The Linley Group

MICROPROCESSOR REPORT

www.MPRonline.com

THE INSIDER'S GUIDE TO MICROPROCESSOR HARDWARE

AMD's Fusion Finally Arrives

Integrated CPU/GPU Chips Strengthen AMD's Low-Power Play

By Tom Halfhill {12/6/10-01}

Tablet computers are the latest craze, making netbooks so...*2009*. So why are AMD's first integrated CPU/GPU Fusion chips intended mainly for netbooks?

For one, Fusion processors for desktop PCs aren't ready yet. Wait until next year. Second, the new processors aren't only for netbooks. If OEM customers want to use these low-power processors to build large-screen notebooks or even desktop PCs, AMD is happy to sell them the chips, no strings attached. And third, despite the hype over smartphones and tablets, netbooks remain a profitable market segment in which AMD has no presence whatsoever. If the struggling company can capture its usual 10% to 20% of the market, its share will be infinitely better than it is now.

In addition, staggering the launch of Fusion processors for notebooks and desktops will give AMD time to educate consumers about this new breed of microprocessors and about a new marketing effort to rate PC performance. By the time Fusion desktop processors appear in mid-2011, AMD hopes users will understand the difference between a Fusion chip with "Vision Premium" performance and a Fusion chip with "Vision Black" performance. Frankly, most users probably won't get it, but flashy point-of-sale placards and little stickers on the computers will be their guideposts. (See the sidebar, "AMD's Blurry Vision.")

Regardless of such gimmicks, integrating a GPU and CPU on the same chip is the biggest leap in microprocessors since their invention almost 40 years ago. As Figure 1 shows, a single Fusion chip now has all the essential elements of a PC: the CPU, GPU, memory controller, and north-bridge system controller. AMD refers to these chips generically as accelerated processing units (APUs). They are an important step beyond Intel's year-old "Arrandale" processors, which package separate CPUs and GPUs in multichip modules, and they match the integration of Intel's "Pineview" netbook processors, which combine the CPU and GPU on a single die.

Previous feats of integration were smaller steps. On-chip L1 caches were the norm by the 1980s, but integrating small amounts of local memory isn't as significant as integrating a peer processor. Microprocessors absorbed FPUs in the 1990s, but most users scarcely noticed floating-point performance back then (or even today). Microarchitecture innovations such as superscalar execution, instruction

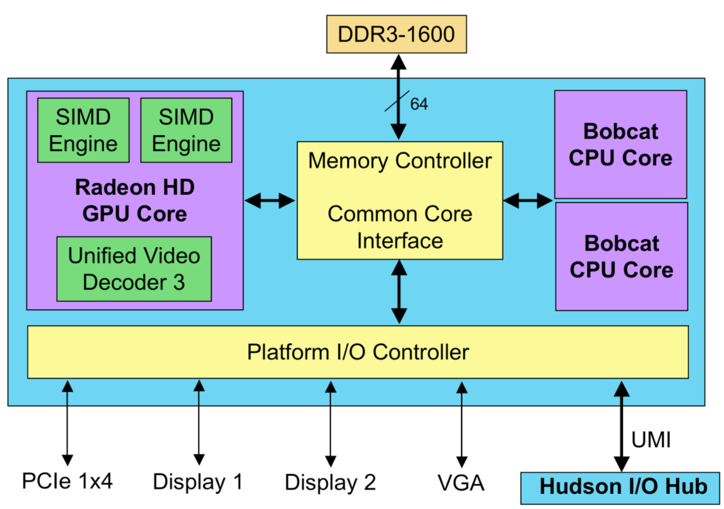

Figure 1. AMD Fusion block diagram. Dubbed an accelerated processing unit (APU), this highly integrated chip has two CPU cores and a GPU. The platform I/O controller provides additional north-bridge functions, linking the APU to external devices and to the "Hudson" south-bridge I/O hub. (Source: AMD)

reordering, and multithreading are a rising tide that lifts all software boats, but an integrated GPU can boost actual and perceived performance much higher, even with some non-graphics workloads. For most PCs other than professional workstations and enthusiast gaming rigs, integrated processors are the future.

Better Late Than Never

CPU vendors have labored to integrate GPUs for more than 15 years. In the 1990s, Cyrix enjoyed some success with its MediaGX processor, but that chip was born with poor graphics and never matured beyond low-performance applications. Later, AMD acquired Cyrix and sold MediaGX and derivatives for embedded systems. (See *MPR 3/10/97-01*, "MediaGX Targets Low-Cost PCs.") Intel designed a highly integrated CPU/GPU in the late 1990s but aborted the project in 2000. One reason was mediocre graphics—a perennial problem for Intel. Another flaw was a main-memory interface that supported unpopular Rambus RDRAM instead of industry-standard SDRAM. (See *MPR 10/16/00-03* , "Intel Deserts Timna.")

