

ARM Grabs Embedded Speed Lead

Digital's StrongArm CPU Boasts Low Power Consumption

by Jim Turley

Proving that high performance and low power consumption are not mutually exclusive, Digital Semiconductor rolled out its first StrongArm processor, the SA-110. In horsepower, the chip rivals many desktop CPUs, while its power consumption is less than that of chips with only a fraction of its speed. This potent combination gives the SA-110 the potential to revitalize a sluggish market for handheld devices.

Paralleling its track record in the desktop CPU market, Digital has managed to take the lead in embedded performance by nearly any measure. The new chip runs at up to 200 MHz, belting out an impressive 230 Dhrystone MIPS from a 2-V supply. At 1.65 V, the chip delivers 185 MIPS, but power consumption drops by half, to 450 mW—the best MIPS/watt ratio of any high-performance microprocessor shipping or announced. Devices are sampling now, with production scheduled for midyear; prices range from \$29 to \$49.

Digital Pumps Up ARM's Strengths

The SA-110 is the first fruit of Digital's ARM license. The company worked through 1995 on the StrongArm core (see [091504.PDF](#)), a major redesign of the ARM pipeline and the basis of the SA-110. Under the terms of their agreement, Digital turns all StrongArm cores over to ARM, and the two companies are currently negotiating with the first StrongArm sublicensee. We suspect this potential second source is NEC.

The SA-110 will begin sampling in three speed grades: 100, 160, and 200 MHz. At top speed, the processor delivers 230 MIPS, based on Dhrystone 2.1. As the chart in Figure 1 shows, this performance puts the SA-110 well ahead of all other embedded CPUs, including the latest MIPS, i960, and SH processors.

Like Digital's other well-known microprocessor design, the SA-110 achieves its performance through fast clock rates and large on-chip caches. As a function of clock frequency, StrongArm's integer performance is not significantly different from other RISC or RISC-inspired architectures. To keep

the embedded chip humming along at 200 MHz, however, large caches are required.

The SA-110 is the first ARM chip to include separate instruction and data caches, each 16K in size. Both are 32-way set-associative (making them almost fully associative) with large 32-byte lines. The write-back data cache is supplemented with a write buffer to help alleviate pipeline stalls. (At such high clock rates, stalls are more than usually disruptive to the processor.) The SA-110's external bus interface operates at a programmable fraction of the internal pipeline clock, from 10 MHz to a maximum of 66 MHz.

Big Caches Aid Performance, Difficult to Flush

The SA-110 is the first ARM processor to have separate caches, and the first to support the write-back update policy,

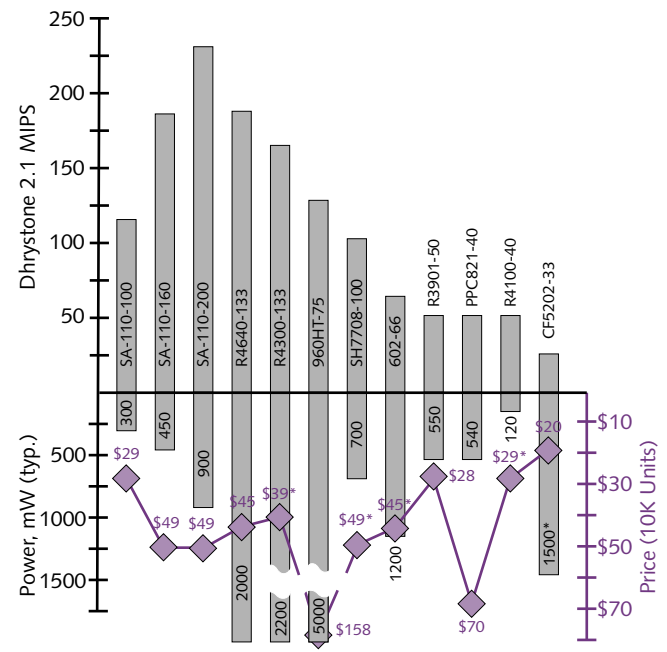


Figure 1. Digital's StrongArms are well ahead of other 32-bit CPUs when comparing the three axes of performance, price, and power consumption, (Source: vendors except *MDR estimates)

