SUPPLEMENT TO

ELECTRONIC DESIGN

FOR ENGINEERS AND ENGINEERING MANAGERS - WORLDWIDE

A PENTON PUBLICATION

MARCH 3, 1997

EMBEDDED SYSTEMS

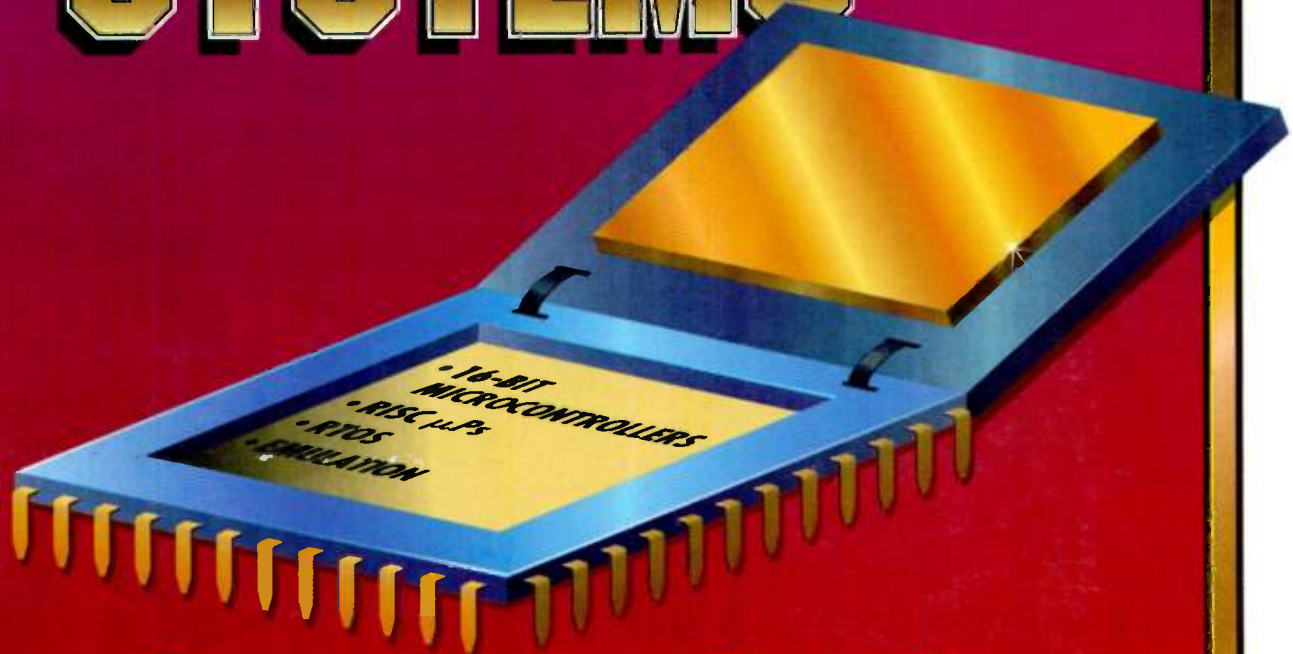

IN ELECTRONIC DESIGN

3.3V ISP™ PLD Performance Breakthrough!

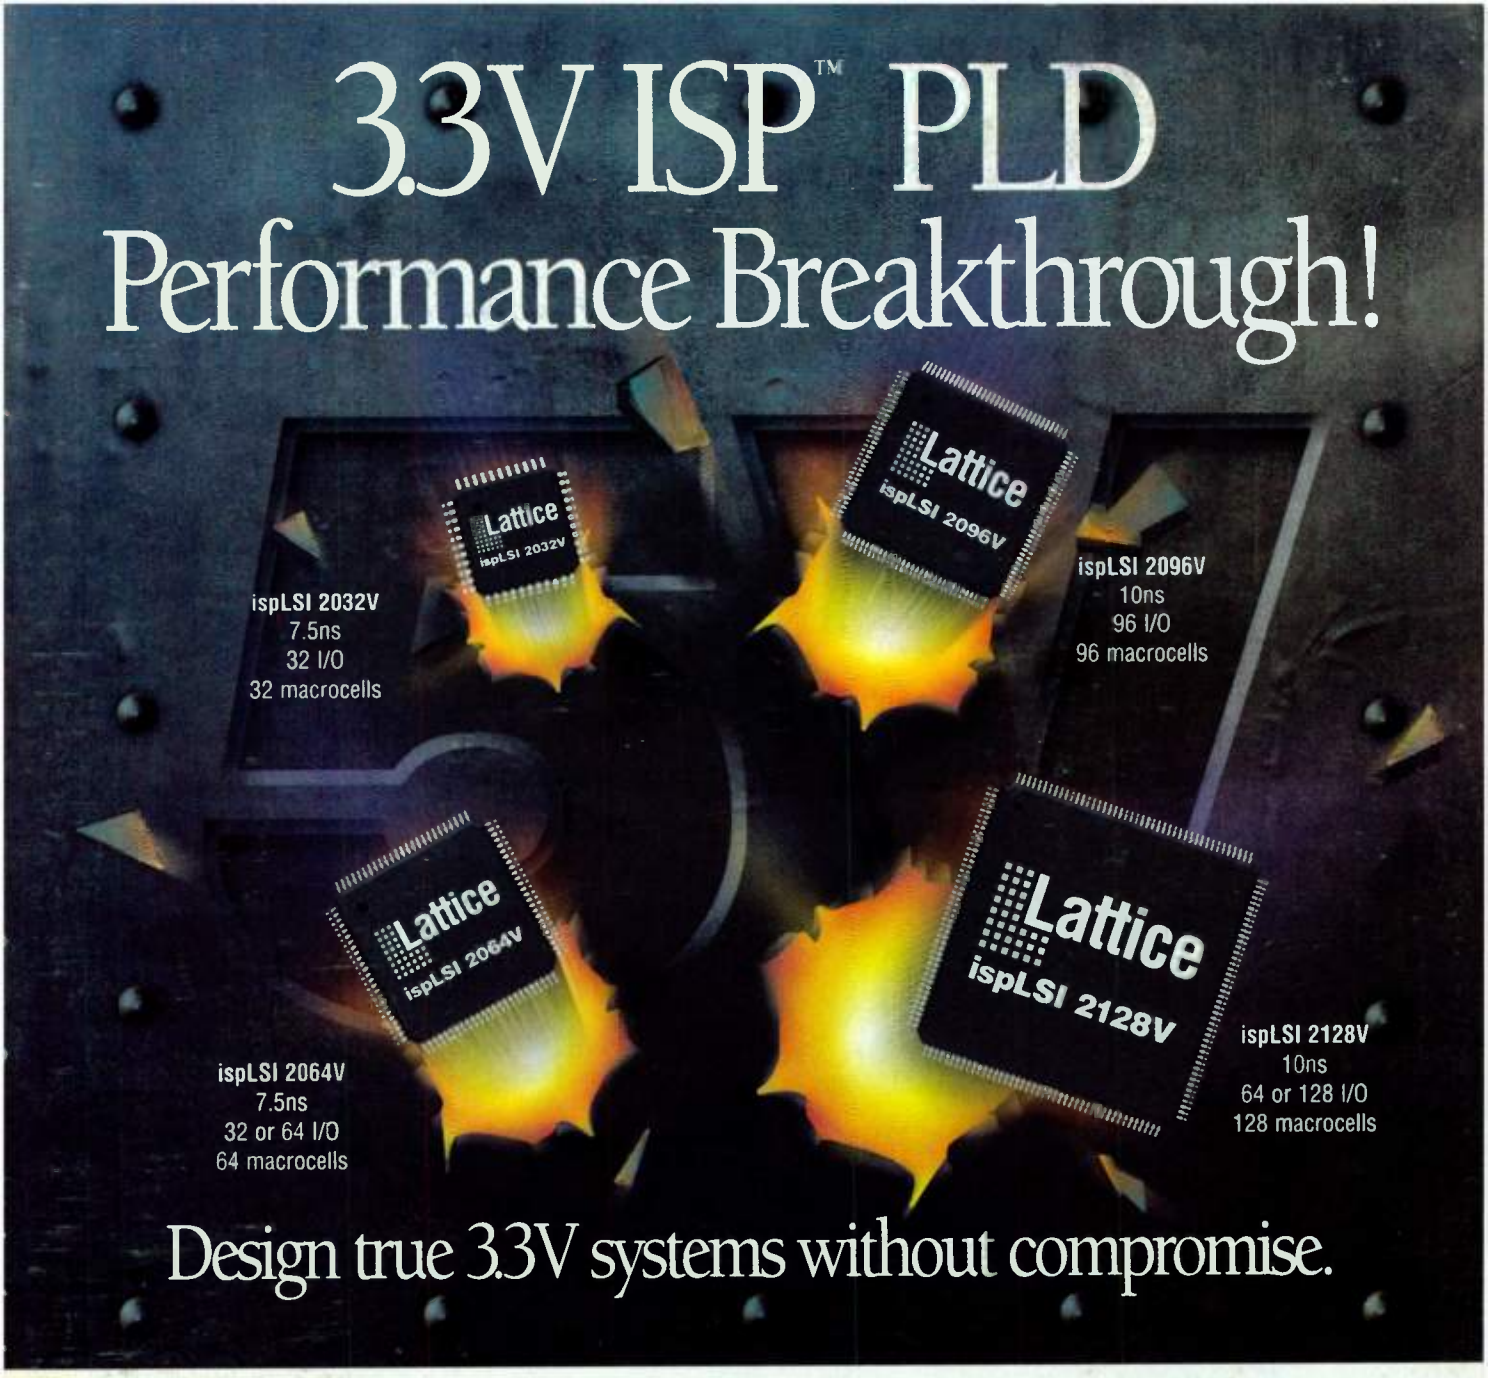

Design true 3.3V systems without compromise.

Smash the 3.3V high-density PLD speed barrier with Lattice's new in-system programmable ispLSI® 2000V Family—the first 3.3V programmable logic devices to deliver 5V performance. Featuring logic densities from 32 to 128 macrocells and I/O options from 32 to 128 pins, this new family gives you all the logic options you need to design 3.3V systems with breakaway speed.

FREE
ISP™ Design Software & Lattice ISP™ Encyclopedia

Includes: Software, Data Book, Handbook, ISP Manual on CD-ROM

All ispLSI 2000V devices are available in space saving thin quad flat pack (TQFP) packages that maximize PCB space. And every ispLSI 2000V device is in-system programmable using only a 3.3V power supply—an industry first! So manufacturing with high-density ispLSI devices saves you both time and money.

Give yourself a break and come up to speed with ispLSI 2000V PLDs. Call us today at **1-888-ISP-PLDS** and ask for information packet 331 or check out our web site at **www.latticesemi.com.**

Lattice™ Semiconductor Corporation

The Leader in ISP™ PLDs

Corporate Headquarters: Tel: (503) 681-0118, Fax: (503) 681-3037 • **France:** Tel: (33) 1 69 33 22 77, Fax: (33) 1 60 19 05 21 • **Germany:** Tel: (49) 089-317-87-810, Fax: (49) 089-317-87 830 • **Hong Kong:** Tel: (852) 2319-2929, Fax: (852) 2319-2750 • **Japan:** Tel: (81) 3-5820-3533, Fax: (81) 3-5820-3531 • **Korea:** Tel: (822) 583-6783, Fax: (822) 583-6788 • **Taiwan:** Tel: (8862) 577-4352, Fax: (8862) 577-0260 • **United Kingdom:** Tel: (44) 1932 831180, Fax: (44) 1932 831181

READER SERVICE 88

From jet planes to Chunnel trains, QNX goes the distance: testing turbines, loading cargo, even controlling air traffic.

QNX helps healthcare applications run faster, and cost less, thanks to its executive-class speed and full x86 support.

The QNX realtime operating system is

With its tiny, full-featured Internet suite, QNX helps web-transaction appliances do big business.

From traffic control to process control, QNX drives thousands of mission-critical applications nonstop, 24 hours a day.

So many applications. So many demands. How does QNX do it?

Start with rock-solid OS technology field-tested for over 15 years. Add in innovative products like the award-winning Photon microGUI™, QNX's embeddable windowing system. Provide a rich, robust toolset so developers hit the ground running. And keep the memory footprint exceptionally small to ensure runtime costs stay exceptionally low.

Most important, make it all fully scalable. So developers can deliver everything from web phones to factory-wide control systems—using a single OS. It all adds up to

Using a QNX-based vision system, shuttle astronauts stay focused on important things. Like launching satellites.

QNX serves up reliable service at the point of sale, whether it's handling fast-food orders or checking credit cards.

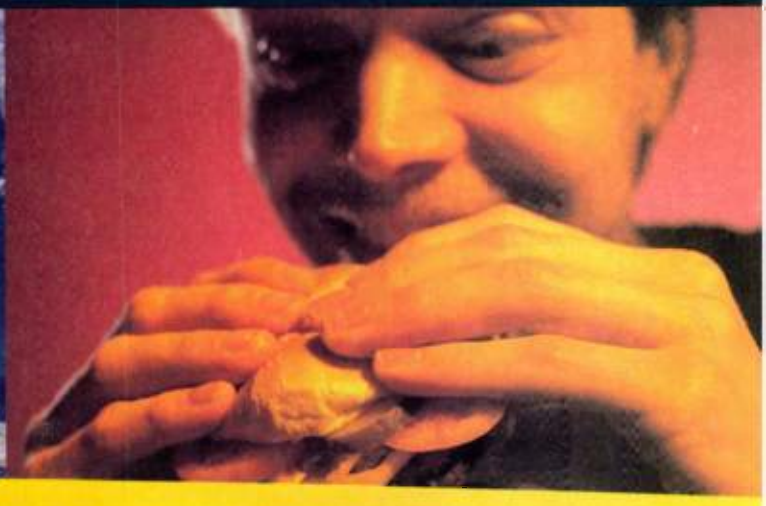

found in these worldwide locations

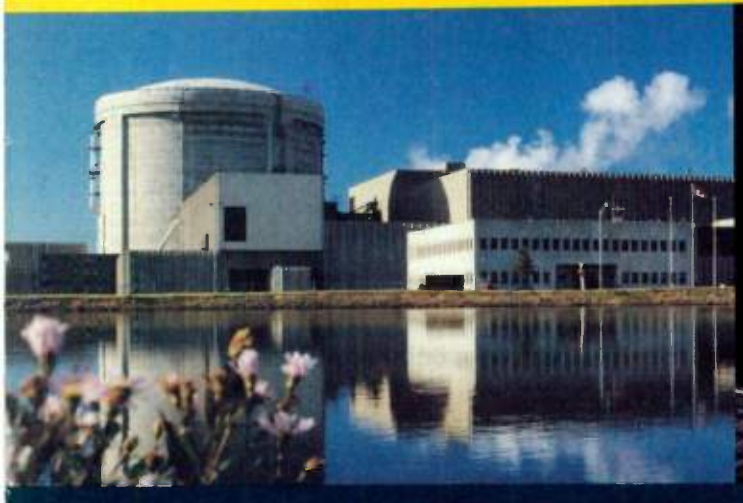

When safety is measured in microseconds, nuclear power stations count on QNX: it's a *real* realtime OS.

Thanks to QNX, the web is coming to your living room faster than you can say "URL."

QNX®

The Leading Realtime OS for PCs

Available Now:

- internet toolkit
- embeddable GUI & browser
- POSIX and Win32 APIs
- embedded filesystems
- memory protection
- fault-tolerant networking
- distributed processing
- multilingual support
- unrivalled x86 support
- embedded OEM pricing

www.qnx.com

call 800 676-0566 ext. 2242

QNX Software Systems Ltd., Voice: 613 591-0931 Fax: 613 591-3579 Email: info@qnx.com
Outside North America: Voice: (44)(0)1923 284800 or 613 591-0931 Fax: (44)(0)1923 285868 Email: QNXeurope@qnx.com

READER SERVICE 94

Palm-sized PC Microcontroller packs industrial I/O power with embedded PC control.

When a proven technological innovator emerges in any industry, it becomes the one to watch. Octagon Systems' new PC Microcontroller™ series blends I/O innovation with embedded PC control. The PC Microcontroller series not only provides a uniquely integrated hardware platform, but includes the industry's largest suite of embedded software in flash memory. You get the best of all worlds–a single card industrial solution that you can adapt to hundreds of applications without the compatibility problems of multiple card variations. Each of the Octagon PC Microcontroller cards includes DOS 6.22, but will run under QNX as well as many other popular real-time operating systems. They also include CAMBASIC, Octagon's fast multitasking control language that supports all the I/O on the card–no drivers to write. Each 4.9"x 4.5" package is rated from -40° to 85°C, plugs into an ISA bus slot or operates stand-alone with a 5V supply.

Octagon PC Microcontroller Series	
CPU	40 MHz
Parallel Port	IEEE1284
Digital I/O	17-65 lines
COM Ports	2-4
Analog I/O	Up to 10
CAMBASIC	✔
Networking	✔
Diagnostics	✔
Flash File	✔
DRAM Supplied	2 MB
Flash Supplied	1 MB
SRAM Supplied	128K
Real Time Clock	✔
KBD/Speaker	✔
Prices Start at	$340/100s

The Octagon PC Microcontroller series. It's loaded with software. There's no memory to buy. You don't need a professional programmer. And it eliminates the headaches of mixing and matching the CPU, I/O and software from different sources. It's an irresistible new product family from Octagon Systems–I.O. with I.Q!–the latest innovation in our sixteen year history of success.

Works with QNX

OCTAGON SYSTEMS®

6510 West 91st Avenue Westminster, CO 80030 USA
TEL: 303-430-1500 FAX: 303-426-8126 www.octa.com

" See us at the Embedded Systems East Conference, Booth #1404 "

READER SERVICE 91

ELECTRONIC DESIGN
FOR ENGINEERS AND ENGINEERING MANAGERS WORLDWIDE

Embedded Systems Features

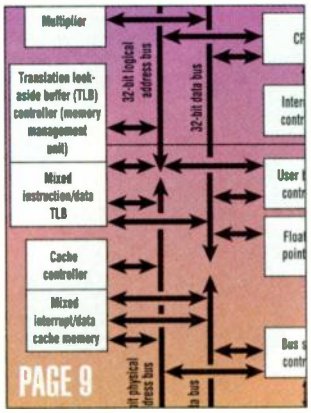

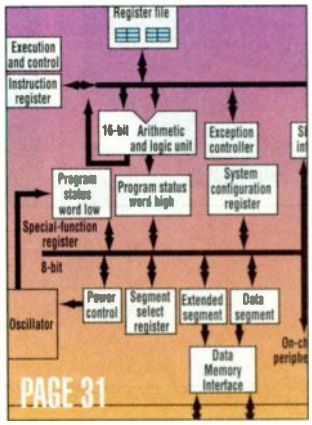

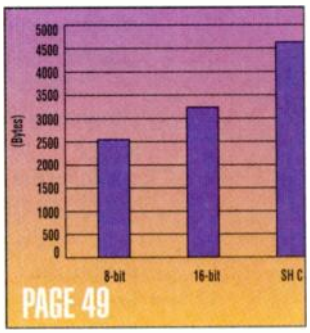

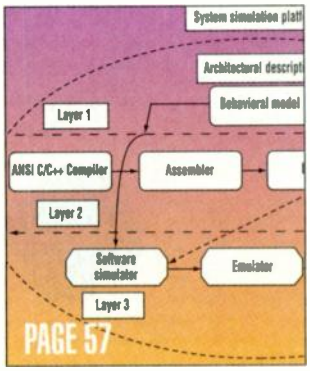

Departments

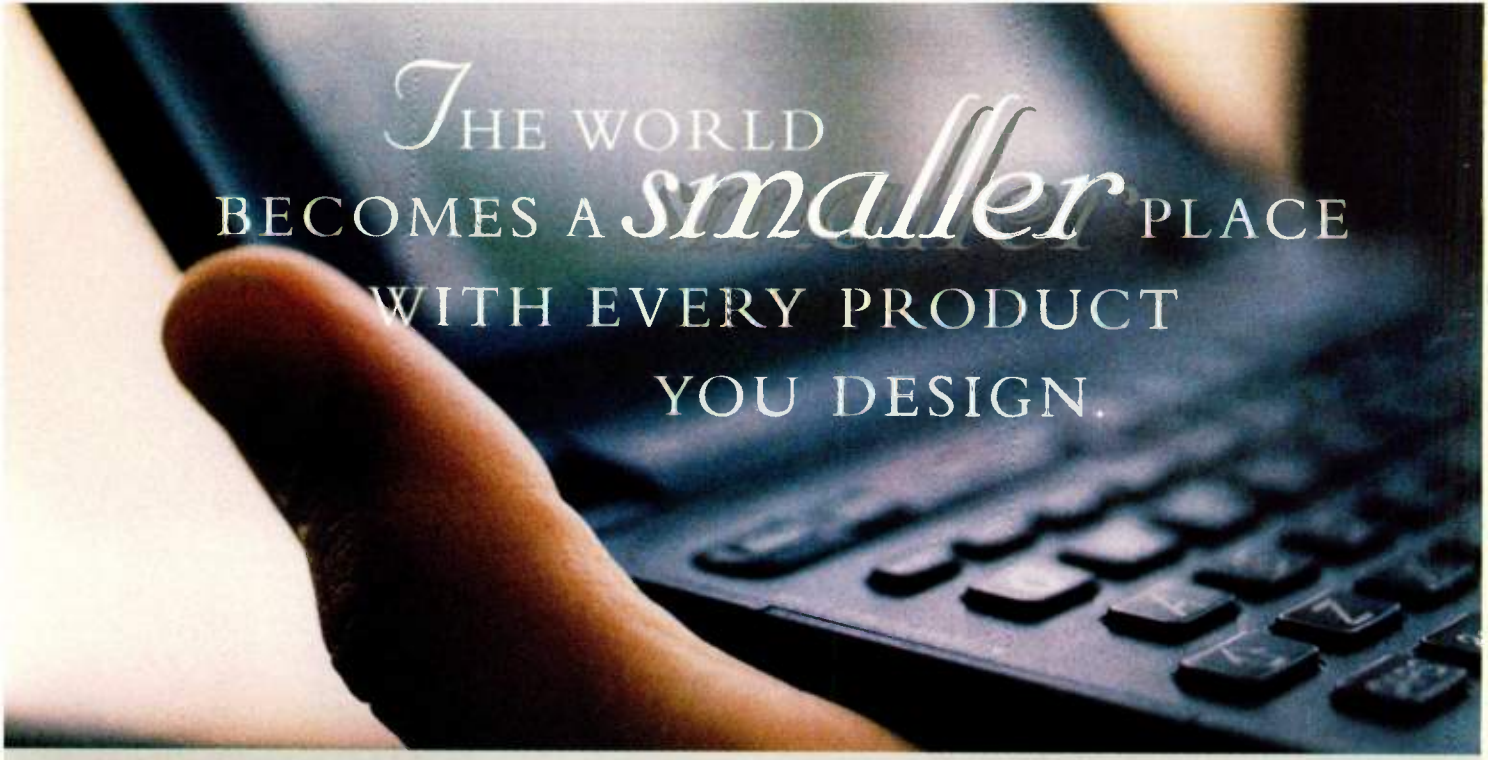





You have a remarkable power.

It is the power to bring people together, by allowing them to be farther apart. And it is a power that, while firmly grounded, is nothing less than nomadic in spirit.

Choose your embedded processor wisely.

PDAs are now reaching their full potential. And you have the power to make the world a smaller place.

But can you really find an embedded processor that can keep up with the shrinking size of reality? The answer, of course, is a resounding "yes." You need look no further than NEC's $V_R4100^{™}$ core technology as your solution of choice.

The V_R4100 Series frees you from the constraints of time and space with a unique, customized "system on a chip," allowing you to create a more streamlined design and send it off to find its place in the world in record speed. And for the more pressing deadline, the V_R4101 offers a built-in memory controller for up to 8Mb DRAM/16Mb MROM.

But is all that enough to change the world? When combined with full 64-bit performance (with a 32-bit interface) and more MIPS/mm² in a low-priced, small die size package, it is very possible indeed. Then, factor in low power consumption, on-board Multiply Accumulate Instruction, plus on-chip management features such as the ability to operate at 3.3 volts, and changing the world of PDAs is virtually assured.

In all, NEC offers significant performance, flexibility and serious power management features in a compact, affordable embedded processor. All of which offers you a chance to bring the world a little closer together.

For more information about the V_R4100 Series, call *1-800-366-9782.* Ask for Info Pack #165. It's a small step to take for something that could go so far.

NEC's V_R4100 processor is uniquely designed to power the next generation of PDA products.

NEC

READER SERVICE 103

Editorial

As electronics pervades our everyday lives, embedded microcontrollers are leading the way–they're all around us, in cellular telephones, in TV sets, in automobiles, in household appliances, in .... They're also pervading industry, where they are the front-line troops in the battle to automate processes and increase worker productivity. Of more importance in the context of Electronic Design and its readers, embedded microcontrollers are causing a profound change in the way electronic systems are being designed.

"Embedded microcontrollers are causing a profound change in the way electronic systems are being designed."

Today, almost every design engineer comes face-to-face, in one way or another, with a microcontroller as part of a system under design. The growth of embedded microcontrollers has had another profound effect: The growth of software as a key element in most new systems, as well as a key skill for EEs that is becoming much more in demand.

This Special Supplement, with its collection of recently published technical articles, focuses on both the hardware and the software elements of the design of embedded systems. Designing an embedded system requires many different types of components and tools, including microcontrollers and microprocessors, memory chips and memory management systems, software tools for development and debugging, real-time operating systems, and others. The articles in this Supplement cover a wide range of topics, from overviews of RISC CPUs and 16-bit microcontrollers authored by the Electronic Design Editorial Staff, to detailed technical explanations authored by experts in industry–on topics such as programming techniques for 32-bit RISC microprocessors, the value of system simulation as a design resource, and the use of flash PC Cards in embedded process control systems. Descriptions of several recently introduced products round out the article package presented here.

This year, to continue our deep focus on this critical subject area, Electronic Design has established a new section: Embedded Software and Hardware, which will appear monthly throughout 1997. The new section will be managed by Electronic Design's Embedded Systems/Software Editor, Tom Williams, who is based in our San Jose, Calif. office. The section will contain news and overview articles authored by Williams, technical articles, and a collection of the most important embedded systems-related products introduced each month.

This Special Supplement thus mirrors many elements of the Embedded Software and Hardware Sections. We hope you find it, as well as the upcoming Embedded Sections, useful references in the future.

STEPHEN E. SCRUPSKI
Editorial Director

And some people still think Microtec is just a tools company.

JIM READY
MICROTEC

Embedded Software Expert

68K, CPU32, x86, PowerPC,™ ColdFire,™ ARM and i960 Supporter

Confidant

RTOS Guru

Networking Expert

Software Debugger

Legacy Software Migrator

Butt-Saver

HW/SW Design Authority

Telecom Specialist

Partner

JavaOS™ Licensee

C++ Wizard

Windows® Developer

Internet Enabler

But we're so much more. Our complete line of products and services can differentiate your designs, add value, and get you to market faster. Products like Spectra®, Microtec's fully integrated—yet remarkably open—embedded development environment. It includes VRTX®, the industry's most powerful and proven RTOS, as well as our renowned XRAY® debugger and C/C++ compilers. Use them together, separately, or with your own choice of RTOS and tools.

Or let us do it for you. Microtec offers a full range of consulting and professional services, so we can provide everything from off-the-shelf point products to completely integrated, turnkey solutions.

So if you're ready to slice serious time and money off your next development project, call Microtec at **1-800-950-5554, Dept. 320**. Mention this ad, and we'll send you your own copy of the Embedded Software Expert badge.

Think of it as a little reminder from your total solutions company.

Microtec®
A Mentor Graphics Company

www.mri.com

See us at the Embedded Systems Conference, Booth #1004

READER SERVICE 104

The latest crop of 32-bit RISC processors integrate many functions needed to trim embedded system costs.

DAVE BURSKY
WEST COAST EXECUTIVE EDITOR

Tuned RISC Devices Deliver Top Performance

AS NEW SYSTEMS ARE BEING DESIGNED THAT REQUIRE MORE INTELLIGENCE, 8- AND 16-BIT CONTROLLERS ARE GIVING WAY TO 32-BIT RISC PROCESSORS. BUT THE RELATIVELY HIGH COST OF first- and second-generation RISC processors has delayed their implementation into cost-sensitive applications such as consumer products, video games, and automotive systems.

But as designers start looking at total system cost rather than the cost of the CPU chip alone, the latest crop of highly-integrated 32-bit processors makes the use of 32-bit architectures more cost-effective than ever. The last year or two has also seen increasing interest in the "system-on-a-chip" approach offered by various ASIC and CPU suppliers. In this approach, designers start with the basic CPU core and add standard and custom megacells. This allows them to craft a processing solution optimized for the system and gives them a level of uniqueness that would be hard for a competitor to duplicate.

Designers can select from among many highly-integrated commodity offerings, or from various cell-based building blocks, which in total, provide a wide range of integration options and cost/performance trade-offs to meet almost any system requirement. The choice of 32-bit embedded controllers includes holdovers from the CISC world, as well as many new challengers based on RISC architectures. CISC choices, in the form of various implementations of the Intel 386/486 architecture, as well as versions of Motorola's 68000 processor, deliver throughputs between 5 and 15 MIPS. The biggest plus of these families is that they can leverage the extensive programming knowledge built up over the years as well as volumes of code libraries and the variety of low-cost software development tools that run on PCs and Macintoshes. These processors will be covered in a future article.