Currently, Intel's Pineview processors for netbooks combine two Atom CPUs with an Intel-designed GPU on a single chip, but the GPU is severely underpowered and widely mocked by reviewers. Intel's "Arrandale" mobile processors link a somewhat more powerful GPU with various Intel x86 CPUs on separate die in a multichip module—a compact but costlier alternative to single-chip integration. (Intel markets Arrandale processors under a variety of brand names: Core i3, Core i5, Core i7, Celeron, and Pentium.) At the January 2011 Consumer Electronics Show in Las Vegas, Intel is expected to introduce its first high-performance CPU/GPU chips based on the new Sandy Bridge microarchitecture. (See *MPR 9/27/10-01*, "Sandy Bridge Spans Generations.")

AMD has wrestled with CPU/GPU integration, too. Fusion is two years late, squandering an opportunity for the company to open a strategic lead over Intel. In 2006, after acquiring GPU vendor ATI for $5.4 billion, AMD said Fusion would debut in late 2008 or early 2009. Instead, the first production chips emerged from AMD's Singapore test and packaging factory on November 8, 2010. (The die are fabricated in Taiwan by TSMC, which also manufactures AMD's discrete GPUs.) OEMs will begin shipping Fusion-based systems in 1Q11, thereby missing the holiday shopping season.

Despite their tardiness, the initial Fusion chips are competitive low-power processors for netbooks and notebooks. All use AMD's new "Bobcat" x86-compatible CPU core, which AMD hopes will be an Atom-smasher. Like Intel's low-power processor, Bobcat is smaller and more power efficient than the CPU cores in mainstream x86 desktop/server processors. Both Atom and Bobcat are dual-issue superscalar machines with 16-stage pipelines, so they should have similar throughput rates and similar clock speeds in comparable fabrication technology. Bobcat, however, uses instruction reordering instead of multithreading, so it probably has a slight edge in single-thread performance. (See *MPR 8/30/10-02*, "AMD's Bobcat Snarls at Atom.")

AMD derived the GPU core in the first Fusion chips from a discrete notebook GPU, the Radeon HD 5430, which made its debut at the Consumer Electronics Show in January 2010. This Fusion GPU is compatible with DirectX 11, Microsoft's latest 3D-graphics application programming interface (API). By contrast, the weak GPU in Intel's Pineview processors is limited to DirectX 9. It's so old that it has trouble running some of the benchmark suites that AMD and independent reviewers are using to measure graphics performance.

	Ontario C-30	Ontario C-50	Zacate E-240	Zacate E-350
CPU Cores	1x Bobcat	2x Bobcat	1x Bobcat	2x Bobcat
CPU Frequency	1.2GHz	1.0GHz	1.5GHz	1.6GHz
L1 Cache (I + D)	32KB + 32KB per CPU			
L2 Cache	512KB	1MB	512KB	1MB
GPU Core	Radeon HD 6250		Radeon HD 6310	
GPU Frequency	280 MHz		500 MHz	
Memory Control	1x 64-bit DDR3-1066, 1.35V–1.5V, 2 DIMMs			
Video Interfaces	2xDisplayPort/1xHDMI/1xDVI/1xLVDS, + VGA with integrated DAC			
PCI Express	PCIe 4x1 or 1x4 Gen2			
Other I/O	LCD backlight, I^2C			
Die Size	75mm^2			
Package	413-ball BGA, 19mm			
IC Process	TSMC 40nm bulk CMOS			
Power (TDP)	9W		18W	

Table 1. AMD's first Fusion processors. All are based on the same die and began shipping to OEMs in November. AMD's power-consumption ratings are similar to Intel's thermal design power (TDP) specification. Note that Ontario and Zacate are AMD's first PC processors manufactured by TSMC. (Source: AMD)

One Die, Four Products

As Table 1 shows, AMD is initially shipping four Fusion processors, all based on the same die. CPU and GPU clock frequencies determine whether the product code-name is Ontario or Zacate and whether the GPU is named the HD 6250 or HD 6310. (To decipher other related AMD code-names, see the sidebar, "AMD & Intel Code-Name Glossary.") Each chip is available as a single-CPU or dual-CPU processor, allowing AMD to save money by selling die with a defective CPU that's been disabled at the factory.

Ontario processors, at 9W TDP, are the lower-power choice for netbooks and the newer category of thin-and-light subnotebooks (inspired by Apple's Macintosh Air). Zacate processors, at 18W TDP, are for small notebooks and nettop PCs. The faster chips are evidently operating at a higher voltage, because the difference in power

AMD's Blurry Vision

In the never-ending quest to explain computing performance to mere mortals, AMD is promoting Fusion with a new marketing campaign that dispenses with clock frequencies, benchmark scores, and abstract numerical speed ratings. Instead, AMD has created five performance labels that correspond to common PC applications. The following figure illustrates the labels and aligns them with their applications.

Associating PC performance with applications is worthwhile, but the labels are mysterious. To begin with, they're inconsistent. AMD calls the lowest performance level "HD Internet," whereas all the others are some variation of "Vision": just plain Vision (which AMD calls "Vision Basic"), Vision Premium, Vision Ultimate, and Vision Black.

Another problem is that Microsoft uses similar monikers (Premium, Ultimate) to distinguish versions of Windows 7. That similarity could mislead people into thinking they need a Vision Premium processor to run a Windows Premium operating system. Also, without a cheat sheet like this figure, it's not apparent that Vision Ultimate is really Vision Penultimate, because Vision Black is higher. In any case, few users will remember which applications correspond to each sticker without referring to a point-of-sale placard.

