

Updated CPU Cheatsheet - Seven Years of Covert CPU Operations

Date: Aug 28, 2004
Type: CPU & Chipset
Manufacturer: Various
Author: Jarred Walton
Page 1

Introduction

Update: 8/27/04 - The charts have all been revised. Thanks go out to all the people that posted corrections in the comments section as well as sending them via email. In addition to the corrections, some further information and commentary has been added to the pages. For anyone that actually comes back to this article for reference information, enjoy the changes!

Foreword by Kristopher Kubicki:

From time to time we stumble upon some truly gifted and patient people here at AnandTech. Some weeks ago I wrote a [CPU codename cheatsheet](#) as just something to do in an airport terminal to kill time. Very soon after, an extremely diligent Jarred Walton showed me his rendition of the CPU family tree that he was keeping just for fun!? Knowing I was bested, I offered Jarred a chance at writing a pilot for AT, and here it is! Please enjoy the second, extremely thorough CPU Cheatsheet 2.0.

*But loud! what lurks in yonder chassis, hot?
A CPU, my programs it will run!
O Pentium, Pentium! wherefore art thou Pentium?
Obscure thy benchmarks and refuse thy name.
What's in a name? that which we call a chip
By any other name would run as fast.*

My sincere apologies to Shakespeare, but that mangled version of *Romeo and Juliet* is an apt description of the world of computer processors. Once upon a time, we dealt with part numbers and megahertz. Larger numbers meant you had a faster computer. 80286 was faster than 8088 and 8086, and the 80386 was faster still, with the 80486 being the king of performance. Life was simple, and life was good. But that is the distant past; welcome to the present.

Where once we had a relatively small number of processor parts to choose from, we are now inundated with product names, model numbers, code names, and features. Keeping track of what each one means is becoming a rather daunting task. Sure, you can always try Googling the information, but sometimes you'll get conflicting information, or unrelated web sites, or only small tidbits of what you're trying to find out. So, why not put together a clear, concise document that contains all of the relevant information? Easier said than done; however, that is exactly what is attempted in this article.

In order to keep things even remotely concise, the cutoff line has been arbitrarily set to the Pentium II and later Intel processors, and the Athlon and later AMD processors. Anything before that might be interesting for those looking at the history of processors, but for all practical purposes, CPUs that old are no longer worth using. Also absent will be figures for power draw and heat dissipation, mainly because I'm not overly concerned with those values, not to mention that AMD and Intel have very different ways of reporting this information. Besides, Intel and AMD design and test their CPUs with a variety of heatsinks, motherboards, and other components to ensure that everything runs properly, so if you use the proper components, you should be fine.

So what will be included? For this first installment, details on clock frequencies, bus speeds, cache sizes, transistor counts, code names, and a few other items has been compiled. The use of model numbers with processors is also something people will likely have trouble keeping straight, so the details of processors for all Athlon XP and later AMD chips and Pentium 4 and later Intel chips will follow. The code names and features will be presented first, with individual processor specifics listed on the later pages. As a whole, it should be a useful quick reference - or cheat sheet, if you prefer - for anyone trying to find details on a modern x86 processor.

With that said, on to the AMD processors. Why AMD first? Because someone has to be first, and AMD comes before Intel in the alphabet.

Page 2

AMD Cheat Sheet

AMD Processors										
Argon (K7)	Athlon	Slot A	500-700	512K	22 + cache	250	184	100		
Pluto (K75)	Athlon	Slot A	550-850	512K	22 + cache	180	102	100		
Orion (K75)	Athlon	Slot A	900-1000	512K	22 + cache	180	102	100		
Spitfire	Duron	462	600-950	64K	25	180	100	100		
Morgan	Duron	462	900-1300	64K	25.2	180	106	100		
Thunderbird	Athlon "B"	462	650-1400	256K	37	180	117	100		
Thunderbird	Athlon "C"	462	1000-1400	256K	37	180	117	133		
Palomino	Athlon XP/M	462	850-1733	256K	37.5	180	129	100/133		
Palomino	Athlon MP	462	1000-1733	256K	37.5	180	129	100/133	1-2	
Thoroughbred A	Athlon XP	462	1467-1833?	256K	37.5	130	80	133		
Thoroughbred B	Athlon XP/M	462	1200-2133	256K	37.5	130	84	133		
Thoroughbred B	Athlon XP	462	2083-2250	256K	37.5	130	84	166		
Thoroughbred B	Athlon MP	462	1667-2133	256K	37.5	130	84	133	1-2	
Barton	Athlon XP/M	462	1467-2133	512K	54.3	130	101	133		
Barton	Athlon XP	462	1833-2167	512K	54.3	130	101	166		
Barton	Athlon XP	462	2100-2200	512K	54.3	130	101	200		
Barton	Athlon MP	462	2133	512K	54.3	130	101	166	1-2	
Applebred	Duron	462	1400-1800	64K	25.2*	130	84*	133		
Thorton	Athlon XP	462	1667-2067	256K	37.5*	130	101*	133		
Thoroughbred B	Sempron	462	1500-2000+	256K	37.5	130	84	166		
Sledgehammer	Athlon FX	940	2200-???	1024K	105.9	130 SOI	193	200	Y	
Sledgehammer	Opteron	940	1400-2400	1024K	105.9	130 SOI	193	200	Y	1-8
Sledgehammer	Athlon FX	939	2400-???	1024K	105.9	130 SOI	193	200	Y	
Clawhammer	Athlon 64	754	1800-2200(?)	512K	105.9	130 SOI	193	200	Y	
Clawhammer	Athlon 64	754	2000-2400(?)	1024K	105.9	130 SOI	193	200	Y	
Newcastle	Athlon 64	754	1800-2600(?)	512K	68.5	130 SOI	144	200	Y	
Newcastle	Athlon 64	939	2200-2600(?)	512K	68.5	130 SOI	144	200	Y	
San Diego	Athlon FX	939	2600-???	1024K	105.9(?)	90 SOI	114(?)	200	Y	
Paris	Sempron	754	1800-???	256K	~50(?)	130 SOI	118	200	N	
Venus	Opteron 1xx	940				90 SOI		200?	Y	
Troy	Opteron 2xx	940				90 SOI		200?	Y	1-2
Athens	Opteron 8xx	940				90 SOI		200?	Y	1-8
Odessa	Athlon 64 M?	754?		512K		130 SOI		200?	Y	
Winchester	Athlon 64	939		512K	68.5(?)	90 SOI	83(?)	200	Y	
Dublin	Athlon XP-M	462			37.5	130 SOI	128	200?	N	
Newark	Athlon 64-M LP	754?				90 SOI		200?	Y	
Lancaster	Athlon 64 M	754?				90 SOI		200?	Y	
Georgetown	Athlon XP M	462/754?				90 SOI		200?	N?	
Sonora	Athlon XP-M LP	462/754?				90 SOI		200?	N?	
Denmark	Opteron 1xx	940				90 SOI		200?	Y	
Italy	Opteron 2xx	940				90 SOI		200?	Y	1-2
Egypt	Opteron 8xx	940				90 SOI		200?	Y	1-8
Toledo	Dual Core FX	939				90 SOI		200?	Y	2C
Palermo	Sempron (?)	939 (?)		256K?	~50(?)	90 SOI	62(?)	200	N?	
Oakville	Athlon 64 Mobile	754?		512K?		90 SOI		200?	Y	

Victoria	Sempron (?)	754 (?)	256K? ~50(?)	90 SOI	62(?)	200	N?
* Die Size and/or transistor count is based off a larger CPU core with a portion of the die disabled.							
** Various steppings/sources listed different die sizes.							
*** The bus speed all Athons/Durons is double-pumped, but the CPU multiplier is based off the listed speed.							

A few notes to clarify the information. The stated die sizes and transistor counts for the Applebred and Thorton reflect the fact that these processors are Thoroughbred and Barton cores, respectively, with half of the L2 cache disabled, which is why they have a single asterisk next to them. There have been reports of hacking the Thorton processors and turning them into full Barton CPUs, but considering the insignificant cost difference these days, it's probably not worth worrying about. AMD plans on discontinuing the Barton soon anyway, and will use the old Thoroughbred core for the Socket A Sempron chips.

