

THE EDITOR'S VIEW

Apple's Macintosh Strategy

Apple to Remain Focused on System 7, Proprietary Systems

By Michael Slater

Last month, Apple invited a few dozen analysts and reporters to a two-day briefing on its system software directions. Apple previewed several exciting new additions to System 7, including the Open Collaboration Environment (a collection of tools to support rich electronic mail), a new imaging model to supplement QuickDraw, a scripting language called AppleScript, and a new microkernel layer that will sit between System 7 and the hardware to provide preemptive multitasking, multiple threads, memory protection, and other capabilities.

These new features look great, but many were in a "pre-alpha" state and will not appear as products for some time. Apple would not give any hint about when any of these features would be available to users. The microkernel features, in particular, appear to be years away—a disappointing fact for those of us who would like our Macs to be more reliable and better at multitasking. The fact that Apple is pitching these capabilities so far in advance of their availability is indicative of its need to convince computer users that the Macintosh has a strong future. Apple apparently wanted to counter a common perception that the "Pink" operating system under development at Taligent was key to its future; if there was one consistent message throughout the two days, it was that Apple is focused on System 7 for the long term.

Indeed, Taligent was barely mentioned until numerous questions were raised about how it fit in, prompting John Sculley to devote the first part of his lunch-time speech to explaining it. Taligent, he said, is aimed at providing corporate computing customers with an effective way to implement customized applications. Oddly enough, this is exactly the explanation Steve Jobs has been giving for where NextStep has been successful. Next will have a five-year-plus lead in this market by the time Taligent ships its first product.

Apple also made it clear that while it will port certain elements of its system software technology to other platforms, System 7 itself is unlikely to be ported to other computers or licensed to other manufacturers. In essence, the Macintosh platform will remain as proprietary as ever. This dooms the Macintosh to a continued minority role in the computer business, and it is the best news Microsoft could hope for.

In addition to Taligent, there is one other vehicle that will allegedly allow other system developers to

make systems that can run Macintosh applications—PowerOpen, the merger of IBM's AIX and Apple's A/UX, hosted on the PowerPC architecture. This received little attention at the briefing, and indeed, very little has been said about PowerOpen. It is not clear whether all the System 7 features will be available in PowerOpen or whether it will be truly open (i.e., available to any system maker that wants it). Apple clearly thinks of it as a way for UNIX users to run Macintosh applications, not as a competitive platform to the Macintosh, and we suspect Apple will ensure that it stays that way. One difference will be in the memory requirements, which are likely to be 16 Mbytes for a minimum usable system. If this were the only limitation, it might still be a viable alternative for Macintosh power users, but we suspect there will be other issues, including application compatibility problems.

Apple described its PowerPC migration strategy, essentially repeating the information it presented at the Microprocessor Forum last November. Apple did, for the first time, demonstrate System 7 running on an 88100 processor (not an 88110), in a fully emulated environment—i.e., none of the operating-system or ToolBox code was native. Apple gave enough of a demonstration to show that emulation does work with a variety of applications, but not enough to show how well it performs.

All of this left us wondering whether the PowerPC architecture will achieve significant volume outside of Apple's own use. If PowerPC systems from everyone other than Apple remain targeted at the UNIX market, Apple will be the dominant user of PowerPC chips—exactly the situation it was trying to escape by abandoning the 88110. The more we hear, the more it seems that Apple would have been better off staying with the 88110, which would have gotten them to market a year earlier with a chip that is twice as fast as the initial PowerPC implementation will be. One Apple manager commented that the software is the critical path in any case, so the switch won't really delay them. If this is the case, we have to wonder about Apple's software development capability. After all, they've been working on this for years. A small startup, Quorum Software, has already shown software that can run Macintosh applications on a variety of RISC systems, using a completely reimplemented ToolBox. If a small company operating without Apple's resources and without access to its source code can do this, why can't Apple? ♦

Most Significant Bits

Ford to Switch to PowerPC

According to sources close to the company, Ford has changed its plans to use an embedded version of the 88000 as its next-generation engine control processor and will instead use a special implementation of the PowerPC architecture. This change is likely to mark the end of the 88000 architecture's use in mainstream embedded control applications, although the 88110 may be of interest in some high-end embedded systems. Motorola had been developing at least two 88300 family members—one for Apple for future portable computers, and one for Ford for engine control. Both of these projects are now dead.