In many cases even these CISC-based processors are rapidly giving way to the latest-generation RISC CPUs. These RISC CPUs have been optimized for high code efficiency, small chip areas, and much higher throughputs. And they're proliferating from two different roots. On one side, designers can select from familiar RISC processors that are extensions of families originally designed for desktop computers. On the other side, designers have available a wide choice of new RISC architectures designed from the ground up to tackle embedded applications.

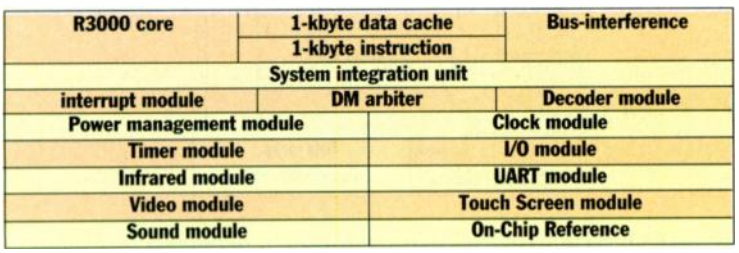

1. ALMOST ALL THE *functionality needed for a personal digital assistant is integrated into the PR30100 RISC processor developed by Philips. At the heart of the processor is a power-reduced version of the 32-bit R3000 core.*

The "traditional" RISC options include the R2000/3000/4000 architectures from the MIPS Technology Division of Silicon Graphics, the SPARC architecture from Sun Microsystems, and the Power PC from IBM and Motorola. And, perhaps not as widely recognized, the PA-RISC architecture, licensed by Hewlett-Packard to OKI Semiconductor and Winbond Electronics, has also been reshaped by those licensees to provide a cost-effective solution for embedded applications. Similarly, the Alpha RISC processor family from Digital Equipment includes one family member targeted for embedded applications.

ARMS FROM THE UK

To a lesser extent, at least here in the U.S., the ARM processor did have desktop roots in its country of origin, the United Kingdom, where Advanced RISC Machines crafted the CPU for the Apricot family of personal computers. But in other regions of the world, most designers view the latest versions as designed from the ground up for embedded applications.

The ARM is now widely licensed to more than half-a-dozen companies that are producing "standard" CPU chips. These chips are aimed at designers who want either an off-the-shelf solution, application-specific circuits that include an ARM core along with a set of features optimized for a particular application, or CPU cores with a uniquely-optimized set of features. The ARM partners include Asahi Kasei Microsystems, Cirrus Logic, Atmel/ES2, GEC Plessey, LG Semicon, NEC, OKI Semiconductor, Samsung, Sharp, Symbios Logic (formerly NCR), Texas Instruments, and VLSI Technology. This grouping of companies represents one of the broadest industry CPU partnerships, and makes the ARM the highest-volume RISC CPU (in all its variations) to date.

An early 32-bit embedded entry, the Am29000 family developed by Advanced Micro Devices, is one of the first to reach a sort of end-of-life status. AMD has indicated it will not develop any new versions of its 29000 family, but it will continue to manufacture and support existing family members and produce the recently-released enhanced versions of the 29200 series.

The other early RISC entry targeted at embedded applications from the start, the i960 family of RISC processors from Intel, has maintained a strong presence, with some significant inroads into applications such as network bridges and routers. The latest versions, the 80960HA/HD/HT and the 80960RP, have targeted high-performance systems. For instance, the HA/HD/HT versions have a superscalar architecture can be clocked at 1X, 2X, or 3X the bus clock speed, respectively, delivering up to 150 MIPS. The RP version is more application optimized and has been crafted to bring intelligence to I/O support in client/server environments.

The Thumb variation from ARM, the StrongARM from DEC, the SH series from Hitachi, the ColdFire family from Motorola, the V800 series from NEC, and the Compact RISC from National Semiconductor, are some of the latest ground-up RISC architectures that designers can select from. The list continues to grow with the release of a 32-bit processor by Mitsubishi Electronics America Inc. (See "Combo RISC CPU and DRAM solves data bandwidth issues," Mar. 4, 1996, p. 67). SGS-Thomson has just released a totally revamped version of its Transputer architecture called the ST20. Next month (April, 1996), Sun Microsystems will release the first Sun-created version of the MicroSPARC II architecture, the IIe, targeted at embedded control applications.

In addition to the traditional CPU suppliers, the embedded RISC field is drawing in new suppliers who feel they can offer some unique capabilities. For example, CSEM (the Centre Suisse d'Electronique et de Microtechnique), Neuchatel, Switzerland, has crafted a very-low-power scalable RISC core dubbed CoolRISC. The core can be scaled from an 8-bit data-word width up to a full 32-bit architecture, yet consumes just a few milliwatts of power. Another company, Patriot Scientific, has developed a novel 32-bit processor called ShBoom. Capable of running at 100 MHz, the 32-bit ShBoom CPU is actually a dual-processor architecture consisting of a RISC CPU based on a zero-operand dual-stack architecture, and an I/O processor that performs timing, time-synchronous data transfers, bit outputs, and DRAM control.

Raw throughput of these latest RISC processors ranges from about 20 MIPS on the low end to well over 100 MIPS for the fastest RISC processors. About a dozen different architectures are now competing in this cost-sensitive arena to form the heart of products such as communication controllers, printers, video games, and industrial

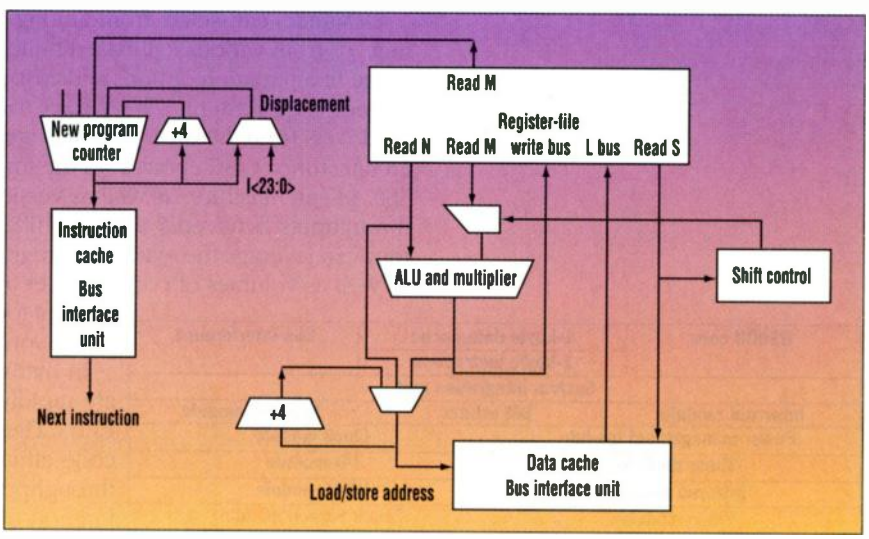

2. AT THE HEART OF THE STRONGARM PROCESSOR *developed by Digital Equipment is a high-performance Thumb-compatible core that can operate at clock rates as high as 200 MHz. The core's data path includes a simple five-stage pipeline, a five-port register file, and a high-performance multiplier that retires 12 bits every cycle.*

robots.

The surplus of choices makes the task of selecting the best processor for the job quite challenging, to say the least. Trade-off matrices can end up including many features, performance numbers, pricing data, software development tools, and so on. And when those matrices still don't come up with an answer, designers can "roll their own" CPU through the use of the available RISC cores and megacell libraries.

One high-interest area in which highly-integrated RISC chips look very attractive is in the hand-held personal digital assistant (PDA) marketplace, in which high throughput, low power and high integration must all come together. And as long as a good C compiler exists, it may not make much difference to the designer which CPU architecture (from among all qualifying candidates) is actually selected, since hardware compatibility and software compatibility are not as important in mostly-closed systems.

For that reason, the novel approach such as employed by Mitsubishi—integrating DRAM main memory onto the CPU chip—could provide a high-performance alternative to today's integrated CPUs. Unlike the high-integration processors now being offered by many suppliers, the M32R/D places lots of memory along with the CPU on a chip. Peripheral support functions are relegated to a separate peripheral chip, since most support functions don't require close coupling to the CPU, and can often run at clock speeds of less than 25 MHz. Memory, both DRAM main memory and SRAM caches, must be closely coupled to the CPU to let the processor deliver top performance. By integrating 16-Mbits of DRAM and 2-kbytes of SRAM cache along with a 32-bit RISC processor, wide, high-speed data buses can be used on the chip to provide high-bandwidth paths to move data and instructions.

Staying with a more traditional approach, the recent release of a MIPS R3000-based design of a PDA on a chip by Philips Semiconductors—the PR30100—provides designers with a low-cost system building block (less than $20 apiece in lots of 100,000 units) (ELECTRONIC DESIGN , Nov. 20, 1995, p. 55). To keep the chip's cost low, Philips' designers integrated only 1 kbyte each of data and instruction cache, and did not include hardware support for multiply-and-accumulate operations in the R3000 core (*Fig. 1*) . One area the designers didn't skimp on was the chip's 4-Gbyte address space. By keeping the address bus to a full 32 bits, the processor will be able to interface and control high-density CD and magnetic storage devices.

Rather than push for supercomputer-like speeds, designers at Philips optimized the chip for low power, with the CPU and bus interface operating at 18.432 MHz (a value derived from the 32-kHz low-cost watch-type crystal used for the PLL-based clock generator). Most instructions execute in a single clock cycle (except loads, stores, and branches), allowing the chip to deliver a throughput of about 15 MIPS.

Internal power-management logic helps keep the active power to just 150 mW when powered by a 3.3-V supply. Moreover, various operating modes allow power savings when all features aren't immediately needed—at its lowest in the "coma" mode, standby current can drop to as little as 30 μA.

Preceding the Philips chip, of course, was Apple Computer Inc., Cupertino, Calif., who developed the Newton PDA based on the ARM RISC processor. The Newtons, however, use a standard ARM processor that's complemented with several custom chips that perform the rest of the system functions. ARM itself, though, has developed several highly-integrated processors. The ARM7100, for example, includes a power-reduced version of the ARM710 RISC core (the 710a), a large 8-kbyte four-way set associative cache (large for an ARM-family processor), an LCD interface, a DRAM interface that controls four 256-Mbyte banks of memory, and various I/O functions. Standby power for the chip is very low—with everything shut down, the current drops to less than 10 μ A, while running at full speed the chip consumes just 66 mW.

Some of the I/O functions included on the chip consist of a DMA controller for high-speed data transfers, a UART-style serial port, an infrared serial port (IrDA-compatible), a codec interface with 16-byte FIFO buffers, and multiple 8-bit ports for general-purpose I/O support. The internal LCD controller delivers half-VGA screen images—640 by 240 pixels—and up to 16 levels of gray-scale resolution. Even with all these highly-integrated features, the small core area of the 710a allows the chip to sell for less than $25 apiece in large volumes. ARM also offers another high-integration processor, the AMR7500, which includes multimedia support features—telephony or CD-quality sound, a video output port capable of 120 MHz pixel data rates and direct drive of CRTs or LCDs, keyboard/mouse/joystick ports, and other I/O functions. This chip is based on the AMR710 core.

On the drawing boards at Advanced RISC Machines is the next-generation CPU, the ARM 8. Although not all features have been defined, this will be the first version to include some aspects of superscalar parallelism and a more heavily pipelined data path (five stages vs. three in previous chips) to further increase CPU performance. This processor—the ARM810—will deliver almost double the performance of the best ARM 7 (up to 80 MIPS), but at the same time, it requires double the chip area due to the more-complex logic. That will cause the ARM 8 to consume about twice the power.

STRENGTHENING THE ARM

As one of the ARM licensees, Digital Equipment has undertaken the challenge of reimplementing the ARM instruction-set architecture on its high-performance 0.35-μm triple-level-metal CMOS process. The result, described at the 1996 IEEE International Solid State Circuits Conference (ISSCC) in San Francisco, Calif., is a chip that delivers a throughput of over 200 MIPS when clocked at 200 MHz—close to a ten-fold performance improvement over the standard ARM processors. Though the advanced process may drive up processing cost, the resulting small chip will be very economical.

Dubbed StrongArm, the chip implements the recently-released Thumb instruction set (version 4 of the instruction-set architecture definition) (ELECTRONIC DESIGN , March 20, 1995, p.163). To get the high throughput, DEC's designers started with a

DREAM TEAM

EPF10K10	EPF10K20	EPF10K30	EPF10K40	EPF10K50	EPF10K70	EPF10K100
TYPICAL GATES	TYPICAL GATES	TYPICAL GATES	TYPICAL GATES	TYPICAL GATES	TYPICAL GATES	TYPICAL GATES
10,000	20,000	30,000	40,000	50,000	70,000	100,000
USABLE GATES	USABLE GATES	USABLE GATES	USABLE GATES	USABLE GATES	USABLE GATES	USABLE GATES
7,000-31,000	15,000-63,000	22,000-69,000	23,000-93,000	36,000-116,000	46,000-118,000	62,000-158,000

ALTERA FLEX 10K

THE DESIGNER'S DREAM TEAM WITH UP TO 100,000 GATES.

The FLEX 10K Dream Team of programmable logic devices is a championship roster from Altera. With densities from 10,000 to 100,000 gates and up to 15 times greater memory efficiency than FPGAs, FLEX 10K can take on even the most aggressive gate arrays.

FLEX 10K offers all the benefits of programmability and features such as memory, in-circuit reconfigurability, and built-in JTAG support. And, Altera's easy-to-use MAX+PLUS II development tools interface with all major CAE tools, giving you a team that will easily score against mid- to high-density gate arrays.

AN EMBEDDED ARRAY GAME PLAN.

The unique FLEX 10K architecture adds an embedded array to a logic array, so you can have up to 24K bits of RAM in a single chip, and megafunctions* such as microprocessors, microcontrollers, DSP and PCI functions, and others.

THE ALTERA ADVANTAGE.

You need every advantage to compete in this league. Value, competitive pricing, high-performance technology, and comprehensive technical support — that's The Altera Advantage.

It's time to put the Dream Team to work for you. Call Altera today for your free FLEX 10K information kit. 800-9-ALTERA (800-925-8372), Dept. A165. Or, find us at http://www.altera.com on the world-wide web.

The ALTERA Advantage

*Developed through AMPP (Altera Megafunction Partners Program) Note: Individual family members available over the next 12 months.

SIEMENS

Hard Wired.

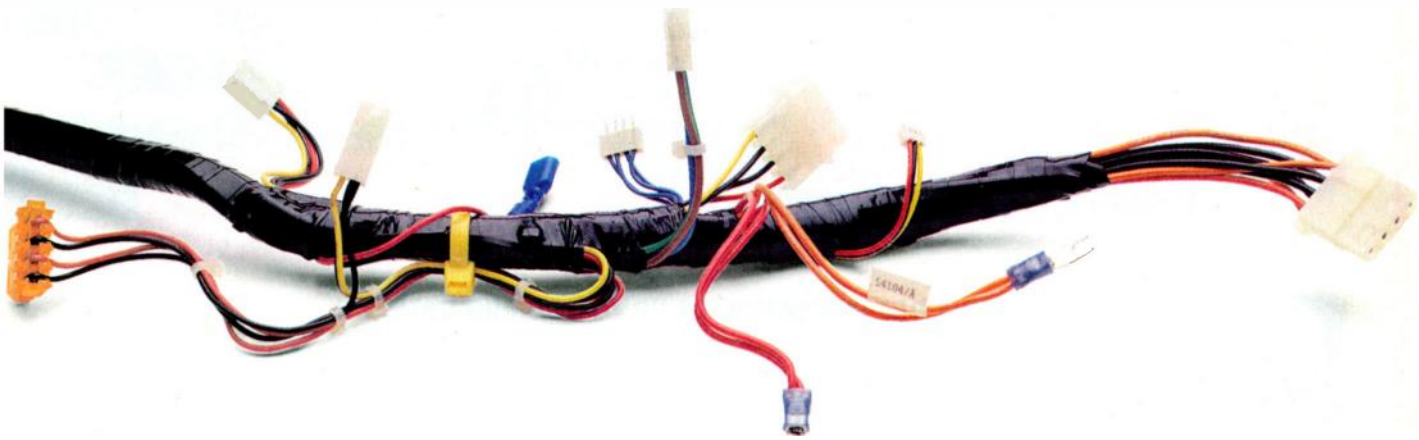

Simplify your life. Network your applications with Siemens microcontrollers and CAN 2.0B

Engineers around the world are discovering the power of Controller Area Network (CAN). And simplified wiring schemes are just the beginning. CAN also offers a host of other benefits, including 1 Mb/s speed over standard twisted pair, rock-solid data integrity in electrically noisy environments, even the ability to send data over AC wiring (great for factory automation). With so many advantages, it's no wonder CAN is rapidly gaining acceptance beyond the automotive market. And no other microcontroller manufacturer offers more CAN 2.0B solutions than Siemens.

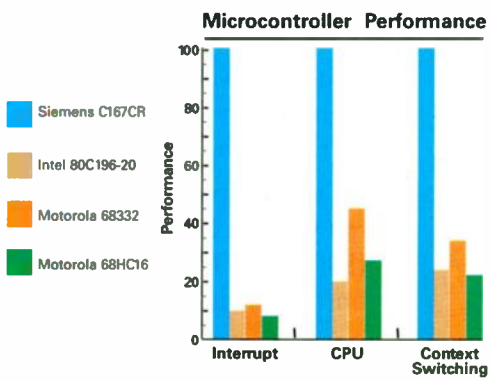

Easy Wired.

COMPLETE CAN CAPABILITY™

16-bit and 8-bit microcontrollers with CAN 2.0B. Plus, stand-alone solutions.

Siemens is your one-stop shop for CAN controllers. For 16-bit, we offer the SAB C167CR LM featuring the world's fastest 16-bit microcontroller architecture (see chart). Or choose the 8-bit SAB C515C LM. Both have CAN Version 2.0B, extending your system capacity to over 500 million different messages. We also have two stand-alone full CAN 2.0B passive products: the SAE 81C90 and the SAE 81C91 — both of which can work in CAN 2.0B active networks.

Volume deliveries today.

With fabs around the world, we have the capacity you need. You also get local support and great third-party development tools.

To see how you can use CAN, call for a 16-bit Evaluation Kit. Ask for Ext. 4, Lit. Pack #M14A050. Plus, visit the CAN section on our Web site.

1-800-77-SIEMENS
http://www.sci.siemens.com/CAN.html

Harvard-style architecture with large instruction and data caches (16 kbytes each) that have independent buses (*Fig.* 2). That replaces the unified cache employed in previous ARM family members and provides more parallelism for read and write operations.

For faster matching, the caches are 32-way set-associative. Although that is a much higher degree of associativity than most CPUs, it is only half that of the 64-way set associativity used in the ARM 610 CPU. Supporting the caches are 32-entry memory-management units and an 8-entry write buffer (16 bytes per entry). The write buffer that allows the CPU to write results and then go on to another task without stalling if the memory subsystem is busy.

The CPU core is a single-issue design with a "classic" 5-stage pipeline. It can perform single-cycle branches and conditional execution of every instruction. Additional resources in the core include an in-line barrel shifter for shift/add and multiply/add operations, and a 32-word by 32-bit register file.

A multiplier-accumulator unit on the chip can perform 32- and 64-bit multiplication in three to six cycles (4.5 ns per cycle), retiring 12 bits at a time. That results in a fairly robust DSP-like processing capability for 12-, 24-, or 32-bit data words. In addition, the MAC includes leading 0 detection to allow for early termination of the multi-cycle computations. The fast computations suit the chip well for DSP algorithms typically encountered in PDAs—data communications (modems), handwriting recognition, and speech recognition and output are some of the key functions the processor can support.

Although high-speed CMOS typically has a reputation for consuming a lot of power, the StrongARM processor was designed for low-power operation, with an idle mode allowing the chip to dissipate just 20 mW, and a sleep mode that drops power to less than 200 μW.

In the normal operating mode, the device consumes about 900 mW when clocked at 200 MHz and running from a 2-V supply. When the core is powered by a 1.65-V supply, it can still run at up to 160 MHz while the power drain drops to just below the 500-mW mark. At that speed and power, the chip delivers a MIPS/W ratio of 411.

Conditional clocking allows sections of the chip not being used during an operation to be turned off, thereby minimizing power. In addition, edge-triggered latches were used throughout most of the design to minimize gate loading on the clock lines. Power can also be further reduced by running the internal clocks off the slower bus clock during cache fills.

In addition to the highly-integrated ARM and StrongARM processors, many of the recent license agreements provide the various semiconductor partners with access to the ARM7TDMI Thumb processor core. One of the partners, Texas Instruments, has embedded the core into a TGC3000 series gate array that packs 100,000 gates in which users can integrate application-specific logic. The core that TI implemented will be part of the company's TMS 470 microcontroller series. The fully-static core delivers 36 MIPS when clocked at 40 MHz.

Another partner, VLSI Technology, not only offers the core CPUs as part of its megacell libraries, but has crafted several standard products that employ the processor core. Some of those products include the Ruby I and II communications controllers, and a two-chip GSM telephone. GEC Plessey has also used the ARM core at the heart of several communications controllers, while Sharp has developed a family of microcontrollers—the LH77790 series that includes on-chip LCD control.

Going after the PDA market with a PowerPC-based solution, Motorola has crafted a chip with a somewhat similar array of features to the ARM 7100 or Philips' PDA chip, but offers a much higher throughput than either of those CPUs—53 MIPS at 40 MHz. The MPC821 is the first result of Motorola's own use of the PowerPC core technology to create a highly-integrated processor targeted at portable computing/communications systems. When clocked at 25 MHz, the processor consumes less than 300 mW. Power consumption drops to less than 10 μW in the low-power stop mode.

In addition to the PowerPC core, the chip includes instruction and data caches (4-kbytes each), a second RISC-based (proprietary) controller that manages many of the I/O support functions, a dedicated multiplier-accumulator, and bus interface logic. A system interface block included on the chip provides the host-system interface, memory-control support for DRAM, ROM, and other memory

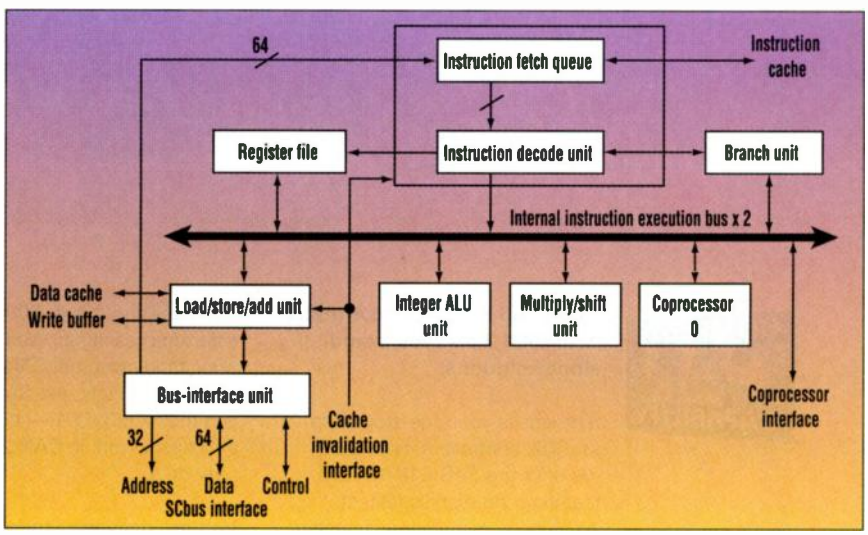

3. THE FIRST SUPERSCALAR PROCESSOR CORE *compatible with the R3000 and R4000 32-bit instruction sets employs an enhanced hardware architecture to provide a higher throughput than other MIPS-based cores. Developed by LSI Logic, the CW4010 core has 64-bit memory and cache interfaces and can issue and retire two instructions per cycle thanks to the use of five independent execution units..*

FOR EMBEDDED APPLICATIONS, THIS 486 COMES FULLY LOADED.

GO AHEAD—KICK THE TIRES.

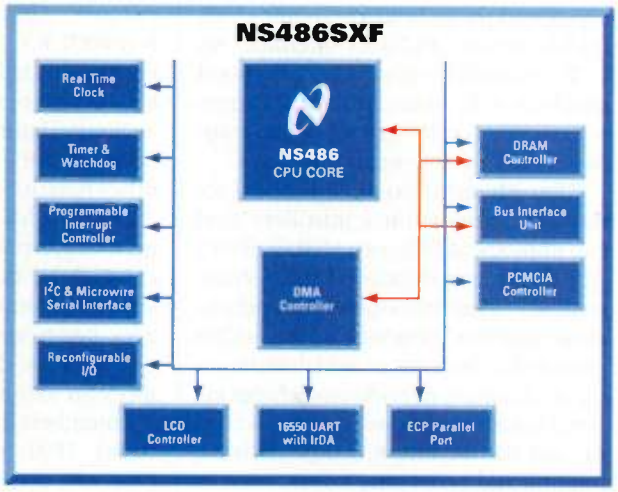

To reduce costs and time to market for your embedded 32-bit application, the **NS486SXF** comes standard with all the features you could want.

With the industry's most complete set of integrated peripherals and on-chip service elements, the NS486SXF is the first true embedded 486 system on a chip. By eliminating costly desktop features, we've created the only 486 CPU core that is optimized for embedded applications.

The familiarity of the 486 architecture means you can develop with confidence. And NS486SXF is supported by the best compilers in the industry, and by tools and kernels from leading real-time operating system vendors.

The NS486SXF Evaluation Kit includes everything you need to generate and debug code and run benchmarks, including evaluation copies of industry-leading development tools and kernels from several world-class software vendors.

ENTER TO WIN A FREE EVALUATION KIT.

VISIT US: At the Embedded Systems Show East.
March 10-12 at the Hynes Convention Center in Boston; Booth #810.

For more information contact us at:

INFO CARD: Mail or Fax

WEB: http://www.national.com./see/NS486

CALL: 1-800-272-9959 Ext. 757

National Semiconductor

types, and packs both a real-time clock and a two-slot PCMCIA interface. An LCD controller supports bit-mapped graphics, with monochrome (4/16 gray scale levels) or 16-color thin-film-transistor (TFT) active-matrix displays.

There are also two multiprotocol serial communication controllers that can implement Ethernet, HDLC/ SDLC, AppleTalk, serial infrared (IrDA-compatible), synchronous or asynchronous receiver/transmitters, and other protocols. Two more serial-management channels provide asynchronous serial communications and can be connected to the internal time-division multiplexed serial channels.

The same basic technology used in the MPC821 went into the creation of the MPC860, also known as the PowerQUICC. This chip is an enhanced version of the 68k core-based QUICC multi-channel data-communications controller (ELECTRONIC DESIGN , Sept. 18, 1995, p. 175).

The PowerPC controller core, as in the MPC821, delivers 53 MIPS when clocked at 40 MHz and controls the activities of up to four on-chip Ethernet channels and HDLC communications support. Supporting the core are 4 kbytes each of instruction and data cache. Also included on the PowerQUICC chip is a second RISC engine—a dedicated 32-bit controller for the system interface functions such as timers, interrupt controllers, virtual DMA channels, parallel I/O lines, and other functions.

Motorola has also crafted a family of general-purpose embedded PowerPC controllers, the MPC-500 series, but since their introduction about a year ago, has not released any new general-purpose chips. The second half of this year should see several new family members unveiled. On the other hand, IBM, the originator of the PowerPC architecture, has been busy developing new family members for embedded applications, both in the form of embedded cores and highly-integrated CPU chips.

PORTABLE MULTIMEDIA

IBM is also trying to attract designers of PDAs, set-top boxes, and other portable systems with a highly-integrated processor in its PowerPC 600 family, the 602. With a power consumption of 18 mW/MHz, and an upper clock limit of 66 MHz, the processor has a host of features that suit it well for PDAs, embedded multimedia support, video games, set-top boxes, and other systems.

Implemented with a 0.5-μm four-level metal CMOS process, the superscalar processor packs 1 million transistors into a small, 7-by-7-mm chip. Dual 4-kbyte caches connect to a high-speed 64-bit internal bus, and separate execution units—floating-point, integer, branch, and load-store—allow two instructions to be issued every clock cycle.

Furthermore, the on-chip floating-point unit can assist in multimedia applications such as audio editing, since its single-precision 32-bit computations can readily provide a wide dynamic range when editing up to 12-bit audio data. For higher-resolution 16-bit audio, the 72-dB dynamic range of the floating-point unit can be extended if software can handle the overflow conditions. Such support results in a 96-dB dynamic range (CD-quality) and only slightly slower throughput than when manipulating the 12-bit audio.