Transistor counts on Paris, Victoria, and Palermo are likely off, but it remains to be seen how AMD actually configures these chips. Early Athlon 64 512K cache chips for socket 754 were Clawhammer cores with half the cache disabled, but the newer models (i.e. 3200+ at 2.2 GHz with 512K, 3400+ 2.4 GHz 512K, and 3700+ 2.6 GHz with 512K) appear to be actual Newcastle cores. The same could very well happen with the Paris cores, where initial shipments are "downgraded" Newcastle cores, and later versions may physically remove the ~18.7 million transistors used in the L2 cache. Regardless, values on these cores should be taken with a grain of salt.

Unreleased processors will likely change from these current estimates, and question marks indicate best guess data at present. If you notice any errors or if you have additional information on forthcoming processors, let us know in the comments section or email.

Take note of the Toledo, Denmark, Italy, and Egypt cores; the 2C next to it stands for dual core. All four models use the same basic core and should come out around the same time in early 2005. Whether they launch as planned remains to be seen, and precise details about the internal layout are not yet clear - recent news suggests that each core will have its own L2 cache. Dual core is best described as SMP on a single chip, and while on the subject of SMP, please note that all of the Athlon XP processors could support multi-processor configurations *unofficially*. 2-way SMP was almost a certainty, but none of the CPUs were verified to function in such a configuration by AMD. While it would not be prudent to take such a risk as a business, quite a few enthusiasts saved themselves a lot of money by putting XP chips into SMP motherboards instead of spending the extra money on MP chips.

The basic core of the Athlon, from the Pluto all the way through the latest Newcastle and Paris processors, changed very little since its inception. It has a 10 stage integer pipeline and 15 stage floating point pipeline, with three identical Arithmetic/Logic Units (ALUs), Address Generation Units (AGUs), and Floating Point Units (FPUs). The FPUs also handle the MMX, 3DNow!+, and SSE/SSE2 support. Opteron increased the length to 12/17 stages, in addition to bringing 64-bit support. Future versions of the Athlon 64 will likely increase the length of the pipeline past the current 12/17 stages in order to increase clock speeds, but I doubt that AMD will ever show the *hubris* of Intel by creating a 31 stage pipeline - at least, not on any iteration of the Athlon architecture. This is especially a problem with the increasing power leakage of high clockspeeds and increasingly small process technology. Until those issues are resolved, I think it's safe to say that pipeline lengths will stay in the 10 to 15 stages (for integers) range with AMD.

Update: One reader was good enough to send a link to AMD's site where they actually list the Opteron as being a 12/17 design. (Thanks Tom!) Finding any good details on the Intel and AMD sites can be a major chore, most likely due to the level of competition between the companies as well as their size. There's a rule somewhere that the larger a company gets, the less informative and helpful their web site becomes! For those that want the link, here's the [Opteron information](#). That means that all Athlon 64 designs are also 12/17, of course. The Denmark, Italy, and Egypt CPUs are also dual core, it appears, and their entries have been updated to reflect this. (The old roadmap didn't include that information.)

Page 3

Intel Cheat Sheet

										Bus Speed (MHz)		
										L1 Cache	L3 Cache	
Intel IA32/EM64T Processors												
Covington	Cel	Slot 1	266/300	8K+8K	7.5	350	118	66				
Mendocino	Cel ("A")	Slot 1	266-433	16K+16K	128K	19	250	154	66	1-2		
Mendocino	Cel ("A")	370	233-533	16K+16K	128K	19	250	154	66			
Coppermine-128	Cel ("A")	370	533-766	16K+16K	128K	28*	180	106	66			
Coppermine-128	Cel ("A")	370	800-1100	16K+16K	128K	28*	180	106	100			
Klamath	P II	Slot 1	233-333	16K+16K	512K	7.5+37.2	350	203+L2	66	1-2		
Deschutes	P II	Slot 1	266-333	16K+16K	512K	7.5+37.2	250	118+L2	66	1-2		
Deschutes	P II	Slot 1	350-450	16K+16K	512K	7.5+37.2	250	118+L2	100	1-2		
Deschutes	P II Xeon	Slot 2	400-450	16K+16K	512K	7.5+37.2	250	118+L2	100	1-2		
Deschutes	P II Xeon	Slot 2	400-450	16K+16K	1M	7.5+74.4	250	118+L2	100	1-2		
Deschutes	P II Xeon	Slot 2	450	16K+16K	2M	7.5+148.8	250	118+L2	100	1-2		
Katmai	P III	Slot 1	450-600	16K+16K	512K	9.5+37.2	250	131+L2	100	1-2		
Katmai	P III B	Slot 1	533-600	16K+16K	512K	9.5+37.2	250	131+L2	133	1-2		
Tanner	P III Xeon	Slot 2	500, 550	16K+16K	512K	9.5+37.2	250	128+L2	100	1-8		
Tanner	P III Xeon	Slot 2	500, 550	16K+16K	1M	9.5+74.4	250	128+L2	100	1-8		
Tanner	P III Xeon	Slot 2	500, 550	16K+16K	2M	9.5+148.8	250	128+L2	100	1-8		
Cascades**	P III Xeon	Slot 2	600-1000	16K+16K	256K	28.1	180	106-90	133	1-2		
Cascades	P III Xeon	Slot 2	700	16K+16K	1M		180	210?	100	1-4		
Cascades	P III Xeon	Slot 2	700, 900	16K+16K	2M		180	385	100	1-4		
Coppermine**	P III	Slot 1	550-1000	16K+16K	256K	28.1	180	106-90	100	1-2		
Coppermine**	P III B	Slot 1	533-1000	16K+16K	256K	28.1	180	106-90	133	1-2		
Coppermine**	P III E	370	500-1100	16K+16K	256K	28.1	180	106-90	100	1-2		
Coppermine**	P III EB	370	533-1133	16K+16K	256K	28.1	180	106-90	133	1-2		
Tualatin	Cel ("A")	370	1000-1400	16K+16K	256K	28.1	130	80	100			
Tualatin	P III	370	1000-1333	16K+16K	256K	28.1	130	80	133	1-2		
Tualatin	P III S	370	1133-1400	16K+16K	512K	45.9	130	110?	133	1-2		
Willamette	Cel-128	478	1700-1800	12Ku+8K	128K	36.5	180	217*	100			
Willamette	P 4	423	1300-2000	12Ku+8K	256K	42	180	217	100			
Willamette	P 4	478	1500-2400	12Ku+8K	256K	42	180	217	100			
Foster	Xeon DP	603	1400-2000	12Ku+8K	256K	42	180	217	100	1-2		
Foster	Xeon MP	603	1400, 1500	12Ku+8K	256K 512K	42+37?	180		100	1-4		
Foster	Xeon MP	603	1600	12Ku+8K	256K 1M	42+74?	180		100	1-4		
Northwood	Cel	478	1400-2800	12Ku+8K	128K	36.5	130	131?	100			
Northwood	Mob. Cel.	478	1400-2800	12Ku+8K	256K		130		100			
Northwood**	P 4	478	1800-2600	12Ku+8K	512K	55	130	146-131	100			
Northwood**	P 4 "B"	478	2267-2800	12Ku+8K	512K	55	130	146-131	133			
Northwood**	P 4 HTT	478	3067	12Ku+8K	512K	55	130	146-131	133			
Northwood**	P 4 "C"	478	2400-3400	12Ku+8K	512K	55	130	146-131	200			
Gallatin**	P 4 EE	478	3200-3400	12Ku+8K	512K 2M	55+123	130	231-237?	200			
Prestonia	Xeon DP	603	1600-3000	12Ku+8K	512K	55	130		100	1-2		
Prestonia	Xeon DP	604	2000-3067	12Ku+8K	512K	55	130		133	1-2		
Prestonia	Xeon DP	604	3067-3200	12Ku+8K	512K 1M	55+61	130		133	1-2		
Gallatin	Xeon MP	603	1500-2800	12Ku+8K	512K 1M	55+61	130		100	1-4		
Gallatin**	Xeon MP	603	2000-2700	12Ku+8K	512K 2M	55+123	130	231-237?	100	1-4		
Gallatin	Xeon MP	603	3000	12Ku+8K	512K 4M	55+246?	130		100	1-4		
Prescott 256?	Cel D	478/775	2400-3200	12Ku+16K	256K		90		133			
Prescott	P 4 "A"	478	2400-2800	12Ku+16K	1M	125	90	112	133			
Prescott	P 4 "E"	478	2800-3400	12Ku+16K	1M	125	90	112	200			
Prescott	P 4 "E"	T/775	2800-???	12Ku+16K	1M	125	90	112	200			
Prescott	P 4 "E"	T/775	???	12Ku+16K	2M		90		200/266			
Nocona	Xeon	T/775?	2800-3600+	12Ku+16K	1M	125	90	112?	200	1-2		
Irindale					2M		90		200?			
Banias	Cel M	478M	1300-1500	32K+32K	512K		130		100			
Banias	P M	478M	900-1800	32K+32K	1M		130		100			
Dothan	Cel M	478M	900-1500	32K+32K	1M		90		100/133			
Dothan	P M	478M	1000-2400	32K+32K	2M		90		100/133			
Potomac							65					
Smithfield										2C		
Jonah	P M?					65?						