The good news for Motorola is that both companies are switching to PowerPC, enabling Motorola to hold on to the business. For Data General and other system companies using the 88000 architecture, however, it is but one more blow to the architecture's prospects.

Ford is rumored to have been unhappy with Apple's switch to PowerPC, in large part because Ford was counting on Apple's use of the 88110 to provide desktop systems for software development and to ensure the availability of a range of compilers and other development tools. From Ford's perspective, the architecture itself is relatively unimportant, as long as it provides the minimum performance level they need and can be implemented in a reasonable amount of chip area.

Motorola Outlines 68000 Future Strategy

In an attempt to clear up the confusion resulting from press reports about Motorola's cancellation of the 68050 project, Motorola has laid out its plans for the next steps in the 68040 evolution. The 68050 was initially planned to be a relatively minor upgrade to the 68040, increasing the cache size and tuning speed paths to enable a higher clock rate. Motorola has now canceled this program to allow more resources to be devoted to products that were to follow the 68050, which Motorola says will now ship about six months sooner than originally expected. Motorola's 68000-family marketing manager Jim Reinhart said that simply doubling the cache size didn't provide a big enough performance boost to interest customers in the chip.

The two next-generation products are the LP040, a low-cost, low-power implementation of the 68040 (see next item), and "Q," the code name for a product that is likely to be called the 68060. The next-generation "Q" processor will be a superscalar, superpipelined design that will borrow some technology from Motorola's 88110 processor. The performance target is three-to-four times 68040 performance. While other 68000-family

processors have been reimplemented as static designs, this is the first 68000-family processor that will be fully static from the start. It will also be a modular design, making it easier to create derivatives that might, for example, omit the FPU and the MMU for embedded applications. Samples are expected in late '93.

As for the 68040, Motorola says it is now delivering production quantities at 33 MHz. It took much longer than expected to get the 68040 up to this speed, primarily due to difficulties in finding and tuning critical timing paths, but Reinhart says he is confident their yield problems at 33 MHz are over. The 33-MHz version, like the 25-MHz version, is fabricated in a 0.8-micron process. A 40-MHz version, fabricated in a 0.65-micron process, is promised for production in the third quarter, with limited quantities of a 50-MHz version by the end of the year. Pricing in 1000s is \$390 for the 25-MHz version and \$499 for the 33-MHz version; prices for the faster clock speeds have not been released.

Motorola's decision to drop the 68050 may be a good one, if indeed it accelerates the development of Q, but the last-minute change in strategy may confuse some customers and raise questions about how well thought-out Motorola's plans are. Motorola has thoroughly alienated its major U.S. workstation customer, Hewlett-Packard, which lost significant market share due to Motorola's delays in shipping the 68040, and HP is busy trying to migrate its customer base to PA-RISC. In the process, however, HP risks losing these customers to another RISC workstation vendor.

Apple is, of course, the key customer for 68000-family processors in general-purpose computers. The 68000-to-PowerPC transition Apple has planned puts Motorola in an odd position. If Motorola does a great job with Q, it could significantly delay the migration of Apple's customers to PowerPC, since the early PowerPC chips won't have spectacular performance. Motorola gets the business in either case, but this situation must make for some difficult strategic planning meetings at Motorola. (Cynics might wonder if there *are* any strategic planning meetings at Motorola's high-end microprocessor operation.)

Low-Power '040 Targets '020 Price Space

The other next-generation 68040 project is the LP040, a complete reimplement of the 68040 design in Motorola's most advanced process—0.5-micron, three-layer-metal. Unlike the 68040, it will be a fully static design, and it will be designed for 3-V operation. Samples are expected in the fourth quarter of 1993, including 33-MHz, 3.3-V and 25-MHz, 3.0-V versions. Pricing

is described as being "in the 68020 price space," which today means under \$30 in volume. The LP040 core will support the inter-module bus used in the 68300 embedded products, and 68300-family devices incorporating this core are expected in late '94.

The LP040 sounds like a great part for pen-based computers and other portable applications, but, in Motorola's grand tradition, it may be too late to get much of this market. The AT&T Hobbit and ARM processors are poised to capture a large part of the hand-held computer business, with Apple committed to ARM and GO targeting Hobbit as the RISC processor for use with its Pen-Point operating system. The LP040 would be a fantastic chip for a next-generation Macintosh notebook computer, however, and this is likely to be one of its major applications.