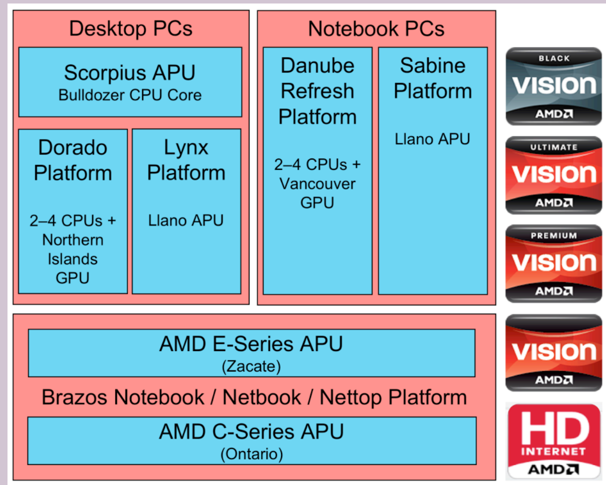

The above figure aligns the five labels with the code-names for AMD's 2011 platforms. Ontario is the lowest-ranking AMD processor, meriting only the HD Internet label. AMD compares it with mobile PCs using Intel's Atom and Celeron processors. Moving up the ladder, AMD compares Zacate (Vision Basic) with mobile PCs using Intel's Pentium processors, Vision Premium with Intel's Core i3 processors, and Vision Ultimate and Vision Black with Intel's Core i5 and Core i7 processors. (To decipher AMD code-names, see the sidebar, "AMD & Intel Code-Name Glossary.")

Marketing microprocessors to consumers has never been easy (see *MPR 12/14/09-01*, "CPU Marketing: The Next Frontier"). AMD's latest attempt to move beyond inscrutable numbers is admirable, but the campaign needs terminology that's more intuitive and consistent.

consumption is not linear with the difference in clock frequency. All these processors integrate the same north-bridge features, including a DDR3-1600 memory controller, three video interfaces, and four PCI Express (PCIe) lanes. Figure 2 shows a die photo of Ontario/Zacate.

AMD has also introduced a companion south-bridge chip code-named Hudson and officially called the M1 Fusion Control Hub (FCH). It's packaged in a 605-pin 23mm BGA and manufactured in TSMC's 65nm CMOS process. It has 14 USB 2.0 interfaces, two USB 1.1 interfaces, 4×1 PCI Express Gen2 lanes (in addition to the ones on the processor chip), and six Serial ATA interfaces (SATA-III, 6Gbps). In addition, the hub supports high-definition audio, performs some power-management functions, controls the fans, and generates clock signals.

Hudson is almost identical to AMD's SB810 I/O hub, an existing south-bridge chip for notebooks. It's also similar to Intel's CG82NM10 platform controller hub (PCH)—the south-bridge chip for Pineview. Intel's I/O hub has eight USB 2.0 interfaces, 4×1 PCI Express Gen2 lanes, and two SATA interfaces. It also supports Intel's HD Audio and AC97 audio. Intel manufactures this hub in older 130nm technology and specifies 2.1W maximum thermal design power (TDP).

Together, an AMD Fusion APU with Hudson I/O hub comprises a nearly complete mobile-PC platform code-named Brazos. The I/O hub consumes 2.7W to 4.7W, depending on the platform configuration and usage model, so the total Brazos TDP (excluding external memory) ranges from 11.7W to 22.7W.

Brazos Embraces Multiple Options

Greater integration can limit flexibility by imposing chip-design decisions on system designers. It's unavoidable to some extent, but AMD has taken pains with Brazos to support multiple system configurations—even if system designers prefer a discrete GPU over the integrated GPU.

As the architecture diagram in Figure 3 shows, a discrete GPU (code-named Vancouver) can connect directly to a Fusion chip over two or four PCIe lanes. These lanes support either PCIe Gen1 (250MB/s per lane) or PCIe Gen2 (500MB/s per lane). They are separate from the four PCIe Gen1 lanes in the Unified Media Interface (UMI) to the Hudson I/O hub. Although a discrete GPU makes the integrated GPU redundant—AMD doesn't plan to split Fusion chips into CPU-only Fission chips—it's a good option for notebooks and nettops needing higher performance.

Ontario has a Radeon HD 6250 GPU and Zacate has a Radeon HD 6310 GPU, but both chips are based on the same die, so their GPUs are identical. The only difference, as mentioned above, is their clock speeds: 280MHz (Ontario) or 500MHz (Zacate). Despite their misleading 6000-series names, these GPUs are derived from AMD's Radeon HD 5430, a discrete mobile GPU introduced at the Consumer Electronics Show in January 2010.

The Ontario/Zacate graphics engine has two SIMD blocks, each with 40 of AMD's "stream cores"—once known as pixel shaders but now known as programmable ALU/FPUs. Each function unit can perform two 32-bit integer or two 32-bit floating-point operations per clock cycle, so the GPU's maximum theoretical throughput is 44.8 gigaflops at 280MHz (Ontario) or 80 gigaflops at 500MHz (Zacate). When rendering 3D graphics, the GPU can process one polygon per clock cycle.