[illegible]

Notes on the Intel side of things are similar to the AMD side. There are again a couple cores that have an asterisk, indicating that the core was a "downgraded" version of a faster core, mostly with the Celeron processors. The double-asterisks are for chips that had varying die sizes in the various steppings. This probably occurs to a small degree in most chips, but in the Cascades, Coppermine, and Northwood cores, the changes were well documented and rather drastic. Thoroughbred A to B in AMD was only a 4 mm2 die size increase, while Coppermine fluctuated between 106 mm2 to 90 mm2, and Northwood went from 146 mm2 to 131 mm2. My guess is that it was due in part to hand-optimizing the layouts of the cores, but if anyone has precise details on the hows and whys of the decreases, I would like to hear them.

You can see that Intel also has 2C (dual core) designs in their roadmap, as well as a highly speculative 16C (sixteen core!) Itanium. Whether or not Tukwila will ever see the light of day is anyone's guess - it could simply be a mythical design that some hardware sites fantasize about. Transistor count on such a chip would likely be several BILLION transistors. (On a different note, I was recently up in Tukwila, WA purchasing a mountain bike from a pawn shop. They didn't have any processors for sale, unfortunately.)

In contrast to AMD, Intel has had several major architecture revisions during the past seven or so years. AMD pretty much stuck with the K7/Athlon core for all their processors, which was admittedly a very good design. Intel, with its deeper pockets, attacked on numerous fronts. First was the Pentium III line, which more or less ended in a draw with their rival AMD. Prompted by marketing - because "clockspeed sells" - Intel came up with a radical new architecture dubbed NetBurst, the basis of the Pentium 4. NetBurst was a success on the desktop, but it really was too power hungry for laptops, so Intel decided to pursue a completely separate architecture for its mobile processors, which is now also penetrating Blade and other low voltage markets. Finally, shortly after the launch of the Athlon 64, Intel countered with their reworked NetBurst architecture and the Prescott line of processors. Add to this the long-awaited launch of IA-64 (roughly ten years in the making!) which was a completely new architecture that was even more radical than NetBurst. Intel has been busy, needless to say.

For their desktop chips, SMP was available both officially and unofficially. The Celeron chips were not intended for SMP use and were never validated (by Intel) to work in such configurations. However, enterprising motherboard makers like Abit with their BP6 board allowed users to run early Celerons in dual CPU configurations. Intel put a stop to that with Coppermine-128 and Tualatin-256 (if you can call it that) Celerons. The P3 Xeon chips were all "multi-processor" configurations, capable of up to 8-way SMP. Such support was more dependent on the motherboard and chipset, though, so most setups topped out at 4-way SMP. Intel had a chipset that linked two 4-way buses together for their 8-way configuration, while ServerWorks created a chipset and motherboard that supported 8-way directly. In theory, they could have even followed Intel's example and linked two buses together to have a 16-way SMP setup, although at that point motherboard size becomes a difficult issue.

Itanium and SMP is a special case that needs further clarification. SMP is not always listed in the above chart, but all Itaniums are said to be capable of 512-way SMP. This is really more of a factor of the motherboard(s) and system design than the chip itself. For example, special high-end clustered systems have been built using AMD Athlon MP and Opteron CPUs as well as Xeon chips that have as many as 128 chips in a "single" system. Itanium is a similar case with SMP. Motherboards with up to eight sockets exist for Itanium, but 512-way SMP requires special hardware beyond the motherboard. (Please feel free to correct me if that's wrong, but I'm pretty sure this is the case. I can't imagine what a motherboard for 512 Itaniums would even look like if it were to exist - 8x8 feet in size?)

Update: A couple people pointed out issues with the naming of the Celeron processors. At the time, Intel used "A" to designate processors that overlapped an existing model. So there were cacheless Celeron 266/300 processors, and the 266/300 with 128K L2 cache had an "A" suffix. This occurred again with the Celeron 533, and once more with the Celeron 1000/1100. In a similar vein, the Klamath core was only 350 nm, while Deschutes was 250 nm. It was initially listed as 350/250 as there were certain Deschutes cores that were released as a pseudo-Klamath, for instance the P2 300 MHz SL2W8. There was not any way to actually tell (other than word of mouth) which P2 chips had the Klamath core and which had the Deschutes core. The chart has now been corrected by putting in a 250 nm 266-333 Deschutes line.

Before we get to the actual charts, I want to take a minute to make clear how the charts are organized. Due to the number of features involved with modern processors, it can become difficult to determine which CPU is actually faster when comparing different models. For example, do you go with the 2250 MHz Athlon XP using the Thoroughbred core, which has a 2800+ model number, or should you go with the 2000 MHz Athlon XP that uses the Barton core, which also has a 2800+ model number? With Intel, it can be even more difficult: you have different cache sizes, bus speeds, and even architectures.

Since I figure a lot of people may actually find some sort of relative sorting useful, I have attempted to do this. How you wish to rate the various factors is of course a topic that could be debated *ad nauseum*. What I am presenting is by no means a definitive answer on which model is faster, but it should give a rough estimate. Below are the various families of processors and the weighting values that I used. I then took the weight factor and multiplied that by the actual clock speed to come up with a final performance ranking.

Since this is simply a rough estimate on my part, I am not including these ranking values in the actual charts, but they are how I sorted the data. Really, the reason for their existence was to get a sorting function that more or less agreed with my own personal opinion, so if I happen to have missed a processor, or if a new processor is released, I can simply add in the processor(s) to the chart and resort it. I'm open for suggestions on how these ratings might be improved, but please realize that there will never be a definitive formula, as relative performance depends on what specific code you are running.

If you don't like math or don't really care to know precisely *how* the charts are sorted, feel free to just skip to the next page. This is only for people that really want to know details. Also, the weighting factors are within each family - they have no correlation with other processor families. (So don't get upset that the Dothan has a 1.6 weighting and Athlon FX only has 1.15!) With that said, here are the weighting factors that I used.

Duron, Athlon, Athlon XP and Sempron

64K	L2 + 100 MHz	bus = 0.7
64K	L2 + 133 MHz	bus = 0.75
256K	L2 + 100 MHz	bus = 0.8
256K	L2 + 133 MHz	bus = 0.85
256K	L2 + 166 MHz	bus = 0.9
512K	L2 + 133 MHz	bus = 0.95
512K	L2 + 166 MHz	bus = 1.0
512K	L2 + 200 MHz	bus = 1.05

Athlon 64

256K L2 + single-channel	(Socket 754)	= 0.9
512K L2 + single-channel	(Socket 754)	= 0.95
1024K L2 + single-channel	(Socket 754)	= 1.0
512K L2 + dual-channel	(Socket 939)	= 1.0

512K L2 + dual-channel (Socket 939) = 1.04
1024K L2 + dual-channel (Socket 940) = 1.11
1024K L2 + dual-channel (Socket 939) = 1.15

Celeron 2 and Pentium 4

128K L2 + 400 FSB = 0.6
256K L2 + 400 FSB = 0.75
256K L2 + 533 FSB = 0.80
512K L2 + 400 FSB = 0.84
512K L2 + 533 FSB = 0.91
1024K L2 + 533 FSB = 0.93
1024K L2 + 800 FSB = 0.98
512K L2 + 800 FSB = 1.0
512K L2 + 800 FSB + 2048K L3 = 1.15
2048K L2 + 1066 FSB = 1.2

Mobile Celeron, Mobile P4, Celeron M and Pentium M

128K L2 + 400 FSB = 0.6
256K L2 + 400 FSB = 0.75
256K L2 + 533 FSB = 0.80
512K L2 + 533 FSB + Northwood = 0.91
1024K L2 + 533 FSB + Prescott = 0.93
512K L2 + 400 FSB + Dothan = 1.25
512K L2 + 400 FSB + Banias = 1.3
1024K L2 + 400 FSB + Dothan = 1.35
1024K L2 + 400 FSB + Banias = 1.4
2048K L2 + 400 FSB = 1.5
2048K L2 + 533 FSB = 1.6

Page 5

Duron and Athlon

I won't bother going into details of the early Athlon and Duron processors. They were great in their day, but they're getting to be rather long in the tooth. If there is a strong demand for more details on these processors, I will add them at a later point, but for now I simply recommend that you bite the bullet and upgrade.