IBM Previews 386SLC Follow-On

IBM has introduced several new systems using its proprietary 386SLC microprocessor, which is pin-compatible with Intel's 386SX but offers 50–85% higher performance at the same clock rate by incorporating an 8K, two-way set-associative on-chip cache and a more efficient processor core (see *μPR* 10/16/91, p. 5). At the same time, IBM demonstrated a next-generation chip, described as "486-like" but not officially named, that is pin-compatible with the 386SLC but provides 16K of four-way set-associative on-chip cache. The new chip uses a similar core design with some minor optimizations, reimplemented in a faster process. It also incorporates a clock-doubler circuit, similar to Intel's 486DX2, allowing it to run internally at twice the external clock rate. IBM will market the chip in the fourth quarter as an upgrade for its existing 386SLC systems; it is likely to appear in new systems as well.

One problem with an on-chip cache in a 386SX pin-out is that there are no snooping control signals. IBM solved this problem by making the cache tags run at twice the processor speed, enabling them to snoop the external bus continuously without requiring any control signals. Another limitation is that the bus doesn't support burst transfers, retarding cache line transfers.

Even though the new chip has twice as much cache, the die size is only 9.0×7.7 mm (354×303 mils), as compared to 12.7×12.7 mm (500×500 mils) for the 386SLC. The smaller size is due to the more advanced process and a more carefully packed design. It is also designed for 3.3-V operation, and IBM claims that it will easily support 66-MHz internal operation with a 3.3-V supply. The combination of the lower supply voltage and special circuit design techniques slashes the power consumption, which is only 1 watt at 50 MHz with a 3.3-V supply. The 386SLC, in contrast, consumes more power than this with only half as much cache, operating at 20 MHz (with a 5-V supply).

IBM claims a Dhrystone performance of 21,600 in a 20-MHz system (operating at 40 MHz internally), compared to 11,500 for the 386SLC and 4,650 for Intel's 386SX. On application-level benchmarks, the new chip operating at 20 MHz external (40 MHz internal) is claimed to be faster than a 486SX-25.

IBM says its agreement with Intel places no limit on the number of these chips it can produce; apparently, the contract clause limiting the percentage of chips IBM can manufacture internally has expired, at least with respect to 386 derivatives. The chip has a system-management mode similar to Intel's 386SL, and it is designed to support suspend/resume modes for portable systems, so it is likely to appear in an IBM notebook computer. IBM appears to be well on its way to phasing out Intel as a supplier of 386-family processors, and the high performance of these chips makes Intel's low-end 486 chips redundant.

SGS and Siemens Offer Flash Microcontroller

SGS-Thomson and Siemens have each announced a 16-bit microcontroller with 32 Kbytes of on-chip flash memory. The device is an extension of Siemens' 80C166, a fast 16-bit microcontroller introduced in early 1990 (see *μPR* 3/7/90, p. 8). Later that year, SGS-Thomson and Siemens entered into an agreement for SGS to alternate-source Siemens' 80C166 line and for Siemens to supply SGS's ST9 microcontroller family. As part of this agreement, both companies have now introduced the new device with 32K of flash program memory. Siemens calls it the 88C166-5S, while SGS-Thomson calls it the ST10F166.

The flash memory is divided into four blocks, each of which can be erased and programmed separately. Erase time is about one second, and word or double-word writes take about 100 μ s. Endurance is initially specified as 100 erase/write cycles, but greater endurance versions are expected. With an endurance of only 100 cycles, it is useful for firmware updates and infrequent calibration information, but not for general data.

Like the original 80C166, the new device includes a 1K on-chip RAM, a 10-channel, 10-bit A/D converter, two USARTs, a 16-channel compare/capture unit for high-speed I/O timing, two general-purpose timers, and a watchdog timer. With a 20-MHz clock, the basic instruction time is 100 ns; it can perform a 16×16 multiplication in only 500 ns, or a division in 1 μ s. The chip is packaged in a 100-pin PQFP and provides 76 I/O lines.

Pricing from Siemens is \$95 in thousands, with samples promised in the second half of this year; a 32K mask-ROM version is priced at \$35 in thousands. SGS-Thomson prices the flash version at \$110 in hundreds, with pricing projected to drop to \$53 in thousands in 1993; the ROM version is \$30 in quantities of 10,000. ♦