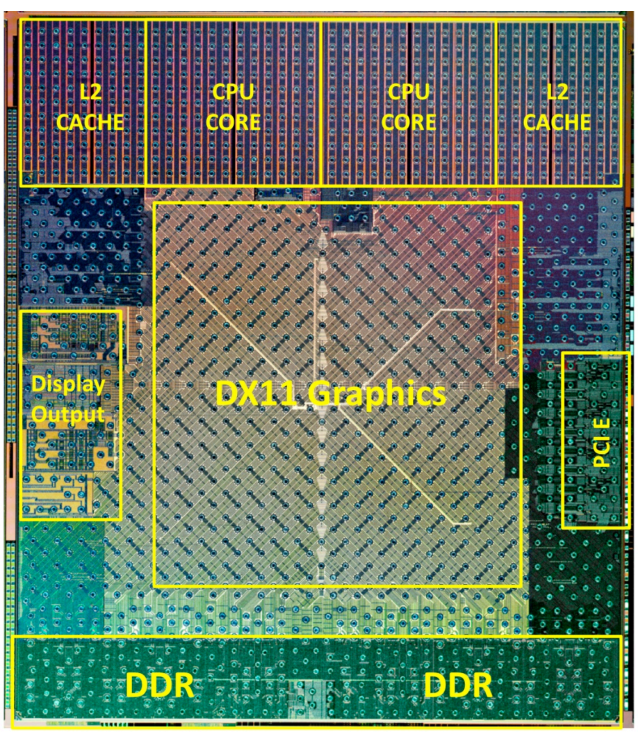

Figure 2. AMD Ontario/Zacate die photo. Both Fusion processors have the same die, but Ontario's CPUs and GPU run more slowly than Zacate's. The GPU core dominates this 75mm2 die, occupying about 25mm2 (34%). By contrast, each CPU core is only about 4.5mm2 (6%). (Source: AMD)

In addition, the GPU has AMD's third-generation Unified Video Decoder (UVD3), which accelerates decoding for MPEG-2, MPEG-4 Part 2 ASP (DivX/Xvid), and Multiview Video Coding (MVC). Dedicated hardware completely offloads decoding from the CPU and supports resolutions up to 1080p. The addition of MVC to the third-generation version of this engine supports Blu-ray 3D playback by decoding two stereo views encoded in an H.264 video stream. Note that the discrete GPU from which AMD derived this integrated GPU has the second-generation UVD2 engine without MVC.

Although Fusion's integrated GPU gets all the buzz, the north-bridge integration is impressive, too. Usually, only consumer SoCs and cell-phone application processors have on-chip display interfaces. To satisfy the varying needs of netbooks, laptops, and desktop PCs, Ontario and Zacate have three such interfaces: two that support high-resolution display standards, plus a VGA interface with an integrated digital-to-analog converter (DAC). Both of the high-resolution interfaces support DisplayPort, HDMI, and DVI, and one of them also supports LVDS (low-voltage differential signaling).

One disappointment is the omission of SuperSpeed USB 3.0, which offers 10 times the I/O bandwidth of Hi-Speed USB 2.0 (4.8Gbps versus 480Mbps). But then, no existing south-bridge chips integrate USB 3.0—not even the newest Intel chips for desktop PCs. AMD says the USB 3.0 standard wasn't quite ready for Hudson but will appear in future I/O hubs. Although separate USB 3.0 controllers are available—some notebook PCs already have them—the

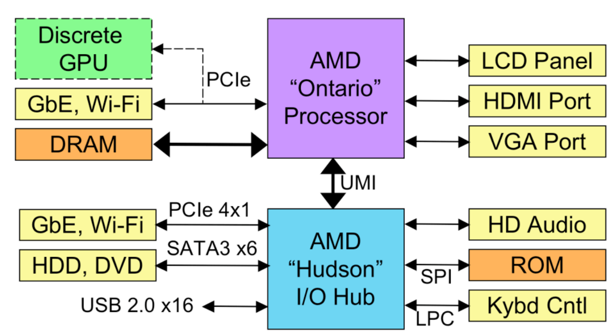

Figure 3. Brazos system-architecture diagram. Both the Fusion APU and Hudson I/O hub have PCI Express interfaces, so the APU isn't wholly dependent on the hub for system I/O. The APU's PCIe interfaces support an optional discrete GPU, bypassing the integrated graphics engine. They support PCIe Gen1 or Gen2, giving system designers the choice of faster throughput or lower power.

additional power consumption is undesirable in smaller-battery netbooks. USB 3.0 peripherals are only beginning to appear.

Netbooks Versus Notebooks

In a typical netbook configuration, using only two primary chips, Brazos supports an LCD screen, HDMI port, VGA port (for external video monitors), Gigabit Ethernet, Wi-Fi, and a SATA-III hard drive or solid-state drive (SSD). Internally, some of the 16 USB interfaces can support a 3G wireless-network adapter and a Secure Digital (SD) card slot, which are popular features in netbooks. Externally, they support user-accessible USB ports. Hudson also integrates the USB physical-layer (PHY) interfaces, eliminating the need for external PHY chips.