For those interested in some historical information, here are a few more tidbits. The early Argon, Pluto and Orion Athlon chips had L2 cache chips contained within the Slot A cartridge. This cache could run at 1/2, 2/5, or 1/3 of the core clock speed - the faster the core, the lower the ratio. This led to situations where, for example, a 700 MHz Athlon with 350 MHz L2 would outperform the more expensive 750/300, or the 850/340 would beat the 900/300 due to the slower cache. Generally speaking, performance comparisons between the Athlon and Pentium III chips of the day were neck-and-neck affairs, with each side winning some benchmarks. Athlon had better x87 floating point performance, while Intel generally won out with features like MMX and SSE - at least in applications that were properly optimized.

The socket A processors switched to an integrated full-speed L2 cache, but the cache was half as large. The increased speed and reduced latencies, however, more than made up for the decrease in cache size. At this time, AMD was able to actually surpass Intel in raw performance for a period of time. The Athlon Thunderbird eventually reached 1.4 GHz, while the Pentium III tried for 1.13 GHz and failed. Later versions of the Pentium III dubbed Tualatin would eventually reach 1.4 GHz, but those only came after the introduction of the Pentium 4. Athlon during these times was *the* chip for gaming systems.

One other item worth noting is that all of the Athlon and Duron systems used the EV6 bus protocol acquired from DEC/Alpha. This was a double-pumped system bus, which improved performance relative to older buses like that used in P6 motherboards. The bus speeds listed in the charts are the base bus speed, which is then multiplied by the CPU multiplier to arrive at the final CPU speed. However, due to the double-pumping, many motherboards will list the bus speed as the doubled value. The actual performance increased gained from the doubling of the bandwidth is not as large as some might expect, but it probably accounts for somewhere between 5 to 15 percent of the total performance of the architecture, depending on the application.

The Athlon 64 and Opteron processors, meanwhile, have switched to a HyperTransport bus running at 800 MHz on the early chips and 1 GHz on socket 939 chips. The main benefit of the HT bus is that it doesn't require as many traces (wires), so it makes motherboard layouts somewhat easier to design. This also allows for multiple high-speed bus connections when used in SMP systems without resorting to designs with more layers.

Athlon XP and Sempron Processors

					Multiplier
Athlon XP (Desktop) & Sempron (Desktop Value)					
Athlon XP 1500+	1333	Palomino	256	133.3	10.0X
Athlon XP 1600+	1400	Palomino	256	133.3	10.5X
Athlon XP 1700+	1467	Palomino/TBA	256	133.3	11.0X
Athlon XP 1800+	1533	Palomino/TBA	256	133.3	11.5X
Sempron 2200+	1500	Thoroughbred B	256	166.7	9.0X
Athlon XP 1900+	1600	Palomino/TBA	256	133.3	12.0X
Athlon XP 2000+	1667	Palomino/TBA	256	133.3	12.5X
Athlon XP 2000+	1667	Thorton	256	133.3	12.5X
Athlon XP 2000+	1533	Barton	512	133.3	11.5X
Athlon XP 2100+	1733	Palomino/TBA	256	133.3	13.0X
Sempron 2400+	1667	Thoroughbred B	256	166.7	10.0X
Athlon XP 2200+	1800	TBA/TBB	256	133.3	13.5X
Athlon XP 2200+	1800	Thorton	256	133.3	13.5X
Sempron 2500+	1750	Thoroughbred B	256	166.7	10.5X
Athlon XP 2200+	1667	Barton	512	133.3	12.5X
Sempron 2600+	1833	Thoroughbred B	256	166.7	11.0X
Athlon XP 2400+	2000	Thoroughbred B	256	133.3	15.0X
Athlon XP 2400+	2000	Thorton	256	133.3	15.0X
Athlon XP 2400+	1800	Barton	512	133.3	13.5X
Athlon XP 2500+	1867	Barton	512	133.3	14.0X
Sempron 2800+	2000	Thoroughbred B	256	166.7	12.0X
Athlon XP 2600+	2133	Thoroughbred B	256	133.3	16.0X
Athlon XP 2500+	1833	Barton	512	166.7	11.0X
Athlon XP 2600+	2083	Thoroughbred B	256	166.7	12.5X
Athlon XP 2600+	2000	Barton	512	133.3	15.0X
Athlon XP 2600+	1917	Barton	512	166.7	11.5X
Athlon XP 2700+	2167	Thoroughbred B	256	166.7	13.0X
Athlon XP 2800+	2250	Thoroughbred B	256	166.7	13.5X
Athlon XP 2800+	2083	Barton	512	166.7	12.5X
Athlon XP 3000+	2167	Barton	512	166.7	13.0X
Athlon XP 3000+	2100	Barton	512	200	10.5X
Athlon XP 3200+	2200	Barton	512	200	11.0X

*** All system buses for Athlon XP, Sempron, Athlon 64, and Opteron are "double pumped", so their data rate is twice the bus speed. The multiplier is based off the listed speed.

Many of the processors listed in the charts were not commonly available, so they may not be well known. Some of these parts were shipped to OEMs who had special requirements, for example they might want to use cheaper PC2100 RAM with a Barton core. Some of the listed chips might also have been mobile parts which were mistakenly listed in the wrong table. However, the majority of these chips actually do exist in various PCs. Note also that some parts were likely to be seen more in overseas markets than in the US. If you are *sure* that a part is incorrect or doesn't exist, feel free to post a comment or send an email.

Athlon XP tweaked some of the finer details of the Athlon architecture to improve performance. Since XP was also going up against Pentium 4 instead of Pentium III, AMD (re)introduced model numbers and began their "clock speed isn't everything" campaign. According to AMD, the XP line was rated in terms of performance relative to the Thunderbird core, but few people actually believe that. It was almost surely market driven, as the Pentium 4 was scaling rapidly in clock speed, and the Athlon cores couldn't possibly keep up in raw MHz. And of course, AMD is correct that clock speed isn't everything - average instructions executed per clock (IPC) multiplied by clock speed would give you the real instruction throughput. Unfortunately, coming up with a precise measurement of IPC is virtually impossible - it varies depending on the code executed. Still, clock-for-clock, Athlons are definitely faster than P4 chips, and the PR ratings were relatively accurate, at least in the beginning.

As the "processor wars" continued, both companies released tweaked designs. Thoroughbred was a process shrink that brought higher clock speeds, but not as high as initially desired. A reworked Thoroughbred B core - which added an extra layer to the core, among other things - helped raise the clock limit a bit more and allowed Athlon XP to eventually reach 2250 MHz. Note that Thoroughbred B cores can often overclock to 2.3 to 2.4 GHz with sufficient cooling, while the A versions are often limited to ~2.1 GHz.

After Thoroughbred, AMD added more cache with the Barton core, and readjusted their model numbers accordingly, since more cache brought more performance. This was really where the model numbers started to become suspect, though, since Intel had also added more cache and increased bus speeds without "adjusting" any model numbers. The 2500+, 2600+ and 2800+ tended to struggle a bit in keeping up with their Intel counterparts, but the real problem came when Intel released the 200 MHz (800 FSB) "C" version of their Pentium 4. The jump to 3200+ with the 200 MHz FSB really only kept the Athlon XP competitive with the P4 2.8C in overall performance comparisons. Of course, here the model names were a stroke of genius, as many people simply assumed that a 3200+ really was the equivalent of the 3.2C.

