

ALCHEMY TRANSMUTES MIPS32

One Integrated Processor Delivers Both High Performance and Low-Power

By Steve Leibson {7/10/00-01}

The microprocessor equivalent to beer advertising's "Great taste, less filling!" is "Low power, high performance!" Intel's StrongARM processor family (originally developed at Digital Semiconductor) has owned the price/performance pinnacle in embedded processors for

several years. Now, a development team formed by many of the original StrongARM designers led by Rich Witek and Greg Hoepfner, has revealed that the MIPS processor architecture is just as amenable to their low-power design wizardry as was ARM. Witek and Hoepfner participated in the design of many groundbreaking microprocessors: MicroVAX, Alpha, PowerPC, ARM, and StrongARM. First technical details of Alchemy Semiconductor's Au1000 highly integrated processor, revealed at last month's **Embedded Processor Forum**, indicate that this chip should break new ground for the MIPS architecture. The chip should run as fast as 500MHz while dissipating 900mW, or at 200MHz while dissipating less than 200mW. These speed and power figures are for a fully functional MIPS32 processor core augmented with a hardware 32 x 16 MAC, separate 16K instruction and data caches, and a long list of peripheral devices, including two memory controllers, two 10/100 Ethernet controllers, separate host and device USB ports, an eight-channel DMA controller, four UARTs, two real-time clocks, parallel I/O ports, and several serial peripheral ports. The target for the Au1000 is low-power applications—specifically, battery-powered devices.

MIPS unveiled the MIPS32 ISA last year as a way of rationalizing the proliferation of MIPS 32-bit architectures into a unified definition more suited to the embedded market (See [MPR 5/31/99-05](#), "Jade Enriches MIPS Embedded Family"). MIPS32 starts with the MIPS-II (R3000) instruction set and adds 19 new instructions, including several

multiplication and multiplication/addition/subtraction instructions. Other added instructions include CLZ and CLO, which count leading zeroes and ones, respectively. These instructions are well suited to specific embedded applications like normalization and cryptography. MIPS32

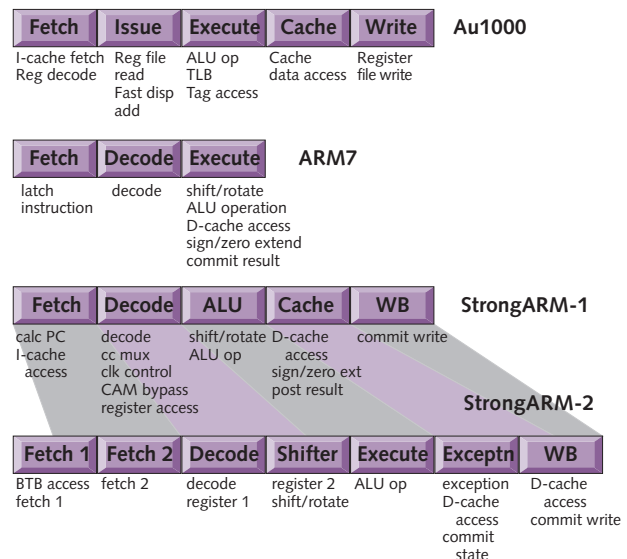


Figure 1. The Au1 core employs a five-stage pipeline, as did the original StrongARM design. StrongARM-2 employs a seven-stage pipeline and therefore achieves higher clock rates at equivalent lithography levels.

also includes memory-management and a privileged-instruction architecture resembling the R4000's. MIPS introduced two Jade cores based on MIPS32 last year. At about the same time, the StrongARM design wizards exited Digital Semiconductor and formed the Alchemy Microprocessor Design Group at Cadence. Before spinning off from Cadence, Alchemy licensed the MIPS32 ISA, but not the Jade core designs.

Instead, Alchemy designed its own MIPS32 core, dubbed the Au1 and built from a custom cell library, also developed by Alchemy. As Figure 1 shows, the Au1 pipeline has five stages, as did the original StrongARM design. Alchemy's cell library is portable across three different foundries, but initial fabrication of the Au1000 employs TSMC's 0.18-micron low-voltage process. In that process, the Au1000 die is expected to measure less than 60mm². Alchemy plans to sell chips, such as the Au1000 based on the Au1 core, but the company can also license its Au1 core design to other MIPS licensees.

Much of the original StrongARM processor's price/performance prowess stemmed from Digital Semiconductor's custom circuit design and advanced (for the time) fab and process capabilities. As a fabless semiconductor vendor shooting for design portability, Alchemy cannot tune a process to its core design. Even so, Alchemy's custom cell library, running on TSMC's standard process, produces competitive results. The Au1000 clocks at 500MHz running at a core voltage of 1.8V. At that clock rate, the chip dissipates 900mW. Lowering the core voltage to 1.25V and the

clock frequency to 200MHz drops power dissipation below a very cool 200mW. The Au1000's I/O is 3.3V compatible.