so some new decisions faced Digital's engineers. The first was how to maintain the consistency of the data cache. Unlike the ARM610 or ARM710, the only way dirty data can be forced out of the cache is to force new data in. A simple—though somewhat inelegant—way to accomplish this is to execute a program loop that reads a large block of uncached but cachable memory. To guarantee a reloaded data cache under all circumstances, the processor must load a contiguous block twice the size of the cache, forcing out any stale data.

There are three conditions that could necessitate a cache flush: when the logical-to-physical MMU mapping has changed, when the processor is about to enter sleep mode, or when the system is turned off. In sleep mode, all processor state is lost and the chip must be reset to restart operation, so flushing the data cache will ensure that all dirty data is written to external storage.

Write Buffer Keeps Performance Up

With a large speed mismatch between the SA-110's 200-MHz pipeline and its 66-MHz bus, a write buffer is essential to preserve performance. The write buffer holds 128 bytes, organized as an eight-deep FIFO with 16-byte-wide entries. Each of the eight entries is tagged with its target address.

Pending write operations are always executed in the order they were committed; the SA-110 does not reorder writes to take advantage of any reference locality. The chip does examine the address of each new entry as it is posted to the write buffer to see if it matches the address of the next-most-recent entry. If both write operations are to the same 16-byte-aligned address range, the two items are combined. Thus, the SA-110 is able to reduce bus traffic somewhat when software performs sequential (or at least, neighboring) stores to memory, as long as they are issued consecutively.

Like most CPUs, the SA-110 gives reads priority over writes. Loads or instruction fetches that miss their respective caches are forwarded directly to the bus interface, cutting ahead of any pending stores in the write buffer. To avoid consistency problems, the SA-110 includes a content-addressable memory (CAM) that checks the load address against all eight addresses in the write buffer. If the load does not fall within the 16-byte-aligned address of any of the pending store addresses, the load executes immediately.

If the load address does match one of the store addresses, the load stalls until the conflicting store operation completes; the chip does not read from the write buffer. The core stalls while the write buffer is drained, flushing pending writes to external memory before the load can proceed. In the best case, this delays the load by a single write transaction, if the conflicting store was already at the head of the FIFO. In the worst case, when the load address matches the most recently posted store, the processor drains the entire write buffer—up to 128 bytes—before the load can proceed.

Core Slows When Not in Use

Unbuffered write cycles and all read cycles that miss the cache

force the processor pipeline to stall until the bus transaction is complete. Seeking to take advantage of this situation, Digital designed the SA-110 to drop its internal pipeline clock to the frequency of the bus clock during bus transfers. This technique lowers power consumption and reduces potentially wasted CPU cycles. The core may not stall completely; it may still be able to execute with cached instructions and data for several cycles, albeit at a reduced clock rate.

In some cases, slowing the pipeline clock will cause a loss of performance; it depends on what is in the instruction cache and what data dependencies the program might have. If the CPU is waiting for a critical data item, no performance is lost and power consumption is significantly reduced. Conversely, if there are no data dependencies, slowing the CPU by a factor of nine or more could seriously impact the performance of cachable loops. Moreover, the effect is unpredictable because it depends on a complex relationship between the contents of the two caches and the write buffer.

MMU Upgraded from ARM610

The SA-110's MMU is nearly identical to the ARM610's, the design of which was influenced by Apple's engineers to support Newton's object-oriented operating system. The new chip includes 64 TLB entries, 32 each for instructions and data, and performs a hardware tablewalk on a TLB miss. The two-level page-table hierarchy still supports 4K, 64K, and 1M pages. Memory can be declared cachable or uncachable on a page-by-page basis. The MMU supports the concepts of domains, clients, and managers, allowing object-oriented software to control access to areas of memory dynamically.