Athlon XP-Mobile Processors

					Multiplier
Athlon XP-M (Mobile)					
Athlon XP-M 850	850	Palomino	256	100	8.5X
Athlon XP-M 900	900	Palomino	256	100	9.0X
Athlon XP-M 950	950	Palomino	256	100	9.5X
Athlon XP-M 1000	1000	Palomino	256	100	10.0X
Athlon XP-M 1100	1100	Palomino	256	100	11.0X
Athlon XP-M 1200	1200	Palomino	256	100	12.0X
Athlon XP-M 1400+	1200	Thoroughbred	256	133.3	9.0X
Athlon XP-M 1500+	1300	Palomino	256	100	13.0X
Athlon XP-M 1600+	1400	Palomino	256	100	14.0X
Athlon XP-M 1500+	1333	Thoroughbred	256	133.3	10.0X
Athlon XP-M 1600+	1400	Thoroughbred	256	133.3	10.5X
Athlon XP-M 1700+	1467	Thoroughbred	256	133.3	11.0X
Athlon XP-M 1800+	1533	Thoroughbred	256	133.3	11.5X
Athlon XP-M 1900+	1600	Thoroughbred	256	133.3	12.0X
Athlon XP-M 1900+	1467	Barton	512	133.3	11.0X
Athlon XP-M 2000+	1667	Thoroughbred	256	133.3	12.5X
Athlon XP-M 2000+	1533	Barton	512	133.3	11.5X
Athlon XP-M 2100+	1600	Barton	512	133.3	12.0X
Athlon XP-M 2200+	1800	Thoroughbred	256	133.3	13.5X
Athlon XP-M 2200+	1667	Barton	512	133.3	12.5X
Athlon XP-M 2400+	1800	Barton	512	133.3	13.5X
Athlon XP-M 2500+	1867	Barton	512	133.3	14.0X
Athlon XP-M 2600+	2000	Barton	512	133.3	15.0X
Athlon XP-M 2800+	2133	Barton	512	133.3	16.0X

*** All system buses for Athlon XP, Sempron, Athlon 64, and Opteron are "double pumped", so their data rate is twice the bus speed. The multiplier is based off the listed speed.

There's not really a whole lot to say about the Mobile AMD processors. They are identical to their desktop counterparts, except they run on lower voltages and can run at reduced clock speeds to save power. Later on, the Athlon XP-M processors gained tremendous popularity due to their unlocked multipliers, which allowed them to overclock very well, as you could keep the bus speed close to the standard 200 MHz.

There are some OEM parts as well in the Mobile Athlon market which use a different socket than the standard 462 pin socket A. For the Athlon XP, there is a 563 pin version, and for Athlon 64 there is a 638 pin version. Further details and information on these parts is, at present, lacking.

Athlon 64 and Opteron Processors

					Multiplier
Athlon 64 & "Performance" Sempron					
Sempron 3100+	1800	Paris*	256	200	9.0X 754
Athlon 64 2800+	1800	Clawhammer	512	200	9.0X 754
Athlon 64 2800+	1800	Newcastle	512	200	9.0X 754
Athlon 64 3000+	2000	Clawhammer	512	200	10.0X 754
Athlon 64 3000+	2000	Newcastle	512	200	10.0X 754
Athlon 64 3200+	2000	Clawhammer	1024	200	10.0X 754
Athlon 64 3200+	2200	Newcastle	512	200	11.0X 754
Athlon 64 3400+	2200	Clawhammer	1024	200	11.0X 754
Athlon 64 3400+	2400	Newcastle	512	200	12.0X 754
Athlon 64 3500+	2200	Newcastle	512	200	11.0X 939
Athlon 64 3700+	2400	Clawhammer	1024	200	12.0X 754
Athlon 64 FX-51	2200	Sledgehammer	1024	200	11.0X 940
Athlon 64 3700+	2600	Newcastle	512	200	13.0X 754
Athlon 64 3800+	2400	Newcastle	512	200	12.0X 939
Athlon 64 FX-53	2400	Sledgehammer	1024	200	12.0X 940
Athlon 64 FX-53	2400	Sledgehammer	1024	200	12.0X 939

					Multiplier
Opteron**					
Opteron x40	1400	Sledgehammer	1024	200	7.0X
Opteron x42	1600	Sledgehammer	1024	200	8.0X
Opteron x44	1800	Sledgehammer	1024	200	9.0X
Opteron x46	2000	Sledgehammer	1024	200	10.0X
Opteron x48	2200	Sledgehammer	1024	200	11.0X
Opteron x50	2400	Sledgehammer	1024	200	12.0X

* The Paris core does not support 64-bit computing. It is included with the Athlon 64 because of the socket and because the integrated memory controller puts it ahead of the Athlon XP in performance.

** All Opterons are available in 1xx, 2xx, and 8xx versions. x=1 is for single processor systems, x=2 is for up to dual processor systems, and x=8 is for up to octal processor systems.

*** All system buses for Athlon XP, Sempron, Athlon 64, and Opteron are "double pumped", so their data rate is twice the bus speed. The multiplier is based off the listed speed.

With the Athlon 64, as the name suggests AMD added support for 64-bit addresses and integers. This was done by widening their pathways and registers, but it wasn't a radical redesign of the core Athlon architecture. It has a pipeline that was increased to 12/17 stages, it got SSE2 support added, and the system bus was switched to a HyperTransport bus. The longer pipelines allow it to scale to somewhat higher clockspeeds, and the HyperTransport buses - there are three in the Opteron - allow for better SMP, but the core remains essentially the same. The addition of x86-64 support has garnered a lot of attention, but so far it's pretty much marketing hype. It has *potential* to improve performance once 64-bit support arrives, but that *potential* has not yet been realized in the mainstream market. The scientific and academic community, however, has greeted the introduction of affordable 64-bit processing with open arms. Most consumers, meanwhile, are stuck waiting for Windows XP-64.

The reason for the superior performance of the Athlon 64 - in current 32-bit code as well as in 64-bit code to a lesser extent - lies mostly with the integrated memory controller, which *dramatically* reduces memory latencies. In effect, it helps to turn system RAM into a very large but still relatively slow L3 cache. It also continues to reduce memory latencies as clock speeds increase. Memory latencies on the Athlon XP were roughly 81 ns at 3200+ speeds, and the P4 3.2C was around 77 ns latency. Meanwhile, the Athlon 64 3400+ comes in at an astonishingly low 48 ns. As mentioned before, those latency figures are getting somewhat close to L3 cache values - for example, the L3 cache in a 3.06 GHz Xeon is about 10 ns. It's still four times slower, but it's also twice as fast as RAM on a P4 system.

No better example of this can be found than the newly introduced Paris core, a.k.a. Sempron 3100+. At 1.8 GHz, it is substantially slower than the fastest Athlon XP in core speed, and yet in typical use it outperforms even the Athlon XP 3200+. This from a part that has half as much cache as the Barton and Newcastle cores! The only area where it fails to keep up is in tasks that generally fit within the L1/L2 cache of the CPU, i.e. certain encoding tasks. In that case, the lack of raw clockspeed is a hindrance.

Of course, reduced latency isn't the entire story of the Athlon 64. In 64-bit mode, the number of useable registers for both integer and floating point operations has been doubled. Depending on the code being run, this could potentially bring 10 to 20 percent more performance. Certain applications that make heavy use of 64-bit integers can also benefit from the added 64-bit support - for example `crvtoaranhv` and encoding tools.

However, MMX and SSE have provided alternative means of improving 64-bit integer performance for many years now - they just require more programming effort to realize.

Page 6

Celeron, Pentium II and III Processors

I'm going to forego listing the various models of these processors for the time being. If anyone has a real desire to see them listed, feel free to let me know. If you're running one of these processors still, I can feel for you. I still use one at work, and I only upgraded from my P3 back in March. Given the price of upgrading, though - \$225 will get you a decent motherboard, 512 MB RAM, and an Athlon XP 2500+ - you really should upgrade if at all possible.

The old Pentium Pro P6 architecture was a 12-stage pipeline, more or less concluding with the Pentium III (more on that later). It had three specialized AGUs, two ALUs - one that handled simple instructions and a second one for the more complex instructions - and one FPU. The FPU also added support for SSE (which AMD lacked until the Athlon XP, but by then Intel was pushing the P4) and MMX - and they were generally faster on these instructions than AMD. That's not too surprising, considering that they created the technologies and AMD had to license them from Intel.