As Figure 2 shows, 200mW powers a pretty capable processor with a large peripheral package. The Au1000 starts with the five-stage Au1 MIPS32 core augmented with a 32 x 16 MAC. The multiplier/accumulator handles one 32 x 16 MAC per cycle and can be double-pumped to produce a 32 x 32 MAC every other cycle. Divide instructions require a maximum of 35 cycles. The Au1 processor core is scalar, but the Au1000 achieves some small measure of parallelism, because the MAC pipeline is independent of the core pipeline. Instructions that use the MAC pipeline exclusively can execute in tandem with other instructions.

Beyond the custom cell library, Alchemy employs a variety of approaches to minimize the Au1000's power dissipation. The core design makes aggressive use of conditional clocking, a technique that has become a favorite of processor designers targeting low-power applications. The Au1000 automatically powers down the MMU, data cache, execution unit, and MAC when they are not in use. The processor has three reduced-power operating modes: idle 1, idle 2, and sleep. In the idle-1 mode, the CPU continues to snoop the external bus and maintains data-cache coherency. Power dissipation is correspondingly higher. In the idle-2 mode, snooping ceases and coherency is lost, which can be dealt with using software in multiprocessor systems. Exiting either idle mode requires fewer than 10 cycles. Exiting sleep mode requires 200ms, because the PLL must relock and a full processor reset occurs during the transition from sleep

to normal operation. The company has yet to characterize the part so power dissipation figures for these low-power modes are not yet available.

Trading Off Speed for Power

Alchemy's processor architects decided to omit some of the more common performance-enhancing but power-hungry structures used in current RISC designs: speculative execution and branch prediction. In lieu of the branch-prediction hardware, the Au1000 incorporates into the pipeline's issue stage a load/store adder that allows the address calculation to occur one cycle early. The processor then fetches the next instruction from the target address during the next cycle, effectively short-circuiting the pipeline's execution stage, as shown in Figure 3. This approach reduces the branch delay to one cycle. The issue-stage adder can also modify base registers so that recomputed base addresses are immediately available for subsequent instructions without incurring a pipeline stall.




Figure 2. Alchemy's additions to the MIPS32 architecture include a 32 x 16 MAC, larger instruction and data caches, and a bevy of high- and low-speed peripheral devices.

One of the features available to MIPS processor designers that Alchemy did not use is the MIPS-16 instruction set, originally developed by LSI Logic and MIPS for LSI Logic's TR4101 TinyRISC core (see *MPR 10/28/96-10*, "LSI's TinyRISC Core Shrinks Code Size"). MIPS-16 is a subset of the MIPS instruction set that employs 16-bit opcodes instead of the standard 32-bit MIPS instruction-set opcodes. The MIPS-16 instruction set uses an entirely different set of opcodes and requires a predecoder that maps the MIPS-16 instructions into the 32-bit MIPS instructions. The smaller instruction set is far less capable than the 32-bit set but the 16-bit instructions can reduce code size by approximately 40%, and therefore may reduce the amount of memory traffic (which can save power). Alchemy says some customers have requested MIPS-16, so it may be added in future core designs. Also missing are multimedia extensions to the ISA. MIPS has defined such extensions for the MIPS64 ISA but not for MIPS32, and Alchemy says it prefers to wait for MIPS to create a standard set of MIPS32 multimedia extensions before adding such extensions to the Au1 core.

MIPS's initial Jade implementations of the MIPS32 architecture included configurable cache sizes, with a recommendation for 8K instruction and data caches. Alchemy's Au1000 incorporates 16K instruction and data caches. Larger caches improve performance and further reduce power dissipation by minimizing traffic on the external memory bus. Both of the Au1000's caches are four-way set-associative caches with 32-byte cache lines. Each I and D cache line can be locked independently. The data cache is a write-back cache. For reads, the data cache allows one outstanding miss. It continues to service read and write requests until a second miss occurs. Only then does the Au1000's data cache stall, until the memory controller satisfies the first missed request. Cache logic snoops the system bus to allow automatic data-cache coherency in multiprocessor systems. Software must take responsibility for maintaining instruction-cache coherency. This should be an issue only in rogue system designs running self-modifying code, or where another processor (such as an I/O processor) is loading or otherwise modifying the local processor's code space.


Two Memory Controllers for the Price of One

The Au1000 incorporates two memory controllers for managing external memory. An SDRAM controller provides a glueless interface to as many as three SDRAM or SMROM (synchronous, masked ROM) banks through a dedicated synchronous-memory port. This port operates at half the clock rate of the Au1000's internal system bus, which, in turn, typically operates at half the processor core's clock rate

(it can operate as slowly as one-fifth of the processor clock rate). Thus, a processor running at 400MHz operates the synchronous memory port no faster than 100MHz. Three programmable chip-select pins associated with the synchronous-memory port permit the design of contiguous system-memory arrays using memory modules of varied size. All memory modules connected to the synchronous-memory port must be 32 bits wide.

