE5 specification, although MPSADBW looks like it'll make it. AMD will eventually add full SSE4 support to its processors but not until the 2009/2010 time frame from what we've heard.



Our second benchmark from Intel is an MPEG-2 encode of an HD video using TMPGEnc 4.0.


	TMPEGEnc 4.0
Intel Core 2 Extreme QX9650 (3.0GHz)	103 seconds
Intel Core 2 Extreme QX6850 (3.0GHz)	135 seconds

The performance difference is a little less significant here, with the SSE4-less QX6850 taking about 31% more time to encode the input file than the QX9650. Both of these are very real-world implementations of SSE4; unfortunately, it's tough to say how long it will be before we see widespread use of the new instructions.

I've Got the Power: 45nm vs. 65nm

Since we're dealing with the same clock speeds as Intel's 65nm processors, power consumption has definitely gone down with the move to Penryn. Let's look at this thing at idle and under load running our WME9 test:






At idle, the QX9650 draws an impressive 34W less than the QX6850 - there's 45nm high-k + metal gate transistors in action for you.

Under load the power advantage is even more impressive, with a 47W delta the QX9650 under load uses only 11W more than its predecessor at idle. If you weren't dazzled by the performance improvements of Penryn, the reduction in power consumption is worth getting excited about.

Page 7

Overclock Me Baby

Naturally we wanted to see how much was left on the table with Intel's 45nm process. We can't help but think that had Phenom been out sooner, we'd be seeing a > 3GHz Penryn launch, but how easy would it be for Intel to ramp up clock speeds?



Our unlocked QX9650 had no problems hitting 333MHz x 12.0, for a final clock speed of 4.0GHz. We had to increase the stock voltage of 1.25V up to 1.40V to achieve it, but the overclock required no additional cooling beyond the standard Intel heatsink/fan. At lower voltages, 3.66GHz should be an easy target to reach.

Page 8

Another Price Drop? A Competitive Update

Although this couldn't be further from the matter at hand, AMD has quietly dropped its prices since the last time we looked at upper mainstream performance. The new/old AMD pricing structure is as follows:


CPU	Old Price	New Price
AMD Athlon 64 FX-74	\$599/pair	\$599/pair
AMD Athlon 64 X2 6000+	\$178	\$167
AMD Athlon 64 X2 5600+	\$157	\$146
AMD Athlon 64 X2 5200+	\$136	\$125
AMD Athlon 64 X2 5000+	\$125	\$115
AMD Athlon 64 X2 4800+	\$115	\$104
AMD Athlon 64 X2 4400+	\$94	\$89
AMD Athlon 64 X2 4000+	\$73	\$68
AMD X2 RF-2350	\$91	\$96

AMD X2 BE 2300	\$73	\$91
----------------	------	------

The price changes are fairly minor, but they do change the way we have to compare these processors. For example, the Athlon 64 X2 6000+ used to be priced close to the Core 2 Duo E6750, but now it's an E6550 competitor, improving AMD's competitive stance. For some reason, AMD's 45W Athlon X2 processors actually went up in price, possibly due to fluctuation in yields.

The other major change is that below the 6000+, all of AMD's chips compete with Intel's E4000/E2000 series, not the E6000 line. We'll be working on an update to our [Midrange CPU Roundup](#) to take some of these changes into account, but for this review we'll do a quick update looking at the 6000+ compared to its new price competitor, the E6550.

The chart below shows the percent difference in scores between the 6000+ and the E6550; the blue bars mean that Intel won that test and the green bars indicate that AMD won:




AMD definitely gets more competitive with its price drop, but Intel still holds on to the competitive advantage. Not only is the E6550 faster overall, it is also a cooler running processor (keeping in mind that the 6000+ is still a 90nm core) with more overclocking headroom.


Page 9

The Phenom Question

The CPU wars aren't over for the year, despite there being only 64 days left in the calendar. AMD is committed to releasing Phenom in 2007, and with it will come a definite change in the balance between AMD and Intel. No one is expecting AMD to take the overall performance crown away from Intel - AMD simply won't be able to hit the clock speeds necessary - but will it be any more competitive at mainstream price points?


