Embedded Processors Focus on Integration

Application-Oriented Processors a Fertile Area for 1992

By Brian Case

To help keep the microprocessor marketplace in perspective, it is worth noting that all of the early microprocessors were what we now call "embedded" processors. While some in the industry doubted that the microprocessor would ever have a sizable market of any kind, almost no one even contemplated the emergence of a market for microprocessors in general-purpose computers.

It may now seem that the microprocessor has made its biggest impact on society as the CPU for personal computers, but that is arguably not the case. The embedded market got a five-year head start on the PC market and has dominated volume shipments ever since. In most homes or offices are at least ten times as many embedded as general-purpose processors: at least one each in a computer keyboard, laser printer, disk drive, and modem, as many as ten in a car, and others sprinkled about in copiers, fax machines, modems, microwave ovens, VCRs, TVs, remote controls, video cameras, still cameras, answering machines, CD players, calculators, intelligent thermostats, etc.

While the bulk of embedded processors is still simple, dirt-cheap four-or eight-bit CPUs, the trend for future growth has been made clear over the past 20 years. Embedded processors are getting faster, integrating more capabilities, and moving (if slowly) to 16 and, most recently, to 32 bits. Excitement over 32-bit embedded processors may be premature, but those who say that 32-bit embedded control will not amount to much sound suspiciously like those who, 15 years ago, said that personal computers would not amount to much.

It is a testament to the vibrancy of the embedded marketplace that the most powerful CPUs are not in personal computers but in personal computer peripherals, such as a 10-MIPS PostScript-compatible laser printer for $1300.

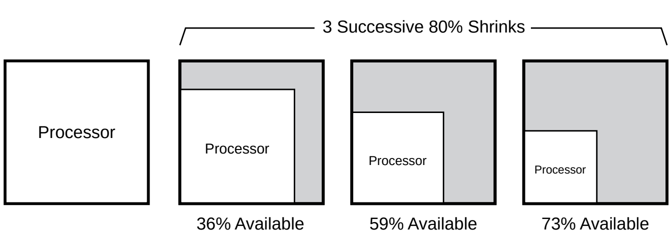

Figure 1. Effect of die shrinks in making area available.

Sometimes it seems that predicting the direction of the microprocessor industry more than a few minutes into the future is fruitless. Yet one prediction can be made with near certainty: IC feature sizes will decrease and more layers of metal will become commonplace. The result will be a dramatic increase in circuit densities, and this means that it is becoming less important to strive for the absolute optimal implementation of a given peripheral function. Even the size of the processor core will become much less important in a few years. Assuming an 80% process shrink every two years, Figure 1 shows graphically how much room is made available on a fixed-size processor die after only six years of such shrinks; the processor shrinks to about one-fourth of its original size.

Developments in the mid-range to high-end embedded processor arena are significant now and set the direction for the future of all segments of the embedded marketplace. As has been the case in the past, the features of high-end embedded processors eventually migrate to the low end. By the end of the decade, 32-bit processors will be rich with peripherals and cheap enough for use in even the most mundane embedded applications. As more memory and peripherals are added to embedded processors and begin to occupy the majority of the silicon area, the incremental cost of a 32-bit processor core over an 8-bit core will be small.

Recent Embedded Developments

At least two trends in high-end and mid-range embedded processor design emerged in 1991. The first trend is characterized by simply integrating a general-purpose system on a single chip, while the second trend calls for implementations that are specialized for a particular application. A third trend, where additional general-purpose processing capabilities are added to an existing processor, seems to be taking shape.

Examples of the first trend—general-purpose integration—include most of the MIPS-derivative embedded processors and the 29030 from AMD. MIPS-based processors from IDT, LSI, and Performance Semiconductor integrate an R3000 processor core and some cache. (See p. 14 of this issue for more on embedded MIPS processors.)