A second memory port, called the "static" port and connected to a second memory controller, supports other word widths and other memory types. This port has separate 32-bit address and data buses and accommodates 16- and 32-bit memories. The port also has four associated programmable chip-select pins that, once again, permit the conglomeration of variously sized memory blocks into one contiguous array. The static port shares address and data lines (but not control lines) with a PC Card/Compact Flash controller that allows glueless connection to removable memory devices. Systems based on the Au1000 can have devices connected to both ports, because only the address and data lines are shared between the static port and the PC Card/Compact Flash interface—the control lines are separate.

Peripherals on the Au1000 are categorized as either high or low speed. High-speed peripherals connect directly to the internal system bus that runs at one-half to one-fifth of the core speed. The roster of high-speed on-chip peripherals includes an eight-channel DMA controller, two Ethernet controllers, a USB host controller, a fast IrDA port, and an EJTAG (enhanced JTAG) controller. Low-speed peripherals, devices that need less attention from the processor, include a long list of devices typically used in embedded systems such as two real-time clocks and four 16550-compatible UARTs. Some of the more unusual low-speed peripherals are an AC97 codec interface, a USB device controller, and several serial ports for connection to peripheral chips such as A/D and D/A converters. The two interrupt controllers on the low-speed peripheral bus represent an often overlooked but critical aspect of embedded systems design: insufficient numbers of interrupts. Each of the two interrupt controllers handles



MICHAEL MUSTACCHI

Alchemy architect Greg Hoepfner describes the low-power features of the Au1000 at the Forum.




Figure 3. Addition of a fast displacement adder to the second stage of the Au1 pipeline reduces the branch delay to one cycle.

32 sources, so it's unlikely that a system based on the Au1000 will run short of interrupts. Low-speed peripherals in the Au1000 connect to a peripheral bus that operates at half of the system bus's clock speed (one-quarter of the processor clock rate). A peripheral-bus interface module links the two on-chip buses.

Because of the Au1000's lineage, some comparisons with StrongARM, particularly the SA-1110, are unavoidable. Although the mix of peripherals is similar on both chips, and the Au1000 certainly has more ports, the SA-1110 has a color LCD controller whereas the Au1000 requires an external controller. Alchemy claims that its initial customers could not agree on a common LCD controller spec, so the designers didn't put one on the chip. However, lack of an LCD controller in the first version of the Au1000 doesn't preclude adding one to the next device in the family.

Know When to Design It; Know When to Buy It

Most of the peripheral devices on the Au1000 are purchased IP. Alchemy's skill is in processor core design, and the Au1000 designers wisely decided to avoid wasting time where they could add little value. Selecting MIPS32 for an ISA is another way Alchemy avoids unnecessary reinvention; the MIPS32 ISA offers instant access to a broad array of well-regarded software development tools and several important operating systems, including Windows CE, Linux, VxWorks, pSOS, and QNX. Windows CE takes the Au1000 into the handheld PC arena; the VxWorks, pSOS, and QNX RTOS products dominate the embedded space; and Linux straddles both the PC and embedded markets.

Price & Availability

Alpha samples of the Au1000 should be available in September. Production ramp is slated for 4Q00. Alchemy expects to charge less than \$50 (in 10,000-unit lots) for the Au1000. More information about the Au1000 is available on Alchemy's Web site at www.alchemysemi.com.

The Au1000's price/performance ratio surpasses the category-leading SA-1110, but that's not much of a surprise—the SA-1110 is built on Digital Semiconductor's (now Intel's) highly tweaked but very old 0.35-micron, three-layer-metal process, while the Au1000 is built with TSMC's most advanced 0.18-micron, four-layer-metal process. The Au1000 does not attain the performance expected from the StrongARM-2, which Intel announced at last year's Embedded Processor Forum. SA-2 jumps from Digital's old 0.35-micron process to Intel's 0.18-micron P858 process and adds two pipeline stages, for a total of seven. These changes produce an expected clock rate of 600MHz, exceeding the Au1000's expected clock rate by 20% while running on only half of the power. Simulations put the Au1000's performance at 569 Dhrystone 2.1 mips, while SA-2 simulations suggest that the processor will achieve more than 700 Dhrystone 2.1 mips at 600 MHz. However, SA-2 is late. According to Intel's 1999 announcement, the company expected to be shipping SA-2 processors by now. With samples expected in September, Alchemy just might deliver its golden Au1000 before Intel can ship SA-2 silicon. ♦

TOP PC VENDORS ADOPT CRUSOE

Transmeta Reveals Roadmap; New TM5600 Has 512K L2 Cache

By Tom R. Halfhill {7/10/00-02}

Four top-tier vendors at PC Expo announced their intention to make notebook computers based on Transmeta's Crusoe processors. Some of these systems will use a new version of Crusoe that has twice as much on-chip L2 cache. Transmeta has also revealed a two-year

roadmap of processors with higher clock speeds, greater integration, lower power consumption, and new VLIW cores, as Figure 1 shows.