To minimize chip power consumption, the execution units are dynamically powered down. That allows the standby power to drop to less than 2 mW. When running at maximum speed, the typical operating power is about 1.2 W. In the envisioned appli-

Manufacturers Of RISC Chips

Advanced Micro Devices Inc.
Austin, Texas
(512) 602-6237

Advanced RISC Machines Ltd.
Los Gatos, Calif.
(408) 399-5195

Asahi Kasei Microsystems
San Jose, Calif.
(408) 436-8580

Atmel Corp.
San Jose, Calif.
(408) 436-4243

Centre Suisse d'Electronique et de Microtechnique SA
Neuchatel, Switzerland
(41) 38-205-111

Cirrus Logic Inc.
Fremont, Calif.
(408) 623-8300

Digital Equipment Corp.
Hudson, Mass.
(508) 568-5856

ES2 (European Silicon Structures)
Rousset, France
(33) 42 33 41 50

Fujitsu Microelectronics Inc.
San Jose, Calif.
(408) 922-9000

GEC Plessey
Scotts Valley, Calif.
(408) 451-4700

Hitachi America Ltd.
Brisbane, Calif.
(415) 589-8300

IBM Microelectronics Corp.
East Fishkill, N.Y.
(914) 894-2121

Integrated Device Technology Inc.
Santa Clara, Calif.
(408) 492-8623

Intel Corp.
Chandler, Ariz.
(602) 554-8080

LG Semicon Co. Ltd.
San Jose, Calif.
(408) 432-5000

LSI Logic Corp.
Milpitas, Calif.
(408) 433-8000

Mitsubishi Electronics America Corp.
Sunnyvale, Calif.
(408) 730-5900

Motorola Inc.
Austin, Texas
(512) 891-8704

NEC Electronics America Inc.
Mountain View, Calif.
(415) 960-6000

National Semiconductor Corp.
Santa Clara, Calif.
(408) 721-5000

OKI Semiconductor
Sunnyvale, Calif.
(408) 720-1900

Patriot Scientific Corp.
Poway, Calif.
(619) 679-4428

Philips Semiconductors
Sunnyvale, Calif.
(408) 991-2000

SGS-Thomson Microelectronics
Lincoln, Mass.
(617) 259-0300

Samsung Semiconductor Corp.
San Jose, Calif.
(408) 954-7008

Sharp Corp.
Camus, Wash.
(206) 834-8700

Siemens Corp.
Cupertino, Calif.
(408) 777-4500

Silicon Graphics Corp.
MIPS Technology Div.
Mountain View, Calif.
(415) 960-1980

Sun Microsystems Computer Co.
Mountain View, Calif.
(415) 960-1300

Symbios Logic
Fort Collins, Col.
(970) 226-9550

Teknema Inc.
Menlo Park, Calif.
(415) 833-7910

Texas Instruments Inc.
Stafford, Texas
(713) 274-3704

Toshiba America Electronic Components Inc.
Irvine, Calif.
(714) 455-2000

VLSI Technology Inc.
Tempe, Ariz.
(602) 752-8574

Winbond Electronics Corp.
Santa Clara, Calif.
(408) 982-0381

This list is only a guide; it is not a complete listing.

ANY MORE INTEGRATED AND THIS 8-BIT MICROCONTROLLER WOULD BE DOING YOUR WORK FOR YOU.

INTRODUCING THE COP8SA SINGLE CHIP SOLUTION.

Oh, sure, you'll still have to connect the chip to the ground and the Vcc. But after that, you're pretty much done.

That's because the new **COP8SA** mid-range OTP microcontroller requires zero external components. Zip. Which means the R/C oscillator, power-on-reset (POR), pull-up resistors, schmitt-triggers, and protection diodes are all built in. Which means your life is instantly easier. There's even patented EMI technology, a latchup capability that meets stringent industry standards, and an ESD rating that's over twice the industry norm.

The result? The quality and reliability of a single-chip solution in an industry-standard footprint. Add easy-to-use development tools and an Evaluation and Programming Unit for under $99, and your product will be in the market before you know it.

The COP8SA. You're almost done already.

FREE INFO KIT—FAST.

To test the COP8SA, contact your local distributor. While you're waiting, give National a call and we'll get you started.

CALL: 1-800-272-9959 Ext. 738.

WEB: http://www.national.com/see/cop8sa

In Europe, fax us at +49 (0) 180-5-12-12-15; in Japan, call 81-43-299-2308; in Southeast Asia, fax us at 852-2376-3901.

National Semiconductor

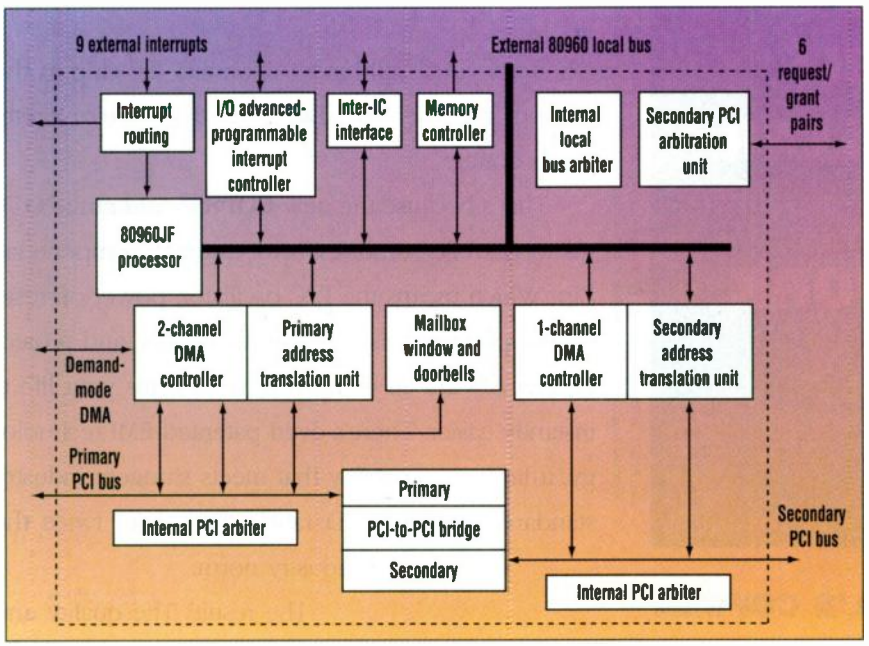

4. INTENDED TO SUPPORT HIGH-PERFORMANCE I/O *operations, the 80960RP from Intel is also the first 32-bit embedded RISC controller to include dual PCI interfaces. Targeted at server network and storage support, the 80960RP provides an intelligent bridge to multiple network controllers or to a RAID subsystem.*

cations, processor operation tends to be very bursty, thus average power will be much lower.

To tackle the low end of the embedded market, IBM designers have trimmed back on the complexity of the PowerPC 403 and 405 cores and created the 401, which occupies an area of just 5.5 mm². The core can execute all PowerPC code, includes hardware support for unaligned data accesses, and includes big- and little-endian support. Such a core will be able to deliver embedded solutions that can cost as little as $10 in large volumes.

Included in the basic core are robust debug capabilities--a critical feature as custom chips are integrated. Also embedded are both dynamic and static power-management facilities, a 32-word by 32-bit register file, and a dual-level interrupt structure. The reduced core leaves many features as options--instruction and data caches, coprocessors, memory management, interrupt control, and other functions. The power-management control allows the core to trim power consumption to about 50 mW when clocked at 25 MHz and powered by a 3.3-V supply. With a 2.4-V supply, the core's power drain drops to just 26 mW--some of the lowest power consumption levels for a 32-bit RISC core.

The growing availability of cores or licensable design files allows system designers and silicon suppliers to craft application-specific solutions. For instance, many core licensees are using the CPU cores to create commercial products that include significant intelligence and control capabilities.

Some prime examples of this trend can be seen with the various CPUs developed as part of Sun's SPARC family. Targeting the embedded control market based on the original SPARC I architecture, designers at Fujitsu Microelectronics created the SPARClite family, which includes over half-a-dozen CPUs, including the MB86934, 86933H, and 86936.

Both the 86934 and 936 are targeted at high-performance applications and pack floating-point coprocessors large on-chip caches, and Harvard-style architectures. The 86934 packs double the instruction cache of the 86936 (8 kbytes versus 4 kbytes), can access 247 different address spaces, with each containing up to 4 Gbytes, and delivering a sustained throughput of 55.5 MIPS when clocked at 60 MHz. The 86936 has a reduced address space and can access only 16 address spaces that each address up to 256 Mbytes. When clocked at 50 MHz, it delivers a peak throughput of 50 MIPS.

In some areas, the 86936 is a superset of the 86934, but in other areas, the 86934 is a superset of the 86936. For example, the 86936 device includes a video interface that can be used to drive printer/digital-copier engines/rasterizers, a pair of 24-bit timers, and a high-performance interrupt controller to list just a few of the distinguishing features. In contrast, the 86934 offers memory support for synchronous DRAMs and half-a-dozen on-chip FIFO buffers to help decouple the floating-point computations from the rest of the chip to achieve a floating-point throughput of 60 MFLOPS (peak).

Targeted at lower-throughput and lower-cost applications, the 86933H is a reduced-functionality CPU that can be clocked at just 20 MHz. The reduced clock rate limits the throughput to about 20 MIPS peak and 18 MIPS sustained. To trim the complexity and cost, the 86933H does not include a floating-point unit or extensive caches—only a 1-kbyte instruction cache is included on the chip.

In addition to extensions of original SPARC core, the MicroSPARC I design has been licensed by several companies that have released or will shortly release products based on the MicroSPARC I core. Some of those companies have crafted application-targeted chips that include the core. Companies that fall into this category include Matra-Harris, a division of Temic, Nantes, France, Hyundai Semiconductor, San Jose, Calif., and C-Cube Microsystems Corp., Milpitas, Calif.

Temic decided to employ the RISC core the same way as Motorola uses the PowerPC core in the PowerQUICC chip to implement a multichannel communications/network controller—the SPARCLET family (TSC701). Targeting multimedia applications, Hyundai employs the MicroSPARC core for use in products such as MPEG-2 decoders for set-top boxes, home theaters, and other applications. And in late 1995, Sun inked an agreement with C-Cube, which will

ONE BOOKMARK GETS YOU 40,000 TECHNICAL PAGES.

WWW.NATIONAL.COM

So you're trying to find the right part. You can hunt through databooks and app notes till you finally get somewhere. But now there's a fast, efficient way to get everything you need to work with: Datasheets. Application notes. Samples. Price/availability. **24 hours a day.**

Just bookmark **www.national.com,** and go straight to our web site. Where you're never more than 4 clicks from the exact information you need. And to give you powerful, customized access to over 14,000 products, we've built in a parametric search engine. It's fast, simple, and totally up to date.

- **PARAMETRIC SEARCH ENGINE**
- **4,000 DATASHEETS**
- **FREE SAMPLES**
- **1,500 APP NOTES**
- **FREE ANALOG DATA BOOKS**
- **SOFTWARE-BASED DESIGN TOOLS**
- **PRICING ON 8,000 PRODUCTS**

No wonder design engineers have called this one of the best sites in the industry. If you haven't tried it yet, we think you should pay us a visit immediately. It's the best way there is to get some work done.

FREE LINEAR SEMINAR HANDBOOK.

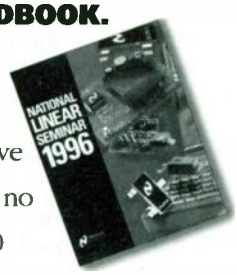

As an extra reason to bookmark this site, we're offering our 350-page **Linear Seminar Handbook**. To prove how simple this site is to use, there's no business reply card to fill out. No 800 number to call. Just go to **www.national.com/design**, and the book is free. Offer limited to first 10,000 visitors, so hurry.

National Semiconductor

Moving and shaping the future.™

WRH

license MicroSPARC core to be embedded in a forthcoming MPEG-2 encoder chip set and other products.

The interest in embedded versions of the SPARC has prompted Sun's SPARC Technology Business unit to create an embedded version of the second-generation MicroSPARC processor, the MicroSPARC IIe. Slated for release next month (April, 1996), the processor will include both integer and floating-point units, a DRAM controller, a memory-management unit, a ROM controller interface, and programmable "hooks" for the chip to tie into many industry-standard buses.

Focusing most of its efforts on variations of the MIPS R3000 core, LSI Logic, as a licensee of the MIPS architecture, has targeted several key application areas such as communications (network control, bridges, routers, ATM systems, etc.), set-top boxes, and industrial systems. Soon to be released is the next turn of the screw, the MiniRISC CW4010, the first superscalar core that will part of the design library.

Based on the MIPS II superscalar core (R4000 32-bit-mode compatible), the CW4010 offers much higher performance than the company's previous MIPS-compatible miniRISC CPU core, the CS4001. The basic core includes the arithmetic and logic unit, a system control coprocessor, a bus interface unit, a load-store unit, and an instruction scheduling unit (*Fig. 3*). To complement the core, designers can add blocks from the CoreWare design library—blocks such as direct-mapped or two-way set-associative instruction and data caches, an MMU, a hardware multiplier-accumulator, and writeback buffers.

The core issues two instructions per clock and can achieve a peak throughput of 160 MIPS (110 MIPS sustained), when clocked at 80 MHz and powered from a 3.3-V supply. The core requires little chip area—in fact, less than many of the ground-up RISC designs—just 3 mm^2. And the small area also translates into low power consumption—just 5 mW/MHz.

Another licensee of the MIPS architecture, Integrated Device Technology, has taken the MIPS architecture it licensed from MIPS and proliferated several families of embedded controllers—the 3050 and 3080 families—targeted at communications control, industrial applications, and computer peripherals. In addition to those families, the 64-bit R4600 CPU, better known as the Orion, has also been optimized for embedded applications, in several new versions released by IDT in the last quarter of 1995. The chips include some enhancements over the R4600, providing more of a system solution at lower cost points than the original R4600.

Toshiba is also offering the 64-bit Orion processor, but has not released any offshoots focusing on the embedded control arena. As an alternative, Toshiba is offering the 32-bit R3900 processor core as a megacell that designers can access through the company's ASIC design tools and standard-cell processing. When clocked at 50 MHz, the core consumes about 400 mW when powered by a 3.3-V supply. At that speed, the core can deliver about 52 MIPS. In addition, designers at Toshiba have crafted a slightly-higher-integration option—the TMPR3901-F—that takes the previous R3900 core and adds to it a 4-kbyte instruction cache, a 1-kbyte data cache, a write buffer, and some additional logic.

Also playing in the MIPS camp, NEC has several high-performance MIPS-compatible processors targeted at embedded applications such as set-top boxes, arcade games, network controllers, laser printers, and other systems. Able to clock at a top speed of 133 MHz, the V R4300 delivers 64-bit throughput of 170 Dhrystone MIPS. That high throughput is made possible thanks to a 16-kbyte instruction cache and an 8-kbyte data cache, a 32 double-entry translation look-aside buffer, a 4-word deep write buffer, and a five-stage pipelined arithmetic unit.

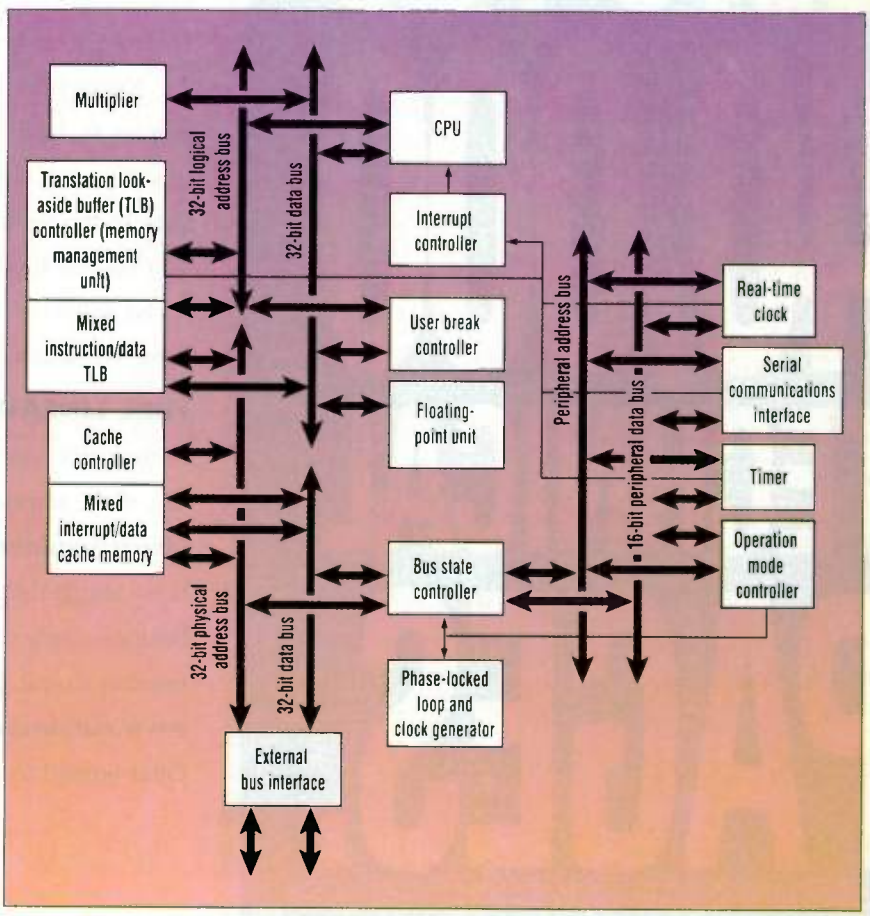

5. THE ADDITION OF A FLOATING-POINT *coprocessor to the SH-3 RISC processor allows designers at Hitachi to use the SH-3E in compute-intensive 3D graphics systems. Also included on the chip are an integer multiplier, serial communications interfaces, a real-time clock, and 2 or 8 kbytes of cache.*

Control Everything . . .

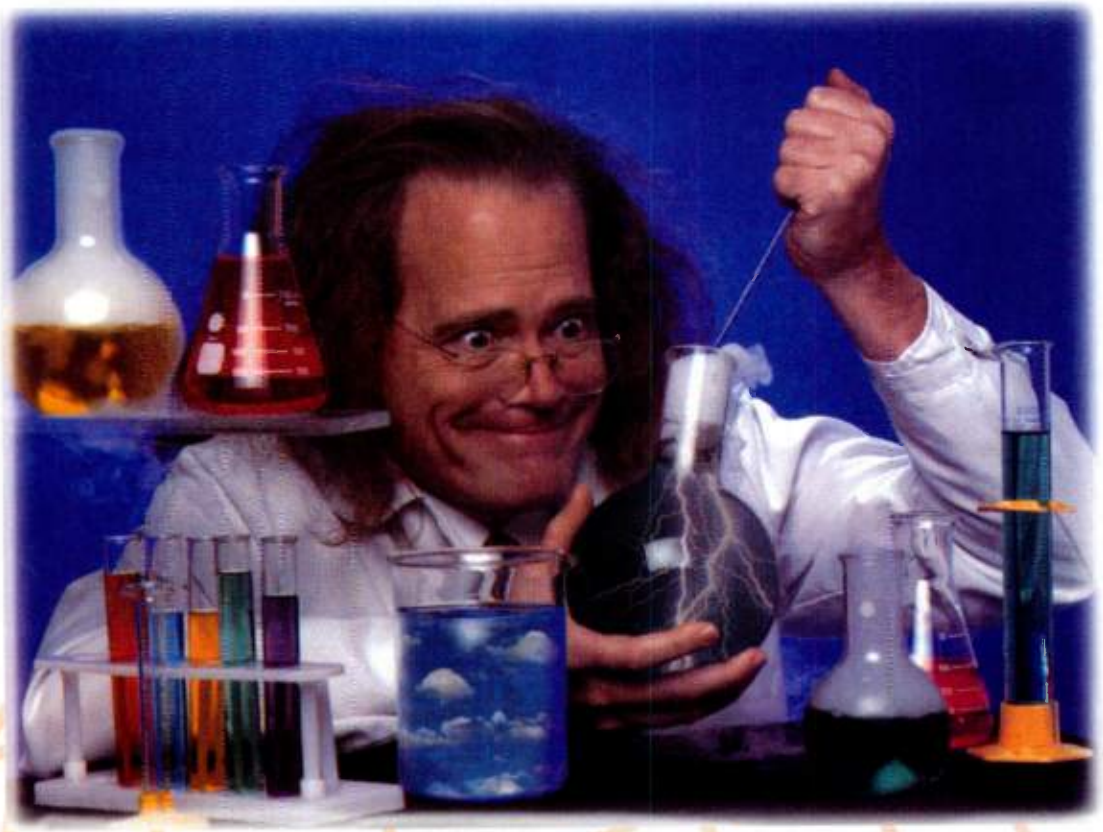

Well Almost.

If you think predicting the weather is tough, try controlling it! Z-World's versatile and powerful controllers can control most anything in your world—like the climate in your building. We control everything from laser surgical devices, to airframe riveting machines. And, we do it in 45 nations. Z-World's controllers interface easily with thousands of sensors, actuators and instruments. They can be networked using standard data communication protocols. Z-World's controllers are easy to program. Simply use your PC with our integrated Dynamic C™ software development system. If you need to get your product to market fast, we offer cost-effective, low-risk solutions for control applications. So, don't be left in the cold. **Call Z-World now.**

NEW! PK 2300 From $179

Quantity one price

The **PK 2300**. 19 I/O (11 are user-configurable). Protected inputs. High-current outputs.RS-232/RS-485 serial ports. Panel or DIN-rail mount. Resistance measurement input. Rugged ABS enclosure.

INNOVATION IN CONTROL TECHNOLOGY

Z-World, 1724 Picasso Avenue, Davis, CA 95616 USA

Telephone 916-757-3737 • FAX 916-753-5141 • To place an order call 1-888-EMBEDUS (USA)

For immediate information, use our 24 hour AutoFax 916-753-0618 or visit our Website http://www.zworld.com

READER SERVICE 100

For 32-bit processing needs, designers can also turn to NEC's V_R4100, which when powered by a 2.2-V supply can deliver a peak of 815 MIPS/W. The 4100 can also operate at 3.3 V. At that voltage level, it delivers a MIPS/W ratio of 375.

NEC's proprietary RISC family, the V800 series, consists of over half-a-dozen members that deliver throughputs ranging from about 10 MIPS on the low end with the V805 up to about 100 MIPS for the best unit currently available. Plans are already in place to up the performance still further, with the goals of 1000 MIPS and beyond. The V810 series delivers about 18 MIPS at 25 MHz, consumes about 500 mW at 5 V and 25 MHz, and has options for operation at lower supply voltages—2.2 V for the V805 and 3.3 V for the V810. At the lower supply-voltage level, the V805 delivers a throughput of about 11.5 MIPS and consumes 100 mW.

Additional family members increase the functional integration on the chip from the basic CPU—the V820 includes a 16-input interrupt controller, dual serial communication ports, a four-channel DMA controller, and a three-channel counter-timer unit. Containing many similar functions to the V820, the V821 adds a DRAM controller, but reduces the bus width to 16 bits, thus reducing the package pin count. Currently at the top of the performance curve is the V830, a 32-bit processor that delivers 118 MIPS when clocked at 100 MHz. At that clock frequency, the processor consumes about 500 mW when powered by a 3.3-V supply. Aimed at closed-loop control and signal-processing applications, the V830 also includes a high-speed multiplier-accumulator that can perform single-cycle (10-ns) multiplies.

Nestled between the 810 and 830 is the V850 series, which are more like single-chip MCUs and include on-chip RAM and ROM/PROM. The V851 and 852 also have 16-bit hardware multiplier-accumulators that can deliver a product in as little as 30 ns when clocked at 33 MHz. Overall throughput for the processors is about 38 MIPS for the 851 (at 33 MHz), and 29 MIPS for the V852 (25 MHz).

One of the longest-surviving 32-bit RISC families that has not lined up with an alternate supplier, the i960 from Intel provides designers with a wide range of CPU chips, from low-cost general-purpose devices that deliver 10 to 20 MIPS, to high-performance I/O controllers that have 40 to 150 MIPS of processing horsepower. One of the newest chips, the 80960RP, combines an 80960JF RISC core and is the only embedded controller to include a PCI-to-PCI bridge interface on the chip. Additional functions integrated on the chip include three chaining DMA controllers, a memory controller that supports DRAM, ROM and flash memories, many other features to support server networking and storage subsystems (*Fig. 4*).

Thanks to the on-chip PCI interface, I/O cards based on the controller can support multiple PCI peripherals on the board, while the board itself only occupies one PCI slot. That greatly reduces board complexity. Or, if added to a motherboard, it can offload the host CPU from control tasks and provide a secondary PCI bus that allows Ethernet communications cards or RAID storage subsystems to be connected.

Yet other proprietary RISC families—the SH series from Hitachi and the ColdFire family from Motorola—have taken cues from the other RISC families and offer unique solutions to both system integration and time-to-market pressures. Both families were developed with high code density in mind. Hitachi accomplishes that on the SH series by using a fixed, 16-bit instruction, while Motorola designers opted for a variable-length instruction approach to pack the code and minimize off-chip memory needs.

Hitachi has already released three generations of SH processors, the SH-

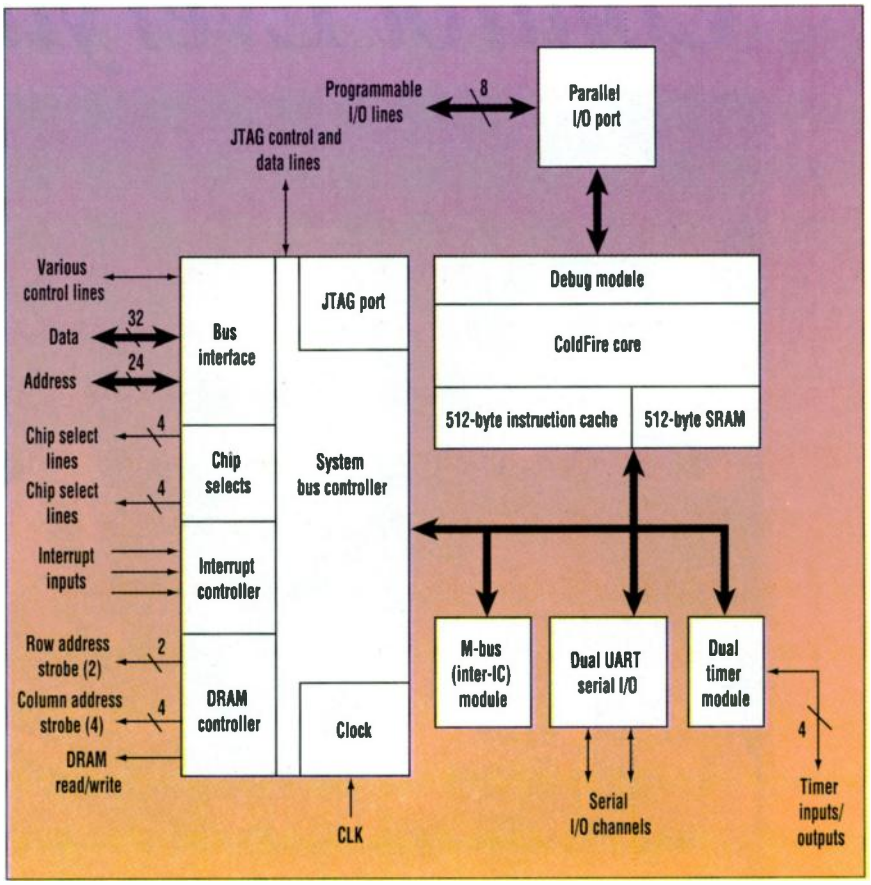

6. VARIABLE-LENGTH INSTRUCTIONS *are used by the ColdFire processor core to minimize the amount of off-chip storage needed for the firmware. The first integrated version of the ColdFire-based processor developed by Motorola includes a 512-byte instruction cache, a 512-byte SRAM, dual asynchronous serial ports, an inter-IC port, a pair of timers, a DRAM controller, an interrupt controller, and eight parallel I/O lines.*

EN GARDE!

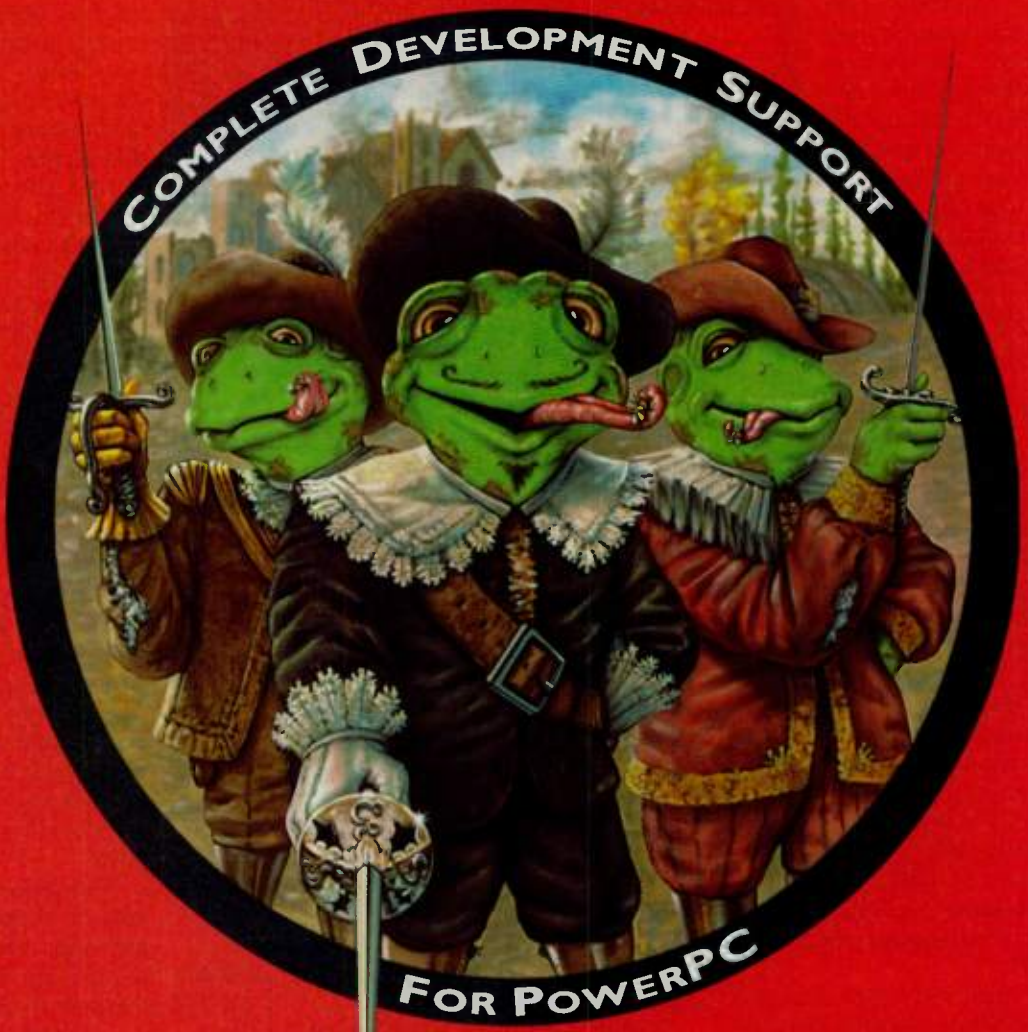

INTRODUCING A POWERFUL TEAM THAT CAN HELP YOU SLASH MPC 860 DEVELOPMENT TIME AND COSTS.