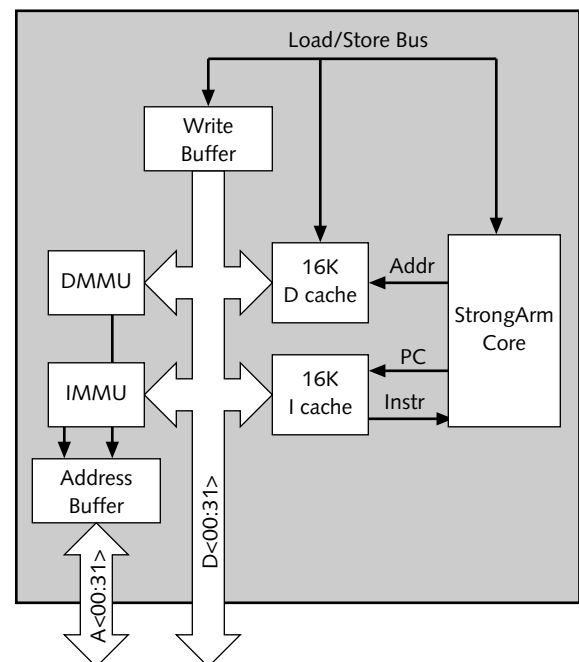


Figure 2. The SA-110 merges the 32-bit StrongArm core with dual 16K caches and dual MMUs similar to those in the ARM610 microprocessor used in the Apple Newton.




Figure 3. The majority of the SA-110's 50 mm² of silicon is devoted to caches. The basic StrongArm core accounts for only about 115,000 of the chip's 2.1 million transistors.

Although the SA-110 has separate TLBs for instructions and data, as Figure 2 shows, both translation tables share the same root pointer, so separate logical-to-physical mappings for data and instructions are not allowed. If the application's code and data do not overlap (share logical addresses), the chip can cache 64 different page-table entries.

Digital's Manufacturing Process Holds the Key

Power consumption is StrongArm's other strong point, helped along by Digital's manufacturing processes. The chip is fabricated in the company's 0.35-micron three-layer-metal CMOS process. The process technology for this device is a crucial factor in achieving its performance and power-consumption goals. It also, not incidentally, gives Digital a headstart over most other ARM licensees in producing StrongArm chips at the most attractive performance/watt levels. Of the list of announced ARM licensees, only NEC, Samsung, and Cirrus (through Digital's fab) have access to a

| | SA-110 | SA-110 | R4100 | MPC821 |
|---------------------|--------------------|--------------------|--------------------|----------|
| Vendor | Digital | Digital | NEC | Motorola |
| Architecture | ARM | ARM | MIPS | PowerPC |
| Frequency | 100 MHz | 200 MHz | 40 MHz | 40 MHz |
| I/D cache | 16K/16K | 16K/16K | 2K/1K | 4K/4K |
| MIPS | 115 | 230 | 40* | 52 |
| Voltage | 1.65/3.3 V | 2.0/3.3 V | 3.3 V | 3.3 V |
| Power | 300 mW | 900 mW | 120 mW | 540 mW |
| Die size | 50 mm ² | 50 mm ² | 25 mm ² | n/a |
| Process | 0.35µ 3LM | 0.35µ 3LM | 0.5µ 3LM | 0.5µ 3LM |
| Production | 2Q96 | 2Q96 | Now | 1Q96 |
| Est mfg cost | \$18* | \$18* | \$8* | \$15* |
| Price (10K) | \$29 | \$49 | \$28* | \$70 |

Table 1. Against other recent PDA microprocessors, the SA-110 has much higher performance with only moderately higher power consumption. (Source: vendors except *MDR estimates)

comparable 0.35-micron process; building StrongArm chips on a less aggressive process would be largely pointless.