The four PC vendors throwing their weight behind Transmeta's unusual x86-compatible processors (see *MPR 2/14/00*, "Transmeta Breaks x86 Low-Power Barrier") are Fujitsu, Hitachi, IBM, and NEC. All the vendors plan to introduce notebooks in the ultralight class, ranging in weight from 2.8 to 3.5 pounds, with TFT screens ranging in size from 10.4 to 12.1 inches. The notebooks are scheduled to ship this fall.

Crusoe processors are well suited for lightweight notebooks, because their low power consumption eliminates the need for cooling fans and large heat sinks. Furthermore, ultralight notebooks don't compete in the same performance class as larger, heavier laptops, where Intel's mobile processors have a speed advantage.

Some of the new notebooks will use the Crusoe TM5600, which has 512K of on-chip L2 cache—twice as much as the TM5400 announced in February. In other respects, the TM5600 is identical to the TM5400. It will be manufactured by Transmeta's foundry partner, IBM Microelectronics, in a 0.18-micron copper process and packaged in a 474-pin ceramic BGA. Doubling the L2 cache increased the die size to 88mm², which is 20% larger than the TM5400's die (73mm²).

According to Transmeta, the TM5600 is 5–15% faster than the TM5400 and consumes 2–17% less power.

Although doubling the size of the L2 cache and enlarging the die would normally increase power consumption, Transmeta says the TM5600 actually uses less power when running typical Windows software, because it makes fewer accesses to main memory over the 3.3V I/O bus.

However, Transmeta still has not released any results of common industry benchmark tests, such as Ziff-Davis

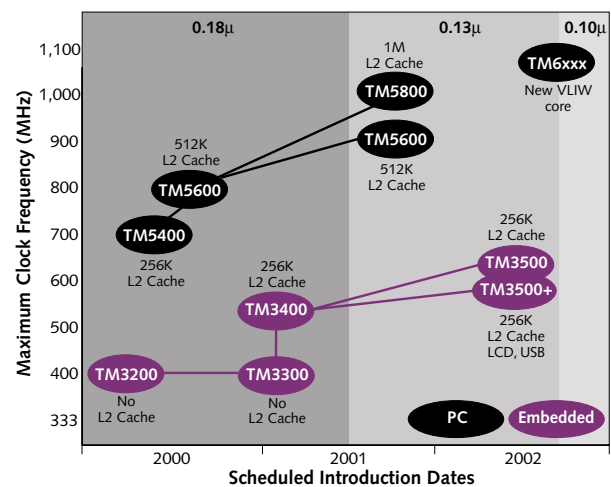


Figure 1. Transmeta's two-year roadmap calls for several new chips based on the existing TM3200 and TM5400 CPU cores, followed by an entirely new core in 2002.

Media's Winstone. Transmeta claims—with some credibility—that existing benchmark programs are misled by the unusual caching and dynamic-recompilation behavior of Crusoe's code-morphing (x86 emulation) software. According to Transmeta's estimates, a TM5400 Crusoe running at 700MHz delivers about the same raw performance as an Intel Pentium III at 500MHz.

Surprising Power Measurements

To back up its claims that Crusoe processors typically consume only 500mW to 1.5W—including the integrated north-bridge controller—Transmeta showed MDR a test system that graphically displays a constant measurement of minimum, maximum, and average power consumption. We experimented with several Windows applications, including Microsoft Word, Internet Explorer, an MP3 audio player, and a DVD movie player. As Figure 2 shows, average power consumption is indeed in the range promised by Transmeta and rarely spikes above 6W. And the test system revealed startling differences in power consumption among applications. For example, merely selecting a paragraph of text in Word briefly gobbled more power than decoding the MPEG-2 stream of a DVD movie.

Faster Crusoe processors are coming next year, according to Transmeta's roadmap. In 2H01, Transmeta will migrate its cores to IBM's 0.13-micron copper process, which offers the option of silicon-on-insulator (SOI) technology (see *MPR 5/1/00-01*, "IBM Paving the Way to 0.10 Micron"). The process shrink will reduce core voltage to 1.2V and boost clock speeds about 25%, even if Transmeta doesn't use SOI. The TM5600 would move into the 700–900MHz frequency range at this geometry.




Figure 2. Transmeta's instrumented test system continuously plots power consumption while other software is running. This trace plots the playback of a DVD movie, with regularly spaced spikes marking the decompression and display of each video frame. The large but brief power surge was caused by launching another Windows application.

To take further advantage of the 0.13-micron process, Transmeta plans to introduce another new processor, the Crusoe TM5800, which will have 1M of on-chip L2 cache. It will use the same core as the TM5600 and TM5400 and retain the 474-pin CBGA package.

In 2002, Transmeta plans to revamp its Crusoe line with an entirely new CPU core based on an enhanced VLIW architecture. This core will have a faster FPU and will target a 0.13- or 0.10-micron process. Transmeta expects this unnamed chip (probably TM6xxx) to have twice the performance of existing Crusoe processors while reducing typical power consumption below 500mW.

The code-morphing software will also get an overhaul. Because that software translates x86 instructions into native VLIW instructions on the fly, Transmeta has virtually unlimited freedom to change the inner workings of the CPU core without breaking compatibility with operating systems and applications.