Applied Microsystems Corporation products and services cover the full development cycle for PowerPC. Applied expertise helps you pull it all together with seminars, training, consulting and support.

Design: Speed up system partitioning with *Eaglei*™ and *VSP-TAP*.

Debugging: Accelerate software and hardware development with *SuperTAP*,™ *CodeTAP*,® and *NetROM*.™

Testing: Improve software quality and performance with *CodeTEST*.™

READER SERVICE 82

WRH

CALL NOW 1-800-426-3925
for **FREE** White Papers on designing with the MPC 860, and on choosing design, debug and test tools. Or visit our home page at www.amc.com.

Applied Microsystems Corporation

European Distributor: Applied Microsystems UK and Europe. Tel: +44 (0)1296-625462. Fax: +44 (0)1296-623460. Germany Tel: +49 (0)89-427-4030. Fax: +49 (0)89-427-40333. France Tel: +33 (0)1644-63000. Fax: +33 (0)1644-60760. PowerPC is a registered trademark of International Business Machines Corp and is used under license therefrom. Designated trademarks and brands are the properties of their respective owners.

1, -2 and -3, and within each generation there can be multiple iterations. The latest release, the SH-3, includes two versions, the SH-DSP and the SH-3E. These two chips share the same base architecture with a CPU that can run at 100 MHz and deliver a throughput of 100 MIPS when powered by a 3.3-V supply. The base SH-3 processor comes in two variations—one version packs an on-chip unified cache of 8 kbytes (the SH7708), while the other includes 2 kbytes of unified cache (SH7702). The cache can be directly addressed and thus can also serve as general-purpose RAM if needed.

At 100 MHz, the processors consume about 1 W. By backing off on the clock to 60 MHz, the MIPS rating and power consumption are proportionally reduced. A further reduction in the power-supply level to 2.5 V and slowing the clock to 30 MHz trims the throughput to 30 MIPS, but more significantly, reduces the power consumption to just 180 mW.

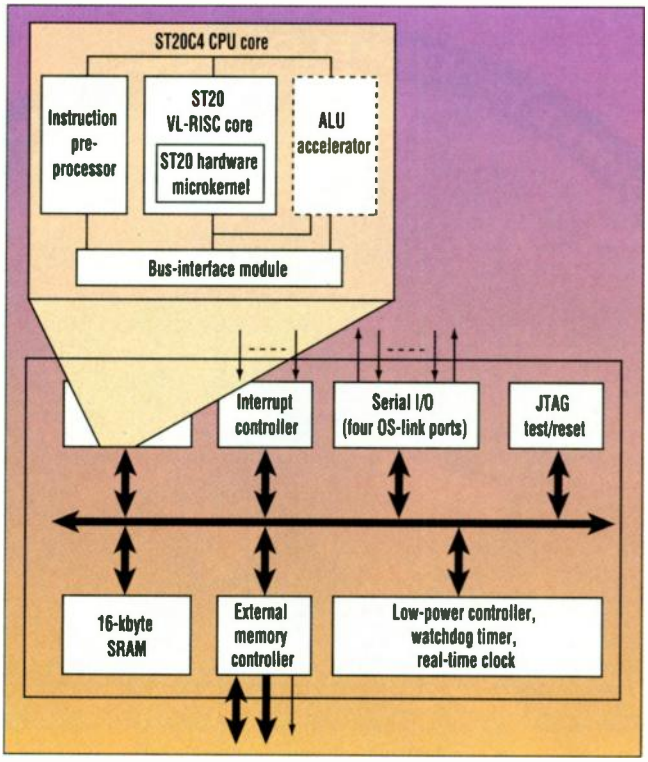

7. BEARING LITTLE *resemblance to the transputer, the ST20-based RISC controller developed by SGS-Thomson contains a completely revamped data path and includes a 16-kbyte SRAM for cache and data storage. Four 20-Mbit/s serial OS-link ports provide high-speed communications to the outside world.*

A very small CPU core is at the heart of the SH processors--the core area is just 2.1 mm^2 when fabricated with 0.35-μm design rules, thus keeping the total chip area to a minimum. On-chip peripheral support includes a real-time clock, three 32-bit timers, a serial asynchronous, full-duplex communications port, an 8-bit general-purpose I/O port, an interrupt controller, and a bus controller that can access seven physical address spaces of 64 Mbytes each, and can interface to DRAMs, SDRAMs, PSRAMs, ROM, and SRAMs. The bus controller also supports PCMCIA interfaces and includes a bus switch for big-endian/little-endian data formats.

The SH-3E processor has been optimized for consumer 3D graphics support and is the first SH family member to include single-precision floating-point math capability on the chip (*Fig. 5*). It is basically an SH-3 with an integrated FPU. Calculations are done in a single cycle and the floating-point unit has a two-cycle latency.

Targeting DSP applications, the SH-DSP includes dual 4-kbyte RAM blocks for parameter storage and a program ROM to store constants or algorithms. In addition to the SH integer processor, the chip includes a DSP unit that packs a high-performance multiplier-accumulator and other support logic. A three-address-bus architecture allows the integer unit and the DSP unit to simultaneously access two operands and one instruction every cycle, thus permitting sustained single-cycle operations for on-chip memory accesses.

Currently in definition, the SH-4 generation will take advantage of superscalar-architectural improvements to achieve throughputs of 300 MIPS when clocked at 200 MHz and powered by a 2.5-V supply. The processor will initially be targeted at multimedia and graphics applications and will be implemented using the company's 0.35-μm CMOS process.

Targeting many of the same applications, from PDAs to communication systems, the ColdFire family from Motorola provides a scalable architecture based on a new RISC core. There are currently three versions of ColdFire available. The MCF5202 and 5203 are just CPUs with a 2-kbyte unified 4-way set-associative cache, a debug module, a bus controller, and a JTAG test port. The MCF5206 adds many peripheral support functions to the core to form a highly-integrated solution (*Fig. 6*).

Through the use of the variable-length instruction-set architecture, very dense code can be developed, thus reducing external memory requirements, allowing the use of slower and less expensive memories, possibly speeding up system throughput. When clocked at 33 MHz, the 5202/5203 can deliver a throughput of 25 MIPS.

The processor core consists of a simple arithmetic and logic unit with 16 user-visible 32-bit-wide registers. The MCF5202 supports dynamic bus sizing for 8-, 16-, or 32-bit data interfaces, while the 5203 is limited to 8- and 16-bit data widths since it uses a reduced bus interface that is only 16-bits wide.

The more feature-rich MCF5206 adds a DRAM controller, timers, parallel and serial interfaces, and all the benefits of the high level of integration. Like the 5202, the 5206 includes dynamic bus sizing for 8-, 16-, or 32-bit data widths and provides a glueless interface to DRAMs, SRAMs, ROMs, and I/O devices.

When clocked at 33 MHz, this version delivers a throughput of about 17 MIPS. The lower throughput could, in part, be due to the reduced cache size—this processor only contains a 512-byte direct-mapped instruction cache and a 512-byte SRAM.

Internally, the ColdFire core consists of two independent, decoupled pipeline structures—an instruction-fetch pipeline (IFP) and the operand execution pipeline (OEP). The IFP is

Don't Limit Your Product's Capabilities—Use Web Technology.

Visit us at http://smallest.pharlap.com to see how you can develop embedded applications that take full advantage of the web.

Making The Most Out Of Your Products

Phar Lap's TNT Embedded ToolSuite,® Realtime Edition, complete with Realtime ETS™ Kernel, now comes with our robust networking protocol—ETS TCP/IP! This cutting-edge feature allows your customers' mainframes, workstations, or PCs to communicate with products on the factory floor, in the lab, or at remote sites—all using Web technology!

As a result, electronic OEMs can use this technology to create intelligent machines and instruments for an unlimited number of applications such as medical instruments, robotics, avionics equipment etc.

In addition, we support a variety of network protocols as well as industry standard tools such as Visual C++, Borland C++, CodeView, and Turbo Debugger.

Never Underestimate The Power Of The Web

Because Phar Lap's TNT Embedded ToolSuite comes with development tools, a Realtime ETS Kernel and ETS TCP/IP, it is a one-stop shopping product with enormous power!

To find out more, catch us on the "World's Smallest Web Server" today at URL http://smallest.pharlap.com or call a Phar Lap sales representative. Let your products realize the power of Web technology!

The 32-Bit x86 Experts

PHAR LAP

Embedded Development — Simply on Target™

Phar Lap Software, Inc. 60 Aberdeen Avenue, Cambridge, MA 02138 • Tel: (617) 661-1510 • Fax: (617) 876-2972 • http://www.pharlap.com

READER SERVICE 93

WRH

a two-stage pipeline for prefetching instructions. Prefetched instructions are gated into the two-stage OEP, which decodes the instruction, fetches the required operands, and then executes the required function. The IFP and OEP are decoupled by an instruction buffer that serves as a FIFO queue. Therefore, the IFP can prefetch instructions in advance of their actual use by the OEP. That minimizes the time the CPU would be stalled waiting for instructions.

The DRAM controller on the chip supports up to 128 Mbytes of DRAM operating with either page-mode or extended-data-out interfaces. The serial interfaces include both a full-duplex dual UART and a separate Inter-IC-compatible Motorola bus (M-bus) interface.

For system debugging, all ColdFire processors include a debug interface that permits background mode debugging and real-time tracing.

Several other companies have developed proprietary 32-bit RISC processor cores that they plan to license. One such core is the CoolRISC processor developed by CSEM.

Designed to execute one instruction per clock cycle, the processor was optimized for low-power operation and can also be scaled for 8-, 16-, 24-, or 32-bit word sizes. Based on a three-stage pipeline, the core includes an instruction set of 20 to 30 generic operations such as branch, call, return, load, store, and many others. With all the variations, the assembly-level command set contains a total of about 150 instructions.

Very efficient instructions give the CoolRISC much of its throughput by allowing various operations to take place in parallel. For example, two operands can be simultaneously addressed, whether they are stored in two registers, or one in a register and the other in the local RAM. Call and Return instructions can be executed with a hardware stack that contains as many program counters as desired (as determined by the ASIC designer). To avoid the growth of the chip's area in the event that many counters are needed, another software call operation can be used with a Branch and Link mechanism to form a stack in the integrated RAM.

The processor core can be compiled with the design software tools for implementation in cell libraries with minimum features of 2, 1, or 0.7 μm, and in the 1-μm process can achieve a throughput of about 20 MIPS when clocked at 20 MHz. The circuit is also very power efficient—an 8-bit version that employs an eight-word register file consumes just 0.3 mW/MHz when powered by a 3-V supply. Therefore, a 32-bit processor core might consume three to four times as much power, about 1 to 1.2 mW/MHz, depending on the additional logic added to the core.

CONTROL IN REAL TIME

Additional RISC cores such as the ST20 from SGS-Thomson provide designers with a processor capable of delivering 40 MIPS with a 50-MHz clock. In addition to the basic processor functionality, the core has been designed with modularity in mind to achieve various levels of throughput (*Fig. 7*). Closely coupled to the RISC core is a hardware microkernel for real-time operations. The kernel provides multipriority process scheduling, trap/exception handling, I/O, DMA, interrupt, and timer support, as well as fast context switching (500 ns at 50 MHz).

Surrounding the ST20 core and microkernel to form the ST20-C4 core are an instruction preprocessor and an ALU accelerator. The accelerator includes an integer multiplier (up to 32-by-32-bit multiplications in three cycles), a single-cycle barrel shifter, and single-cycle adder.

Another version of the core, the ST20C2, drops the ALU accelerator and merges some of the arithmetic functions into the ST20 core, thus reducing overall core area. SGS designers added a small register cache that makes up for some of the throughput lost by eliminating the accelerator. The register cache allows program variables to be cached, thus speeding up access to the variables as the executing software requires them.

The instruction preprocessor allows the core to use variable-length instructions that range in size from 8 to 32 bits. The variable-length commands are assembled by the preprocessor from the basic 8-bit commands, allowing programmers to minimize the amount of memory required to hold the application code, thereby reducing overall system cost.

Based on a VHDL model, the core and a companion VHDL-based macrocell library allow designers to configure an integrated processor very easily. An internal bus on the ST20—the OMI324 bus—allows simple interconnections of the blocks. The bus has been adopted by SGS-Thomson, Siemens, ARM, Matra, and Philips as a "standard" on-chip interconnect bus. The OMI324 bus permits high-speed communications between the core and peripheral blocks, with a latency of just two machine cycles and a 200-Mbyte/s bandwidth for access to on-chip and off-chip memory.

Testability has also been integrated into the design process. The core includes a test access port that supports the IEEE 1149.1 JTAG test standard for boundary scan testing. The boundary scan capability can be used for board-level testing as well allowing the parallel scan-path testing of each block within the chip.

For even lower-cost applications, the company has stripped the core even more by eliminating many of the instructions, thus simplifying the logic and further reducing core area. The ST20-C1 is targeted at applications such as smart cards and other deeply embedded systems.

Testability has also been integrated into the design process. The core includes a test access port that supports the IEEE 1149.1 JTAG test standard for boundary scan testing. The boundary scan capability can be used for board-level testing as well allowing the parallel scan-path testing of each block within the chip.

Based on the ST20 core, the ST20450 processor is a more fleshed out CPU chip that designers can purchase. In addition to the ST20 core, the chip includes 16 kbytes of SRAM, a programmable memory interface, four high-speed serial communication links and the hardware microkernel. The integral microkernel reduces application development time, reduces memory requirements, and can eliminate the royalties typically paid for software kernels.

Originally published in the March 18, 1996 Electronic Design.

SOME OF THE BEST NAMES IN THE BUSINESS WILL BE THERE . . .

EXHIBITORS: AMD-Logic Products Division • Advanced RISC Machines • AER Energy Resources • Analog Devices • Annabooks • Annasoft • AVX • Battery Engineering • Battery Technologies • Benchmarq Microelectronics • Berg Electronics • Bourns • Chips and Technologies • Crystal Semiconductor • CSEM SA • Dallas Semiconductor • Duel Systems • EE Product News • Electronic Design • Electronic Design China • Energizer Power Systems • Fujitsu Microelectronics • Hadco • HEI • Hewlett Packard, Optical Comm. Div. • Jbro Batteries • LCO Tech Rep • Linear Technology • Linfinity Microelectronics • Lucent Technologies • Megatel Computer • Micrel Semiconductor • MicroModule Systems • Microwaves & RF • Miniature Card Implementers Forum • Motorola, HPESD Div. • National Semiconductor • NEC Electronics • Nexcom Technology • Opti • PCMCIA • Phihong USA • Rayovac • S-MOS Systems • SAFT America • Samsung Semiconductor • Sanyo Energy • Sharp Electronics • Temic Semiconductors • Texas Instruments • TriTech Microelectronics • Unitrode Integrated Circuits • USAR Systems • VESA • Wireless Systems Design • WSI • **CONFERENCE:** 1-800-Batteries • Aavid Engineering • ACTiSYS • Advanced Micro Devices • Advanced Risc Machines • AER Energy Resources • Ampro Computers • Anadigics • Analog Circuit Design • Analog Devices • Battery Engineering • Battery Technologies • Benchmarq Microelectronics • Boulder Technologies • B-Tree Systems • California Micro Devices • Counterpoint Systems Foundry • CPS • Crystal Semiconductor • Duracell • Elantec • Energizer Power Systems • Exar • Fluid Dynamics • Fujitsu Microelectronics • Genoa Technology • Gore • Gould Electronics • Harris Semiconductor • Hewlett-Packard • IBM • Intel • Intellon • JKL Components • Kimmel Gerke Associates • Linear Technology • Lucent Technologies • Maxim Integrated Products • Microchip Technology • Motorola • M-Systems • National Semiconductor • NEC Electronics • Opti TriTech • Polytechnic University of Catalunya • Powerdex • Questra Consulting • Rayovac • S3 • Samsung Semiconductors • Sandia National Labs • Seagate Technology • Sensory • SiRF • Symbol Technologies • SystemSoft • Temic • Texas Instruments • Thermacore • Toshiba America Electronic Components • Unitrode Integrated Circuits • USAR Systems • Vadem • Zilog •

ALSO: Jack Kilby • Bob Pease • Tom Beaver • Philip Wennblom • Robin Saxby • Vaughn Watts, and more!

. . . WILL YOU?

Keynote Luncheon sponsored by BENCHMARQ MICROELECTRONICS

Industry Reception sponsored by intel

For more information on the Conference and Exhibition, call 201/393-6075; Email: portable@class.org

Portable by Design

CONFERENCE & EXHIBITION

MARCH 24-27, 1997

SANTA CLARA CONVENTION CENTER • SANTA CLARA, CA

FREE.

NEW.

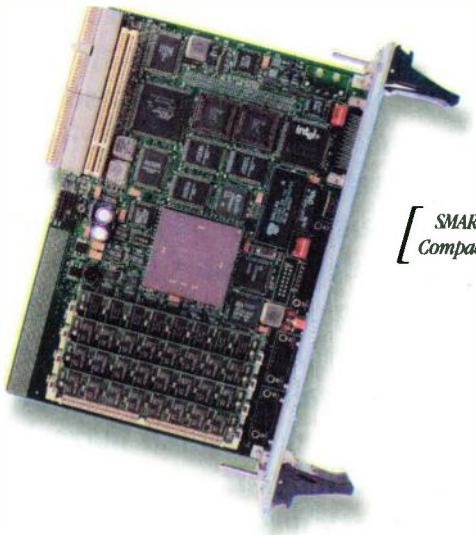

[SMARTengine/50PCI CompactPCI Computer]

NOW. 510-624-8221

The *first* R5000 CompactPCI computer.

It's fast, high-bandwidth, cool-running, and versatile. And it's available now.

Powered by a 200 MHz 64-bit MIPS R5000 processor, the SMARTengine/50PCI is designed to provide top performance in embedded applications such as internetworking, telecommunications and imaging. The CPU, memory and logic all run at 3.3 volts to minimize power consumption. And the dual-issue superscalar R5000 CPU provides programs with a generous on-chip 32 KB of instruction and 32 KB of data caches for blazing performance.

Fax 510.623.1434
http://www.smartm.com

A full complement of subsystems—including capacity for up to 128 MB of RAM, serial interfaces, watchdog timer, nonvolatile memory, ethernet, and SCSI—completes the high-performance picture. Software support for third-party real-time operating systems is comprehensive and complete.

The SMARTengine/50PCI is the *first* R5000 CompactPCI computer on the market—not surprising when you consider it comes from SMART Modular Technologies, Inc., a technology leader in modular memory, input/output and computing products for OEMs. With state-of-the-art manufacturing facilities in Fremont, California, Puerto Rico, and Scotland, SMART is uniquely qualified to be your long-term embedded computer partner.

FREE EVALUATION.

Now SMART offers qualified companies the chance to evaluate the SMARTengine/50PCI for a limited time at no cost.* In parallel, evaluate our ability to tailor a system solution to meet your specific needs, our software support, and our ability to function as an integral extension of your development team.

Contact us—now—since this offer can be extended for a limited time only.

SMART®
Modular Technologies
Embedded Systems – The SMART Way

See us at the Embedded Systems Conference, March 10 through 12, 1997, Boston, Massachusetts, Booth 2005.

EMBEDDED SYSTEMS CONFERENCE

CompactPCI™

READER SERVICE 98

Higher levels of integration and throughput in the newest 16-bit MCUs help trim system cost.

DAVE BURSKY
West Coast Executive Editor

16-Bit Embedded Controllers Open Up New Markets

IT'S GETTING HARDER TO DEFINE WHAT A 16-BIT MCU SHOULD CONSIST OF. IS IT A CHIP WITH A 16-BIT DATA PATH AND 16-BIT INTERFACE TO THE OUTSIDE WORLD? OR A DEVICE THAT USES 16-bit instructions and employs an 8- or 32-bit arithmetic and logic unit (ALU)/data path? Or a circuit that has an 8-bit external data path and includes 16-bit internal data paths? Or a processor with a 16-bit data path that employs a combination of variable-length instructions? Well, the latest crop of 16-bit embedded controllers provides designers with mid-range processors that deliver better performance than the popular high-end 8-bit microcontrollers (MCUs), but without the high cost typically associated with the newer full 32-bit controllers. These 16-bit controllers now achieve 20-MIPS throughputs, and provide close to single-chip logic solutions for many system designs, thanks to higher levels of integration that place more of the system functions on the CPU chip, along with advanced processes that allow the chips to clock faster and add new memory options.

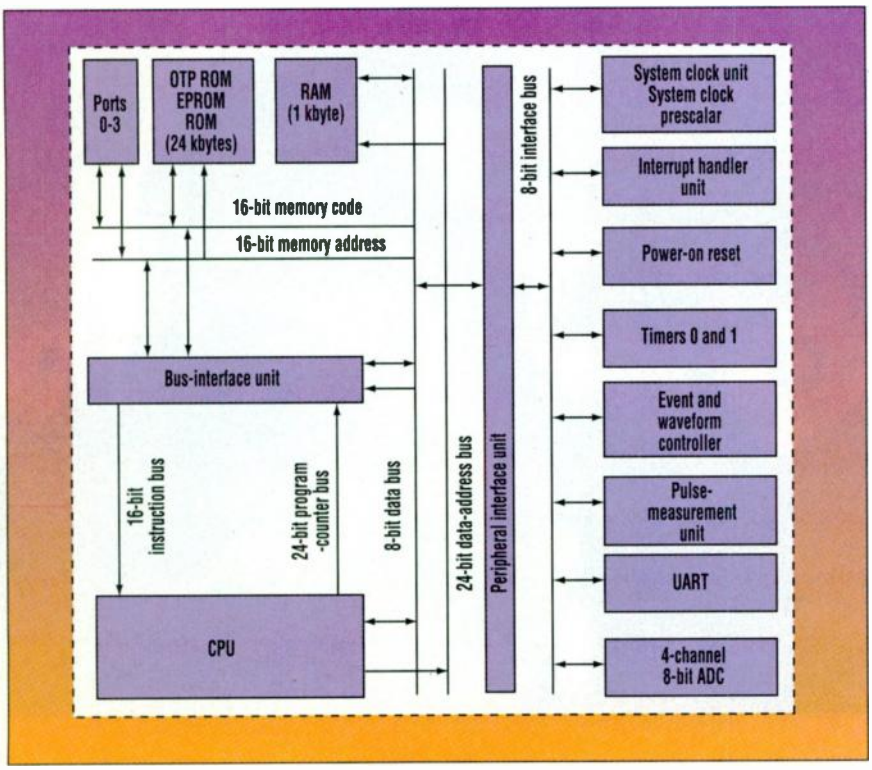

1. BY ADDING A MULTICHANNEL ADC AND *more memory to Intel's 80251 extension to the 8051 architecture, designers at Temic provide a real-world interface and larger program storage.*

For this report, we will define a 16-bit processor as a chip with a CPU core that includes a full 16-bit internal data path that can use any length instruction (one to four bytes). It may also include an 8- or 16-bit data-bus interface.

Variable-length instructions usually result in very compact code, but at the same time, some throughput inefficiencies occur because most memory accesses begin on a word boundary. For commands with an uneven number of bytes, a filler or blank byte must typically be added. This may cause a second memory fetch cycle, which

16-BIT CONTROLLERS

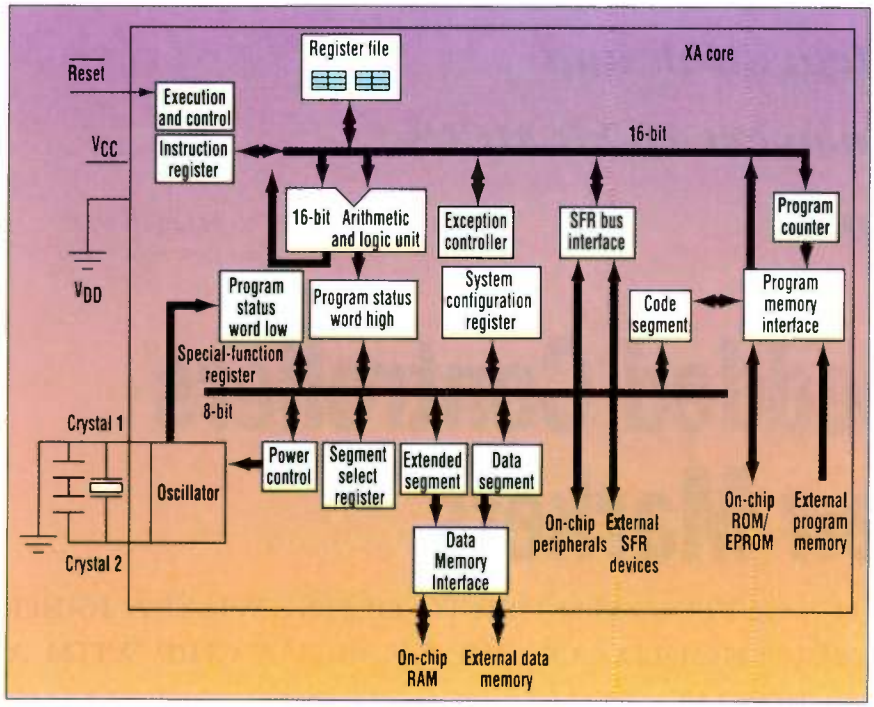

2. AN ALTERNATIVE TO the 80251, the 80C51XA *family developed by Philips Semiconductors contains a full 16-bit data path. Hardware-assisted multiplication in the ALU allows a 16-by-16-bit multiplication to be done in 12 clocks.*

would slow program execution, or increase the amount of memory needed to hold the program code. Some of the latest 16-bit MCUs overcome this fetch slowdown, but it is a key area to evaluate for software efficiency.

Most 16-bit processors that use variable-length instructions have roots in the 8-bit world, and are enhanced versions of older 8-bit architectures that have not reoptimized the instruction set. Newer processors avoid this problem by taking cues from the RISC world using single-word, 16-bit instructions for easier software compilation.

Controllers with 32-bit data paths and 16-bit host interfaces or 16-bit instruction sets can compete cost-wise with 16-bit MCUs. Companies like Advanced RISC Machines, Hitachi, Mitsubishi, Motorola, and Sharp have crafted their 32-bit CPUs to do just that. But according to designers at Hitachi, who examined the problem and analyzed the performance issues, just trimming bus width isn't enough (see the box, "How Do You Measure The

Performance Of An Architecture?").

Based on the 16-bit MCU's definition, digital signal processors may also fit the description. Most DSP chips, however, do not have a usable instruction set in control applications. But that's changing as enhanced versions of the DSPs' packaging features and instructions allow them to handle control applications.

For example, the latest 16-bit DSP ICs from Analog Devices, Motorola, Texas Instruments, and Zilog each include enhancements that efficiently handle embedded-control operations. Further adding to the confusion are many new MCUs that pack multipliers or full multiplier-accumulators that perform more DSP-type operations.

The need for higher performance has led Intel, Philips, and Temic to develop extensions of Intel's 8051 8-bit MCU. The Intel 80251, the Temic TSC80251, and the Philips 8051XA all provide 16-bit extensions to the 8051 instruction set and architecture. The 80251 performs many 16- and 32-bit data-transfer, logical, and arithmetic operations, and provides extended addressing modes that support indirect, relative displacement, and bit addressing. Initial versions of Intel's 80251 retain an 8-bit ALU and include 8 or 16 kbytes of ROM or OTP EPROM, 512 bytes or 1 kbyte of RAM, a full-duplex UART, on-chip power management, and two 16-bit counter-timers. One offshoot, the 82930A, adds a Universal Serial Bus (USB) port to the 80251, making the chip usable as a controller in USB-compatible computer peripherals.

The architectural enhancements—a three-stage pipeline, a 40-byte register file that allows access to its contents as bytes, words, and double words, and a 64-kbyte extended stack space—enable 8051 code to run unchanged, achieving up to a five-fold performance throughput improvement. If the new instructions are used, performance kicks of up to 15X are possible. A code reduction of 30 to 40% also can be gained if 8051 code is recompiled for the 80251.

In addition to Intel's basic 80251 definition, designers at Temic have defined a derivative product that packs hardware multiplication support which performs fast 16-by-16-bit multiplication or 16-bit division. The TSC80251A1 contains 24 kbytes of one-time programmable or UV-erasable PROM or mask ROM (a 50% increase of the Intel chip), 1024 bytes of RAM (expandable to 256 kbytes externally), and a 24-bit linear address space (16 Mbytes) for code and data (*Fig. 1*). For control operations, designers added a programmable counter array, a four-channel 8-bit ADC (not available on the Intel chip), and three pulse-width measurement units.

Going to a full 16-bit data path, designers at Philips increased the 8051 performance to the level of 16-bit controllers, yet maintained source-code compatibility with the 80C51 (*Fig. 2*). The 80C51XA-G3 includes support for multitasking operating systems and high-level languages such as C. Designers used a Harvard-style archi-

tecture to separate the code and data, giving programmers more flexibility in handling code either on- or off-chip. A linear, unsegmented 16-Mbyte address space is available for off-chip data and code access.

The 80C51XA-G3 provides a 20-bit address range with 1 Mbyte each for program and data space, 32 kbytes of on-chip EPROM/ROM, 512 bytes of RAM, three 16-bit counter-timers, a watchdog timer, two enhanced UARTs, four 8-bit I/O ports, and various power-management features.