Intel could have certainly stuck with the design for a lot longer, as the last gasp Tualatin core offered pretty competitive performance clock for clock with the Athlon up to 1.4 GHz (the last Pentium III-S). In fact, the later 1.0A to 1.4A Celeron processors were very good overclocking chips, and a 1.1A running on a 133 MHz bus gave pretty decent performance. (I have just such a system powering my Home Theater PC.) Newer and better chipsets could have improved speed further, but Intel cut off the line and focused on pushing the Pentium 4 and NetBurst. This appears now to have been more of a marketing driven decision, although for the most part it can't be said that it was the worst idea ever.

Celeron 2 and Pentium 4 Processors

				Multiplier	L3 Cache
Pentium 4 and Celeron (Desktop)					
P4 1.3	1300	Willamette	256	100	13.0X 423
C 1.7	1700	Willamette	128	100	17.0X 478
P4 1.4	1400	Willamette	256	100	14.0X 423
P4 1.4	1400	Willamette	256	100	14.0X 478
C 1.8	1800	Willamette	128	100	18.0X 478
P4 1.5	1500	Willamette	256	100	15.0X 423
P4 1.5	1500	Willamette	256	100	15.0X 478
C 2.0	2000	Northwood	128	100	20.0X 478
P4 1.6	1600	Willamette	256	100	16.0X 423
P4 1.6	1600	Willamette	256	100	16.0X 478
C 2.1	2100	Northwood	128	100	21.0X 478
P4 1.7	1700	Willamette	256	100	17.0X 423
P4 1.7	1700	Willamette	256	100	17.0X 478
C 2.2	2200	Northwood	128	100	22.0X 478
P4 1.6	1600	Northwood	512	100	16.0X 478
C 2.3	2300	Northwood	128	100	23.0X 478
C 2.4	2400	Northwood	128	100	24.0X 478
C 2.5	2500	Northwood	128	100	25.0X 478
P4 1.8	1800	Northwood	512	100	18.0X 478
C 2.6	2600	Northwood	128	100	26.0X 478
C 2.7	2700	Northwood	128	100	27.0X 478
C 2.8	2800	Northwood	128	100	28.0X 478
P4 2.0	2000	Northwood	512	100	20.0X 478
P4 2.2	2200	Northwood	512	100	22.0X 478
C D 320	2400	Prescott	256	133.3	18.0X 478
P4 2.4	2400	Northwood	512	100	24.0X 478
C D 325	2533	Prescott	256	133.3	19.0X 478
C D 325/J	2533	Prescott	256	133.3	19.0X T/775
P4 2.26B	2267	Northwood	512	133.3	17.0X 478
C D 330	2667	Prescott	256	133.3	20.0X 478
C D 330/J	2667	Prescott	256	133.3	20.0X T/775
P4 2.4B	2400	Northwood	512	133.3	18.0X 478
P4 2.6	2600	Northwood	512	100	26.0X 478
P4 2.4A*	2400	Prescott	1024	133.3	18.0X 478
C D 335	2800	Prescott	256	133.3	21.0X 478
C D 335/J	2800	Prescott	256	133.3	21.0X T/775
P4 2.53B	2533	Northwood	512	133.3	19.0X 478
C D 340	2933	Prescott	256	133.3	22.0X 478
C D 340/J	2933	Prescott	256	133.3	22.0X T/775
P4 2.4C	2400	Northwood	512	200	12.0X 478
P4 2.66B	2667	Northwood	512	133.3	20.0X 478
P4 2.8B	2800	Northwood	512	133.3	21.0X 478
P4 2.6C	2600	Northwood	512	200	13.0X 478
P4 2.8A*	2800	Prescott	1024	133.3	21.0X 478
P4 2.8E	2800	Prescott	1024	200	14.0X 478
P4 520/J	2800	Prescott	1024	200	14.0X T/775
P4 3.06B HTT	3067	Northwood	512	133.3	23.0X 478
P4 2.8C	2800	Northwood	512	200	14.0X 478
P4 3.0E	3000	Prescott	1024	200	15.0X 478
P4 530/J	3000	Prescott	1024	200	15.0X T/775
P4 3.0C	3000	Northwood	512	200	15.0X 478
P4 3.2E	3200	Prescott	1024	200	16.0X 478
P4 3.2C	3200	Northwood	512	200	16.0X 478
P4 3.4E	3400	Prescott	1024	200	17.0X 478
P4 550/J	3400	Prescott	1024	200	17.0X T/775
P4 3.4C	3400	Northwood	512	200	17.0X 478
P4 560/J	3600	Prescott	1024	200	18.0X T/775
P4EE 3.2	3200	Gallatin	512	200	16.0X 478 2048
P4 540/J	3800	Prescott	1024	200	19.0X T/775
P4 570J	3800	Prescott	1024	200	19.0X T/775
P4EE 3.4	3400	Gallatin	512	200	17.0X 478 2048
P4EE 3.4	3400	Gallatin	512	200	17.0X T/775 2048
P4 580J	4000	Prescott	1024	200	20.0X T/775
P4EE 3.46	3467	Gallatin	512	266	13.0X T/775 2048
P4EE 3.73	3733	Prescott	2048	266	14.0X T/775

* Prescott 2.4A and 2.8A processors have HyperThreading Technology (HTT) disabled.

*** Front Side Bus (FSB) Speeds are "quad pumped", so Intel's FSB numbers are four times the actual bus speed on Pentium 4/Celeron, Pentium M/Celeron M, and Itanium processors. Multipliers are based off the base bus speed, not the FSB value.

NetBurst consists of a deep 20-stage pipeline coupled to an 8-stage fetch/decode unit. Due to the time spent fetching and decoding instructions, Intel created a new type of cache called a trace cache. This contained pre-decoded micro-ops, so for a large percentage of instructions, NetBurst runs as a 20-stage pipeline. Certain types of code run very well on NetBurst, while others - specifically branch-heavy code, like that seen in compilers and some games - do not. An incorrect branch prediction on P4 costs about twice as many lost cycles as an incorrect branch prediction on P3 or Athlon, which is why Intel added a more robust branch prediction unit.

The long pipeline allowed clockspeeds to scale very quickly with NetBurst. It was also a bandwidth hungry design, so increasing bus speeds combined with dual-channel memory eventually pushed the P4 beyond the reach of the Athlon XP. On the server front with the Xeon processors, the bandwidth was provided by adding L3 cache.

The Prescott further extended the NetBurst pipeline to 31 stages in addition to the 8 fetch/decode stages. For whatever reason, Intel generally describes the pipeline of the Prescott as 34 stages while only calling the

The Prescott further extended the netburst pipeline to 23 stages in addition to the 6 fetch/decode stages. For whatever reason, Intel generally describes the pipeline of the Prescott as 31 stages while only calling the earlier design a 20 stage pipeline. Besides the additional stages, Prescott doubled the L2 cache of the Northwood, added SSE3 support, and to the best of my knowledge contains deactivated x86-64 support - called EM64T by Intel and AMD64 by its creator AMD. Xeon versions of Prescott with the 64-bit support enabled are now shipping, and likely by the time XP-64 is released we will see 64-bit enabled desktop processors.

The Pentium 4 architecture also saw the introduction of Symmetric Multi-Threading (SMT) for Intel processors - they chose to call it Hyper Threading Technology (HTT). It appears to have been a part of the core from the very beginning, but Intel didn't enable the functionality until the P4 3.06 was launched, at which time it became available in the Xeon platforms as well. Later, it was enabled in all the 800 FSB "C" processors. Due to the length of the P4 pipeline, HTT allows the execution units to stay busy in the event of an incorrect branch prediction. The second thread can continue to run while the other thread recovers. In an ideal scenario, HTT could potentially increase performance by 20 or even 50 percent. In real world tests, however, rarely does it improve performance by more than 5 to 10 percent, and there are even times when it hurts performance.

With the switch to socket 775 LGA, Intel has also adopted model names. This likely has something to do with the recent difficulties Intel has encountered in scaling the NetBurst architecture to higher speeds. However, an even bigger problem is Intel's own Pentium M architecture (which is the next section). Anyway, we now have new model numbers which are supposed to reflect the overall capabilities of the chip, with higher numbers indicating more desirable chips. Comparing between families of chips should not be done based solely off the model number, however - there will certainly be instances where a 5xx chip offers better performance than a 7xx chip, and perhaps we'll also see some 3xx chips outperform their "superiors". For the time being, all of the 5xx chips are Prescott cores with 1 MB of L2 cache and an 800 MHz FSB. Future processors are also listed, and you can see where they will likely fall in the performance spectrum.