Equally Aggressive Embedded Roadmap

Transmeta's roadmap calls for similar improvements to the company's line of Crusoe processors for Mobile Linux information appliances. The current product is the TM3200, formerly known as the TM3120 (renamed because of a trademark conflict).

Next year, Transmeta plans to introduce two new chips, the TM3300 and TM3400. Both will come in smaller packages (360-pin CBGAs instead of 474-pin CBGAs), achieved by eliminating the unused pads reserved for the DDR-SDRAM interface on TM5xxx-series chips, but they will retain their SDR-SDRAM interfaces. The downsized package will allow Transmeta to sell the TM3300 for less than \$50. The TM3400 will be the higher-end model, adding a 256K on-chip L2 cache and LongRun power manager.

In 2002, Transmeta plans to introduce the TM3500, which migrates the existing core to a 0.13-micron process. That will reduce the core voltage to 1.2V and boost the clock frequency to 600MHz. There will also be a more integrated version of the TM3500 that has 256K of on-chip L2 cache, LongRun power management, an LCD controller, a USB interface, and a PCI interface. Those additions will bump the pin count back up to 474 but provide system vendors with a more complete system solution. Table 1 summarizes the features of Transmeta's current and future chips.

Less than six months after its much-hyped public debut, Transmeta has gained the crucial support of some

Feature	TM5400	TM5600	TM5800	TM6xxx	TM3200	TM3300	TM3400	TM3500
CPU Core	5400	5400	5400	6xxx	3200	3200	3200	3200
Core Freq*	500–700MHz	500–800MHz	700MHz–1GHz	>1GHz	333–400MHz	333–400MHz	400–533MHz	400–600MHz
L1 Cache (I/D)	64K/64K	64K/64K	64K/64K	n/a	64K/32K	64K/64K	64K/64K	64K/64K
L2 Cache	256K	512K	1M	n/a	None	None	256K	256K
LongRun?	Yes	Yes	Yes	Yes	No	No	Yes	Yes
SDRAM Ctrl?	SDR+DDR	SDR+DDR	SDR+DDR	n/a	SDR	SDR	SDR	SDR
LCD Ctrl?	No	No	No	n/a	No	No	No	Yes
USB Ctrl?	No	No	No	n/a	No	No	No	Yes
PCI Ctrl?	Yes	Yes	Yes	n/a	Yes	Yes	Yes	Yes
IC Process	0.18 μ 5LM	0.18 μ 5LM	0.13 μ	0.10–0.13 μ	0.18 μ 5LM	0.18 μ 5LM	0.18 μ 5LM	0.13 μ
Core Voltage	1.6V	1.6V	1.2V	1.2V	1.5V	1.5V	1.5V	1.2V
Die Size	73mm ²	88mm ²	n/a	n/a	77mm ²	77mm ²	n/a	n/a
Package	CBGA-474	CBGA-474	CBGA-474	n/a	CBGA-474	CBGA-360	CBGA-360	CBGA-474
Power (typ)	<1.5W	<1.5W	<1W	<500mW	<1.5W	<1.5W	<1.5W	<1W
Price (1K)	\$119–\$329	n/a	n/a	n/a	\$65–\$89	<\$50	n/a	n/a
Availability	Now	Aug-00	2H01	2002	Now	1Q01	1Q01	2002

Table 1. Transmeta has aggressive plans to expand its Crusoe line over the next two years while taking advantage of IBM's latest semiconductor process. (*In this table, "core frequency" refers to the maximum clock-frequency ratings of individual Crusoe processors, not to the range of frequencies supported by LongRun power management. n/a = information not available.)

top-tier PC vendors and major financial backers. In addition to the PC vendors that announced Crusoe-based notebooks at PC Expo, Gateway has committed to using a Crusoe chip in a Linux-compatible information appliance designed for AOL (see *MPR 6/5/00-04*, "Transmeta Lands Gateway-AOL IA"), and several high-profile investors have kicked in \$88 million of funding (see *MPR 5/1/00-06*, "Transmeta Lands \$88M").

Although Transmeta's performance and power-consumption claims haven't been independently verified, this situation is clearly not hurting the company's ability to attract key customers and investors. And Transmeta's early success is exerting pressure on Intel to come up with a comparable line of low-power processors.

Telling Watt's Right From Wrong

No doubt in response to Transmeta's challenge, Intel says its latest mobile PC processors come close to matching the average power consumption of the TM5400. On June 19, Intel announced five mobile Pentium III and Celeron processors and claimed their average power consumption to be 0.8–1.6W for the 500/600MHz Pentium III and 1.6–2.8W for the 600/750MHz Pentium III (see *MPR 7/3/00-02*, "Intel Strikes Back at Transmeta").

Those numbers are far below Intel's own thermal design power (TDP) specifications, which represent worst-case design points for engineers who want to ensure that their systems won't overheat. The disparity between Intel's

average-power estimates and the TDP specifications provoked a quick reaction from Transmeta. The company accused Intel of manipulating the average-power figures by basing them on tests with Ziff-Davis Media's BatteryMark program, which typically spends 80% of its time in idle states.