The chip can perform a 16-by-16-bit multiplication in 480 ns and executes context switches in less than 2 μs at 25 MHz. By using the 16-bit extensions to the instruction set, the XA-G3 delivers a four- to five-fold throughput improvement for 80C51 translated code. With direct execution of the new instructions, it can achieve a 10- to 100-fold throughput improvement.

Derivative versions of the XA also have been defined by Philips. The XA-C3 is optimized for automotive applications and includes a CAN 2.0 (controller area network) interface, replacing one of the two UARTs on the G3. The SRAM space was doubled to 1024 bytes, and one of the four 8-bit I/O ports was removed. Another derivative, the XA-S3, doubles the G3's SRAM to 1024 bytes, includes the dual UARTs, and adds an inter-IC bus, a programmable counter array, a fifth parallel port, an 8-bit ADC, and a dual-channel universal parallel interface.

Although Motorola offers a 16-bit MCU family (the HC16 series), it saw the need for a lower-cost 16-bit family that provides a code-compatible extension of its 8-bit 68HC11 series. The result: The 68HC12 family, which can directly execute code written for the HC11 series, but adds 64 new instructions, a 20-bit ALU with an instruction queue, and enhanced indexed addressing. With the new CPU instructions and paged memory addressing, users have access to over 4 Mbytes of program and 1 Mbyte of data space.

Like the HC11 series, the HC12 includes a background debug capability. The multiline interface to the background debug logic on the HC11 was reduced to a single line. It also was significantly enhanced to consume fewer system resources and provide more

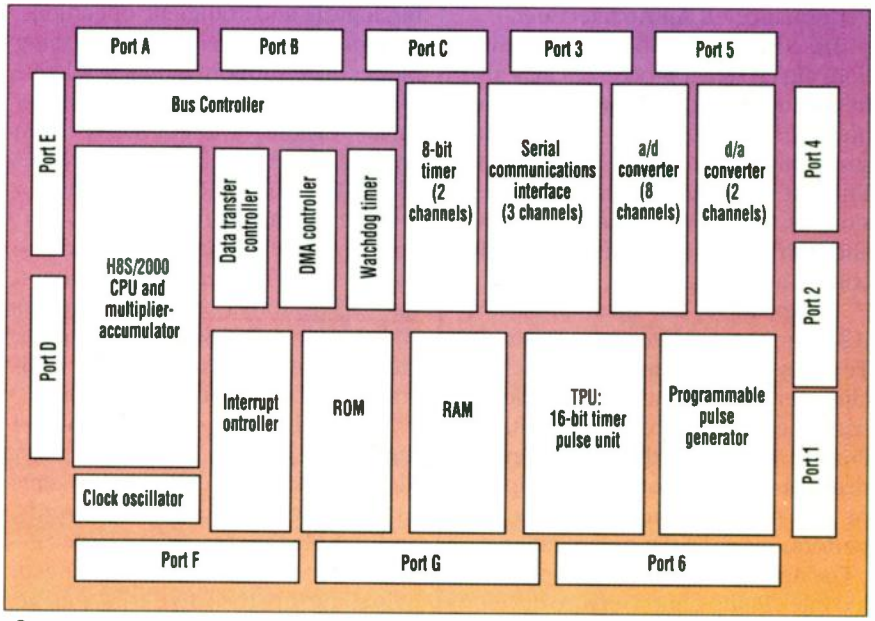

3. THE MOST ADVANCED 16-BIT *microcontroller series to come from Hitachi is based on the H8S-enhanced RISC-like, 16-bit core. The H8S/2000 series includes DSP, thanks to a full 16-bit multiplier-accumulator.*

transparent (nonintrusive) operation. With the background debug mode, users can set up two hardware breakpoints and code patching can be done for short code updates in lieu of redoing the entire program store.

The first members of the HC12 family include the 68HC812A4 and the 68HC912B32, with the former now sampling and the latter slated for release in the next quarter. The A4 chip contains the 16-bit core plus 4 kbytes of byte-erasable EEPROM; 1 kbyte of SRAM; an 8-channel, 8-bit ADC; an 8-channel, 16-bit timer block that allows each timer to perform input capture, output compare, or 16-bit pulse accumulation; and two asynchronous serial ports and one synchronous serial peripheral interface.

The B32 chip is the first 16-bit MCU to include 32 kbytes of flash-erasable program storage and 768 bytes of byte-erasable EEPROM. Other resources include a J1850-compatible data-link communications module for automotive applications; 1-kbyte of SRAM; an 8-channel 8-bit ADC; an 8-channel 16-bit timer array; and an 8-bit, 4-channel pulse-width modulator.

In addition to the new upward-compatible family, Motorola also has the 68HC16 series of programmable and configurable microcontrollers. Based on an upward code-compatible extension of the 8-bit 68HC11 microcontroller series, the HC16 series was designed with a highly modular architecture that surrounds an intermodule bus used to interconnect all of the functional blocks.

The HC16's data path consists of a full 16-bit ALU with dual 16-bit accumulators, three 16-bit index registers, and a 16-bit multiplier-accumulator that supports DSP applications. Various off-the-shelf versions come with memory options ranging from a simple 512-byte boot ROM (the MC68HC16V1) to a unit that packs 48 kbytes of flash EEPROM, 2 kbytes of block-erasable (with byte/word programming) flash, and 2 kbytes of SRAM (the MC68HC916X1). The V1 has a maximum clock speed of 20.97 MHz, while the 916X1's maximum speed is 16.78 MHz.

Also planning to put flash memory on board, designers at Sharp are readying a new family of 16-bit MCUs that are expected to operate from a 2.5-V supply and run at 10 MHz. The family will deliver relatively high throughput, executing a 16-bit multiplication in 1 μs and a 32/16-bit division in 3.9 μs.

Sharing some of the early history with the 6800 family, the 6502 microprocessor created in the mid-1970s has evolved thanks to development work at The Western Design Center. The company has expanded the 8-bit ar-

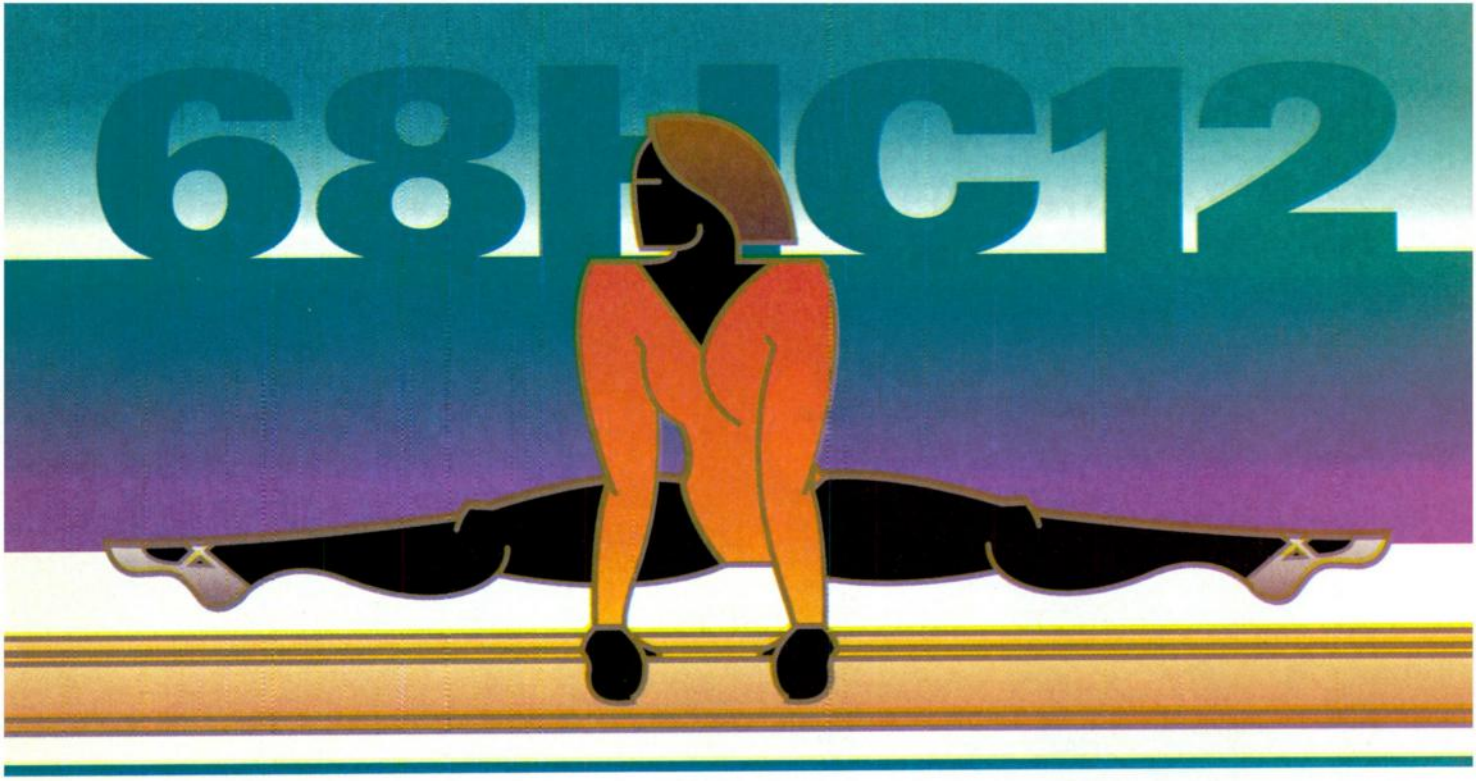

PERFECT BALANCE

The Competitive Edge of Motorola's New 16-Bit Microcontroller.

On a scale of one to ten, it's a 12. The 68HC12.

This new cost-effective microcontroller family offers the perfect balance of low power consumption and high performance. With a CPU12 core that runs at 8 MHz from 2.7 to 5.5 volts, the 68HC12 will launch next-generation designs for portable consumer, automotive, wireless communications, and industrial control products.

A perfect score. The 68HC12 is a highly integrated 16-bit architecture that is completely source-code compatible with Motorola's industry-standard 8-bit 68HC11 microcontrollers. Customers who need increased performance can now migrate designs directly up to the 68HC12 without sacrificing valuable investments in software.

The 68HC12 is packed with features that make it a market first. Low voltage, low power, and low noise at full bus speed. Enhanced Background Debug™ Mode for non-intrusive in-circuit programming and debugging. Flash memory. High-level language optimization. Fast math operations. On-chip fuzzy logic instructions. A modular design that leverages Motorola's wide array of existing, proven peripheral modules. And, the 68HC12 is backed by a complete set of development tools from Motorola and independent suppliers.

One performance you don't want to miss. To experience this winning combination for yourself, call your local Motorola sales office or authorized distributor to order the 68HC12 Evaluation Kit (M68HC12A4EVB). For more information call (800) 765-7795, ext. 942; fax us at (800) 765-9753; or visit our Web site at http://freeware.aus.sps.mot.com

MOTOROLA

Microcontroller Technologies Group

What you never thought possible.™

SGS-THOMSON MICROCONTROLLERS ARE MAKING THEMSELVES RIGHT AT HOME

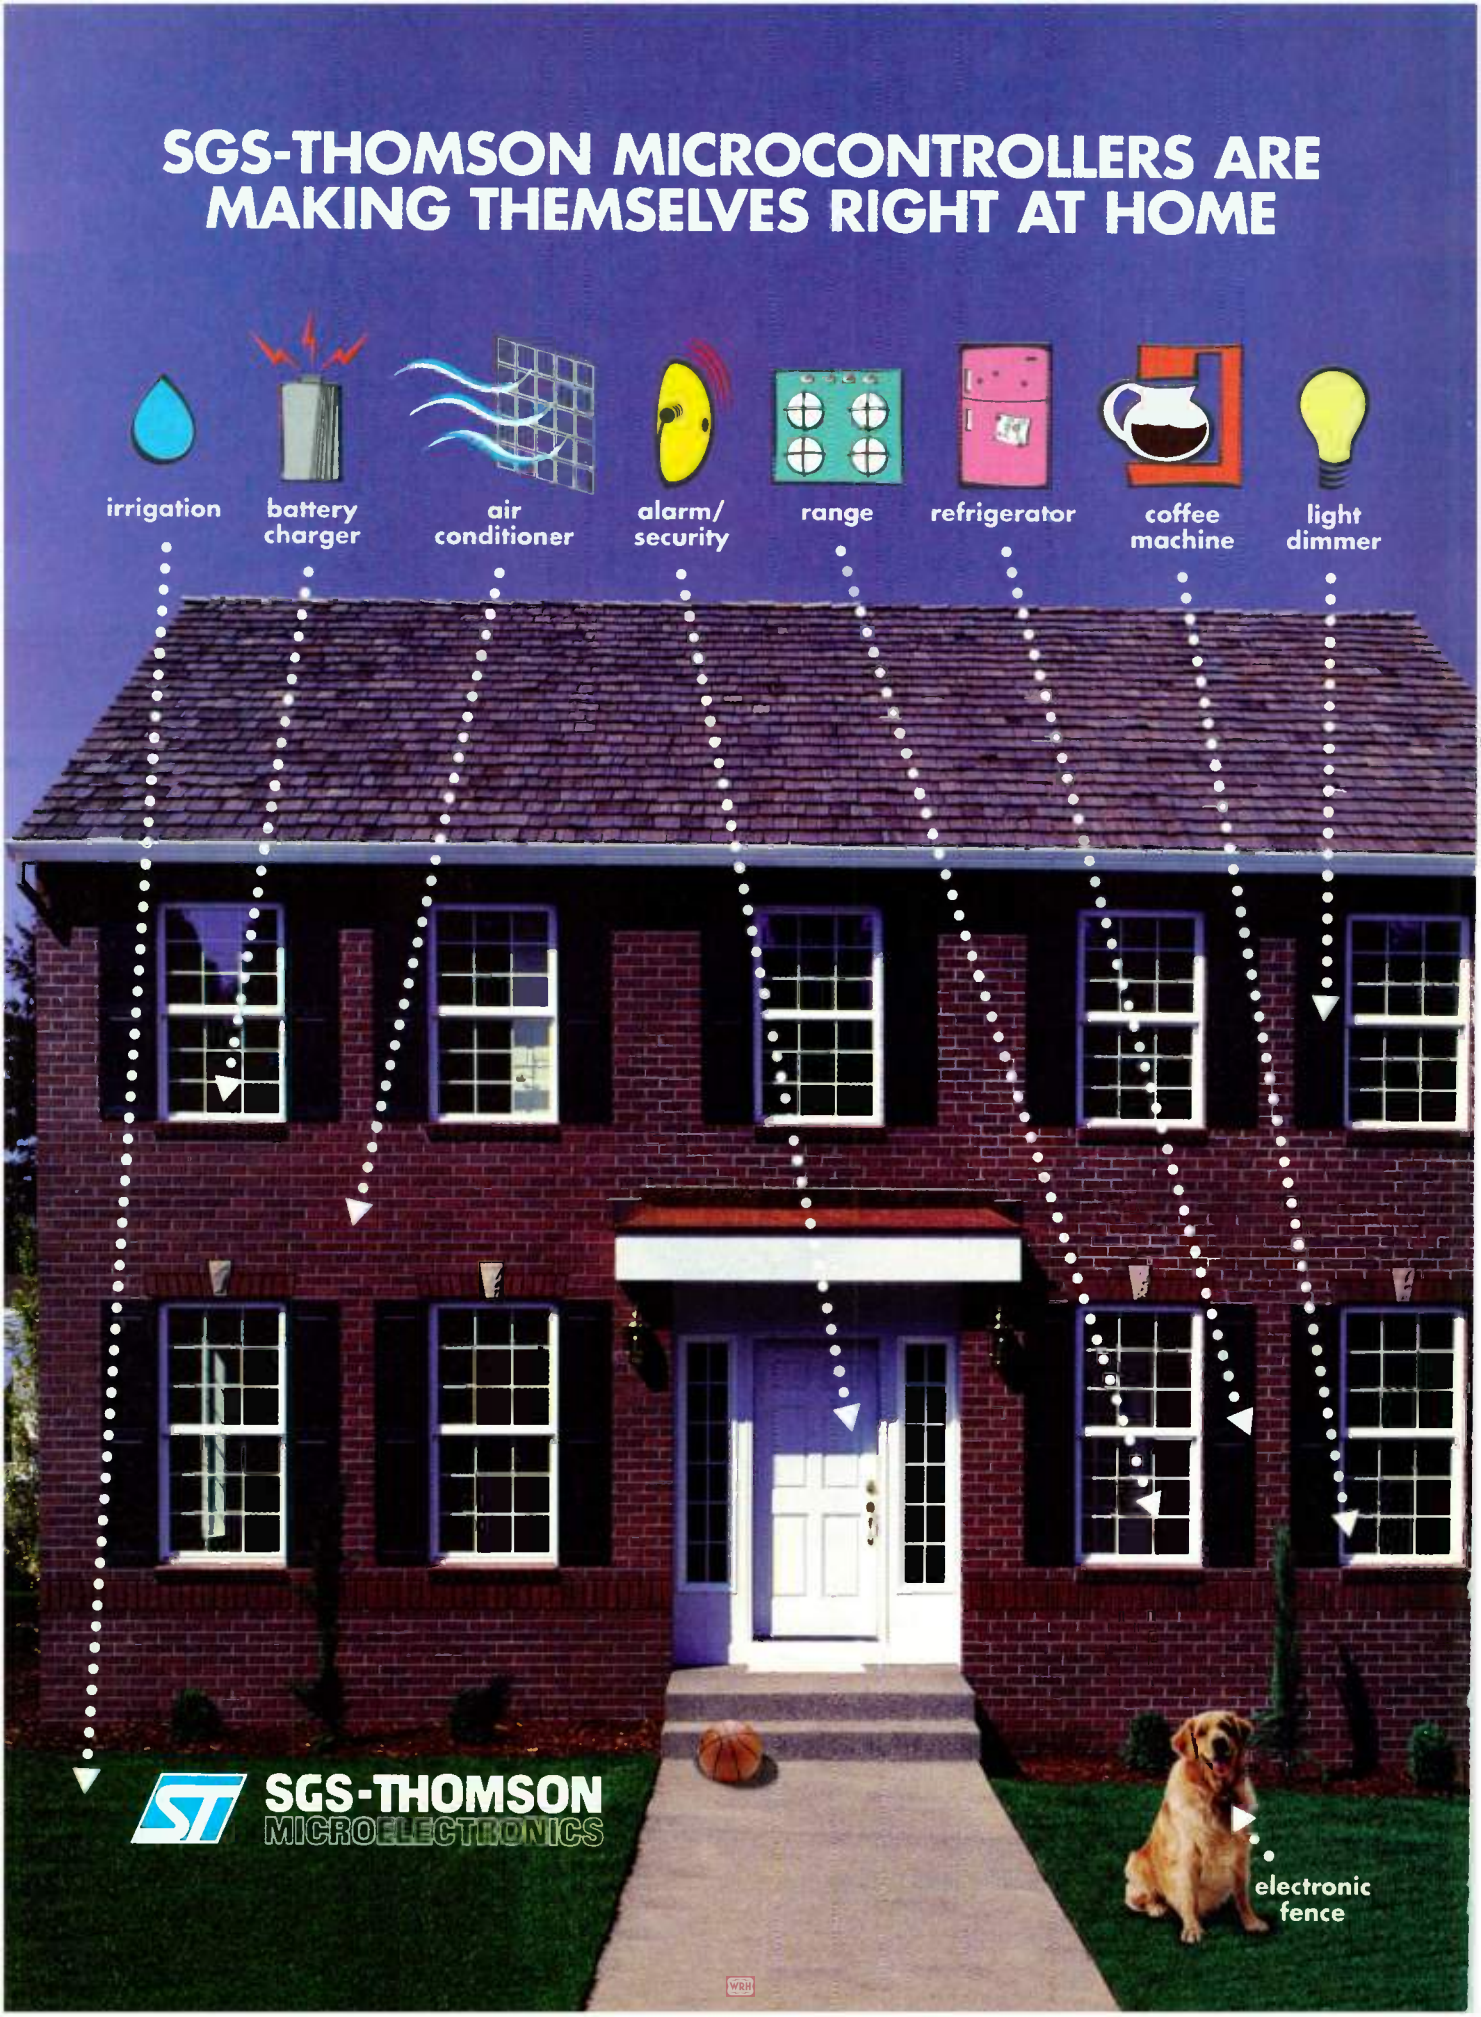

LET OUR MICROS MAKE YOUR PRODUCT A HOUSEHOLD WORD

The reason SGS-THOMSON's 8-bit ST62 family of MCUs is making itself at home in so many household appliances can be summed up in one word: *Value.* ST62 devices deliver more performance in less space for less money. Even the core is optimized for cost-effective operation. Add ESD protection and unmatched noise immunity and you begin to understand why the ST62 is opening doors to consumer applications that remain closed to ordinary MCUs.

Additional ST62 MCU Applications

- washing machine
- power tool
- heater
- UPS
- thermostat
- scale
- programmable timer
- vacuum cleaner
- home bus

All ST62s contain ROM, RAM, an 8-bit timer with 7-bit programmable prescaler and multifunctional individually programmable I/O ports. Also available: Devices with high-current buffers to directly drive LEDs or TRIACs, along with a wide range of peripherals such as PWM and LCD drivers. A wide operating voltage range and robust design allow ST62 microprocessors to be powered directly from a battery or the main with minimum external components.

In addition to the extensive ST62 family, SGS-THOMSON offers other 8-32 bit micro solutions such as ST7, ST9, ST10 and ST20 families. All of these products are fully supported with extensive development tools including C compilers for most families. So why not let our micros help make your product a household word. To find out more fax 617-259-9442, write SGS-THOMSON, 55 Old Bedford Road, Lincoln, MA 01773, or e-mail: info@stm.com. And visit our web site at **http://www.st.com**

ST6
The 8-bit MCUs of choice in automotive and industrial as well as consumer applications. Instruction set and addressing modes maximize code efficiency.

ST7
Powerful industry standard 8-bit core, surrounded by numerous advanced peripherals. 3K to 48K of ROM with different RAM sizes. Choose EPROM and OTP versions for proto-typing. On-chip EEPROM is also available for integrated data storage.

ST9
8/16-bit micro family fills requirements of most advanced computer, consumer, telecom, industrial and automotive applications.

Service and Technology

For More Info Turn Page

READER SERVICE 101

series, which offers a more modular internal architecture for easier customization, and the ability to run at clock rates from dc to 25 MHz. Functions such as serial channels, DRAM refresh control, and power management were added to provide more functionality. Versions of the Ex family run at supply voltages of 3 V and offer more power-saving modes than their previous family.

Both NEC and Advanced Micro Devices (AMD) crafted their own versions of the integrated processors—the V40/50 series from NEC and the Am186/188 family from AMD. NEC now offers several integrated V40/V50 chips, as well as other versions for embedded applications that require x86-instruction-set compatibility. A customized version of the V50 also was developed by Vadem.

The AMD offerings are part of a licensing arrangement with Intel, and the company has recently expanded its family with several versions—the Am186/188 ES and ESLV series. The families are architecturally identical, with the main difference focusing on operating voltage and power consumption (Fig. 4). The ESLV family can operate with 3.3-V supplies at clock rates of 20 or 25 MHz, while the ES family operates from a 5-V supply, but can run at clock rates of up to 40 MHz.

Enhancements to the chip architecture include an improved bus interface that packs 32 programmable I/O pins, two full-featured asynchronous serial ports with independent baud-rate generators (each port performs full-duplex transfers with 7-, 8-, or 9-bit data words), a multidrop 9-bit serial port protocol, DMA transfers to and from the serial ports, and pseudostatic RAM control support .

Another microcontroller architecture created by Intel also survived the test of time—the MCS 96 family. The latest version is the 8XC196. Its register-based CPU core eliminates the accumulator bottleneck and permits fast context switching. The redesign also runs much faster than previous versions, clocking up to 50 MHz. The processor core is a full 16-bit architecture and is flexible enough to handle bit, byte, word, and 32-bit operations.

Three major variants of the 8XC196 family are available—one version has an event processor array, the second version packs high-speed I/O features that include FIFO buffers and other I/O support, and the third family includes motion-control support in the form of a waveform generator and event processor array. The latest addition to the series with the event pro-

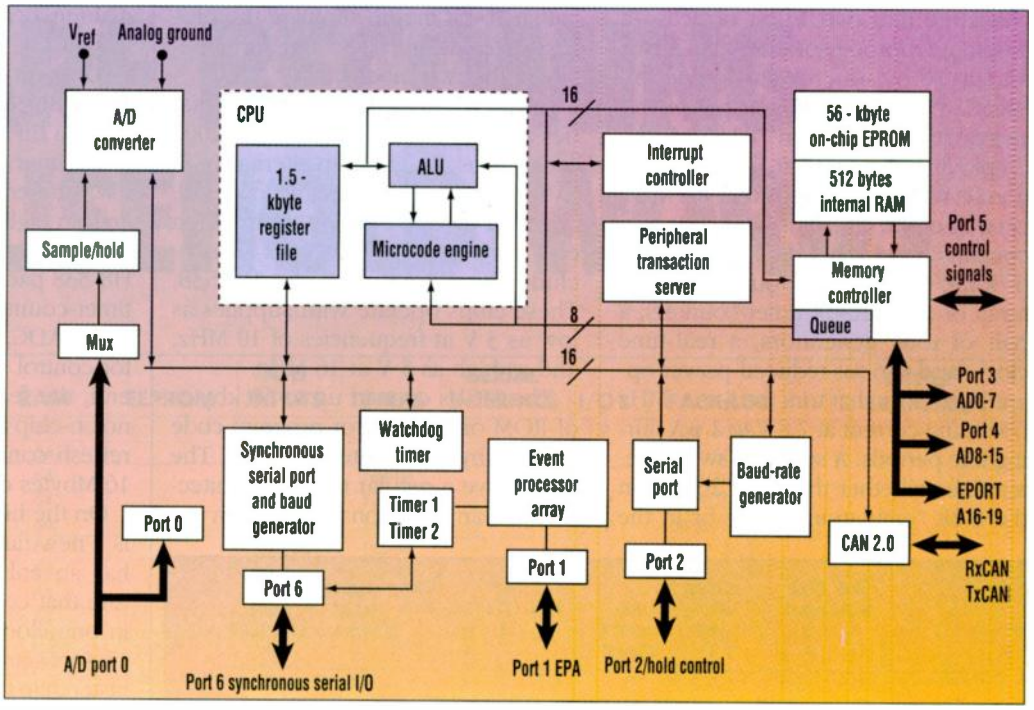

5. IN ADDITION TO THE 16-*bit processor core,* the 8XC196CB *microcontroller developed by Intel includes a CAN 2.0-compliant interface port along with all of the peripheral support required to handle engine control and other automotive needs, including a multichannel ADC, a pair of serial ports, an event processing array, 512 bytes of RAM, and 56 kbytes of EPROM.*

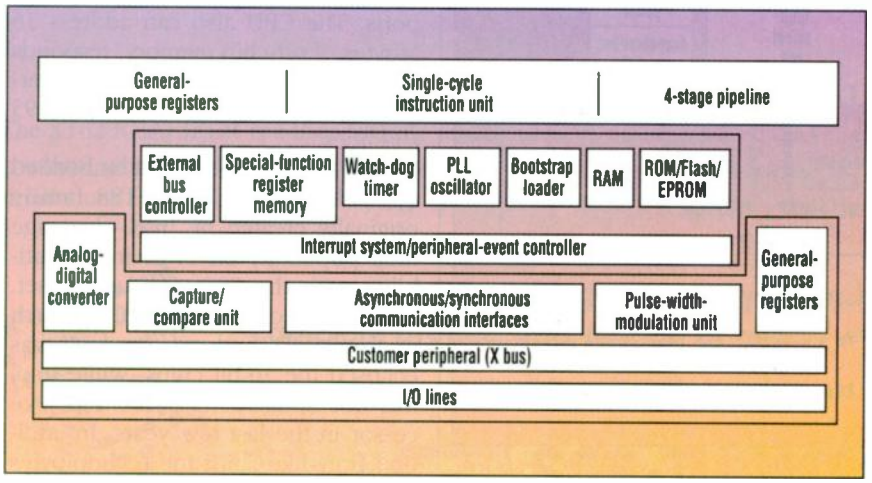

6. A FOUR-STAGE PIPELINE IN THE CPU *allows the SAB-166 series of microcontrollers designed by Siemens/SGS Thomson to deliver high instruction throughputs of 100 ns/instruction. Among the options included are 32 kbytes of flash memory, up to 111 I/O lines, two serial ports, and complex counter-timer support blocks for input capture and compare operations.*

"Stop Noise Two Ways"

Most people know that a pacifier will stop noise, sometimes very quickly, other times it takes a little longer. The method, however, works. Add a little sweetness and it works even faster. And for certain applications the pacifiers come in different shapes and sizes.

Conec Filter Connectors stop noise as well. Noise created by EMI/RFI in today's high speed digital systems. And just like pacifiers, Conec Filter Connectors come in different configurations - industry standard d-sub, high density d-sub, filltered adapters, modular jacks and combo d-sub, with power filtering and signal filtering. We don't add any sweeteners. What we do add is our patented planar filters. The technology with a proven track record, combined with the industry's best connectors—ours. Both technologies fully integrated, for best results in filtering.