Digital's process is nominally tailored for 2-V operation but tolerates supply voltages from about 1.5 to 2.1 V, with clock frequency naturally peaking at the highest voltage. Individual SA-110 chips do not tolerate the entire supply range, so Digital has qualified parts at two discrete voltage levels. At 1.65 V, the chip is available at 100 and 160 MHz; at 2 V, the SA-110 reaches 200 MHz. In either case, the pad ring requires a separate 3.3-V supply.

The die measures just 50 mm², most of which is cache, as Figure 3 shows. The basic execution units (labeled IBOX and EBOX) account for just 3.3 mm² (7%) of the total area.

The MDR Cost Model estimates an \$18 manufacturing cost for the SA-110, allowing Digital a comfortable margin on the part's \$29-\$49 price tag, particularly for the faster devices. Alternatively, this margin could allow Digital to give deep discounts to a high-volume customer (e.g., Apple).

Power Consumption Brings Substantial Benefits

At 100 MHz, the chip typically dissipates less than 300 mW at 1.65 V, as Table 1 shows. Turning up the clock speed increases consumption fairly linearly, to just under 450 mW at 160 MHz. At this rate, the SA-110 delivers an astounding 411 MIPS/watt, easily the best performance/power ratio of any high-end CPU. Its closest competitor, NEC's R4100, coincidentally achieves about the same ratio, but only because it squeaks out a mere 40 MIPS.

The 100- and 160-MHz versions are aimed directly at the handheld market, where absolute performance must be tempered with moderate power consumption. For so-called tethered applications, where battery life is not an issue, Digital offers its "high voltage" 2-V version at 200 MHz. Although its efficiency (in terms of MIPS/watt) is not up to the example of the slower parts, it still burns just 900 mW and could be used in portable devices as well.

What's remarkable is that even the 2-V version is much better than competitive devices. At less than 1 watt, the 200-MHz StrongArm still outperforms anything in its speed class. Intel's top-of-the-line 960HT, for example, fares pitifully against the Digital speed demon; a 75-MHz 960HT turns out 125 MIPS while consuming 4-5 watts, for \$150. Even IDT's impressive R4650-133 falls short: 75% of the performance, at about double the power, for twice the price. Other chips can match the SA-110's power consumption; a few others come close in performance. Nothing can touch it on both.

Even the SA-110's price sets a new benchmark. Unparalleled features and a single source of supply typically spell premium prices. Yet Digital's asking price of \$49 for the 200-MHz chip works out to just \$0.21/MIPS. Compared with about \$0.28 for NEC's R4300, \$0.69 for PowerPC 602, and \$1.25 for the 960HT, the SA-110 is a bargain, besides.

Digital chose to price the 160-MHz, 1.65-V part at \$49 as well. The company feels that power-conscious customers

will buy the slower part while performance-hungry ones will opt for the faster part. The \$29 price for the 100-MHz chip is a throwaway—it's better to sell underperforming chips at a discount than turn them into fobs for key chains.

Bus Specified for Light Loads

The SA-110's power specifications are believable, almost worst-case, measurements. Using a 33-MHz bus with 20-pF loading as a model, they assume the chip is continually toggling 50% of its address and data lines by generating back-to-back burst transactions. Only the pathological case of consistent cache misses, a full write buffer, and alternating I/O lines generates higher power numbers.

Although 50-pF loads have traditionally been the norm when specifying power consumption, Digital's more lightly loaded bus model is generally representative of what modern handheld or highly integrated systems actually experience. Other microprocessor vendors often use 50 pF, or even higher, loading when qualifying their devices, so direct comparisons are not always straightforward.

The chip supports two power-saving modes of operation, one reversible and one not. In idle mode, software can halt all internal pipeline and bus activity, leaving only the PLL running. Power consumption drops by 95% and all internal state is preserved.