With the same two primary chips, the Brazos platform supports a typical notebook-PC configuration. Essentially, it's the same as the netbook configuration, but it optionally substitutes the higher-speed Zacate APU for Ontario and optionally bypasses the integrated GPU in favor of a discrete GPU. In addition, the SATA interfaces can also support an optical drive.

Bypassing the integrated GPU for a discrete GPU may seem illogical. Why use an integrated graphics processor but not the integrated graphics? Actually, this option broadens the market for Ontario and Zacate. OEMs can use the same motherboard or a similar one in multiple designs, simplifying development work and cutting costs. AMD's substitute for the integrated GPU—"Vancouver"—is actually a family of Radeon Mobility GPUs scheduled for introduction in 1H11. All support DirectX 11, and some will be second-generation designs with faster DirectX engines and Blu-ray 3D support.

For mainstream notebooks, AMD will offer two GPUs that hold power consumption below 20W TDP: "Robson" (first generation) and "Seymour" (second generation). For higher-performance notebooks, AMD will offer two GPUs in the 20–30W TDP range: "Capilano" (first generation) and "Whistler" (second generation). Additional Vancouver GPUs exceeding 35W TDP will support gamer notebooks. All these discrete GPUs will outperform the integrated graphics in Ontario and Zacate. One disadvantage of integration is that discrete processors tend to evolve faster than a GPU core frozen early in a chip design.

Despite their features and integration, Ontario and Zacate have little potential in consumer-electronics products beyond PCs. Their power consumption isn't competitive in that market. They still require active cooling (fans), whereas consumer-electronics manufacturers prefer quiet and reliable passive cooling. The same limitation rules out this generation of Fusion processors in smaller computing devices, such as tablets and smartphones. Their rich display interfaces won't be wasted, however, because even some existing netbooks and laptops have HDMI ports for HDTVs.

Atom Chipsets Use Less Power

Intel's Sandy Bridge processors—set to debut in January—will be faster chips intended for higher-performance notebook and desktop PCs. In mobile Sandy Bridge chips, we expect the CPUs to run at 2.2GHz to 2.7GHz—68% faster than the 1.6GHz dual-CPU Zacate processor. We expect Sandy Bridge GPUs to cruise at 650MHz to 850MHz, briefly surging to 1.3GHz using Intel's Turbo Boost technology. Nevertheless, the Sandy Bridge GPU executes only 83 gigaflops at 850MHz, barely edging Zacate's GPU (80 gigaflops at 500MHz).

Of course, the tradeoff for higher performance is higher power consumption. We expect mobile versions of Sandy Bridge to burn 35W to 55W TDP. Even the low end of that range is 54% hotter than Zacate's maximum TDP (22.7W)—and the Sandy Bridge estimate doesn't include an I/O hub like AMD's Hudson. Sandy Bridge aims for a different mobile market segment.

For netbooks, small notebooks, and nettops, Intel's graphics-integrated Atom processors are a closer match for Ontario and Zacate. These Atom processors are code-named Pineview-M for mobile systems (launched in 3Q10) and Pineview-D for desktops (2Q10). On a single chip, Pineview integrates two Atom CPUs with a GPU and north-bridge functions. The Intel-designed GPU is the Graphics Media Accelerator 3150 (GMA3150). The paired south-bridge chip is the CG82NM10 platform controller hub (PCH), code-named Tiger Point. The resulting two-chip platforms are code-named Pine Trail-M and Pine Trail-D.

Intel specifies 10.6W maximum TDP for the Pine Trail-M chipset, which includes the 1.5GHz Atom N550 Pineview-M processor. Intel specifies 15.1W maximum TDP for the Pine Trail-D chipset, which includes the 1.8GHz Atom D525 Pineview-D processor. As Figure 4

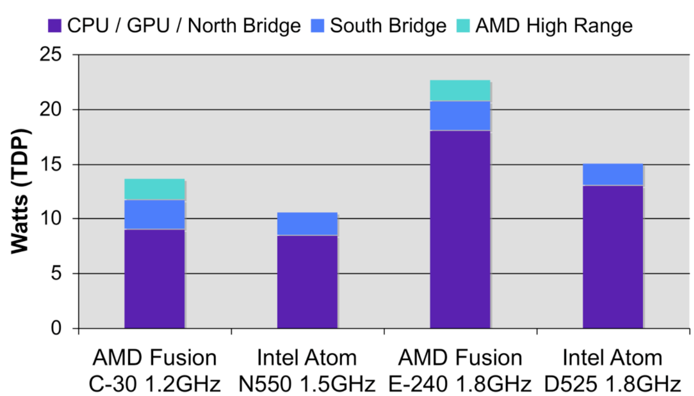

Figure 4. Power consumption for AMD's Brazos platform versus Intel's Pine Trail platforms. AMD's and Intel's power consumption is virtually the same at the low end of the Brazos range, but the Atom CPU runs faster. Despite using the most power of all of these combinations, AMD's 1.6GHz Zacate processor will have trouble competing with the CPU performance of the 1.8GHz Atom D525. (Source: AMD and Intel)

Price & Availability

AMD began shipping the first Fusion chips to OEM customers in early November. Systems are expected to reach the market in 1Q11. AMD hasn't publicly announcing chip pricing. For more information, access *http://fusion.amd.com*.

shows, Pine Trail-M uses virtually the same power as AMD's lowest-power Ontario Brazos chipset, which has only a single 1.2GHz CPU. Pine Trail-M uses about 3.0W less than AMD's 1.0GHz dual-CPU Ontario Brazos chipset. Zacate suffers in this comparison, too. The Pine Trail-D chipset uses significantly less power than either the 1.5GHz single-CPU or 1.6GHz dual-CPU Zacate Brazos chipset. Overall, Intel delivers higher CPU clock speeds and more CPU performance for less power.