Contact us today and put a little sweetness in your life. You can find us at our website:
http://www.conec.com, or
e-mail 24926@ican.net

CONEC QUALITY
"PEACE OF MIND"

• ISO 9001 CERTIFIED •

CONEC®

"TECHNOLOGY IN CONNECTORS"

72 Devon Rd., Unit 1, Brampton, Ontario Canada L6T 5B4
Tel: 905-790-2200 • Fax: 905-790-2201

READER SERVICE 102

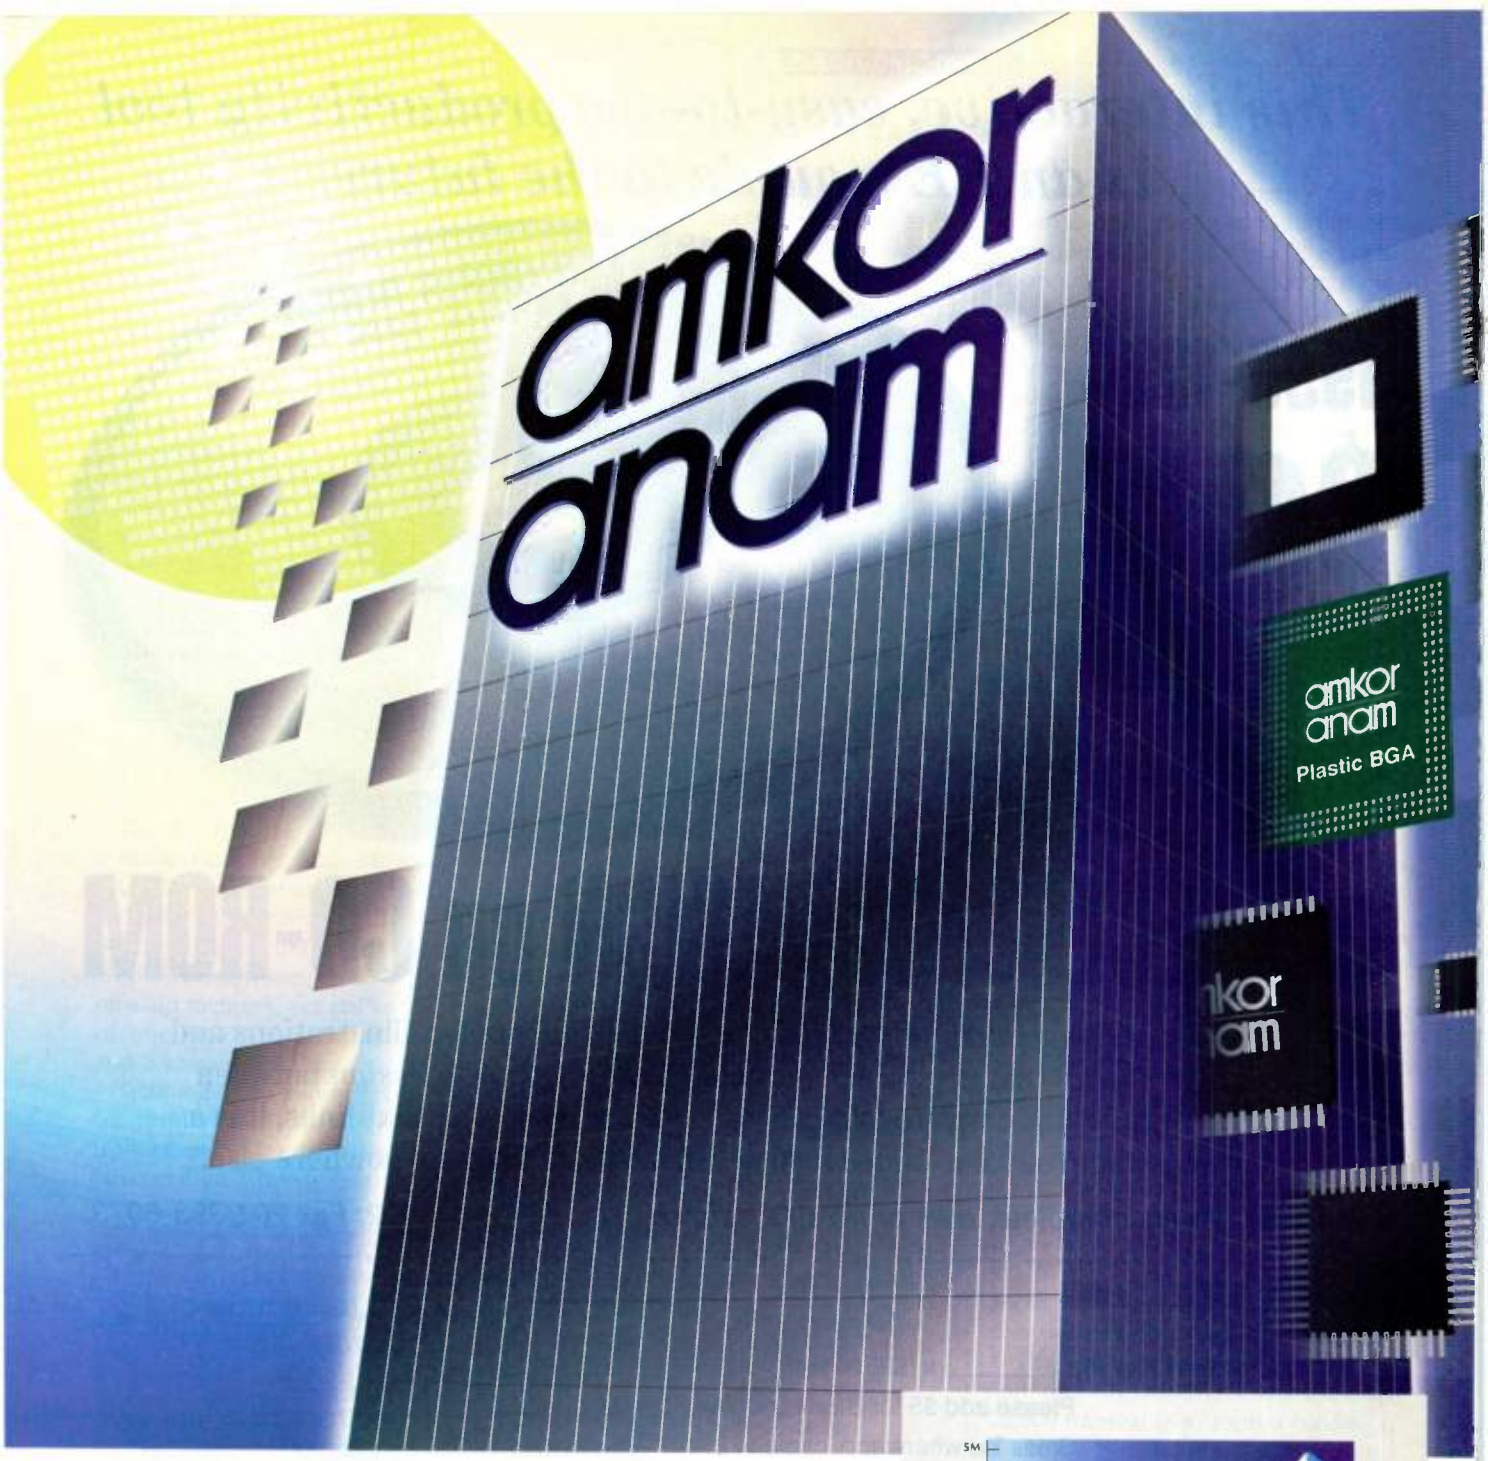

WE'RE PACKAGING THE FUTURE

AMKOR/ANAM has been packaging the future for over 25 years. Today, more companies look to us for IC assembly, test and new package ideas than anyone else. That's because we've *always* been ready to do whatever it takes to earn and keep our customers' business.

What about you? Need more IC packaging capacity right now, next year, five years from now? You can be sure we'll be your *strongest* supporter — regardless of market conditions. In fact, AMKOR/ANAM dedicates more state-of-the-art equipment per square foot of IC factory space

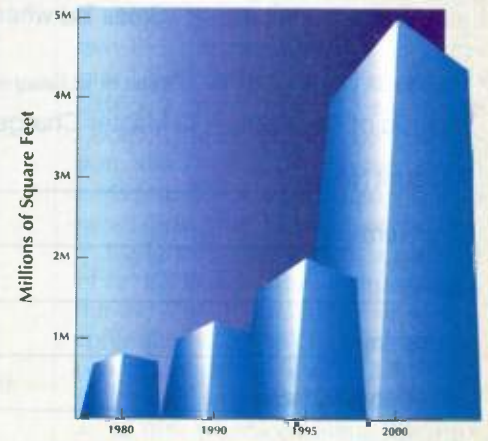

AMKOR/ANAM, already the world's largest IC package foundry will more than double its production capacity by the year 2000.

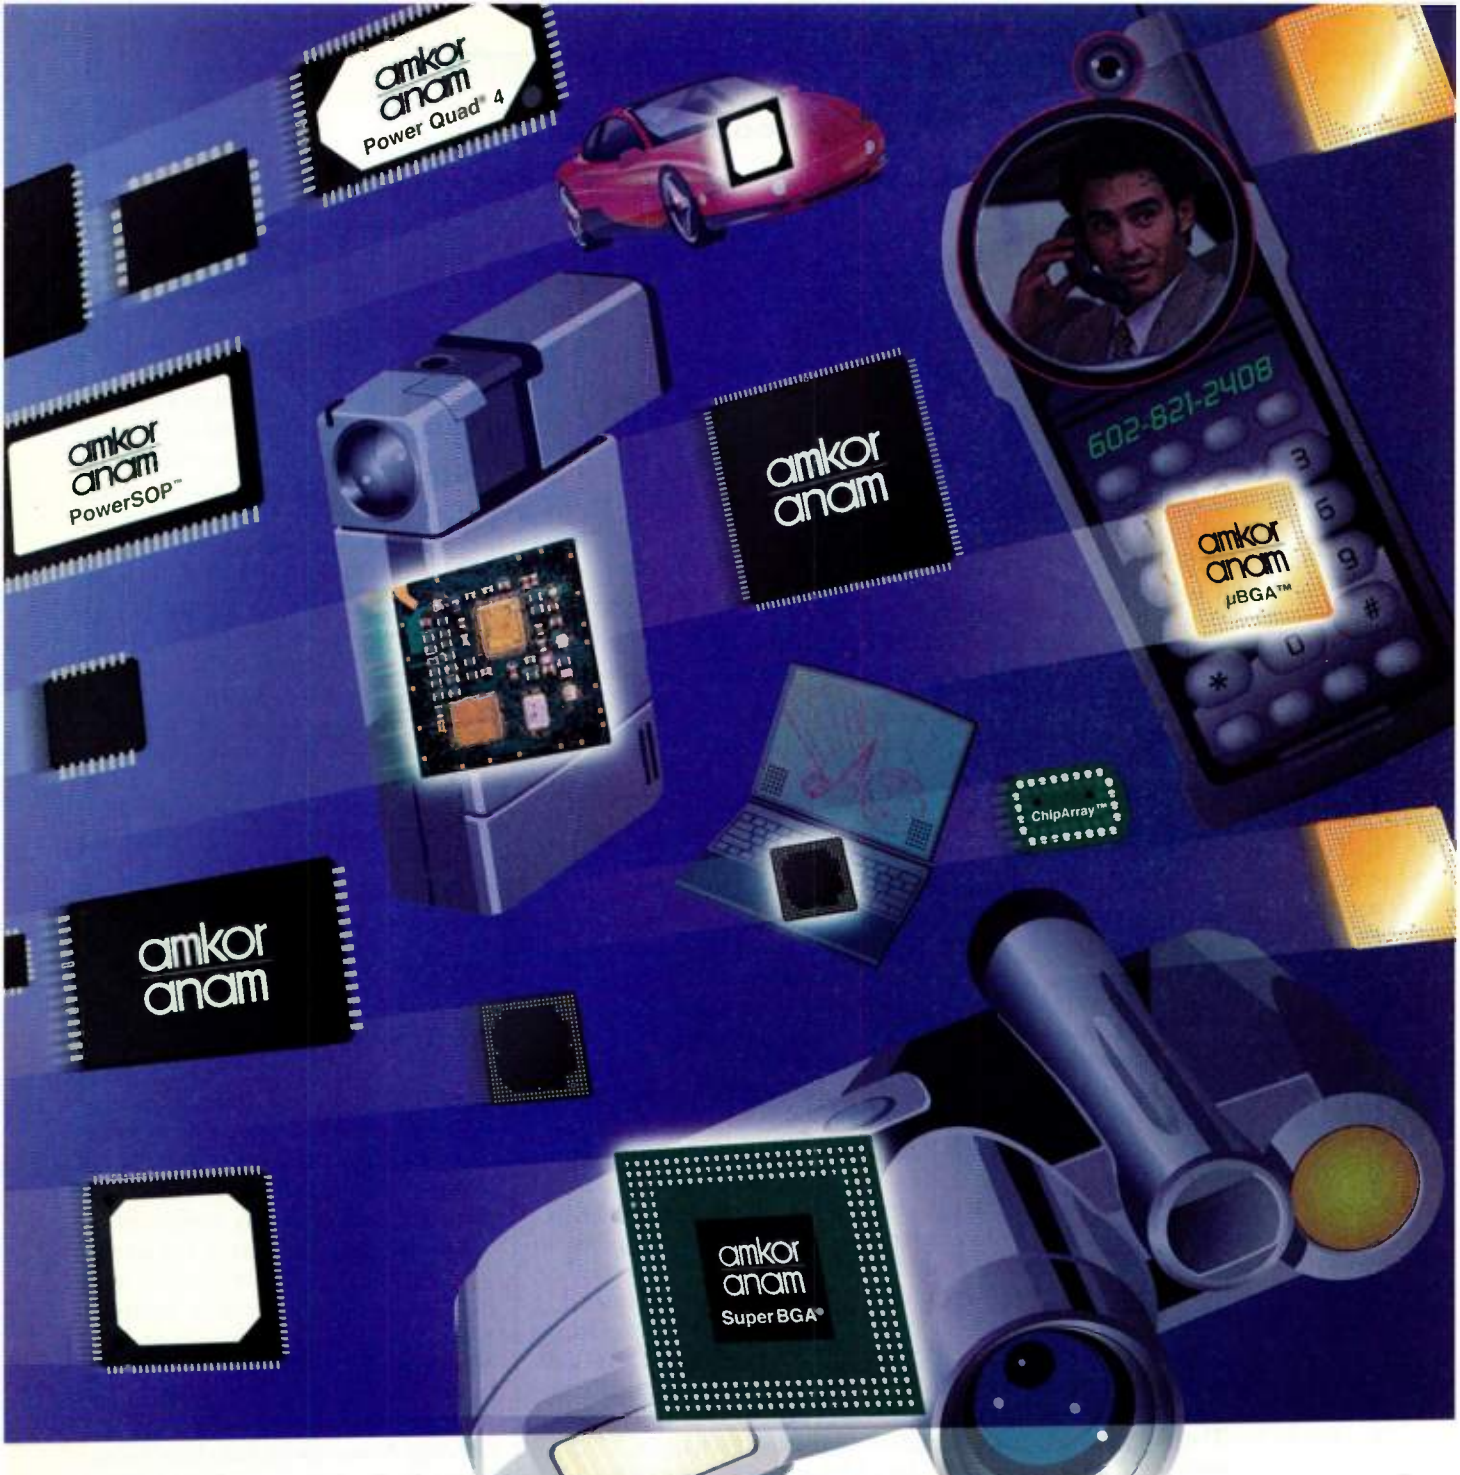

to you than anyone else in the world. And with the addition of new facilities in the Philippines, Korea and the USA, we're growing even stronger.

Need IC test services? From wafer probe to drop shipping — no one is more capable or more thorough than AMKOR/ANAM.

And if you're serious about moving your products into the 21st Century, count on our design team to lead the way. Semiconductor manufacturers and OEMs already recognize powerful, patented *Super*BGA®, ultra-small Chip Scale and other AMKOR/ANAM package innovations as pathways to the future.

Whether you need extra IC production capacity, industry-leading IC test services or cutting edge package designs, get in touch with the company that is packaging the future. AMKOR/ANAM. Visit our web site for the location of the service center nearest you: **http://www.amkor.com** or call: 602-821-2408, ext. 2000.

amkor
anam

READER SERVICE 81

16-BIT CONTROLLERS

dressing modes for maximum flexibility. Most of the instructions are less than 3 bytes long and can execute in less than 1 m s when the chips are clocked at 25 MHz.

One unique command offered on the MCUs is the Block Transfer instruction, which can move up to 64 kbytes at a time anywhere in the MCU's 16-Mbyte address space. Minimal support on the chip is provided for multiplication, but the chip still can produce a 32-bit result from a 16-bit multiply in just 1.92 µs.

The MCU family is subdivided into several subgroups—the M37702/03, the 37704/05, the M37710, the M37730, and the M37732. The first group contains the most core-like versions which are used as the heart of the other family members. These cores include from 512 bytes to 2 kbytes of RAM, up to 60 kbytes of ROM or OTP EPROM, eight channels of 8-bit analog-to-digital conversion, eight counter-timers, two UARTs, and up to 68 I/O lines. Optimized for motor control, the M37704/5 family includes three-phase motor controls and pulse-width-modulated control lines. The M37710 adds an eight-channel 10-bit ADC, while the 37730 and 32 have slightly reduced resources that optimize their use in pulse motor-control applications.

A more RISC-like 16-bit family, the M16C, was released in Japan last year by Mitsubishi. The register-based architecture includes hardware multiplier support so that 16-bit multiplies require just 500 ns to complete when the chip runs at 25 MHz. However, the processor also operates at very low power levels when clock speed is reduced—at a clock of 7 MHz, the processor consumes 22 mW when powered by a 3-V supply.

Most commands in the instruction set require one to four clock cycles for execution, while instruction sets from previous 16-bit MCUs would require the equivalent of 4 to 11 clock cycles to execute. The M16C architecture implements the instructions more efficiently, allowing developers to reduce code complexity by 15%.

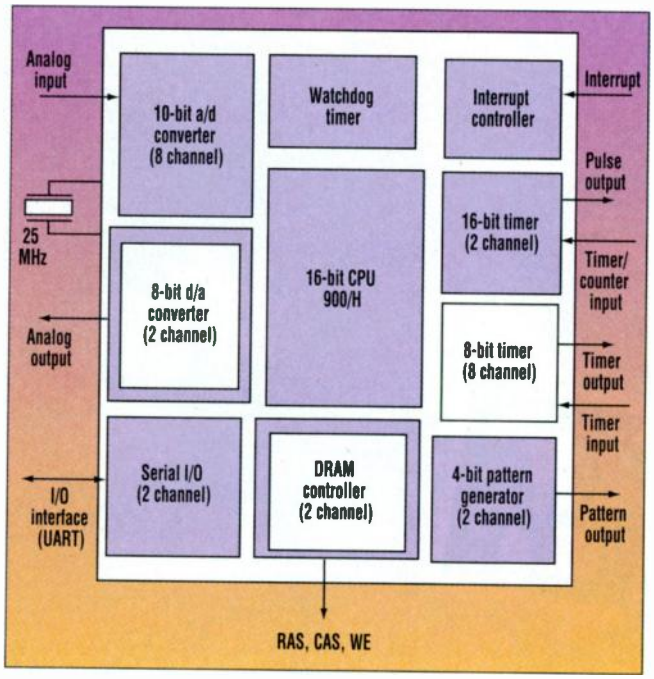

7. THE HIGHEST *performance member of the TLCS-900 series developed by Toshiba, the TMP95C063F, packs a 10-bit, eight channel ADC, a dual-channel DAC, a dual-channel DRAM controller, a pair of serial ports, multiple 8- and 16-bit timers, and a dual-channel 4-bit pattern generator.*

The first chip in the family, the M16C/60, includes the 16-bit CPU and hardware multiplier; 10 kbytes of RAM; 65 kbytes of ROM; a 10-channel, 10-bit ADC; a two-channel, 8-bit DAC; a pair of asynchronous serial ports; an interrupt controller; a DMA controller; and eight timers, plus a watchdog timer. Up to 1 Mbyte of off-chip memory can be addressed by the M16C/60. In addition, many of the industrial/consumer applications will typically see high electrostatic discharges.

In addition to x86-compatible processors described earlier, NEC also offers a family of embedded controllers based on its 78K series of microcomputers. The 78K series starts with low-end, 8-bit MCUs and increases in complexity to two families of 16-bit MCUs, the 78K/III and 78K/IV series. Instructions require a minimum of 250 ns to execute when the chips are clocked at 16 MHz (the highest-performance version), allowing the chips to handle many complex tasks.

Built-in support for pseudostatic RAMs and interrupt handling suit the chips well for real-time industrial applications. The interrupt controller can service vectored interrupts, handle macroservice routines, and initiate context switching. The macroservice support, which is a carryover from the 8-bit microcontrollers, reduces overhead to the CPU during the interrupt-handling process, such as when a counter or timer hits the desired value.

The 78K/III family includes a pulse-output capability for pulsed motor control, 512 to 2048 bytes of RAM, up to 48 kbytes of ROM or EPROM, a 10-bit or 8-bit multichannel ADC, high-speed math support and data-string operations for high-performance data processing, and an integrated DAC. For high software integrity, some versions of the K78/III family include error-checking and correction logic on the ROM interface to ensure that only correct instructions are delivered to the CPU. Some of the MCUs have a sum-of-products capability that can sum 16-by-16-bit products in 10 steps (13.4 µs at 32 MHz).

The newer 78K/IV series can operate over a wider supply range—from 2.7 to 6.0 V. In addition, it has multiple power-management modes, and provides a larger memory address space—1 Mbyte for programs and 16 Mbytes for data. With a 25-MHz clock, the chips will have a minimum instruction-execution time of just 160 ns. On-chip resources include four 16-bit counter-timers, three serial interfaces (two asynchronous, one synchronous), an 8-channel ADC, a dual-channel DAC, dual pulse-width-modulated outputs, up to 64 I/O lines, and a programmable clock output.

Also providing an upgrade to its 8-bit TLCS-90 family, Toshiba has a family of 16-bit microcontrollers that offers upward software compatibility with the TLCS-90 series. The TLCS-900 series is divided into three subfamilies—the general-purpose TLCS-900 devices, the low-power TLCS-900L subset, and the high-performance TLCS-900H series. The 900-series de-

nohau 80C196 In-Circuit Emulators

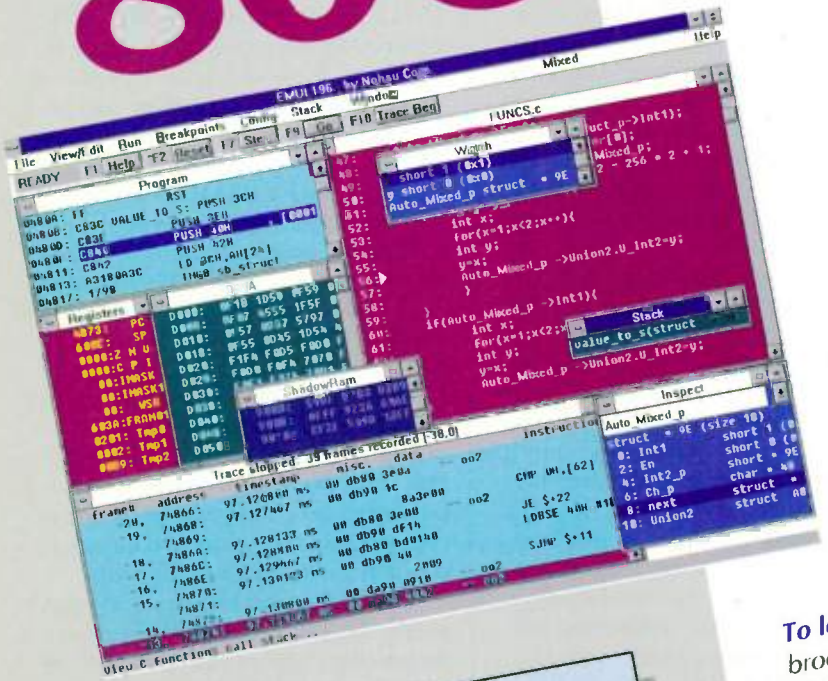

Real-Time Microprocessor Development Systems

EMUL196-PC™

To learn more, call **(408) 866-1820** for a product brochure and a FREE Demo Disk. The Demo can also be downloaded from our web site- **http://www.nohau.com/nohau.**

For more information via your Fax, call our 24-hour Fax Center at (408) 378-2912.

EMUL196-PC FEATURES:

- Support for 80C196: CA/B, EA, JQ/R/T, KB/C/D, KQ/R/S/T, NP/T/U and many more.
- Real-time emulation at maximum chip speeds.
- Uses bond-out chips for accurate emulation.
- Hosted on PC's and workstations.
- High Level support for popular C-compilers.
- Unlimited hardware breakpoints.
- Break in real-time on Internal Access, both on data value and address.
- Trace board up to 512K deep, 104 bits wide, with 40 bit timestamp. Triggering and filtering with full instruction queue decoding.
- Memory contents shown during real-time emulation (Shadow RAM).
- Code Coverage and Program Performance Analysis.
- CCB's controlled from user interface.
- PCMCIA compliant card for use with laptops.

MICROSOFT WINDOWS/'95 COMPATIBLE

nohau CORPORATION

51 E. Campbell Avenue
Campbell, CA 95008-2053
Fax. (408) 378-7869
Tel. (408) 866-1820

NOHAU INTERNATIONAL

Argentina 54 1 312-1079/9103, Australia (02) 654 1873, Austria 0222 27720-0, Benelux (078) 681 61 33, Brazil (011) 453-5586, Canada 514 689-5889, Czech Republic 0202-811536, Denmark 45 43 44 60 10, Finland 90 777 571, France (1) 69 41 28 01, Germany 07043/40247, Great Britain 01962 733 140, Greece +30-1-924 20 72, India 0212-412164, Israel 03-6491202, Italy 02 498 2051, Japan (03) 3405-0511, Korea (02) 784-7841, New Zealand 09-3092464, Norway (+47) 22 67 40 20, Portugal 01 4213141, Romania 056 200057, Singapore +65 749-0870, S.Africa (021) 234 943, Spain (93) 291 76 33, Sweden (0)40-92 24 25, Switzerland 01-745 18 18, Taiwan 02 7640215, Thailand (02) 668-5080.

Support also available for:

8051 P51XA™ MCS®251 68HC11 68HC16 683xx

How Do You Measure The Performance Of An Architecture?

Design engineers have an ability and desire to categorize and compartmentalize things to evaluate and ultimately select the best solution for a particular task, usually defined in terms of performance and cost. In the past, categorizing a microprocessor architecture was easy: It was either vonNeumann or Harvard. However, the proliferation of CISC and RISC microprocessors (MPUs)/microcomputers (MCUs) and DSP chips has made consistency in comparing and classifying devices very difficult. Nevertheless, several methods can accomplish this task.

An often-used measuring stick for processor performance is the instruction word's bit size. But this method quickly runs into problems. Do we mean the size of the op code (operation code) field? Or is it the size of the whole instruction, including the source/destination addresses/registers, as well as any immediate data or offsets that must be added to a register to give the effective address of the data? If the latter, do we use the smallest or largest case as the bit size? For a Harvard-type architecture, this might lead to classifying a 4-bit MCU as a 10-bit, 12-bit, or even 14-bit device, thereby putting a low-end microcontroller in the same class as some high-performance DSP devices.

Another performance measurement ruler involves the size of the data bus (or data path), either internal or external. But using the external data bus width may wind up classifying microcontrollers as 0-bit devices because they usually lack an external data or address bus. And by using this measurement ruler consistently, the same CPU would earn different ratings when used as a core in ASICs that have 8-bit, 16-bit, or 32-bit external data buses for different applications.

Another difficulty arises when using the data bus/data path ruler for digital signal processors. Many newer DSPs employ a modified Harvard architecture that provides two data paths from memory to the ALU to ensure that all data elements get to the ALU. Then these elements can be processed in one CPU cycle. Should the data bus width be considered to be the sum of the two? Also, if you couple a wide (say, 32-bit data bus) to an ALU (arithmetic logic unit) that's 8 bits wide, how efficient will that combination really be?

A third ruler for measuring MPU/MCU/DSP performance is ALU size. This method is superior because it relates directly to how efficiently a processor manipulates the data it receives, rather than merely how fast the processor gets that data. However, even this measure can lead to contradictions if used without discretion. For example, in the 1980s, one particular minicomputer had a 4-bit ALU, but handled its data as 12- or 16-bit words, so it was sold as a 16-bit machine.

Why do we need three (and probably many more) different ways to measure computer performance anyway? Performance, by itself, isn't what drives the decision of what's best for an application, whether from a user's or a chip designer's point of view. Rather, what's best is typically determined by the combination of performance and cost.

Designing a new MPU or MCU requires the chip designer to look carefully at balancing the trade-offs involving the performance that a market (or customer) needs at a cost (system cost) that the market/customer is willing to pay. Only a careful examination of available choices such as vonNeumann vs. Harvard, scalar vs. superscalar, pipelined vs. non-pipelined, RISC vs. CISC, and small data path vs. large data path will lead to an optimum implementation.

Contributed by H. Lyle Supp, microcomputer product marketing manager, Hitachi America Ltd., Brisbane, Calif.

vices operate at clock speeds of 20 MHz and have minimum instruction execution times of 200 ns as well as a large, linear address space—64 kbytes to 16 Mbytes for programs, and up to 16 Mbytes for data. Depending on the version, on-chip resources include 64 kbytes of ROM, 2 kbytes of RAM, a four-channel DMA controller, two 16-bit timer-counters, an 8- or 10-bit four-channel ADC, a two-channel pattern generator, and a DRAM controller.

Targeting low-power systems, the 900L series operates with supply levels as low as 2.7 V; at 3 V can run at 12.5 MHz, drawing 7 mA. During standby, the chips draw 10 μA. On-chip resources include 2 kbytes of RAM, 32 or 64 kbytes of ROM, up to 79 I/O lines, a four-channel DMA controller, an eight-channel 10-bit ADC, two serial ports, a two-channel by 4-bit pattern generator, two 16-bit timer-counters, plus two pairs of 8-bit timer-counters/pulse-width modulators.