Sleep mode, on the other hand, is irreversible. Putting the chip to sleep halts the PLL in addition to all other logic, and all internal state is lost. The processor must be reset to exit sleep mode, so flushing the contents of the data cache beforehand is crucial.

External Bus Interface Compatible with ARM7

The SA-110's external bus interface is a simple one, reflecting a legacy of straightforward ARM designs before it. The chip is hardware compatible with previous generations of ARM processors, although it can run at a considerably faster rate.

The nonmultiplexed bus consists of 32 address and 32 data lines. Simple cycle-start, cycle-end, byte-mask, and direction signals round out the set. All bus cycles are synchronous with the MCLK (memory clock) signal. Unaligned accesses are not supported.

The bus has two optional operating modes that can be enabled either individually or in combination to balance performance and ease of use. First, the external bus clock can operate either synchronously or asynchronously with the SA-110's internal pipeline clock. In asynchronous mode, the bus and the pipeline have no phase relationship, and internal synchronizers handle signals crossing clock boundaries. In synchronous mode, the bus clock is derived from the pipeline clock through a programmable divider, up to a maximum frequency of 66 MHz.

The other option controls cache-line fills. In standard mode, the SA-110 performs a strict linear fill, starting each 8-word burst at address 0 and incrementing addresses sequentially. Enhanced mode loads the line in interleaved

| Missed Word | Fill Order |
|-------------|-------------------------------|
| 0 | 0 → 1 → 2 → 3 → 4 → 5 → 6 → 7 |
| 1 | 1 → 2 → 3 → 0 → 5 → 6 → 7 → 4 |
| 2 | 2 → 3 → 0 → 1 → 6 → 7 → 4 → 5 |
| 3 | 3 → 0 → 1 → 2 → 7 → 4 → 5 → 6 |
| 4 | 4 → 5 → 6 → 7 → 0 → 1 → 2 → 3 |
| 5 | 5 → 6 → 7 → 4 → 1 → 2 → 3 → 0 |
| 6 | 6 → 7 → 4 → 5 → 2 → 3 → 0 → 1 |
| 7 | 7 → 4 → 5 → 6 → 3 → 0 → 1 → 2 |

Table 2. The SA-110 is able to fill cache-line burst requests in an interleaved manner compatible with most burst-EDO DRAMs.

order, somewhat similar to Pentium's (which Intel has patented), as shown in Table 2. Although the interleaved fill option returns the critical item first, yielding better overall performance, the resulting addresses can be difficult to decode using simple logic. Either burst order is compatible with synchronous SRAMs, synchronous DRAMs, and most burst-EDO DRAMs.

Strictly speaking, even in synchronous mode, the bus is not synchronous in the conventional sense because MCLK is not technically even a clock signal. That is, MCLK can have nearly any duty cycle and can have its high or low phase extended for an arbitrary amount of time; the clock may even be stopped. While such clock-stretching techniques are still common with inexpensive microcontrollers, it is very unusual to see this feature in a modern 32-bit microprocessor. Again, the SA-110 shows its roots in the cost-driven ARM design philosophy.

By allowing clock-stretching and asynchronous bus timing, the SA-110 lends itself to economical embedded systems using low-cost components. Rather than miss an entire bus cycle because a peripheral or an SRAM cannot meet the required setup time, it is possible to stretch MCLK by a few nanoseconds to get the required margin.

Bus Doesn't Tolerate Aborted Cycles

Once writes are posted to the write buffer, they cannot be aborted. To simplify the SA-110's error-recovery logic, Digital's engineers stipulated that external write cycles coming from the write buffer must complete without an error. The depth of the SA-110's write buffer and the speed mismatch between its buses would have made backing up the machine state for precise exceptions inordinately difficult.