Although they burn less power, Pineview chips are larger than Ontario and Zacate (87mm^2 versus 75mm^2). Intel manufactures Pineview in an older 45nm process, not its latest 32nm technology. That difference, we estimate, enlarges Intel's die about 10–20% compared with the TSMC 40nm process that AMD uses for Ontario and Zacate. But Intel compensates for its larger die with lower current leakage. Transistors built in Intel's 45nm process have high-*k* metal gates, and we estimate that Ontario and Zacate have twice as many transistors as Pineview's 176 million.

AMD hasn't announced pricing for Ontario and Zacate, but Intel's Pineview pricing is a clue: the Atom N550 costs $86, and the Atom D525 costs $63. Usually, AMD prices its products lower than Intel's unless it has a clear advantage.

AMD's Graphics Crush Pine Trail's

Pineview excels in CPU performance per watt but suffers badly in video flexibility and graphics performance compared with Ontario and Zacate. Intel's GMA3150 integrated graphics processor is only a minor improvement over the GMA950 built into the older Atom N270 north-bridge chip. The GMA3150 supports only DirectX 9, and it has only one LVDS interface (for the system LCD) and a VGA interface. By contrast, Ontario and Zacate support DirectX 11 and have LVDS, VGA, and additional interfaces that support DVI, DisplayPort, and HDMI.

Graphics performance isn't even close. Independent benchmark tests on 1.66GHz Pine Trail netbooks have reported FutureMark 3DMark06 scores in the 60s—and that with difficulty, because some tests in the suite won't run properly on the GMA3150. Using the same benchmark suite, AMD quotes graphics scores of 1,748 for the 1.0GHz Ontario and 2,399 for the 1.6GHz Zacate. A few additional watts are easier to swallow when they buy a graphics processor that's 30 times faster. Intel's graphics are so bad that Nvidia has built a tidy sideline business by selling Ion graphics chips for Atom processors.

AMD says that most PCs will cost less than $500 by 2012 and that graphics performance matters more than CPU performance in the sub-$400 market segment. If that's true, Ontario and Zacate are well positioned. In addition to outperforming Intel's graphics-integrated Atom chips, they also outrun some less-integrated notebook processors, including AMD's own chips.

Previous AMD mobile platforms integrate Radeon graphics in the south-bridge chip. On 3DMark06, Zacate is 10% to 50% faster than AMD's two most popular platforms, code-named Danube and Nile. Danube pairs an Athlon II P320 CPU with Radeon HD 4250 graphics; Nile pairs an Athlon II Neo K325 CPU with Radeon HD 4225 graphics.

Ontario and Zacate also consume less power than AMD's other mobile platforms. One reason is that integration eliminates the high-power signals between the CPU and GPU. At 21W TDP in a typical notebook configuration, the Brazos chipset uses 16% less power than Nile, which was AMD's lowest-power mobile platform when unveiled only seven months ago. On 3DMark06, Brazos uses 40% less power (6.5W versus 10.83W). On MobileMark 2007, Brazos uses 34% less power (2.71W versus 4.14W). AMD says the MobileMark 2007 measurements are low because that benchmark test includes significant idle time, allowing both platforms to power down their CPU, GPU, and north-bridge logic.

AMD has benchmarked the fastest Zacate processor at 8.47 frames per second with Cinebench OpenGL, which measures 3D-graphics performance and is also relevant for video streaming. That's more than twice as fast as a current Athlon II N330 processor (2.3GHz, dual cores) and more than eight times faster than a Pentium P6000 (1.86GHz, dual cores), a 35W TDP Arrandale notebook processor.

GPUs Go Beyond Graphics

CPU-intensive benchmarks tell a different story. On the PCMark Vantage Productivity suite, the Athlon II N330 and Pentium P6000 were almost twice as fast as Zacate, which scores only 2,300. And, as noted above, Pineview Atom CPUs can reach higher clock speeds while using less power, negating the slight advantage that AMD's Bobcat CPU should enjoy by executing instructions out of order.

Productivity-software workloads are primarily integer code. Fusion's real strengths are graphics and floating-point performance—if the software can use both the CPU and GPU. Zacate's maximum theoretical floating-point throughput is 93 gigaflops, which includes 80 gigaflops for the 500Hz GPU and 13 gigaflops for the 1.6GHz CPU. Even Intel's architecturally superior Sandy Bridge processors will have trouble beating that performance, despite using much more power.