For higher performance system needs, the TLCS-900H series removes the on-chip memory and adds a DRAM controller so that high-speed off-chip memory can be used (*Fig. 7*). Separate address and data buses allow the controllers to access data at 100 ns/word. Moreover, they enable the chip, while operating at 25 MHz, to deliver about twice the performance of other TLCS family members. The DMA controller on the chip also was enhanced so that it can transfer data at a rate of two bytes every 640 ns when clocked at 25 MHz (versus two bytes every 1.6 μs at 20 MHz for the TLCS-900).

At the high end of Toshiba's family is the TLCS-9000 series of 16-bit MCUs. These chips trim instruction execution time to just 50 ns when using a 20-MHz clock. Employing a register-based architecture that supports up to 256 register banks, the processor can easily handle multitasking applications and fast context switching. The ALU can handle three-operand instructions and perform bit-field extraction or insertion on fields of from 1 to 16 bits. Hardware-suported multiplication takes 350 ns (16 by 16 bit), and a sum of products operation (16 by 16 plus 32) takes 400 ns.

Originally published in the June 10, 1996 Electronic Design.

Compiler optimizations and programming techniques ease working with embedded 32-bit RISC microprocessors.

YUGO KASHIWAGI
Hitachi Ltd. Semiconductor & Integrated Circuits Div.
Tokyo, Japan

Coping With 32-Bit Code Density

BECAUSE EMBEDDED SYSTEMS ARE BECOMING MORE COMPLEX, 32-BIT RISC MICROPROCESSORS HAVE GROWN IN POPULARITY. AT FIRST, 32-BIT RISC MICROPROCESSORS WERE USED MOSTLY IN computation-intensive systems such as graphic engines. Hitachi's Super-H RISC Engine (SH series) microprocessors, which used simple RISC architecture to reduce die size and power consumption, made it possible to introduce 32-bit power for control-intensive, single-chip applications such as engine control.

Code density is the key issue because on-chip memory is the most important resource in single-chip applications. Instead of the 32-bit fixed-length instruction format of conventional RISC architectures, SH adopts 16-bit fixed length, encoding frequently used instructions into 2 bytes. At Hitachi, we developed an optimizing compiler for SH and programming techniques to reduce code density based on the study of real-world embedded application programs.

As embedded systems move from 8- or 16-bit microprocessors to 32-bit RISC processors, a common problem is an increase in code size. Two factors typically increase code size: The first is instruction length: Conventional 32-bit RISC processors have 32-bit fixed length instructions. The most frequently used instructions such as register-register operation occupy 4 bytes instead of 1 or 2 bytes in 8/16-bit microprocessors. The second factor is address specification: 32-bit RISC processors have a 4-Gbyte memory space. This results in a 32-bit area to store the address specification instead of a 16-bit area.

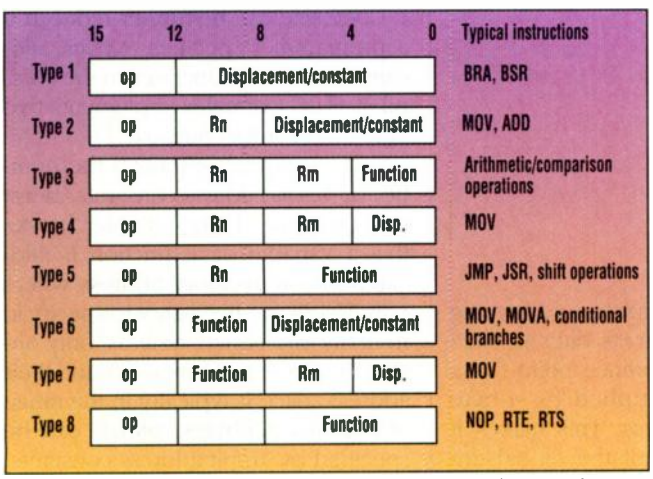

1. THE INSTRUCTION FORMATS *for the SH family show that the architecture has adopted a 16-bit fixed-length instruction format.*

SH architecture solves the instruction length problem by adopting a 16-bit fixed length instruction format. Fig. 1 shows the instruction formats of SH architecture. We focused on the address specification problem in implementing the compiler's code size optimization. In the first place, how can 16-bit fixed length encode 32-bit address values? Our strategy uses PC relative data load. A typical data load is implemented by the following code sequence:

```
MOV.L  LAB_a, R0    (loads address constant a using PC
                     relative addressing)
MOV.L  @R0,R0 (loads data specified by the address a)
...
LAB_a:
.DATA.L a
```

The address constant is placed after an unconditional

branch, such as BRA or RTS instructions. Compared with 8/16-bit microprocessors, this is the major reason for the increase in code size for the SH architecture. Loading or storing global variables and function calls requires 32-bit address constant loading, and these are the most frequent operations in embedded application programs. The following examples show the comparisons:

```
C Code:
   a=b;
8-bit code:
   MOV.L  @a,R0
   MOV.L  R0,@b (typically 8 bytes:
     (2 byte opcode with 2 byte address)x2)
SH:
   MOV.L  LAB_a,R0
   MOV.L  @R0,R1
   MOV.L  LAB_b,R0
   MOV.L  R1,@R0 (16 bytes (including
     address data))
   ...
LAB_a:
   .DATA.L a
LAB_b:
   .DATA.L b
(the label "LAB_x" is placed after the near-
est unconditional branch such as BRA or RTS)
Another example:
C Code:
   f();
8-bit code:
   JSR   @f (typically 4 bytes (2 byte opcode
    with 2 byte address))
SH:
   MOV.L  LAB_f,R0
   JSR   @R0 (8 bytes including address data)
   ...
LAB_f:
   .DATA.L f
```

The above comparison raises the following problems for the SH optimizing C compiler: The 32-bit architecture uses 4 bytes to specify an address constant, instead of 2 bytes in 8-bit microprocessors; being a RISC, SH requires extra instructions to load addresses compared with CISC, where the address loading operation is implicit in an operand. A similar problem can be encountered in the workstation/mainframe world, which is now moving from 32 bits to 64 bits.

Besides existing optimizations and various RISC-oriented optimizations, our optimizing compiler added two optimizations to solve the problem stated in the previous section. First, the compiler shares address constants among load/store code sequences. And second, it allocates address constants to registers, making the most of 16 general-purpose registers.

SHARING 32-BIT ADDRESSES

SH C compiler shares 32-bit addresses among load/store instruction sequences. This reduces the size growth caused by 32-bit addresses. This optimization could not have been done if SH were a CISC, where addresses are specified in operands instead of separate data.

```
C Code:
   a=1;
   ...
   x=a;
Before Optimization:
   MOV #1,R0
   MOV.L  LAB_a0,R1
   MOV.L  R0,@R1
   ...
LAB_a0:
   .DATA.L a
   ...
   MOV.L  LAB_a1,R0
   MOV.L  @R0,R1
   MOV.L  LAB_x,R0
   MOV.L  R1,@R0
   ...
LAB_a1:
   .DATA.L a
LAB_x:
   .DATA.L x
(26 bytes)
After Optimization:
   MOV #1,R0
   MOV.L  LAB_a,R1
   MOV.L  R0,@R1
   ...
   MOV.L  LAB_a,R0
   MOV.L  @R0,R1
   MOV.L  LAB_x,R0
   MOV.L  R1,@R0
   ...
LAB_a:
   .DATA.L a
LAB_x
   .DATA.L x
(22 bytes)
```

The reach of PC-relative addressing to load 32-bit address data is 1024 bytes from the reference point (8-bit displacement multiplied by 4-byte longword alignment). This makes it possible to share all the global addresses in medium-size functions.

A typical 32-bit CISC uses 6 bytes to load global data (2-byte operand and 4-byte address specification). SH uses 8 bytes (two instructions and 4 byte address data). But if the variable is accessed n-times, CISC uses 6n bytes whereas SH uses 4n+4 bytes. SH generates more compact code if a global variable is accessed more than twice within a 1024-byte range.

ALLOCATION OF ADDRESSES

We can reduce the number of address-load instructions by allocating 32-bit addresses to registers. In conventional 32-bit systems, 32-byte addresses are embedded in operands, and usually were considered "cheap" because they did not increase the number of instructions. This is not true. Even in conventional architectures such as 32-bit CISC, register allocation of 32-bit address constants is effective because of the reduction of the size of instructions.

In SH, we can eliminate address load instructions by allocating a 32-bit address to a register. The example in the previous section can be further optimized to the following code:

```
   MOV #1,R0
   MOV.L  LAB_a,R1
   MOV.L  R0,@R1
   ...
   MOV.L  @R1,R0 (R1 holds the address "a")
   MOV.L  LAB_x,R0
   MOV.L  R1,@R0
   ...
LAB_a:
   .DATA.L a
LAB_x
   .DATA.L x
(20 bytes)
```

LANGUAGE EXTENSION

Code size can be further reduced if a programmer specifies explicit allocation of data or functions to the compiler. SH C Compiler implements two #pragmas for this purpose.

SH C Compiler optimizes the number of 32-bit address constants. But it is even better if the compiler knows that a variable or a function is allocated in low memory address where only 16 bits are necessary to specify it. In many embedded systems, only on-chip ROM and RAM are used, and their address space is typically in the range of -32k to 32kbytes, which can be specified by 16-bit address constants.

We implemented a language extension (C #pragma) called "abs16" (absolute 16-bit address), to specify that

Where can you go to meet and converse with two of the electronics industry's most respected personalities?

The Fourth Annual

Portable by Design

CONFERENCE & EXHIBITION

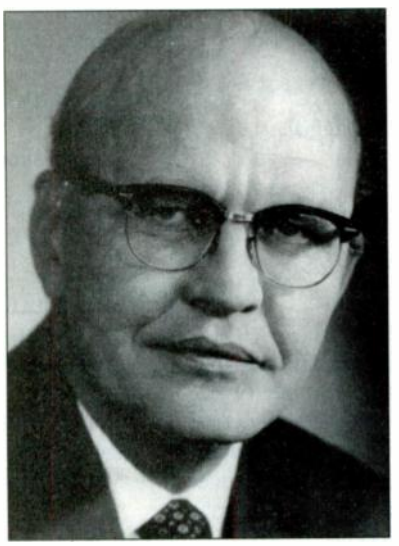

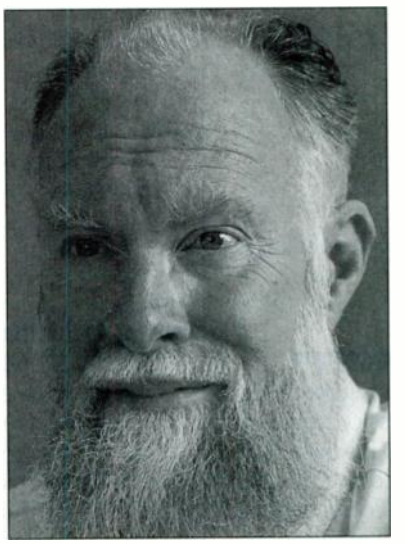

KILBY & PEASE

Come meet Jack Kilby, world-renowned inventor of the integrated circuit, and Bob Pease, analog guru and author of the famous Pease Porridge column in Electronic Design magazine, read by design engineers worldwide.

TUESDAY, MARCH 25, 5-8 P.M.
Mr. Kilby will present the first annual ***Electronic Design* Award For Technical Innovation** at the Portable by Design Industry Reception which will be held in the Exhibit Hall.

Industry Reception sponsored by

intel

WEDNESDAY, MARCH 26, 11 A.M.
Mr. Pease will enlighten attendees with a unique presentation in the Portable by Design Product Demonstration Area. At 1:00 P.M. that same day, Mr. Pease will be on hand to speak with attendees and autograph copies of his ***Electronic Design Compendium*** of Pease Porridge columns.

data or functions is allocated in the range of -32k to 32k. To reduce code size, a programmer can allocate frequently accessed variables and functions in low memory. For example:

```
C source:
#pragma abs16(a, b)
   ...
   a=b;
Without pragma:
   MOV.L LAB_a,R0
   MOV.L @R0,R1
   MOV.L LAB_b,R0
   MOV.L R1,@R0
   .....
LAB_a:
   .DATA.L a
LAB_b:
   .DATA.L b
(16 bytes)
With Pragma:
   MOV.L LAB_a,R
   MOV.W @R0,R1
   MOV.L LAB_b,R0
   MOV.W R1,@R0
   .....
LAB_a:
   .DATA.W a (Word size instead of long)
LAB_b:
   .DATA.W b
(12 bytes)
```

GLOBAL BASE REGISTER

SH has a special base register (GBR) to hold the base address of frequently used global variables. This base register can be used to hold the base address for frequently accessed global flags or I/O ports if the application is I/O-intensive. SH C Compiler implements a pragma ("gbr_base") to allocate global variables to this area. The area is small (128 bytes), but the allocation reduces the number of instructions as well as the address constant.

```
#pragma gbr_base (a)
   ...
   a&=1;
Without Pragma:
   MOV.L  LAB_a,R1
   MOV.B  @R1,R2
   AND #1,R2
   MOV.B  R2,@R1
   ...
LAB_a:
   .DATA.L a
(12 bytes)
With Pragma:
   MOV  #a-$G0,R0    ($G0 is the base of
                     the GBR section, and a-$G0
                     is less than 128)
   AND.B  #1,@(R0,GBR)
(4 bytes)
```

After implementing the above optimizations and pragmas, we further studied programming techniques to improve code size. The study of 8/16-bit embedded application programs has shown that they use global variables extensively, and programming techniques to eliminate them can further improve code size. In 32-bit microprocessors, global variables (or functions) are expensive. Related functions should be put in the same file. Then the compiler knows the relative address of the callee, and calls can be done using BSR (with relative address). This also helps modularize the program structure (a). Example:

```
   f();
When the function "f" is near the caller:
   BSR   f
(2 bytes)
When the function "f" is external:
   MOV.L  LAB_f,R0
   JSR @R0
   ...
LAB_f:
   .DATA.L f
(8 bytes)
```

DATA STRUCTURING

Structuring related global data into a "struct" reduces the number of 32-bit addresses (b) By doing so, several references to a global address constant can be combined into one reference. This also helps data modularization. Example 1: Source code:

```
extern int a, b, c;
   ...
   a=1;
   b=2;
   c=3;
Object code:
   MOV.L  LAB_a, R0
   MOV #1,R1
   MOV.L  R1,@R0
   MOV.L  LAB_b, R0
   MOV #2,R1
   MOV.L  R1,@R0
   MOV.L  LAB_c, R0
   MOV #3,R1
   MOV.L  R1,@R0
   ...
LAB_a:
   .DATA.L a
LAB_b:
   .DATA.L b
LAB_c:
   .DATA.L c
(30 bytes)
Example 2:
Source code:
extern struct s {int a, b, c;} x;
   ...
   register struct s *p=&x;
   p -> a=1;
   p -> b=2;
   p -> c=3;
Object code:
   MOV.L  LAB_p,R0
   MOV #1,R1
   MOV.L  R1,@R0
   MOV #2,R1
   MOV.L  R1,@(4,R0)
   MOV #3,R1
   MOV.L  R1,@(8,R0)
   ...
LAB_p:
   .DATA.L p
(18 bytes)
```

Looking at the source code, the sec-

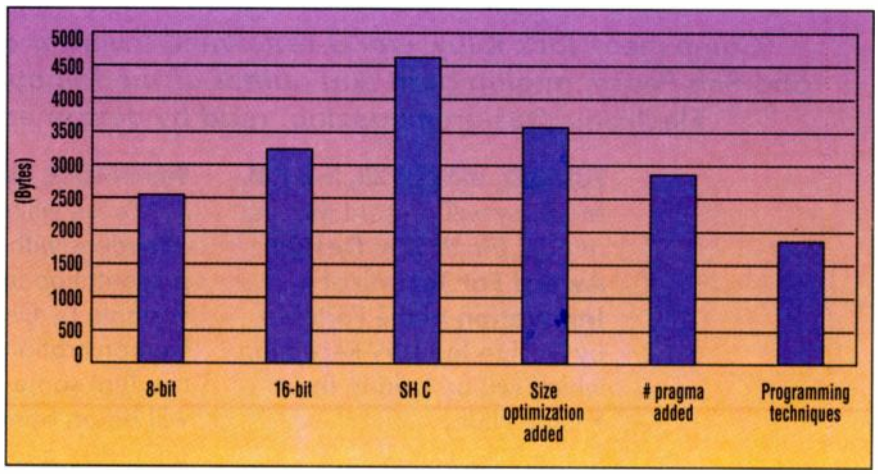

2. THE CHART DEMONSTRATES *how embedded code benefits from the application of optimizations, pragmas, and programming techniques.*

ond example seems to be more complex. But as to the number of 32-bit addresses in the code, the first example includes 3 (a, b, and c), whereas the second includes 1 (only x). The indirect access operation is efficient using register-with-displacement addressing modes. Once data are structured, the remaining 32-bit address in the code appears when we first load the base address of the structure. This can be eliminated by passing the base address as a parameter (c). In the following example, two global address constant references (in "g"and "h") are reduced into one by moving them to the caller "f."

```
f()
{
   g();
   h();
}
g()
{
   register struct xx *p=&x;
   p -> a=p -> b;
}
h()
{
   register struct xx *p=&x;
   p -> c=p -> a;
}
```

Example of passing base register as a parameter:

```
f()
{
   register struct xx *p=&x;
   g(p);
   h(p);
}
g(p)
register struct xx *p;
{
   p -> a=p -> b;
}
h(p)
register struct xx *p;
{
   p-> c=p-> a;
}
```

The programming techniques illustrated above are preferred from the software-engineering standpoint. Programs and data are structured using techniques (a) and (b), and references to global variables are essentially eliminated by the technique (c), improving the locality of program modification. These techniques are effectively embodied by the classes of C++. This suggests that the disciplined use of C++ is a better choice for embedded applications with 32-bit RISC processors. Fig. 2 shows the effect of our optimizations, pragmas, and programming techniques on a real-world embedded program. This shows that SH C compiler generates better code than 16-bit microprocessor C compilers if appropriate pragmas are specified. The compiler generates code as good as that supplied by 8-bit microprocessor C compilers if programming techniques are applied to eliminate 32-bit address constants.

Originally published in the April 15, 1996 Electronic Design.

Characteristics of RISC machines require special considerations in choosing a real-time operating system.

Tom Barrett
Embedded System Products, Inc.
Houston, Texas

How To Mix RTOS With RISC And Come Out A Winner

USING A REAL-TIME OPERATING SYSTEM (RTOS) AS A SOFTWARE FOUNDATION IS A GOOD DECISION FOR TODAY'S COMPLEX EMBEDDED APPLICATIONS. AN RTOS CONSISTS OF A KERNEL dealing with processor specifics such as CPU allocation and scheduling, register context changing, and memory management. Around the kernel is a library of routines, the RTOS services, that perform system level functions intended to achieve certain effects in the operation of the application code. The application is decomposed into a suite of tasks that get control of the CPU according to some multitasking scheduling algorithm managed by the RTOS scheduler. An application task, typically written in something other than assembly language, invokes an RTOS service by calling its corresponding application programming interface (API) function. Assuming use of some language such as C, the RTOS and its API library effectively mask the inner workings and hidden mechanisms of the processor, be it CISC or RISC, 8-, 16- or 32-bit, so that the application software engineer need not be too concerned with the actual processor.

That's the view from the outside. From the inside, the RTOS's view, the landscape is quite different. RISC machines differ from CISC microprocessors and microcontrollers and those differences often require special consideration where operating systems are concerned. RISC machines are built for speed, running at high clock rates and normally performing instructions in a single cycle. They often employ pipelines with multiple levels so that instruction and data prefetching and branch address evaluation can be done by the time the instruction rolls out of the pipeline into the execution unit. Whenever that pipeline flow is dis-

RTOS-On-RISC Care Abouts	
Processor Context •Number of General Purpose registers •Special Function registers •Floating Point registers	**Context Saving/Restoring** •Minimum number of registers •Reserved registers
Stack Frame Conventions •ABI compatibility •Alignment	**Coding** •Pipeline integrity •Choice of language •Can be debugged
Interrupt Processing •Single or multiple interrupt vectors •System Responsiveness requirements •Enable or disable interrupt doing ISR	**Tool Interoperability** •Compatibility with debuggers and emulators •ABI compatibility

rupted, as in some branch or jump instructions, performance takes a hit; so minimizing such occurrences yields the best performance. Simply put, the RISC machine is like a Formula I race car, fast on the straight-aways but it must slow down in the turns.

There are lots of "turns" in an application using a multitasking RTOS resulting from "straight-away" application code or the RTOS causing a change in the normal flow of processing. For instance, in multitasking applications, the processor is shared among many tasks which make calls to RTOS services causing the RTOS to do a lot of processing within itself (not all of which is "straight-away") in order to perform the requested operation and to decide which task is next to run and to give it control of the CPU. Add to all those changes in processing flow the possibility of more caused by the occurrence of external interrupts and it is easy to see the prospect of long stretches of high speed straight-away application code diminishes. So, if the RTOS is slow through the "turns", the application's performance will also suffer. One of the keys, then, to a successful RTOS on a RISC processor is making it perform efficiently through those "turns".

At each change in a task's processing flow resulting from a RTOS service request or an interrupt, the CPU's register context must be managed correctly so that processing can continue without error whenever the same task regains control of the CPU. RISC machines usually employ a large number of registers which requires special consideration by the RTOS designer. How many registers is it necessary for the RTOS to save when an application task makes a call to an RTOS service? To a function in the RTOS? For an interrupt? If too few are saved, the results are quickly obvious in that the RTOS is going to fail; so that's only a consideration early in the design of the RTOS. The converse is handling too many registers, a less obvious problem but one that causes the performance to drop. Fortunately, the RTOS designer can get help when making these decisions from guidelines provides by the processor's application binary interface (ABI) specification.

The ABI, usually written under the aegis of the RISC processor's manufacturer (e.g., the PowerPC Embedded ABI sponsored by Motorola), is a specification of conventions useful to developers of software tools such as compilers, assemblers and debuggers as well as components such as RTOSs. An ABI would contain definitions of register usage for argument passing during function calls, stack frame conventions, reserved registers and other information useful to a compiler or debugger designer. Because an RTOS is a piece of software that exists almost unto itself, most of the specifications in an ABI do not concern the RTOS designer. However, the stack frame conventions and reserved register definitions can be most helpful in designing an efficient RTOS.

Take reserved registers for example. They are non-volatile and need not be considered part of the processors' context that needs to be managed by the RTOS during context save and restore operations. By saving and restoring only the volatile registers, the RTOS saves a few cycles every time it operates on the processor's context. Considering that context management operations can occur tens of thousands of times per second, even those few saved cycles add signigicant time for the CPU to do something else.

The RTOS designer will likely make use of the ABI's stack frame conventions to ensure that task stacks are aligned on proper boundaries, that stack pointers are properly adjusted, and that frames created by the RTOS during context saves reflect the format specified in the ABI. Adhering to the ABI spec yields an important benefit: it becomes possible to view the stacks with an ABI compliant debugger.

There is also an unseen benefit derived by following the processor's ABI specification when implementing an RTOS: interoperability with tools. Instead of having to produce a separate binding of the RTOS for each tool chain as would be the case with non-ABI compliant tools, an ABI compliant RTOS should be compatible with any ABI compliant tool chain. Such compatibility makes the RTOS developer and the user happy, the former because it isn't necessary to keep porting the code to make it work with all of the non-compliant tools, and the latter because the RTOS and runtime libraries and debuggers are going to work together "out-of-the-box". Both parties save time and improve productivity.

CHOOSING A LANGUAGE

The RTOS designer must also be concerned with the issue of the best language for coding an RTOS on a RISC processor. RISC processors are complex and often their assembly languages are best used by people who enjoy self inflicted pain. Those more squeamish often revert to simpler solutions such as compilers that allow code to be produced with greater respect for one's sanity. Regardless of the language, selection of compiler is very important because it can make the RTOS code perform well or badly on a RISC processor. For instance, the way a loops is written may affect the continuity of the instruction pipeline. Breaking the pipeline requires flushing and refilling it with a corresponding loss of efficiency the cumulative effect of which is greatly reduced performance. RTOS designers usually need to write code for a RISC processor differently than an application designer so that the CPU's pipeline is broken only when necessary.

Another way of breaking pocessing flow is through an interrupt or other type of exception to normal operation. Interrupts not only break normal processing flow but a consequence of their occurrence, the required servicing of the interrupt, can adversely affect system performance and responsiveness if done poorly. While that can be said for almost any system, it is especially true for those based on RISC processors because their normally large processor contexts and hardware interrupt designs are diabolically unfriendly to software. An interrupt is an exception to normal processing; the generally accepted RTOS design is to handle them as expeditiously as possible, but it isn't always easy to define what is "expeditious" on a RISC processor running an RTOS.

Whenever an interrupt occurs, the associated interrupt service routine (ISR) must save some or all of the processor context, identify the source of the interrupt, service the interrupting device, and restore a normal processing path. That portion of that ISR pro-

cessing dealing with saving the processor's context is usually done with the processor's interrupt system disabled allowing a recoverable state to be stored without fear of corruption. But, with the interrupts disabled, other devices can't request service and must wait until the ISR re-enables interrupts allowing them to be recognized. When the processor has a large number of registers to be saved, as in a typical RISC machine, the time to save the context can chew up a lot of machine cycles and that increases the system's interrupt latency and decreases its responsiveness. Ideally, the interrupts would be enabled once the processor's context is saved but that brings up the complicating possibility of another device interrupting an active ISR.

Where there are interrupt priority levels or separate vectors, multiple interrupts are readily managed, but if all external interrupts go through one or two vectors, as in a PowerPC or ARM, the prospect of reentrant interrupt servicing becomes very real. While turning the interrupts on can certainly improve responsiveness to other interrupt requests, handling reentrant interrupts is complex and the RTOS must be designed to accommodate the situation so as not to leave an interrupt partially serviced. Is it more expeditious to have a simpler RTOS and keep the interrupts disabled for the life of the ISR and avoid the problem entirely or to enable interrupts and have an ISR and RTOS smart enough to handle reentrant interrupts?

There is probably not a pat answer to some of those questions because there are applications considerations that play a role in determining what is most expeditious. However, consider that the ISR may need to call one or more RTOS services in order to make the operating system aware of the event caused by the interrupt. Do the interrupts in the first scenario remain disabled during the RTOS service as well? Doing so has the obvious effect of extending the interrupt latency to the extent that system responsiveness may suffer noticeably. In short, what is expeditious may end up being what kind of system responsiveness is needed. With the simple RTOS, simpler ISR code is possible at the expense of reduced responsiveness versus. With a smarter RTOS and more complex ISR code, better system responsiveness is possible.

In the final analysis, the needs of the application have a lot to do with the RTOS selected for an embedded application using a RISC processor. But it is useful to remember that an RTOS well-suited to the application as well as the RISC processor will be fast down the straight-aways and keep up a good pace through the turns. On the other hand, if it is ill-suited, it will still go like a greyhound when the path is straight but in the turns it's a real dog.

Originally published in the September 16, 1996 Electronic Design.

System designers have not taken full advantage of system or architectural simulators as design resources.

Navin Govind
Intel Corp.
Chandler, Ariz.

System Simulators Can Speed Time-To-Market

DESIGN PROTOTYPES CAN BE AN IMPORTANT AID TO DESIGN ENGINEERS IN MEETING THEIR TIME-TO-MARKET GOALS. IN PARTICULAR, PROTOTYPES CAN BE OF GREAT USE IN

Generating optimized real-time applicationscode with the shortest development cycle. But in the early stages of design and development, physical prototypes are usually not available. In such instances, software simulation of the prototype model early in the design cycle can make the difference in meeting time-to-market goals. System simulation using high-level languages opens the door to system emulation and the selection of the rest of the tools within the tool chain, including a real-time operating system, compilers, assemblers, and debuggers for various projects. While the concept of using a system or architectural simulator has been around a few years, system designers have not taken full advantage of it as a design resource, particularly when compared to the widespread use of logic or circuit simulation. A key to successful system simulation are the accurate high-level, bus-functional models now available for co-development environments. These models are device models with complete timing accuracy. Tying the bus model interface to a system simulator facilitates software/hardware co-simulation and system validation. A system development environment is shown at a block diagram level (*Fig. 1*). The system architecture is defined at a hierarchical level, beginning with the bus-functional model to the final system integration.