The 29030 from AMD integrates only an instruction cache, but this allows it to eliminate the separate external instruction bus and control signals that gave the original 29000 a high pin count. By integrating an I-cache, the 29030 brings to many systems a performance

Defining Embedded Processors

Traditionally, an embedded controller is a microprocessor that has significant I/O capability, such as parallel and serial ports, pre-decoded memory interfaces, analog ports, etc. Embedded processors are defined by their application, however, rather than by the processor design itself, and any microprocessor can be "embedded." Essentially, an embedded processor is one that is dedicated to running a single program and is not user-reprogrammable. Perhaps the most important aspect of the definition of an embedded processor is simply "inexpensive." Support software is also a key factor: a processor for which a popular operating system and a large suite of applications are not available quickly becomes "embedded."

feature they could not otherwise afford, while at the same time easing system design with a simpler bus interface. AMD wins as well since the die size of the old 29000 was essentially pad-limited; with fewer pins, AMD is better able to reduce the die size and cost.

The second trend—application specialization—is clearly exemplified by AMD's 29200, LSI Logic's 33020, and National's 32AM160. Motorola's 68300 family, with its on-chip intermodule bus, has the potential to be an even more striking example of specialization. So far, however, the members in this family are not really specialized for any one application, with the possible exception of engine control.

AMD's 29200 laser-printer controller is an example of a relatively quick customization of an existing processor for a specific market. The BTC and MMU of the old 29000 core were removed to make room for peripheral functions—DMA channels, DRAM control, video interface, etc.—that reduce the cost of the electronics for a laser printer. The most difficult part of bringing such a product to market is getting potential customers to agree on the features that should be included.

LSI Logic's 33020 X-terminal controller is a similar story. LSI saw an emerging market—X-terminals—with no processor specifically tailored for that application. Since the X-windows software is essentially just a regular program written in C, their existing 33000 chip had most of what was needed—a fast processor core, caches, DRAM control, etc.—but did not offer a compelling advantage over other processors that might seem cheaper. To create a compelling advantage, LSI integrated a small but powerful graphics unit that does not increase die cost much.

While AMD and LSI have recently discovered the wisdom of targeting a specific application area, National has been pursuing this strategy for some time. Their latest offering is the 32AM160, which is designed for voice applications—primarily digital telephone answering machines. While AMD and LSI still seem to feel most comfortable in the high end of the embedded spectrum, National is being more aggressive in low-cost, high-volume applications.

By basing the 32AM160 on the now decade-old 32000 CISC processor, National kept design time and die size down. While the 32000 core's performance is not up to the standards of today's RISC architectures, the 32AM160's autonomous DSP unit gives the chip the performance it needs for the targeted application without requiring a super-fast processor core.

Compared to the 29200 and the 33020, the 32AM160 is much more like a traditional microcontroller. The external bus is 8 bits wide, and the chip includes a large ROM and a small RAM. The on-chip program ROM sets this chip apart from most other 32-bit embedded controllers.

The third trend—functional augmentation—is not such a clear long-term trend and even overlaps the others somewhat, but it currently makes sense for some applications and some processors that have less than leading-edge performance. Processors exemplifying this trend are National's HPC+ and Zilog's Z86C94 with their integrated DSP units. While the DSP units in these processors are philosophically similar to the graphics coprocessor of the 33020 and the DSP unit in the 32AM160, the HPC+ and Z86C94 are not as specifically aimed at a single application.

These two chips are promoted for digital servo-loop applications, such as hard-disk drive control, but with the general set of peripherals—A/D, D/A, timers, serial ports—and on-chip ROM and RAM, they are appropriate for a variety of real-time control applications. National and Zilog are capitalizing on the economical instruction sets of their processor cores—24K bytes of instruction ROM goes a lot farther with the Z86C94 than with a RISC—while improving performance in important applications with the DSP units.

Future Embedded Developments

In the near-term future, vendors of embedded processors can be expected to continue the product strategies that they have started. Application specialization is the most exciting trend in embedded control because it promises to produce the most rapid development of new products and the most competition between vendors. With transistor counts at about 500K for cost-effective die sizes and the pervasive use of high-quality automated design tools, very-quick-turn customization of existing processor cores is becoming feasible.

With the ability to do quick-turn designs, the striking successes of one vendor can be expected to provoke other vendors to introduce products aimed at the same markets. If the 29200 proves to be a money maker for AMD, it is likely that one of the MIPS-based vendors will define a similar chip aimed at laser printers. Such a copy-cat chip would probably need to have a compelling advantage—such as higher performance or lower system cost or both—to overcome the design-in inertia created by the market momentum of the 29200.