If a bus access might generate an error, it must bypass the write buffer (and thus be uncachable as well), which forces the processor to stall until the transaction is complete. This model forces the hardware designer to classify all external resources into one of two categories: reliable memory, which can be cached and buffered, and unreliable peripherals, which are neither. In designing the SA-110's bus interface, Digital felt that in a "closed" embedded system like a PDA, soft DRAM errors are not a concern, and accessing a nonexistent resource is a programming error. In this view, creating a complex bus interface that is tolerant of occasional faults is not a prudent use of silicon.

Price & Availability

The SA-110 is currently sampling to selected customers at 100, 160, and 200 MHz; general sampling begins this month. The chip is expected to be in volume production in July. In 10,000-unit quantities, the 100-MHz SA-110 is priced at \$29; the 160- and 200-MHz are both priced at \$49. For more information, contact Digital Semiconductor (Hudson, Mass.) at 800.332.2717 or 508.628.4724; fax 508.626.0547, or via the Web at www.digital.com/info/semiconductor.

New Hope for Newton

Apple is evaluating the SA-110 for a new generation of Newton handhelds, with an announcement expected at Fall Comdex and deliveries scheduled near the end of this year. With nearly 12 times the integer performance of the anemic ARM610 inside current Newtons, the SA-110 has the capability to help propel Newton out of its current doldrums. The unit's handwriting recognition, long a sore point, would be vastly improved. The extra processing power could also be turned toward NSP-like functions such as a software modem or better audio, including speech recognition. Besides adding more features, this performance could reduce the amount of support logic needed, lowering system cost as well.

Interestingly, the 20-MHz ARM610 in current Newtons and the 160-MHz StrongArm have similar power numbers. (Apple clings to the 5-V version of the 610 for compatibility with Newton's other logic, even though a 3-V version has been available for more than a year.) Thus, Apple could turbocharge Newton's performance without altering its power budget or affecting battery life. The bus protocol is nearly identi-

cal between the two chips, allowing Apple to leverage its existing hardware designs, although turning the SA-110's bus down to 20 MHz would sap some of the potential performance gain. Although Newton would benefit greatly from a StrongArm infusion, the actual performance gain running Newton OS would fall well short of the order-of-magnitude improvement seen with Dhrystone.

StrongArm Overshadows ARM810 Release

The StrongArm core was the first of two ARM processor core upgrades announced at the end of 1995. The other, ARM8, shares a few features with StrongArm, particularly its five-stage pipeline. But the similarities end there.

Unlike StrongArm, the ARM8 core (*see 0917MSB.PDF*) was developed entirely at ARM's Cambridge (U.K.) facility as a booster for its fab partners' current ARM7 designs. As such, the ARM8 enhancements had to be somewhat less aggressive than StrongArm's, so the design would be portable across different fab processes.

The performance of the 810 lies between that of today's ARM7 devices and the SA-110, as Figure 4 shows. However, the 810's portable process rules keep it from achieving the same kind of power efficiency as Digital's chip: at 75 MHz (the current top speed) the 810 dissipates 500 mW at 3.3 V—66% more than the slowest SA-110, which delivers 40% better performance. In other words, the SA-110 squeezes more than 2× the performance from the same amount of juice.

The 810 die measures about 40 mm², or about 80% the size of the SA-110, but with one-fourth the cache. This and other factors place its manufacturing cost at about \$11, according to the MDR Cost Model. Unless the 810's vendors can price the chip below \$20, it will have a hard time competing with the SA-110. ARM has not yet revealed which vendors will be building the 810, except to say that parts will be available in 2H96, a bit later than Digital's shipments.

StrongArm Overtakes Newest MIPS Chips

NEC's R4300 microprocessor is a close match for the SA-110 in many respects. At its top speed, the 133-MHz MIPS chip (*see 0916MSB.PDF*) is the SA-110's equal in performance, clock for clock, but can't match StrongArm's higher clock speeds. Its die size, manufacturing cost, and volume pricing are also competitive. But the R4300 falls down in peak clock speed and power consumption. The Nintendo engine burns around 2 watts, far too much for handheld applications, which, admittedly, were not Nintendo's immediate concern.