Mobile Celeron, Mobile P4, Celeron M and Pentium M Processors

Multiplier					
Mobile Pentium/Celeron Chips**					
MC 1.4	1400	Willamette	128	100	14.0X 478M
MC 1.5	1500	Willamette	128	100	15.0X 478M
MC 1.6	1600	Willamette	128	100	16.0X 478M
MC 1.7	1700	Willamette	128	100	17.0X 478M
MC 1.4	1400	Northwood	256	100	14.0X 478M
MC 1.8	1800	Willamette	128	100	18.0X 478M
MC 1.5	1500	Northwood	256	100	15.0X 478M
MC 2.0	2000	Willamette	128	100	20.0X 478M
MC 1.6	1600	Northwood	256	100	16.0X 478M
CM 353/J	900	Dothan	1024	100	9.0X 478M
MC 2.1	2100	Willamette	128	100	21.0X 478M
CM 333	900	Banias	1024	100	9.0X 478M
PM 900 (ULV)	900	Banias	1024	100	9.0X 478M
MC 1.7	1700	Northwood	256	100	17.0X 478M
MC 2.2	2200	Willamette	128	100	22.0X 478M
MC 1.8	1800	Northwood	256	100	18.0X 478M
CM 373/J	1000	Dothan	1024	100	10.0X 478M
MC 2.3	2300	Willamette	128	100	23.0X 478M
PM 1.0 (ULV)	1000	Banias	1024	100	10.0X 478M
MC 2.4	2400	Willamette	128	100	24.0X 478M
MC 2.0	2000	Northwood	256	100	20.0X 478M
PM 723/J (ULV)	1000	Dothan	2048	100	10.0X 478M
PM 1.1 (LV)	1100	Banias	1024	100	11.0X 478M
CM 350/J	1300	Dothan	512	100	13.0X 478M
MC 2.2	2200	Northwood	256	100	22.0X 478M
PM 1.2 (LV)	1200	Banias	1024	100	12.0X 478M
CM 320	1300	Banias	512	100	13.0X 478M
MC 2.4	2400	Northwood	256	100	24.0X 478M
PM 1.3	1300	Banias	1024	100	13.0X 478M
PM 718 (LV)	1300	Banias	1024	100	13.0X 478M
CM 330	1400	Banias	512	100	14.0X 478M
MC 2.5	2500	Northwood	256	100	25.0X 478M
CM 360/J	1400	Dothan	1024	100	14.0X 478M
MC 2.6	2600	Northwood	256	100	26.0X 478M
CM 340	1500	Banias	512	100	15.0X 478M
PM 1.4	1400	Banias	1024	100	14.0X 478M
PM 713 (ULV)	1400	Banias	1024	100	14.0X 478M
MC 2.7	2700	Northwood	256	100	27.0X 478M
CM 370/J	1500	Dothan	1024	100	15.0X 478M
MC D 325	2533	Prescott	256	133.3	19.0X T/775
MC 2.8	2800	Northwood	256	100	28.0X 478M
PM 1.5	1500	Banias	1024	100	15.0X 478M
PM 705	1500	Banias	1024	100	15.0X 478M
PM 733/J (ULV)	1400	Dothan	2048	100	14.0X 478M
PM 738/J (LV)	1400	Dothan	2048	100	14.0X 478M
MC D 330	2667	Prescott	256	133.3	20.0X T/775
MC D 335	2800	Prescott	256	133.3	21.0X T/775
PM 1.6	1600	Banias	1024	100	16.0X 478M
PM 715	1500	Dothan	2048	100	15.0X 478M
PM 758/J (LV)	1500	Dothan	2048	100	15.0X 478M
MC D 340	2933	Prescott	256	133.3	22.0X T/775
PM 1.7	1700	Banias	1024	100	17.0X 478M
MC D 345	3066	Prescott	256	133.3	23.0X T/775
MP4 2.8	2800	Northwood	512	133.3	21.0X 478M
MP4 2.8 HT	2800	Northwood	512	133.3	21.0X 478M
PM 735	1700	Dothan	2048	100	17.0X 478M
MC D 350	3200	Prescott	256	133.3	24.0X T/775
PM 730/J	1600	Dothan	2048	133.3	12.0X 478M
MP4 518	2800	Prescott	1024	133.3	21.0X 7478M
PM 745	1800	Dothan	2048	100	18.0X 478M
PM 753/J (ULV)	1800	Dothan	2048	100	18.0X 478M
MP4 3.0	3000	Northwood	512	133.3	22.5X 478M
MP4 3.0 HT	3000	Northwood	512	133.3	22.5X 478M
PM 740/J	1733	Dothan	2048	133.3	13.0X 478M
MP4 532	3067	Prescott	1024	133.3	23.0X 7478M
MP4 3.2 HT	3200	Northwood	512	133.3	24.0X 478M
MP4 538	3200	Prescott	1024	133.3	24.0X 7478M
PM 750/J	1867	Dothan	2048	133.3	14.0X 478M
PM 755	2000	Dothan	2048	100	20.0X 478M
PM 760/J	2000	Dothan	2048	133.3	15.0X 478M
MP4 552	3467	Prescott	1024	133.3	26.0X 7478M
MP4 558	3600	Prescott	1024	133.3	27.0X 7478M
PM 770/J	2133	Dothan	2048	133.3	16.0X 478M
PM 765	2400	Dothan	2048	100	24.0X 478M

** There are several chips in the mobile sector. PM is for Pentium M, MP4 is the Mobile Pentium 4, CM is the Celeron M, and MC is the Mobile Celeron (P4 core).

*** Front Side Bus (FSB) Speeds are "quad pumped", so Intel's FSB numbers are four times the actual bus speed on Pentium 4/Celeron, Pentium M/Celeron M, and Itanium processors. Multipliers are based off the base bus speed, not the FSB value.

With the scaling clock speeds of the Pentium 4, not even the specially designed Mobile versions were really suited for use in laptops. (Of course, they were still used, but Intel had other plans.) Higher clockspeeds mean higher power requirements as well as increased heat output, which makes it very difficult to get increased battery life. In response to pressure from companies such as Transmeta, Intel commissioned a design team in Israel to put together a high-performance, low-power processor. The end result was the Pentium M. Where the push for high clockspeeds was the driving force behind the NetBurst design, Pentium M was targeted at reaching specific thermal requirements. While specific details are rather hard to come by, since Intel is trying to protect its lead in the Mobile space, the Pentium M appears to be a modified version of the venerable P6 architecture.

One of the improvements made to the P6 architecture was a large L2 cache, which could be reserved and accessed in 256K sections. This allows large sections of the cache to be in a low-power "sleep" mode at any

One of the improvements made to the P4 architecture was a large L2 cache, which could be powered and accessed in 32K sections. This allows large portions of the cache to be in a low-power "sleep" mode at any given time, so they get the performance benefit of a large cache without incurring as much of the usual power increase. The L1 cache was also doubled from the PIII to 32K+32K data and instruction. Floating point performance was increased with the doubling of MMX/SSE units - although this really only helped with SSE optimized code - and there were a few other architectural changes. Overall, the Pentium M is able to provide performance that's roughly the equivalent of an Athlon processor of the same clock speed, while requiring much less power. Battery life in laptops that use the Pentium M can often be 25 to 50 percent longer than equivalent laptops that use the Mobile Pentium 4, Mobile Celeron or Mobile Athlon XP chips.

The length of the above chart should be an indication of how big the mobile market has become. One of the reasons for this increase in size is likely the cut-throat conditions that exist in the desktop CPU market. Intel charges a hefty premium for most of their mobile processors since, generally speaking, anyone looking for a high-performance laptop has more money to burn. This is what I call the "mobility tax"; you should only buy a laptop if portability is a *primary* concern; otherwise, your money will go a lot further with a desktop system. Certainly, business types that use computers for presentations and work on the road will be willing to pay this so-called tax.