Despite the fact that the bulk of shipments occur in low-end embedded control, some significant opportunities are found at the high end of the application spectrum. When Apple introduced the LaserWriter, for example, its 300-DPI resolution was a compromise reached to reduce processing and memory requirements while still providing compelling output quality. Even at the time of introduction, the engines were capable of 400 DPI, and many engineers knew that the resolution could be increased even more by modulating the laser in the Canon engine at a faster rate.

Now that very high-performance embedded control processors are available and memory is dramatically cheaper, higher-resolution, 600-DPI laser printers (and controller-board replacements for existing printers) are starting to appear. Even 1200-DPI laser printers are available, though currently expensive. In a few years, high-resolution laser printers will account for the bulk of the market, and the vendors who get a toehold in this market today could end up with large profits later.

To tap the high-resolution laser printer market, vendors can be expected to combine a set of peripherals similar to those in the 29200 with a higher-performance processor core. The 29200, without cache or floating-point hardware and with a low maximum clock rate, is clearly aimed at what is becoming the low end of the laser printer market.

As fabrication technology improves, vendors will be looking for ways exploit it other than reducing die sizes and costs. Improving performance by increasing cache sizes and integrating multiple processors are two ways of soaking up the increasing number of transistors, but other ways will evolve as well.

One use for a large number of transistors is to include some field-programmable gate-array logic on the chip. Consider, for example, the Xilinx SRAM-based technology. Already, Xilinx chips big enough to accommodate the implementation of a rudimentary processor are sampling. If a standard processor core were supplemented with a generous amount of SRAM-based logic tied into the internal processor buses, some very useful, application-specific (or even real-time) customization of processing capability could be done. For example, an instruction could be added that performs a simple yet non-standard logic function appropriate for fuzzy logic or neural-network computations.

As another option, the FPGA logic could be used to provide a peripheral function not present in the standard configuration. If the chip provides one serial port but two are needed, the FPGA could be used to provide an additional UART.

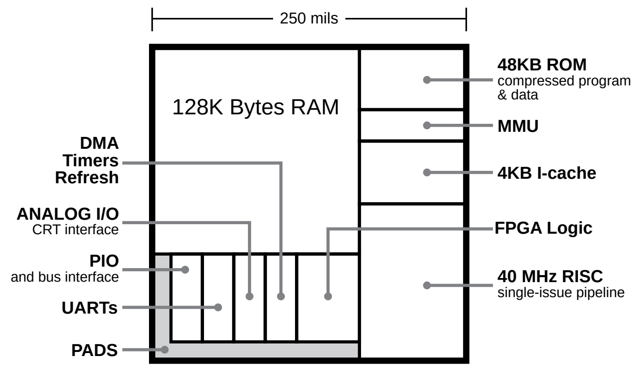

Figure 2. One possible embedded processor, circa 1998.

The problems with on-chip, reconfigurable logic are many, however. There is the cost differential: reconfigurable logic requires an order of magnitude more die area than equivalent custom logic. Also, changing the behavior of a processor or even a peripheral could create a debugging nightmare. Still, in five years, fabrication technology will reach a point where the cost for a significant amount of reconfigurable logic will only be a small percentage of a typical chip.

Yet another way of soaking up excess transistors is to simply integrate a large chunk of main memory onto the die with the processor. The first uses of this strategy could be in the display controller area. A laptop or palmtop display needs less than 40K bytes, while a large-screen, black-and-white monitor requires only about 100K bytes of frame buffer memory. Thus, in the near future, it will be possible to include an embedded graphics processor and a one-bit-deep frame buffer on the same die. Already, TI's SuperSPARC processor integrates 40K bytes of static cache memory. Note that frame buffer memory need not be as fast as cache and need not be static. Even with dynamic RAM, the performance of rudimentary graphical-user-interface functions could be dramatically improved by this integration of the graphics processor with the frame buffer.

Figure 2 shows the partitioning of a 1998 fantasy embedded controller. This chip would be about 250 mils on a side and would sell for perhaps $50 in volume. Most of the die would be dedicated to RAM; the program would be stored in compressed form in a ROM and decompressed into RAM upon power-up. While the specifics of this device aren't necessarily meaningful, it illustrates a key point: the processor core will be a small part of the chip.