When AMD's Barcelona launched we attempted to [simulate desktop performance vs. its K8 architecture in a handful of applications](#); we came up with the following chart:





Using those numbers, we took our Athlon 64 X2 6000+ results and scaled them according to where we expect Phenom to perform. We then put this simulated Phenom head-to-head with an identically clocked Core 2 Duo E6850 as well as the price-competitor to the 6000+, the E6550.

Note that this is a very rough comparison, first because the scaling values we have were taken from a quad-core Barcelona vs. two dual-core K8s and we're applying it to a dual-core K8 here. Second, AMD isn't going to be launching a 3.0GHz dual-core Phenom until sometime next year - most likely after Intel has already started shipping mainstream Wolfdale parts - which means that AMD will be competing against a completely different beast when that happens. Regardless, the comparison is still an interesting one to make; let's see if we can get any sort of expectations for what is to come.



*Denotes time in seconds, lower bars mean better performance

Clock for clock, Intel still holds on to the lead, but note that there are some situations where our simulated 3.0GHz Phenom X2 outperforms the E6550, which could be its price competitor. Again, we're assuming that a dual-core desktop Phenom will scale as well as our quad-core Barcelona and we're also assuming that AMD can price Phenom this competitively and that Intel doesn't respond with even more aggressive pricing. That's a lot of assumptions for AMD to be able to pull ahead of Intel at the same price point, but then we need to factor in the other P: Penryn.

In a couple of the areas that will be close between AMD and Intel, such as DivX and 3dsmax, Penryn happens to do really well. That means any gains AMD could make there with Phenom may be negated by the performance boost we're seeing from Penryn, leaving us in much of the same situation that we are in today. The only saving grace for AMD is that there are some areas where Intel just doesn't get a big performance boost from Penryn (e.g. SYSMark, Lightwave), and in those benchmarks Phenom will make AMD more competitive.

The end result is that we expect Phenom to make AMD more competitive, but because of Penryn and aggressive pricing - and if our assumptions are correct - it doesn't look like we'll see the sort of upset that Intel pulled on AMD with the launch of the original Core 2. (Or that AMD pulled on Intel with K8 vs. NetBurst.)

Page 10


Final Words


It's almost a bit disheartening to see Penryn launch at only 3.0GHz, as we know the architecture is capable of so much more. Unfortunately, what it looks like we're seeing here today is an artificial slowing of Intel's roadmap in response to the competition. While the tick-tock model is still very much in play and Nehalem is on track for a late '08 release, we're getting a 1333MHz FSB 3.0GHz QX9650 today instead of a 1600MHz FSB 3.2GHz Q9770 because there's no competitive reason for anything more. AMD's Phenom still isn't out and all indications are that we won't see 3GHz from AMD until sometime next year, meaning that Intel definitely doesn't have any real incentive to push the performance envelope this year.

Penryn really offers a wide range of performance; at worst it's no faster than its predecessor, but at best we're looking at low double digit performance gains. The biggest improvements seem to come in the encoding, 3D rendering and image manipulation applications, which makes sense when you look at the nature of the architectural improvements in Penryn.

The tremendous reduction in power consumption is very exciting and arguably one of the more desirable features of Penryn. The initial SSE4 results do look promising, but as we've often seen with new instructions, it's anyone's guess as to when we'll see widespread adoption of them from developers (and how much they will truly help in everyday use). The fact that we're able to bring you real world DivX performance results today though is promising.

Given that we're expecting a seamless price transition, it's tough to really find fault with Penryn. While the performance bumps can vary from minor to major, the cost should be nothing to the end user, in which case we can't complain. It's not the revolutionary step that Conroe was, but Penryn is merely making everything a little better all while driving power consumption down. So long as Intel keeps its prices the same, we're happy.





So what are the final words on it? As an evolution of Conroe, Penryn gets our vote. We just want it at lower price points, please.