Scaling down the R4300 produces an R4100, a chip that's better suited to power-sensitive applications (*see 090403.PDF*). At 3.3 V, the R4100 yields a MIPS/watt ratio every bit as good as the SA-110's, but at 40 MHz, its performance cannot begin to reach StrongArm's. The R4100 and the SA-110 are on the same performance/power curve, but at different ends.

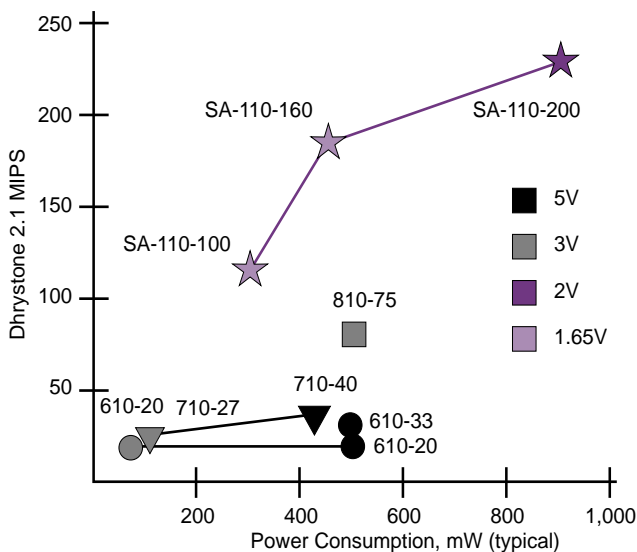


Figure 4. Even compared with other recent ARM implementations, Digital's SA-110 chip offers superior power/performance.

New Opportunities Open Up for ARM

Certainly there are other targets of opportunity for the SA-110 than just PDAs. On performance alone, the chip opens possibilities in television set-top decoders, video games, the fabled Internet browser, and other consumer-priced applications. Oracle President Larry Ellison, for one, has stated his intention to use the integrated ARM7500 device (*see* [080905.PDF](#)) for that company's first Internet terminal, with StrongArm chips powering a follow-on version. Factoring in the chip's low price just makes it more attractive, whether power consumption is an issue or not. For the time being, Digital has the privilege of trawling some of these markets by itself.

The downside is that the StrongArm chip is currently sole sourced, a situation that might make some high-volume OEMs think twice before committing to a new microprocessor. Digital has a short track record supplying CPUs on the merchant market, and Alpha experience is not relevant here. If anything happens to Digital's fab capacity, or if market conditions force a change of plans, the only source for StrongArm chips could conceivably dry up.

On the other hand, Digital has no shortage of fab capacity for StrongArm. Although Digital has touted the SA-110 as a means to fill its fab, apart from Alpha chips and some PCI peripherals, the company's Massachusetts factory is underex-

ploited, even after selling half of its capacity to Cirrus. Assuming an optimistic run rate of 500,000 processors per year, the little 50-mm² chips would fill scarcely 1,200 wafers per annum, hardly enough to keep Digital's fab lines full. As before, the company must find a way to amortize its investment in manufacturing soon, before all that expensive equipment goes to waste—or gets sold.

An additional high-volume vendor for StrongArm microprocessors would appease many OEMs without necessarily cannibalizing Digital's business. If the alternate source also offered ASIC capability—an ARM tradition—it could broaden StrongArm's appeal in the market and provide royalties for Digital. A second set of engineers could also hasten the development of application-specific derivatives of the part for individual markets, again benefiting Digital, ARM, and the architecture in the long run.

Technically, the SA-110 is a remarkable part and a milestone for both Digital and ARM. It offers a nearly unbeatable combination of performance, price, and power consumption. Working together, the two firms have soundly whipped any remaining doubts about the ARM architecture's growth path and future prospects. Competitors in the high-performance and low-power fields have some catching up to do. With any luck, the new systems that are enabled by this new microprocessor will be as exciting as the chip itself. ■