Conclusions

The direction for the embedded market in the near future is likely to be an extension of the trends already

Continued on page 21

only NFL quarterbacks with at least 1500 passing attempts—a closed system growing slowly with time. In contrast, there is no such barrier for entry to consideration for inventor of the microprocessor.

As you see from this chart, there wasn't much interest in who invented the microprocessor during the 70s. The ratings in the chart for the 70s mostly show the effect of working on integrated circuit design and publishing some of the early technical papers. The 57 engineers actually doing the work at ten companies probably thought it was obvious. You can see the increased interest beginning in the early 80s. The figure is incomplete, since I was able to collect and plot only a small fraction of the recently identified contenders, but it does show some interesting features. (Perhaps Miami Dolphin's quarterback Dan Marino used a similar plot of his ratings improvement and its projection in negotiating his recent $25 million five year contract.) For example, the line beginning in 1989 that appears to be rising faster than the number of instructions in the MIPS architecture, shows the effect the patent office and a large public relations staff might have on a contender's rating (see line *a* in Figure 1). This line may represent the only candidate on the list with the equivalent of the NFL's 1500 attempts requirement, except in this case all of them are correspondence with the patent office.

Line *b* on the chart, showing a preliminary peak in the early 80s and tapering off later, may indicate the possible ill effects of over-aggressive self promotion. Other lines show consistent long-term or belated efforts to improve ratings either by individuals or, in at least one case, by an aggressive legal department at a large company which may be embellishing the patent claims of one of its former employees (see lines *c* and *d*, for example).

Unfortunately, I don't have any final results yet. I'm still collecting data and tuning coefficients, but progress is slow because the project is unfunded and, therefore, must be relegated to "hobby" status. I suppose I should be filing grant proposals with DARPA or NSF to get funding to complete the project, but I'm not from a big name university and this doesn't look like the beginning of a fad, so its chance of being funded is nil. It's too bad, because I think this kind of work is important to computer science and society.

If the intense acrimony built up over the last fifty years in the competition for recognition as inventor of the computer is any indication of what is in store for the candidates aspiring to recognition as the inventor of the microprocessor, we owe it to society and the giant candidate pool to head off the coming street fight. We owe it to ourselves and our profession to establish rules for friendly competition and a definitive, quantitative result. And may the best candidate win! ♦

Embedded Processor Trends

Continued from page 13

established. Some very-high-end embedded processors will continue to simply collapse more of a general-purpose system onto a single chip. Others will be designed with a specific application area in mind; these will have a set of peripherals peculiar to their intended use. Still others will consist of older processor cores augmented with additional processing units that make them appropriate for modern applications and let them exploit the cost advantages of simpler implementations.

Increasingly dense IC processes combined with better automatic design tools will lead to less time spent "hand tweaking" logic and circuits. In turn, a new processor derivative can proceed from definition to implementation in less time. Automatic design tools might not give the optimal implementation, but they help get silicon functioning correctly in less time. Consequently, the pace of innovation in the embedded market can be expected to accelerate.

The ability of vendors to produce customized embedded processors will have to improve, given the fact that the number of embedded applications will also increase at an accelerating rate. In ten years, personal calculators went from being expensive tools for engineers and students to being disposable "freebies" available to anyone with a promotional need. In the future, small televisions, telephones, and even personal computers may be relegated to such disposable status.

The golden rule of embedded control is that low cost wins. The most successful members of the 960 and 29000 families are not those that have higher performance but those that have reduced costs. In the 29000 line, the newest family members show performance moving sideways while cost moves down.

Over the next several years, as fabrication technology permits the integration of several million transistors on a low-cost die, embedded processors will provide a tremendously rich field for innovation. Innovation in general-purpose processors will consist mainly of increasing cache sizes, more pipelining, instruction decode and execution units capable of processing multiple instructions per clock cycle, and possibly multiple processors on a single chip. Embedded innovation will include all of the above plus increasingly complex mixtures of processors and peripherals, support for fuzzy-logic or neural-network processing, perhaps some field-programmable logic, and ever-expanding on-chip RAM and ROM.

In short, while general-purpose processors must evolve in ways that are compatible with existing application software and long-lived operating systems, embedded processors will evolve with the products into which they are embedded. ♦