Sorting out these claims and counterclaims isn't easy. Intel's TDP specification is 9.5W for the 500/600MHz mobile Pentium III and 15.8W for the 600/750MHz mobile Pentium III. However, those numbers don't include a north-bridge controller, which is integrated on Crusoe chips and included in their power ratings. Intel does offer mobile

Feature	Transmeta Crusoe TM5400	Transmeta Crusoe TM5600	Intel Mobile Pentium III	Intel Mobile Pentium III
CPU Core	5400	5400	Coppermine	Coppermine
Freq Range*	200–700MHz	200–800MHz	600/750MHz	500/600MHz
L1 Cache (I/D)	64K/64K	64K/64K	16K/16K	16K/16K
L2 Cache	256K On-Chip	512K On-Chip	256K On-Chip	256K On-Chip
North Bridge	Yes	Yes	No	No
SDRAM Ctrl	SDR+DDR	SDR+DDR	No	No
PCI Ctrl?	Yes	Yes	No	No
MMX?	Yes†	Yes†	Yes	Yes
SSE?	No	No	Yes	Yes
IC Process	0.18 μ 5LM	0.18 μ 5LM	0.18 μ 6LM	0.18 μ 6LM
Metal Layers	Copper	Copper	Aluminum	Aluminum
Core Voltage	1.1–1.65V	1.1–1.7V	1.35/1.6V	1.1/1.35V
Die Size	73mm ²	88mm ²	106mm ²	106mm ²
Power Modes	LongRun	LongRun	SpeedStep	SpeedStep
Power (typ)	0.5–1.5W	0.5–1.5W	1.6–2.8W	0.8–1.6W
Power (max)	~6W	~6W	13.9/20W	12.2/16.6W
Price (1K)	\$329	n/a	\$562	\$316
Availability	Now	Aug-00	Now	Now

Table 2. Transmeta's Crusoe processors have an advantage in power consumption, whereas Intel's mobile Pentium III processors have an advantage in performance: Transmeta says a 700MHz TM5400 is about as fast as a 500MHz Pentium III. (*In this table, "frequency range" refers to the variable clock rates of an individual chip in its different operating modes. The TM5400 and TM5600 can vary their frequencies within these ranges in increments of 33MHz, while the mobile Pentium III chips are limited to the two frequencies shown. †Not including north bridge. n/a = information not available.)

Price & Availability

Transmeta will begin shipping production volumes of the Crusoe TM5600 with 512K L2 cache in August. The price has yet to be announced. The TM5400 and TM3200 are shipping now. TM5400 prices range from \$119 at 500MHz to \$329 at 700MHz, and TM3200 prices range from \$65 at 333MHz to \$89 at 400MHz (all for 1,000-unit quantities). For more information, go to www.transmeta.com.

Pentium III processors in MMC-2 cartridges with a north bridge, but that feature increases the TDP to 12.8W for the 500/600MHz mobile Pentium III and 19.1W for the 600/750MHz mobile Pentium III.

There are two clock-frequency numbers for each mobile Pentium III processor because they automatically reduce their core voltage and clock rate when unplugged from AC power, a feature Intel calls SpeedStep. For example, the 500/600MHz mobile Pentium III normally operates at 600MHz and 1.35V on AC power, but it steps down to 500MHz and 1.1V on batteries. Likewise, the 600/750MHz mobile Pentium III normally operates at 750MHz and 1.6V on AC power; however, it then steps down to 600MHz and 1.35V on batteries. Table 2 compares the specifications of Intel's and Transmeta's mobile processors for PC notebooks.

The differences between Intel's TDP specifications and average-power estimates are roughly an order of magnitude, which implies the average-power figure is based on a 10% duty cycle. ZD Media's eTesting Labs (formerly ZD Labs) told MDR that BatteryMark does indeed leave a system in




Figure 3. This power-consumption trace shows a 500/600MHz mobile Pentium III processor running the ZD Media BatteryMark 3.0 program, which keeps the CPU idle most of the time. That allows Intel's QuickStart power manager to reduce average power consumption well below 3W.

idle states about 80% of the time, because that's how ZD Media's engineers think real people use notebook computers. So, as Transmeta alleges, BatteryMark's light duty cycle could explain the gap between Intel's average and worst-case power numbers. Figure 3 shows a power-consumption trace of the 500/600MHz mobile Pentium III in battery mode (500MHz).

The Quest for Better Benchmarks

The truth about Intel's and Transmeta's power-consumption claims is more elusive, however. For one thing, Intel didn't base its average-power estimates solely on BatteryMark. The company tested a wide range of desktop PC applications, as did Transmeta.

Not by coincidence, one application Intel tested was a DVD movie player, which is Transmeta's favorite demo. Playing a DVD movie on a 600/750MHz mobile Pentium III at 600MHz consumes an average of less than 2W, says Intel, and the same task consumes less than 1W on a 500/600MHz mobile Pentium III at 500MHz. But the screen photo in Figure 4 (supplied by Intel) traces the latter processor during DVD playback in battery mode, and it appears to show average power consumption in the 2.3W range. And Intel's TDP specifications clearly indicate that, to avoid meltdowns, system engineers must design for brief periods of much higher power consumption.

