AMD Stakes Claim to Fastest Mobile Micro

With a stroke of its pen and very little engineering, AMD has added three new speed grades to its mobile K6-2 lineup, all at near-desktop-processor prices. In working with key OEMs, AMD says it has discovered that the traditionally accepted 11-W thermal power limit for notebook processors is just a myth. Major OEMs are telling AMD that they can actually accommodate up to 16 W.

AMD pounced on the 5 W of extra headroom as if it were manna from heaven, raising the voltage on its mobile K6-2 processors from 1.8 V to 2.2 V. The 22% boost allowed the company to raise frequency by 14%, to 380 MHz. The voltage boost could actually have supported processors up to 400 MHz, but instead AMD decided to apply a portion of the new-found headroom to increasing yields, thereby lowering manufacturing costs.

To differentiate its mobile processors that take advantage of the 16-W envelope from its previous parts that fit within the 11-W envelope, AMD will append a "P" to the part designator of the higher-power parts. AMD also said that it will eventually offer a technology like Geyserville (see MPR 3/29/99, p. 14), giving the K6-2P parts a low-power mode for extending battery life.

Even without the low-power feature, however, battery life will not suffer much from the higher-power "P" parts. According to the Ziff-Davis Battery Mark, the K6-2P/380 depletes a battery charge only 7% (12 minutes) faster than a K6-2/333 in the same system. The diminished effect is due to the fact that the processor consumes only 10–15% of the power in a notebook; the rest goes to backlighting, disks, and other electronics.

Admittedly, using the extra 5 W of headroom is somewhat of a sleight of hand. AMD rationalizes that its customers are clamoring for higher-frequency and lower-priced processors for the low-cost notebook market, which is now growing. Because some notebooks can dissipate the extra heat, AMD is simply meeting customer demands. Indeed, as the table below shows, the new K6-2P parts are priced more like desktop parts than mobile parts. The new K6-2P/380 at $169 (quantity 1,000) is 10% less expensive than the $187 Mobile Celeron-300, even though it outperforms the Celeron part substantially on Winstone 99, according to AMD.

Compaq is on board with the concept, offering the new K6-2P in its Presario 1675 (380 MHz) and 1670/1270 (350 MHz) notebooks. If other mobile OEMs really do want high frequency and low price more than they need low power, the new K6-2P parts could increase AMD's share of the U.S. retail notebook market beyond its current 27%. But if Intel follows AMD's lead and engages in a price war, it could further erode both companies' average selling prices—something AMD can ill afford. —*K.D.*

	K6-2		Celeron or Dixon (Dx)	
Frequency	Desktop	Mobile	Desktop	Mobile
400 MHz	$133	–	$133	–
380 MHz	–	$169 (P)	–	–
366 MHz	$93	$149 (P)	$93	$696 (Dx)
350 MHz	–	$119 (P)	–	–
333 MHz	$73	$299	$73	$465 (Dx)
300 MHz	–	$187	$63	$187

Celeron Accelerates to 433 MHz

Keeping the pressure on AMD, Intel has introduced the next step in its Celeron line, a 433-MHz version of Mendocino. Like earlier Celerons (see MPR 1/25/99, p. 18), the new part continues with the 66-MHz bus, distinguishing it from the Pentium II-400 and -450, which use a 100-MHz bus. Staying with the slower bus stretches the clock multiplier to 6.5×, which clearly isn't good for performance on any memory-bound application. Most PC applications, however, perform reasonably well on the Celeron parts.

The Celeron-433 carries a 1,000-piece list price of $169. Intel reduced the prices of the slower Celeron parts just last month (see MPR 2/15/99, p. 4), so those prices did not change with the introduction of the 433-MHz part. The new Celeron matches up against AMD's K6-2/450, preventing AMD from charging much of a premium for that part and helping keep Intel customers in the fold. Early figures indicate AMD will not gain market share in 1Q99, showing that Intel's new Celeron strategy is having some success. —*L.G.*

PA-8600 Due in Early 2000

Now that HP has begun shipping systems using the PA-8500, the company has turned its sights to the next generation. HP recently announced more details about the PA-8600, which will tape out in the next few months and is due to ship in systems in 1Q00. The biggest performance gain will come from an increase in clock speed: the new chip will ship at clock speeds ranging from 500 to 550 MHz, a 25% boost from the 440 MHz achieved by the 8500.

HP indicated that the die size of the 8600 will be 467 mm^2, the same as that of the 8500. This implies that the 8600 will continue to use a 0.25-micron-generation process, perhaps with a transistor shrink to increase speed. The enormous die size is the result of including 1.5M of primary cache on the die, eliminating the need for any external cache.

The company has not disclosed the foundry for the 8600 (or the 8500, for that matter), although the metal pitches it published at ISSCC '99 are identical to those in P856, strongly suggesting that Intel is building the chips. Given the chip's ship date, the failure to move to a 0.18-micron process indicates that HP's outsourcing strategy is not providing access to leading-edge process technology.

With most of HP's CPU designers now working on McKinley and other IA-64 processors, there are few functional

changes in the 8600. HP has improved the cache-prefetch algorithm and moved from round-robin set selection to LRU (least recently used). Finally, HP added a lockstep mode for fault-tolerant designs, presumably at the behest of Stratus, which builds highly available PA-RISC systems.