With the release of the Dothan core Pentium M chips, Intel has also switched to model numbers. Here, however, there are many factors that influence the overall number. Ultra-Low Voltage processors running at lower clock speeds can end up rated higher than faster processors that require more power. This is supposed to reflect the relative desirability of certain features, as an increased battery life could be more important to some people than raw performance. Of course, Intel specifically states that the model numbers are not measures of performance, but only the technically literate are likely to know this. In their own words: "Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. See http://www.intel.com/products/processor_number for details."

Itanium and Itanium 2 Processors

				Multiplier		L3 Cache	
Itanium (Server)							
Itanium	733	Merced	96	66	11.0X	PAC-418	2048
Itanium	733	Merced	96	66	11.0X	PAC-418	4096
Itanium	200	Merced	96	66	12.0X	PAC-418	2048
Itanium	200	Merced	96	66	12.0X	PAC-418	4096
Itanium 2	900	McKinley	256	100	9.0X	PAC-611	1536
Itanium 2	900	McKinley	256	100	9.0X	PAC-611	3072
Itanium 2	1000	McKinley	256	100	10.0X	PAC-611	1536
Itanium 2	1000	McKinley	256	100	10.0X	PAC-611	3072
Itanium 2 LV	1000	Deerfield	256	100	10.0X	PAC-611	1536
Itanium 2 LV	1500	Deerfield	256	100	15.0X	PAC-611	1536
Itanium 2	1300	Madison	256	100	13.0X	PAC-611	3072
Itanium 2	1400	Madison	256	100	14.0X	PAC-611	4096
Itanium 2	1500	Madison	256	100	15.0X	PAC-611	6144

*** Front Side Bus (FSB) Speeds are "quad pumped", so Intel's FSB numbers are four times the actual bus speed on Pentium 4/Celeron, Pentium M/Celeron M, and Itanium processors. Multipliers are based off the base bus speed, not the FSB value.

Itanium processors are likely one of the least understood CPUs by most computer enthusiasts. Given that the cheapest models still cost over \$1000, that's not really surprising. These processors are meant to target the high-end corporate world. They are often used in massively parallel processing situations, and Itaniums are capable of working in up to 612-way SMP systems. Of course, that doesn't really explain *what* the Itanium is.

For starters, Itanium is the way that Intel envisioned 64-bit computing, and it is built on a new instruction set dubbed IA-64 (Intel Architecture 64).. IA-64 was a clean break from x86 legacy code, and it was designed for the future. Really, its competition isn't the Xeon or Opteron CPUs, although some mistakenly compare it with these processors. Itanium is meant to compete in the high-end corporate 64-bit computing world, going up against servers based on the IBM Power4/5, HP PA-RISC, Sun UltraSparc-III, and DEC Alpha. If none of those names ring a bell, that's not very surprising. The quad-processor IBM Power4 system that was used as the main server at a company I worked for (and they had two units for redundancy) cost somewhere in the neighborhood of \$500,000, and the RAID-5 array that provided data storage was another \$500,000. Perhaps more important than the hardware was the service contract with IBM that helped guarantee everything stayed running. The cost of the support contract with IBM (for dozens of such setups) was supposedly around \$300 million dollars a year!

The Alpha technology, interestingly enough, was purchased by Compaq, who merged with HP, and HP worked with Intel on the design of the Itanium, with the intention of using it in place of PA-RISC once it was complete. I believe that some (all?) of the Alpha technology was later transferred to Intel, most likely for use in furthering the design of the Itanium processors. Compaq/HP has continued to support this chip for the past several years, but they haven't invested a lot of money into researching new iterations of the design. This makes sense, since HP is encouraging its enterprise customers to switch to their Itanium platforms. Recently, HP announced that the 1.3 GHz (I think that was the speed) EV7 version of the Alpha chip will be the last.

These systems are often referred to as "Big Tin" systems, and they're in a league of their own. They are frequently used in systems that process *huge* amounts of data - their 64-bit addressing allows the use of many gigabytes of physical RAM - and they are usually optimized for input/output functions. Of course, reliability and up-time are far more important than actual performance numbers, and often once a system has been built around a specific architecture, large corporations will stick with that hardware unless there is *tremendous* incentive to switch to something else. Switching usually consists of several years of coding, testing, debugging, and validation - a task not to be undertaken lightly, to be sure.

For the processor design, Intel continued with their radical departure from accepted norms. Instead of a RISC or CISC approach, Intel went back to a technology that had been used in old mainframes and other computers of yore, VLIW (Very Long Instruction Word). Itanium is not a strict VLIW machine, though, as VLIW has some well known drawbacks that Intel worked to overcome, and Intel chose to call their new approach EPIC, "Explicitly Parallel Instruction Computer". In contrast to designs such as the Xeon and Opteron, which can issue up to three instructions per cycle, the Itanium 2 (forget Itanium 1 for a minute) can issue *eight* instructions per clock, and unlike VLIW designs, future Itanium chips could further increase the issue width without needing to recompile the code. In theory, then, a 1 GHz Itanium chip could perform roughly as fast as a 2.66 GHz Xeon/Opteron, or the 1.5 GHz Itanium 2 would be roughly as fast as a 4 GHz Xeon/Opteron. That's just theoretical performance, of course, and the overall system design will play a large role in determining how much of the potential of any system is actually realized.

To help reach that potential, Itanium chips run off a 128-bit quad-pumped system bus, using standard SDRAM (for the time being). The lower clock speeds combined with the wider bus make the SDRAM less of an issue than with high-speed desktop systems. The initial Itanium design, Merced, had four integer units (ALUs), two floating point units (FPUs), and three branch units (BRUs), two SIMD (i.e. MMX/SSE) units, and two load/store units - also called address generation units (AGUs) in other CPUs. The modified McKinley (and later) designs have six ALUs, three BRUs, two FPUs, one SIMD, two load units, and two store units - sort of like having 4 AGUs, except that they're more specialized. In addition, the McKinley has roughly three times the cache bandwidth as Merced. Merced was also a six issue design with a deeper pipeline (10 stages) and less memory bandwidth - a rather problematic design. McKinley and later designs are eight issue designs with shorter pipelines (8 stages) and more memory bandwidth. While Merced rarely made full use of its six issue design, McKinley's enhancements help it come a lot closer to issuing the maximum eight instructions per clock.

That doesn't really tell a whole lot about the architecture, and I don't really want to go much deeper than that right now. Suffice it to say that Itanium depends in a large part on compiler technology in order to reach its potential, and Intel has apparently had more difficulties in that area than they initially anticipated, but lately this seems to be less of a problem. The initial Merced design was also flawed, if you couldn't tell from my above description of the architecture, but Itanium 2 goes a long way toward rectifying the problems.

Many have called the Itanium a failure - coming up with such names as Itanic to describe the processor - especially now that AMD has launched Opteron and Intel is following suit with x86-64 support. However, they're really very different goals, and in the target market segment, Itanium is still managing to compete. Needless to say, it helps that Intel has very deep pockets thanks to the income generated from their desktop and mobile processor divisions. Itanium may or may not live in the long term, but short term Intel has plans to keep it around at least another three or four years, and they will likely keep it around longer to support existing clients. Honestly, though, I doubt any of us will ever be running an IA64 processor on our desktop systems.

Conclusion

Obviously, that's only a brief glimpse at the processor histories of AMD and Intel, with a vague picture of the future. Dual core designs should start appearing within the next year, and rumors of quad core processors are also floating around the web. At some point, we will likely reach the limits of current manufacturing technologies, but that day is still a long ways off. AMD and Intel both have technologies in development that should carry us past 45 nm process technologies, and probably down to single digits in our lifetime. That's assuming we don't get quantum computers first, that make all of the current binary systems seem quaint by comparison.

The amount of processing power sitting in front of you right now was beyond comprehension a couple decades ago. Even the "average" computers of today would seem amazing to people even one decade in the past. Ten years ago, 3D was only dreamt about, and professional 3D accelerators cost thousands of dollars while doing far less than a "cheap" GeForce 3 or Radeon 8500. Ten years ago, 32-bit processors were still looking for a real operating system, and 64-bit was only used by governments and research centers. Ten years ago, a 100 MHz processor was as good as it got. Ten years ago, few people had ever used a networked computer at home, and 28.8 modems were amazingly fast. Here's hoping the gurus at AMD, Intel, and other companies can continue to amaze us for another ten years!

Stay tuned for more insider articles from Jarred, including a much anticipated GPU cheat sheet as well!