Unfortunately, Transmeta has been less forthcoming with power-consumption benchmarks than Intel. Transmeta criticizes Intel's BatteryMark score but won't release one of its own, saying that conventional tests like BatteryMark don't yield accurate results, because of the TM5400's unique LongRun technology. LongRun can vary the chip's core voltage from 1.1V to 1.65V and the clock frequency from 200MHz to 700MHz in 33MHz increments—all while running on battery power and in response to soft-

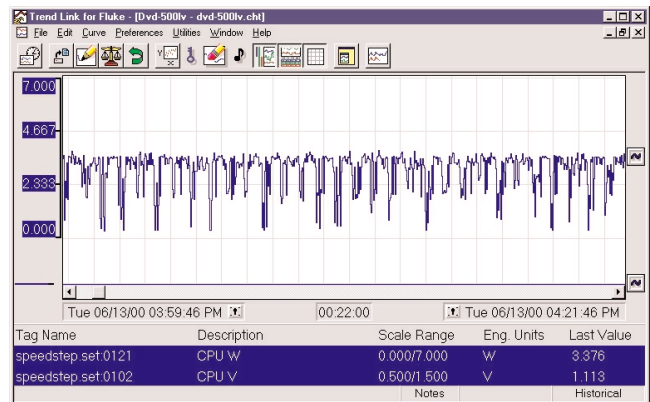


Figure 4. This trace shows a 500/600MHz mobile Pentium III processor playing a DVD movie in battery-optimized (500MHz) mode. Compared with the similar trace from Transmeta in Figure 2, the Intel processor appears to be consuming at least twice as much power—even without a north bridge, which would add about 4W.

Transmeta Explains LongRun

The last mysteries of Transmeta's LongRun technology were explained in a presentation at **Embedded Processor Forum** last month by Marc Fleischmann, the company's director of low-power programs. Although LongRun is somewhat less sophisticated than we imagined after Transmeta's public announcement in February, it is nonetheless effective.

At that time, Transmeta said LongRun scales a Crusoe processor's core voltage and clock frequency up or down in response to software demands, allocating just enough performance to handle the varying workload while conserving power. That sounds like a neat trick, but we wondered how LongRun could tell the difference between a true workload and a tight event loop that's simply waiting for something to happen. Answer: it can't. At least, not without a clue from the operating system in the form of an idle command.

To conserve power, modern operating systems periodically issue idle commands to the CPU during brief periods when the need for processing power is low. Those periods might occur between the frames of an MPEG video stream, or even between a user's keystrokes. The idle command triggers the CPU's power-saving mode. The CPU goes to sleep until it needs to process an interrupt, then wakes up.

It turns out that Transmeta's code-morphing (x86 emulation) software constantly tracks the amount of time the CPU spends in sleep mode and uses that information to help LongRun regulate the chip's core voltage and clock frequency. For example, if the code-morphing software discovers the CPU is asleep half the time (e.g., a 50% duty cycle), it tells LongRun to cut the chip's voltage and frequency by a little less than 50%—say, 366MHz in the case of a 700MHz TM5400. As the figure below shows, that leaves enough performance to handle the software workload while boosting the duty cycle to near 100%.

By itself, reducing the clock rate would save little or no power, because the CPU is now working twice as hard at half

the frequency to do the same amount of work. But LongRun also reduces the core voltage, which significantly cuts power consumption. That's because $W = \frac{1}{2} CV^2 F$, where W is watts, C is capacitance, V is voltage (squared), and F is frequency. Therefore, power consumption drops faster than the linear reduction in frequency.

The CPU still gets a few catnaps after the clock-rate adjustment, because LongRun chooses a clock frequency that leaves some performance headroom. During those naps, the TM5400 enters a special low-voltage sleep mode that reduces power consumption to 40–50mW. The TM5400 consumes even less power in this mode than it does during the normal sleep mode.


LongRun is a clever and innovative approach to power management that does have the effect of scaling power consumption to software demands. And unlike AMD's similar PowerNow feature (see [MPR 5/1/00-05](#), "AMD's K6 Family on the

Move"), LongRun requires no modifications to the BIOS or operating system, because it takes advantage of existing power-management technology.




MICHAEL MUSTACCHI

Marc Fleischmann, Transmeta's director of low-power programs, explains LongRun at EPF2000.



$$\text{Power} = \text{Capacitance} \times \text{Voltage}^2 \times \text{Frequency}$$



ware workloads, unlike SpeedStep (see sidebar "Transmeta Explains LongRun"). The TM5400 would probably fare better with tests that simulate a heavier workload than BatteryMark, because LongRun could scale the core to some intermediate voltage and frequency within its broad range. A mobile Pentium III with SpeedStep would,

however, be limited to its single battery-mode voltage and frequency.