These changes will have little impact on per-clock performance for most applications, although a few may see a bigger benefit from the cache changes. We estimate the 550-MHz 8600 will achieve scores of 37 SPECint95 (base) and 60 SPECfp95 (base). Although it will be reasonably competitive, the 8600 will not keep pace with the Alpha 21264, which should exceed 45 int/80 fp at clock speeds of at least 1 GHz in the same time frame. Without access to leading-edge fab technology, HP's PA-RISC line will be hard pressed to deliver leading-edge performance. —*L.G.*

Metaflow Decimated

Following STMicroelectronics' acquisition in May 1997 of a majority stake in Metaflow (*www.metaflow.com*) (see MPR 6/23/97, p. 4), the relationship between the companies has soured. Sources indicate that the two were unable to agree on which direction to take Metaflow's technology, creating a rift that resulted in ST's buying out Metaflow's last remaining founders: President Val Popescu, Vice President Bruce Lightner, and Director of Engineering Gary Gibson. Gibson has since taken a position with Mosart Systems (*www.mosart.com*), working on some as-yet-undisclosed new processor.

The originator of much of the technology underlying today's superscalar out-of-order microprocessors, Metaflow was never able to reap its just rewards. Begun in 1985, the company's first project, a SPARC-based ECL gate-array processor, was supplanted in 1989 by Lightning, a CMOS design backed by funding and IC-design resources from Hyundai. But the division of labor between the two companies proved stormy, and in 1991 Lightning was discharged, creating Thunder—a 0.8-micron three-chip processor designed entirely by Metaflow (see MPR 2/12/92, p. 9). In 1995, the company successfully demonstrated an 80-MHz Thunder processor—which delivered 2.5 SPECint92/MHz—just as Hyundai decided it wanted x86-based processors instead of SPARCs.

Metaflow had a brush with success in the early '90s when, working under contract to Intel, it convinced the processor giant that out-of-order design offered the best hope for building fast x86 processors. In fact, sources indicate that inside Intel the P6 was initially referred to as the "Metaflow processor." Sources have also revealed that Intel actually tried to buy Metaflow, but Andy Grove, unable to come to terms with Hyundai, had to abort the purchase.

Intel subsequently acquired the IP rights to Metaflow's technology through a patent cross-license with Hyundai for DRAM technology. Today the P6, with its centralized reservation station and instant repair of out-of-order and speculative execution, bears the clear mark of Metaflow's involvement. More than 150 Intel patents cite Metaflow intellectual property as prior art.

In 1997, frustrated by the inability to control its own IP and thus its own destiny, Metaflow convinced Hyundai to sell out to ST. Since then, ST and Metaflow have been jointly developing an x86-based processor. But, already unhappy with ST's plans for that part as well as with its lackluster process technology, Metaflow's leaders were pushed over the brink when ST traded Metaflow's IP to IBM in an x86-for-PowerPC swap (see MPR 8/3/98, p. 10).

Some Metaflow designers continue to work on ST's x86 project, but with the leaders gone and ST more interested in the low end of the x86 market, chances are slim that the x86 chip will ever see daylight—especially now that ST has a Pentium-class x86 project under way in India. ST officials declined to comment, other than to acknowledge that ST has indeed taken full ownership of Metaflow and to express confidence in the design team that remains. —*K.D.*

Somerset Fuels Intel's x86 Efforts

In an ironic turn of events, sources have confirmed that Mark McDermott, former director of Somerset (the once-joint IBM/Motorola PowerPC design center), is putting together a large engineering team in Austin to develop x86 processors for his old enemy Intel. Contrary to some speculation, the new Intel design center is working on neither StrongARM or IA-64 but is instead working on post-Willamette x86-based processors.

McDermott's team has been seeded with some key Intel personnel, including Intel Fellow Uri Weiser, architect of MMX. The team is growing rapidly, with McDermott snatching engineers from companies all over Austin, including AMD, IBM, and Motorola. Not surprisingly, however, much of the growth has come at the expense of McDermott's former staff at Somerset. Unable to resist Intel's compensation packages and desirous of working again for their well-respected leader, several key Somerset engineers and architects have quietly slipped over to the dark side (as PowerPC supporters often refer to the x86).

Although good for Intel, the brain drain on Somerset is another setback for PowerPC. Normal attrition, IBM's pull-out (see MPR 6/22/98, p. 4), and Motorola's internal struggle over whether to optimize Somerset's chips for Apple's computers or Cisco's routers have already chased away some irreplaceable talent. The new defections cannot be welcome news to Motorola's primary PowerPC customer, Apple, which is depending on the strength of Motorola's processors to get back in the game (see MPR 3/29/99, p. 22).

Although Motorola continues to recruit engineers to Somerset, and a solid core of engineers still remains, the company says it has now lost at least 15 PowerPC engineers to Intel. To halt a possible implosion, Motorola was forced to take aggressive action, which it did by filing suit against Intel and McDermott to block further pillage of its engineers (and the trade secrets that inevitably go with them). —*K.D.*