Traditional PC benchmarks test the CPU and GPU separately. These metrics need revising to fully describe the performance of an integrated CPU/GPU chip. Increasingly, programmers are using GPUs for special tasks that run too slowly on CPUs. Some tasks are purely computational and don't involve graphics, whereas others use graphics only to illustrate the results. This trend—general-purpose GPU (GPGPU) computing—started in the scientific community but is spreading to consumer applications.

A popular scientific application is protein folding, which evaluates potential medicines before spending money on lab testing. Energy companies use GPUs to analyze sonar probes of underground oil and gas deposits; investment banks try to exploit small movements in commodities and currencies. In the consumer market, programmers are using GPUs to accelerate Photoshop filters, transcode digital video, and stabilize the shaky home movies shot with handheld camcorders. Typically, these tasks use intensive floating-point math and have lots of data parallelism.

Nvidia has been quicker to pursue these applications, but AMD isn't far behind. (See *MPR 10/5/09-01*, "Looking Beyond Graphics," and *MPR 12/22/08-01*, "AMD's Stream Becomes a River.") Fusion offers new opportunities, because it tightly couples the CPU and GPU with an on-chip I/O interface that's much faster than board-level PCIe. To help software developers use their GPUs, both AMD and Nvidia are promoting OpenCL. (See the sidebar, "OpenCL Tries to Standardize Parallel Programming," in "AMD's Stream Becomes a River.") Integrated processors will change how PC performance is achieved, measured, and experienced.

Integration Threatens Disintegration

CPU/GPU integration also promises (or threatens, depending on your point of view) to reshape the PC industry. By itself, graphics integration gives no advantage to AMD, because Intel is doing the same thing. Graphics *performance* is what matters. In this contest, AMD has a clear advantage, because the ATI acquisition gives the company experienced design teams in both disciplines, and AMD has had four years to meld those teams. Now, all AMD has to do is execute—a frequent problem for the company.

Intel, on the other hand, is infamous for poor graphics performance. One reason is that Intel aims for the bulk of business users and consumers who don't particularly need fast graphics. But another reason, apparently, is that Intel has lacked the ability or urgency to develop competitive high-performance GPUs. Intel's latest attempt to develop a discrete GPU recently suffered a setback when the Larrabee project stumbled (see *MPR 6/14/10-02*, "Intel Adapts Larrabee for HPC"). Simply put, Intel has never excelled in this field.

One alternative is to acquire the necessary expertise and intellectual property. Intel has certainly been acquisitive lately, spending $7.9 billion for malware-masher McAfee and $1.4 billion for Infineon's wireless operations (see *MPR 8/2/10-01*, "Intel Shakes Up Cellular Market"). The most likely acquisition target for world-class GPU technology is Nvidia—another multibillion-dollar gulp that would raise antitrust issues and be difficult to digest even if swallowed.

Licensing a GPU isn't out of the question for Intel, which already licenses PowerVR engines from Imagination Technologies. Intel owns 14% of the British company, inciting rumors that it plans to buy the whole operation. But PowerVR engines are small GPU cores optimized for lower-power systems, such as netbooks and smartphones. They would need major improvements to compete with AMD in the broader PC market. Licensing Nvidia's superior GPU technology would seem to make more sense, except that Intel and Nvidia are engaged in a bitter legal battle over bus licensing, which makes a closer relationship unlikely—unless they get desperate.

Nvidia may indeed get desperate, if integrated processors decimate the discrete-GPU market. Nvidia insists that won't happen. The company says that professional users and gamers are numerous enough to preserve a healthy market and that demand for better graphics and GPGPU computing is growing. Today, millions of users are satisfied with the lackluster graphics in their low-end home computers, business PCs, and notebooks. Integrated processors will give them better graphics without hurting Nvidia's discrete business, beyond the damage already done. Better integrated graphics may even whet their appetites and inspire upgrades.

More likely, Nvidia's PC revenue will continue to erode unless the company can follow the integration trend. Nvidia can keep trying to sell its own integrated processors (ARM-based Tegra chips for smartphones and tablets), license GPU technology to the only large x86 vendor that needs it (Intel), or become an x86 vendor itself (as persistent rumors suggest). The established high-end market of gamers and workstation users is probably large enough to guarantee that Nvidia won't go the way of Weitek, a once-thriving FPU vendor that vanished after CPUs began integrating FPUs. But when integrated CPU/GPUs catch on, the market for discrete graphics may be smaller than Nvidia's current revenue, forcing some retrenchment if efforts to broaden its product lines fail.

Much depends on the evolutionary trajectory of GPUs. Historically, discrete GPUs evolve faster than integrated GPUs, so discrete graphics are always superior. If integrated processors dominate the market, they will attract more-aggressive development, but they will always have trouble matching the performance of discrete GPUs on larger silicon. On the other hand, CPU-GPU cooperation is inherently faster when both processors are closely coupled on the same chip, and the rising number of nontraditional applications for GPUs makes that tight union even more compelling. In any case, this leap in microprocessor evolution will almost certainly mean big changes for the industry, and Fusion positions AMD to make the most of those changes. ♦

AMD & Intel Code-Name Glossary

Arrandale: Intel code-name for a CPU chip and GPU chip packaged together in a multichip module. Product names vary by the CPU and span Intel's product line: Core i3, Core i5, Core i7, Celeron, and Pentium.