On the other hand, a mobile Pentium III is likely to deliver more performance than a TM5400, and not just because the TM5400 has to emulate the x86. At its lowest LongRun-controlled power level of 1.1V, the 700MHz

TM5400 cuts back its clock frequency to a pedestrian 200MHz. At that same voltage, the 500/600MHz mobile Pentium III sprints at 500MHz.

We agree with Transmeta that existing power-consumption and performance benchmarks aren't the best way to evaluate the TM5400's unique abilities. Still, it would be nice to have more data points. Our experiments with Transmeta's power-monitoring system lead us to believe that the brief 6W power spikes we observed roughly correspond to Intel's TDP ratings, and that the TM5400's average power consumption is indeed in the 1.5W range. If our casual observations are confirmed by more rigorous tests on

the Crusoe-based notebook computers soon to hit the market, Transmeta's promise of "all-day computing" could be realized.

One good sign is that Transmeta has joined EEMBC (EDN Embedded Microprocessor Benchmark Consortium) to help define a new suite of power-consumption tests and is working with other benchmarking organizations to refine their existing suites. Those efforts could benefit all CPU vendors, not just Transmeta. As more CPU architects turn to unusual techniques like LongRun to prolong battery life, power-consumption benchmarks will need to get more sophisticated—and now is not too early to start. ♦

To subscribe to Microprocessor Report, phone 408.328.3900 or visit www.MDRonline.com

TIDBITS

{7/10/00-03}

◆ WILLAMETTE NAMED PENTIUM 4

On June 28, in a move that surprised absolutely no one, Intel announced it had decided on Pentium 4 as the brand name for its next-generation microprocessor, code-named Willamette. Although the company could surely have come up with a more innovative name for its next-generation microarchitecture, it simply has too much marketing capital invested in the Pentium name to abandon it. Furthermore, the company is in the midst of trying to establish Itanium as a new brand name for its next-generation 64-bit instruction-set architecture; trying to establish another new name for Willamette at the same time might have created considerable confusion in the marketplace.

The only minor surprise about the new brand name is that Intel decided to eschew the Roman numerals it had used in the previous Pentium II and Pentium III generations for the less pretentious Arabic numeral in Pentium 4. Perhaps the company foresaw a long-term problem for future generations, such as Pentium XXX. —K.D.

◆ CELERON INCHES TOWARD 1GHZ

On June 26, Intel announced the availability of three new Celeron speed grades: 633, 667, and 700MHz. Like all Celerons, the new speed grades operate at a bus speed of 66MHz, bringing the clock multiplier to a staggering 10.5x. Pentium IIIs with twice the amount of on-chip cache operate at only a 7.5x multiplier, even at speeds of 1GHz. With a 10.5x multiplier and a small 128K L2, Intel has taken Celeron well beyond the point of diminishing performance returns for increased clock rate. In the sub-\$1,000 market where Celeron is focused, however, frequency is still more important than performance, so Intel may get away with the slow 66MHz bus for this round of frequency increases. But it will soon have to boost Celeron bus speeds to 100MHz or higher if it expects performance to improve much with frequency.

The three new Celerons are offered in flip-chip PGA packages and, in quantities of 1,000 units, list for \$138, \$170, and \$192, respectively. —K.D.

◆ TSMC TURNS INDUSTRY ON HEAD

In a dramatic reversal of roles, TSMC—the world's largest semiconductor foundry (see *MPR 6/5/00-01*, "TSMC Sets Sights on #1")—has become the first foundry in history to license its semiconductor technology to a large integrated-device manufacturer (IDM): National Semiconductor. Previously, semiconductor foundries like TSMC have always looked to IDMs as the source of IC-process technology.

Under their agreement, TSMC will transfer several 0.25- to 0.10-micron logic and embedded-memory processes to National for implementation in that company's South Portland (Maine) facility. TSMC's processes are restricted to use at the South Portland site. With this move, TSMC expects to gain access to excess capacity at that facility, and it will also receive license and royalties for parts National manufactures using TSMC processes. —K.D.

◆ AMD BOOSTS MOBILE K6-2+ TO 550MHZ

On June 26, AMD announced the availability of 0.18-micron Mobile K6-2+ processors running at 533MHz and 550MHz. Perhaps most noteworthy about the announcement is that HP will use the new processors in its HP Pavilion N3300 notebooks with the PowerNow feature enabled. AMD claims PowerNow can extend battery life in notebooks by up to 30%. Other OEMs, including Compaq, Fujitsu, and NEC, are already using K6-2+ chips but are not using the PowerNow feature. AMD worked with Phoenix Technologies and Insyde Software to incorporate the BIOS changes needed to take advantage of PowerNow.

The new K6-2+ speed grades operate on a core voltage of 1.4V to 2.0V and dissipate less than 3W of power in the battery-saver mode. The parts have a 100MHz Socket 7 front-side bus and are offered in a 321-pin CPGA. List prices in quantities of 1,000 parts are \$99 for the 550MHz part and \$85 for the 533MHz part. —K.D. ◆