APU: accelerated processing unit. AMD's generic term for a Fusion chip that integrates an x86-compatible CPU and an AMD GPU on the same die.

Bobcat: AMD code-name for a new low-power CPU core in x86 mobile processors. The first two chips with Bobcat are code-named Ontario and Zacate.

Brazos: AMD code-name for a mobile-PC platform that uses CPU/GPU integrated chips. Brazos includes Zacate (CPU/GPU for notebooks) and Ontario (CPU/GPU for netbooks).

Bulldozer: AMD code-name for a new high-performance CPU core in x86 desktop and server processors. Supersedes K10.

C-series: AMD part-number prefix for lowest-power mobile-PC processors intended mainly for netbooks. Ontario is the first C-series chip. Known variations are the C-30 and C-50.

Danube: AMD code-name for a notebook-PC platform based on the Athlon-II or V-120 processor, AMD RS880 north-bridge chip, and AMD SB810 south bridge with Radeon HD 4250 GPU. Will be superseded in 2011 by Danube Refresh, probably with Bobcat CPU cores and a Vancouver GPU.

Danube Refresh: AMD code-name for a notebook-PC platform based on a next-generation CPU/GPU chip. Unlike Sabine, a similar platform, it won't use the Llano CPU/GPU processor with K10 core. Instead, it will probably have two to four Bobcat cores and a Vancouver GPU. Spans the performance range from Vision Premium to Vision Black. Scheduled for 2011.

Dorado: AMD code-name for a desktop-PC platform based on a next-generation CPU/GPU chip. Unlike Lynx, a similar platform, it won't use the Llano CPU/GPU processor. Instead, it will probably have two to four K10 cores and AMD's Northern Islands GPU. Spans the performance range from Vision Premium to Vision Ultimate. Scheduled for 2011.

E-series: AMD part-number prefix for low-power mobile-PC processors intended mainly for notebooks. Zacate is the first E-series chip. Known variations are the E-350 and E-240.

Fusion: AMD brand name for integrated CPU/GPU chips.

Hudson-M1: AMD code-name for the Fusion Controller Hub, a south-bridge chip with I/O interfaces for USB 2.0, SATA-III, PCI Express (4×1 Gen2), high-definition audio, SPI, and low pin-count (LPC) interface. Connects to an integrated CPU/GPU Fusion processor via AMD's proprietary Unified Media Interface.

Llano: AMD code-name for an integrated CPU/GPU chip for notebook PCs. Based on the current K10 CPU core.

Lynx: AMD code-name for a desktop-PC platform based on the Llano integrated CPU/GPU chip. Spans the performance range from Vision Premium to Vision Ultimate. Scheduled for 2011.

Nile: AMD code-name for a notebook-PC platform based on the Athlon-II Neo processor (K135 or K325), AMD RS880 north-bridge chip, and AMD SB810 south bridge with Radeon HD 4225 GPU. Will be superseded in 2011 by the Brazos platform.

Oak Trail: Intel code-name for an Atom-based chipset designed for tablets and similar devices. May also be used for ultrathin notebooks and netbooks.

Ontario: AMD code-name for a C-series integrated CPU/GPU chip with one or two Bobcat CPU cores. It's a lower-power version of Zacate intended mainly for netbook PCs. Known variations are the C-30 and C-50.

Pine Trail: Intel code-name for a two-part Atom processor chipset designed for netbooks, subnotebooks, and nettop PCs. The processor chip is code-named Pineview, and the south-bridge I/O hub is code-named Tiger Point. Pine Trail-M is for mobile systems; Pine Trail-D is for desktop PCs. Pine Trail's product name is the Intel NM10 Express Chipset.

Pineview: Intel code-name for a dual-core Atom processor with integrated north-bridge functions and an Intel GMA3150 graphics processor. Pineview-M is for mobile systems; Pineview-D is for desktop PCs.

Sabine: AMD code-name for a notebook-PC platform based on the Llano CPU/GPU chip. Spans the performance range from Vision Premium to Vision Black. Scheduled for 2011.

Sandy Bridge: Intel code-name for a new x86 CPU that will be integrated with an Intel GPU. First products are expected in 1Q11.

Scorpius: AMD code-name for an integrated CPU/GPU desktop-PC processor using the new Bulldozer CPU core. Intended to deliver Vision Black performance. Scheduled for 2011.

Tiger Point: Intel code-name for the south-bridge I/O chip in the Pine Trail platform. Tiger Point's product name is the CG82NM10 platform controller hub (PCH).

Timna: Intel code-name for an integrated CPU/GPU designed in the 1990s but aborted in 2000.

Vancouver: AMD code-name for a family of DirectX 11 discrete GPUs. May be integrated in future Fusion chips. Known members are Blackcomb, Capilano, Granville, Robson, Seymour, and Whistler.

Vision: AMD brand name that broadly defines the performance of Fusion processors.

Zacate: AMD code-name for an E-series integrated CPU/GPU chip with one or two Bobcat CPUs. A higher-power version of Ontario intended mainly for notebook PCs. Known variations are the E-240 and the E-350.