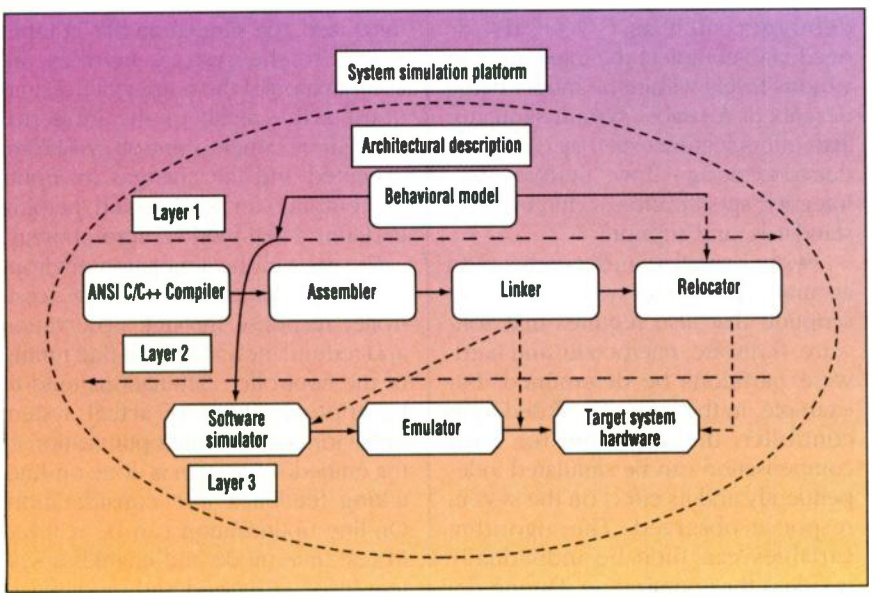

1. THIS BLOCK DIAGRAM *illustrates a typical system development environment. The architecture is defined at a hierarchical level beginning with the bus-finctional model and ends at the final system integration*

Behavioral modeling can be implemented at a high level to validate the design. When using a behavioral model, results are monitored through interactive graphical user interface (GUI) tools such as debuggers that allow the editing of variables, terms, and rule blocks for fine tuning. Model sim-

ulation directly affects the design and development time, as well as verification of a system. For example, complex systems can be analyzed for variable critical margins and system thresholds that would render a system unstable or unrealizable in the final development-cycle stages. In a simulation environment, blocks in a GUI editor are displayed and edited in a well-defined format. Using graphical interface editors, the system input and outputs can be manipulated for defining function types, and output shapes are then outputted to the simulation block for validation.

Despite its advantages, modeling suffered from serious shortcomings. First, it required a time-consuming and complex process of identifying system parameters and dynamics. And as most designers know, real-life embedded systems can interact with more complex systems. In these cases, generic floating point-based behavioral modeling, although useful for algorithm validation, are rarely sufficient. But due to its complexity, hardware and software simulation have been insufficient. So it's not surprising that a demand for effective commercial modeling tools emerged. These tools are becoming popular worldwide, even though most designs are still hand-stitched, and problems are discovered during integration, with the only option being a redesign.

The designer's ability to use off-the-shelf real-time operating systems has been particularly problematic. Significant advances have been made in real-time kernels, with a variety of architecture-specific, highly optimized, user-configurable kernels available. Some come integrated with feature-rich native development environments, GUI interfaces, configuration control, network management and Internet access facilities, all tuned to make software development and system debugging easier and more efficient. Unfortunately, many developers in the embedded world still do not use these integrated development tools because they believe they do not deliver accurate modeling, simulation, automatic code generation, and hardware/software emulation of systems running such microkernels.

SYSTEM SIMULATION

In a simple design such as an adaptive PID controller, time-to-market pressures can be eased if the system software integration is 90% complete beforehardware is targeted. This can be donr with a system simulator that allows system-level integration to be handled over the entire design cycle rather than just the back end. Desirable features in a system simulator include a completely or partially reconfigurable GUI that supports symbolic disassembly with multiple breakpoint and single-step execution. Support for a source-level, high-level language debugger such as C/C++ also is needed to complete the integration of various levels within the model being described. A usable system simulator also offers features such as easily understood design flow, performance, integrity, specification within industry standards, and support.

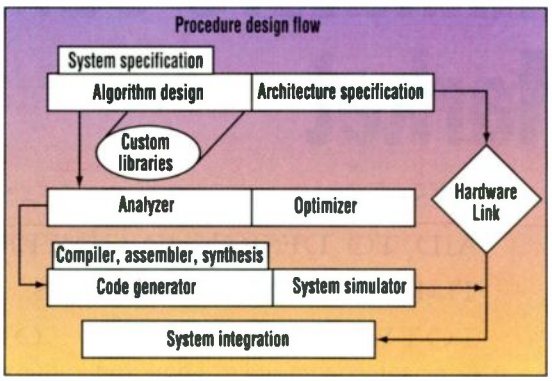

2. THIS TYPICAL SYSTEM SIMULATION MODEL *describes a procedural design flow. The choice of tools and the complexity of the system under consideration play a key role in the possible trading off of cost versus accuracy.*

System simulation depends on an accurate, complete system-level description that also requires that software, firmware, microcode, and hardware partitions be determined. For example, in the case of a closed-loop controller, the algorithm for error compensation can be simulated independently and its effect on the system response observed. The algorithm variables can then be individually tweaked for optimization. During debugging, a variable edit option can be used to specify the range, names, and data type of the variable within the simulation tool in floating-point or integer-type resolution.

While implementing a complex control system, a limitation in defining the variables is the degree of scalability the tool supports during code conversions. Standard practice usually dictates that variables be set to a limit within 10% by observation and are defined as integers that can be represented between 0 to 255 for computation. Values of coefficients and variables are chosen so that excessive undershoot and overshoot due to control are minimized. Values can be chosen by observing the change in the control surface plot and designing for a required system gain and zero overshoot.

CODE GENERATION

In the simulator environment, ANSI C code for computation and I/O handling is generated with 16-bit resolution using a GUI tool. Assembly-language code generation using the GUI is an option if the code is required to be compact and fast. In the optimization stage, off-line optimization allows system performance analysis using model simulation, while on-line optimization allows process hardware to be connected to the host system and optimization of the controller performance during run time. When simulation of a control loop is initialized, the simulation fills in input values to the system, invokes the computation of the output values, and outputs the result of the inference simulation. Single control cycles are executed and the changes in inputs and outputs can be observed, helping to define a real-world control strategy.

Changes in the simulation environment and determination of the controller response through observation and redundancy allows for fine tuning of the controller. Simulations tend to be approximations of actual system behavior. Appropriate optimization of the embedded system is done on-line, taking feedback into consideration. On-line optimization can be realized in real-time mode and enables a system to be visualized and modified in real time while the process is running. The generated code from the GUI tools are recompiled without the on-

line option, and integrated into the C196 hardware system in either ANSI C or assembly. The recompiled code compacts the code, since optimization features are enabled during recompilation. Superior simulator performance is attributed to the higher level of abstraction in describing the model and using accurate and procedural modeling techniques.

Since some controllers are prone to immunity from noise and system parameter variations, simulation tends to be ignored during the design and imple-mentation stages. Simulation speed often is a topic of debate. While logic and circuit simulators tend to be time-consuming, a C/C++ language-based simulator will run faster than EDA-based simulators. This means that a fast PC could run a simulation model quickly and effectively.

When rigid design methods are followed, repeated simulation runs must be observed to accurately predict the per-formance pattern. Problems occurring due to pure time delay can be reduced by including a model that predicts the future output of the system in general. Repeated simulation runs predict the performance and behavior patterns of the system under development and establish observable parameters for stability analysis.

A generic simulator can be used for simulations for similar designs. The compatibility of the tools and the design being implemented must be verified before the tool can be reused. The flexibility of using one set of tools for compatible multiple designs ensures maintaining previous designs, and also speeds up the design cycle.

SYSTEM EMULATION

The next step after system simulation is system emulation to verify software and hardware co-design. This allows system verification long before the hardware is ready for implementation. Software emulation of a target system helps system designers to select the right algorithm for optimization and fine tuning as well as in choosing the right set of tools. The idea of reuse lies within the design and simulation environment. Once the system model is created and verified, it would be possible in most cases to eliminate compatibility issues and a similar simulation platform used on derivative projects.

A visual description of procedural design flow for a system simulation model is shown (*Fig. 2*). Cost-versus-accuracy trade-offs depend on the choice of tools and the complexity of the system. For a design based on derivatives, the trade-offs are minimal. Debugging at the system level provides a process for discovering and correcting design and integration problems early and accelerates time-to-market by weeks or months. Simulation and model verification should be a necessity rather than an option. The parameter selection and code generation process using the right set of tools determines the robustness and efficiency of the controller-based design. By using a complete ensemble of visual graphical tools, a successful embedded system can be implemented and verified in a short design time-frame.

The advanced concepts of system modeling, simulation, automated design and validation are all well understood. Translating them into practical tools, however, has been difficult. But it has been changing recently. An example of the new approach is the joint effort of Intel Corp., and Integrated Systems Inc., Sunnyvale, Calif. They have ported ISI's MatrixX tool chain to Intel's MCS(R) microcontrollers. MatrixX is a full-featured environment for system modeling and visualization, automatic code generation (ANSI C and Ada), virtual and real-time hardware simulation, and automatic documentation.

The growing system complexity and availability of tools such as real-time microkernels, combined with the trend toward expensive, long turnaround, custom components, is tipping the scales in favor of automated modeling, design and simulation tools.

Originally published in the October 24, 1996 Electronic Design.

Computer boards combine with flash-file system software and PC Cards to upgrade a control system.

Kent Tabor
Granite Microsystems
& Raz Dan
M-Systems

Embedded Process Control Gets Boost From Flash PC Card

MANY APPLICATIONS THAT WERE ONCE BASED ON PROPRIETARY HARDWARE AND SOFTWARE ARE NOW BEING UPGRADED WITH OFF-THE-SHELF SINGLE-BOARD COMPUTERS (SBCS). These SBCS are actually embedded personal computers that have been ruggedized. Factories, hospitals, process control systems, and other "mission-critical" applications require reliable solutions that can work in harsh and demanding environments that a normal desktop PC cannot handle.

Rugged and reliable data storage is a key requirement for such critical applications. This application brief describes how combining off-the-shelf SBCs from Granite Microsystems and flash PC Cards (formerly PCMCIA—or, Personal Computer Memory Card International Association—cards) and TrueFFS software from M-Systems can upgrade a 20-year-old process control system in a concrete mixing plant.

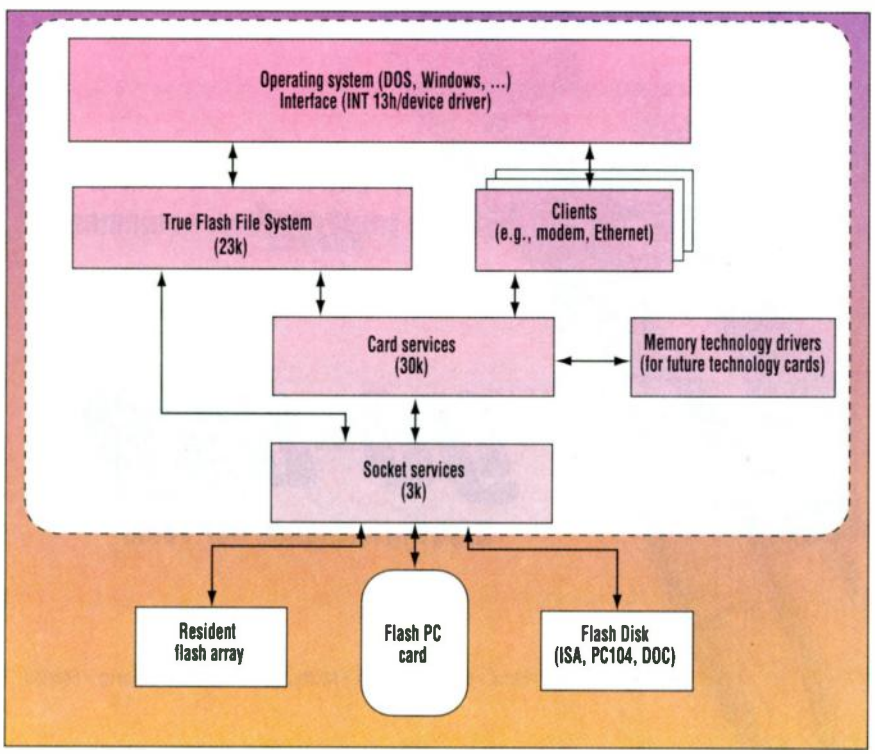

1. DISK EMULATION *using TrueFFS provides for upgrades in functions.*

One challenge Granite faces in providing SBCs to replace any legacy application is the need to develop a general-purpose system that can work with a variety of different operating systems (OS) and interface with existing hardware. In addition, the SBCs must be able to withstand harsh operating environments; provide cost-effective, yet rugged, storage; and fit a form factor compatible with the existing space available at the application site.

Granite has found that hard disk drives are the weakest link in harsh industrial environments, because they often fail when subjected to temperature extremes, vibration, or shock. For industrial and process control, reliability is especially critical: If a plant's

operations fail, tens of thousands of dollars can be lost in a single day. Flash PC Card storage gives several advantages that hard drives cannot match: greater tolerance of temperature ranges, immunity to a rugged environment, lower power consumption, faster execution speed, compact design, and high mechanical and data reliability. A 3.5-in. hard drive takes up 25 times the cubic volume of a flash PC Card, weighs 16 times as much, and consumes 71 times as much power. A flash PC Card can withstand 33 times as much shock as a hard drive, and its average seek time is 100 times faster.

Granite designed a system using two SBCs on a split passive backplane for the concrete batching control system. Both were 486 DX4-100-based machines. One computer ran the application's real-time RT kernel OS and the entire plant's control system, totaling several hundred I/O points. The other board served as the front-end graphical user interface (GUI) for operators, running Windows. Granite's Windows SBC allowed operators to design recipes for concrete mixes on-screen while the plant simultaneously batched concrete, performed diagnostics, generated tickets for truck drivers, uploaded and downloaded scheduling information, and kept track of inventory.

The previous control system consisted of a patchwork of 20 years of software development, none of it based on DOS or the PC architecture. All its features had to be maintained in the upgrade. None of the existing code was portable to a PC-based platform, so all of it had to be rewritten. One of the main challenges of this application was the fact that the RT kernel OS was not DOS-based, but off-the-shelf utilities are all written for DOS. Using off-the-shelf utilities in non-DOS applications requires adding a great deal of verification and testing to the development process.

Granite needed utilities that were compatible with the multiple OS being used in this application, and a file system that could be easily ported to all of them. In addition, a well-designed and thoroughly tested file system increases reliability.

The new control system also had to include a storage medium with the smallest possible size and the highest possible reliability. The new system was far more complex than the previous one, in that it was a deterministic, real-time design with multiple OS and dual processors, and it had to fit physically into a smaller envelope than the previous generation of hardware.

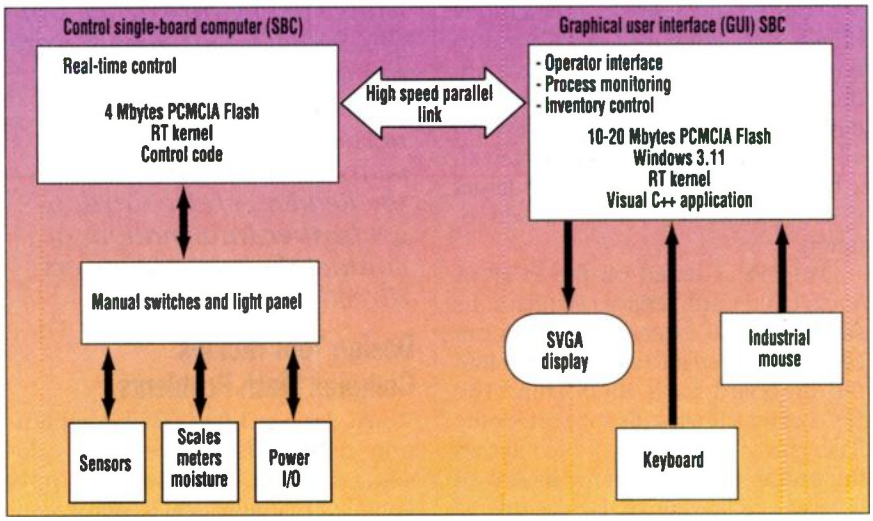

2. CONCRETE BATCHING CONTROL *system uses Granite's single-board computers and M-Systems' Flash PC Cards.*

The control system's complexity was increased because it was designed to be fully operable in the manual mode in case either computer experienced a falilure. Granite designed electronics into the control SBC that could talk directly to the I/O.

Additional electronics were designed to allow the operator to manually override the system if necessary. Both computers were partially redundant to each other, connected via a high-speed parallel port, and equipped with watchdog timers.

If the Windows SBC failed, the control SBC would continue controlling the plant's processes, but it would not perform additional functions such as tracking inventory. If the control SBC failed, then functions such as batching would be interrupted and would have to be done manually. In that case, however, all the events, such as inventory tracking and truck ticketing, would continue to be recorded, so that no data would be lost.

To meet the concrete plant's reliability and form factor requirements, Granite decided to use flash PC Cards from M-Systems. By adopting an industry-accepted standard, Granite was able to deliver a rugged solution in a credit-card size package. By incorporating the PC Card (PCMCIA) interface chip on the SBCs, Granite could also provide additional features for future memory or functionality upgrades, such as adding a modem card (*Fig. 1*). The fact that the flash PC Card was removable meant that capacity requirements could be tailored to the customer's needs; an easier upgrade path to higher-capacity flash cards was possible; data and programs could be either developed or modified directly on the application's SBC, or developed on a standard desktop or laptop PC and transferred to the target SBC; and the overall solution was more cost-effective.

Both SBCs in this application contained M-Systems' flash PC Cards. The SBC controlling the plant's processes used a 4-Mbyte flash card, and the Windows front-end machine used a 10-Mbyte flash card with a 20-Mbyte option (*Fig. 2*).

Besides the memory cards, Granite was also responsible for providing the software that allowed the SBCs to boot from the flash PC Cards. The cards had to be capable of emulating a hard drive while running under Windows and the RT Kernel real-time OS. Granite chose M-Systems' TrueFFS, the de facto industry-standard flash file system software package for working with flash memory. True FFS supports flash cards from multiple vendors such as Intel, AMD, Samsung, and Toshiba.

FLASH PC CARDS

The TrueFFS package includes drivers that allow a flash PC Card to emulate a mechanical hard drive under a variety of operating systems. This emulation allowed Granite to develop all the software on a standard PC with a hard drive and port it to the target system with its flash PC Card-based solution.

TrueFFS is based on the PCMCIA Socket and Card Services standard. Its standard BIOS extension module can be programmed into an EPROM or into the on-board flash BIOS chip. The BIOS extension installs support for the flash card via INT 13h, which allows the system to boot from the flash card instead of a hard drive. Device drivers are also available for systems that do not need to boot from the flash card, such as laptop and desktop PCs. The flash cards' industry-standard package makes them easily interchanged between platforms.

Granite installed M-Systems' TrueFFS BIOS module on the flash BIOS chip, giving a bootable solution that also emulated a hard drive. The BIOS extension with the TrueFFS package did not require customization because it supports the standard, Intel 82365-type interface chip that Granite already used on its board.

TrueFFS guarantees high data integrity under extreme conditions. In case of power failures or card removal while writing to the flash cards, the algorithms ensure that the data structures that map information on the card will not be corrupted.

Error-detection mechanisms built in to TrueFFS can automatically retire flash memory storage blocks if they become worn out, without affecting system operation. The occurrence of bad blocks is minimized through the incorporation of third-generation wear-leveling algorithms. The effects of an error are always localized and do not affect data integrity and the ability to access data globally.

Granite's concrete batching control system is the first in a growing number of off-the-shelf solutions that combine the PC platform and flash PC Cards for industrial and process control applications.

Originally published in the June 24, 1996 Electronic Design.

NEW PRODUCTS

Unlike the feature articles in this Supplement, which were previously published in Electronic Design, the products described here are appearing for the first time. Use the Reader Service Card to get further Information, or contact the manufacturers directly.

Design Tool Tackles Complex Math Problems

Visual Science 1.0 is a 32-bit application that lets users interactively design, simulate, analyze, and apply complex mathematical systems visually. Long available for Unix-based workstations, Visual Science 1.0 now brings the technique to Windows 3.1, 95, and NT personal computers. Using the software, designers can create a visual model of a mathematical system; bring out the important hierarchical structure of a system and hide unwanted details; and gain complete control over executing a mathematical system.

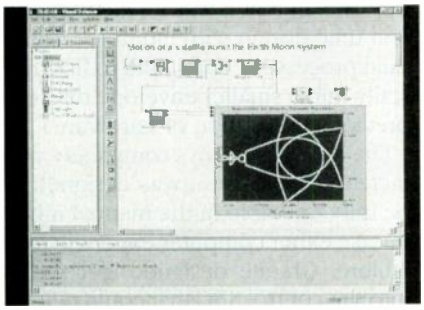

The software also makes it possible to analyze a system using interactive tools; scale from simple problems to complex dynamic systems; simulate parallel execution; and manage large projects.

Key features include the MathCalc matrix/array calculation language, seamless support for MATLAB and IDL, over 100 double-precision (64-bit) real and complex mathematical functions, support for multiprocessor hardware on Windows NT, and extensive on-line help. Visual Science 1.0 is expected to benefit technical disciplines such as computer, engineering, life, mathematical, physical, social, and statistical sciences. The $895 suggested retail price includes two floppy disks and hard-copy documentation. The company also offers a discounted fully functional version for qualified students.Potential users can download a trial version of Visual Science 1.0 by visiting acroScience's home page at: http://www.acroScience.com.

acroScience Corp., *1966 13th St., Suite 250, Boulder, CO 80302; (303) 541-0089 or 1 (800) 600-MATH. e-mail: info@acroScience.com*

CIRCLE 140

Software Adds Interactive Shopping To The Web

With the two latest software packages designed by Altia Inc., consumers can examine, operate, and compare product features on the World Wide Web without going to the store. Altia Design 2.0 is a software package that lets non-programmers simulate the features and behavior of a product prototype in an interactive computer model.

The toolset consists of a graphics editor, animation editor, stimulus editor, and control editor. The open architecture enables electronic prototypes to be linked to external applications developed in C, C++, and Microsoft Visual Basic programming environments. A run-time player allows distribution of electronic prototypes to third parties without royalty fees. Development and run-time platforms are available for Microsoft Windows 3.x, 95, or NT, SGI-Irix, Sun Solaris, Sun OS, IBM-AIX and HP-HPUX.

Altia's ProtoPlay is a Netscape plug-in for Altia Design that lets product developers post electronic prototypes on the Web to allow Internet users to interact with product features on-line. Consumers can provide feedback through on-line links to the manufacturer's Web-server environment. The end-user requires a 486-based PC and a 14.4-kbit/s modem. A complete Altia Design system costs $5900 for PCs and $9900 for Unix workstations. ProtoPlay is added as a plug-in at no additional charge.

Altia Inc., *5030 Corporate Plaza Dr., Suite 200, Colorado Springs, CO 80919; (719) 598-4299; Web: http://www.altia.com.*

CIRCLE 141

Flash-Disk Development Tools Build Embedded Applications

A suite of flash-disk solutions developed by M-Systems allows designers to select the level of complexity needed to integrate flash-memory-based, solid-state storage into their systems.

The tools are TrueFFS software, the LFDC-1016 Linear Flash Disk Controller, a FlashDisk chip set, and a developer's kit.

TrueFFS technology uses a unique block allocation method to provide total flash management and full disk emulation. It requires just 23 kbytes of memory to provide full disk emulation and has a sustained read speed of up to 3 Mbits/s. This makes TrueFFS a natural flash file system for both PC cards and embedded flash-memory solutions.

The single-chip LFDC-1016 controller comes bundled with TrueFFS software and supports the PCMCIA FTL standard. It provides full hard-disk emulation for up to 32 Mbytes of on-board flash, and is compatible with a range of operating systems, including DOS, Windows, QNX, and pSOS. Supporting both 8- and 16-Mbit NOR flash devices, the controller eliminates the need for glue logic and provides an ISA bus interface. No I/O address space is required, and only 10-kbytes of system memory window is needed. The window base address is user-selectable. The controller is housed in a 100-pin PQFP and comes with schematics for a reference design.

LFDC-1016 pricing for OEM quantities is $10 each, which includes a TrueFFS software license. A chip set solution consisting of the controller, TrueFFS software, and a flash memory component is available in capacities of 1 to 32 Mbytes.

The 2-Mbyte chip set goes for under $40 in OEM quantities.

Priced at $800, the Embedded TrueFFS Integrator's Kit (E-TIK) includes the following: one 4-Mbyte PC FlashDisk for the ISA bus, 10 TrueFFS licenses, and all of the information needed to design in and integrate an on-board flash disk.

M-Systems, *4655 Old Ironsides Dr., Suite 200, Santa Clara, CA 95054; (408) 654-5820.*
e-mail: info@ccm.msyscal.com..
CIRCLE 142

Development System Generates Motorola DSP Software

The Link-56K development system created by Domain Technologies eliminates having to use a different emulator for each of Motorola's 16- and 24-bit digital signal processors. The emulator is a PC platform tool with a source-level debugger running under Microsoft Windows. It links the PC and the DSP through the PC's RS-232 port and the DSP's OnCE or JTAG port.

The debugger has a windowed user interface for displaying up to 24 DSP resources. These include program, registers, command, trace, calls, stack, I/O, flags, watch, up to 10 data windows, direct memory access, cache, and view. Pull-down menus are used to configure and operate the debugger, and a toolbar provides quick command execution. A status bar displays the status of the DSP and debugger.

The debugger can read source code files, executable code files, and the information that links the executable code with the source code. Therefore, an assembly-language program can be debugged at the source level and displayed as it appears in the source text files.

True hardware breakpoints may be set on memory read, write, access, or fetch, and on a specific memory location or a range of memory locations. Link-56K supports full C debugging and complies with Motorola, BSO, and Tartan DSP C compilers. A block of memory can be displayed in hexadecimal, decimal, fractional, binary, or ASCII format. Resource windows, toolbar, colors, and fonts are user-customizable.

Software upgrades supporting new Motorola DSP releases are available at no cost through the Internet.

Domain Technologies Inc., *1700 Alma Dr., Plano, TX 75075; (214) 985-7593.*
e-mail: info@domaintec.com
CIRCLE 143

Design Tool Speeds Embedded Apps Development

CARDtools version 4.4 is a multifaceted design tool for solving design issues in real-time embedded systems such as cellular phones, hard disk drives, modems, PDAs, and multimedia set-top boxes. Simulation capabilities make it possible to review design options before implementing flawed or inadequate systems. The CAE-like technology helps reduce system development risks, design costs, and time to market while improving software quality.

Developers can use multitasking and concurrent design modeling techniques for quick application design and implementation. Device behavior modeling techniques are employed to co-simulate hardware with software, as well as between separate external hardware devices for full system-level simulation.

By creating simulation models using CARDtool's Tasking and Timing Application Simulator (TNT Sim), developers avoid timing and design constraint trade-offs such as missed timing deadlines, memory use, capacity overflows, poor scheduling, and race conditions.

TNT Sim also supports mixed-mode simulation, and allows the developer to analyze CPU and RTOS trade-offs on user-specific applications. The developer also can automatically generate standard ANSI C code with a push of a button and generate updated software documentation as the development progresses.

Code-generation capability includes C prototypes, header files, logic, pass-through, and user-selected RTOS calls. CARDtools 4.4 complies with ISO 9000 requirements, including design traceability.

Pricing depends on the configuration and number of copies purchased. An annual maintenance contract and

NEW PRODUCTS

customized technical support (on-site or off-site) are available.

CARDtools Systems Corp., *101 Metro Dr., Suite 250, San Jose, CA 95110-1314; (408) 894-9500; e-mail: cardsb@aol.com.*

CIRCLE 144

Development Tools Span All Classes Of Embedded Systems

The PDOS PowerSuite development environment for embedded systems comprises a complete 32-bit Windows-based development toolset plus the scalable PDOSpro real-time multitasking operating system. For designs with limited memory requirements, the operating system can be configured to a minimum size of 2 kbytes. For systems with larger memory requirements and resources, a robust set of Installable System Modules (ISMs) is available.

Except for the kernel, all components of the PDOS PowerSuite and PDOSpro are ISMs.

Therefore, the end product has no unnecessary operating system overhead depleting valuable resources. With PDOSpro, even the scheduler is considered as an ISM.

Because scheduling requirements differ between applications, PDOSpro offers a number of different scheduling routines. To assist system developers in integrating the various components of the PDOSpro kernel and operating system, the PDOS PowerSuite includes a tool called PDOS Build.

With this tool, users can automatically select and link the required operating-system components by using a mouse button under Windows 95 or Windows NT.

Other tools include PDOS ViewPort control center for file management and terminal emulation, CodeWrite editor, C and C++ compiler links, source-level debugger links, and on-line documentation. Microsoft compiler support will be available for 80x86 and Pentium processor applications. Initially, PDOS PowerSuite will support 80x86 and PowerPC processors.

Pricing starts at $3000 for a single seat license. Multiple seat licenses are available as well.

Eyring Corp., *6912 South 185 West, Midvale, Utah 84047; 1 (800) 937-7367.*
e-mail: pdos-info@eyring.com.

CIRCLE 145

Design Tool Co-Simulates Hardware And Software

CARDtools version 4.4 is targeted at solving design issues in real-time embedded systems, such as cellular phones, hard disk drives, modems, PDAs, and multimedia set-top boxes. Using this CAE-type technology, developers can use multitasking and concurrent design modeling techniques, as well as co-simulate hardware and software for full system-level simulation.

Other functions include the ability to perform CPU and RTOS trade-off analysis on user-specific applications, automatically generate standard ANSI C code, and automatically generate updated software documentation as the development progresses. Code generation includes C prototypes, header files, logic, pass-through, and user-selected RTOS calls.

CARDtools eliminates timing and design constraint trade-offs, such as missed timing deadlines, memory use, capacity overflows, deadlocks, poor scheduling and race conditions. This is accomplished through simulation models that use TNT Sim (Tasking and Timing Application Simulator).

TNT Sim also supports mixed-mode simulation. Furthermore, external C routines can be included in a task behavior model. The technology promotes compliance to ISO 9000 requirements, including design traceability.

CARDtools version 4.4 runs on Sun workstations and operates in a client-server X-Window and Motif environment.

Pricing depends on the configuration and number of copies purchased. An annual maintenance contract and technical support (on-site or off-site) also are available.

CARDtools Systems Corp., *101 Metro Drive, Suite 250, San Jose, CA 95110-1314; (408) 894-9500.*
e-mail: cardsb@aol.com.

CIRCLE 146

ADVERTISERS INDEX

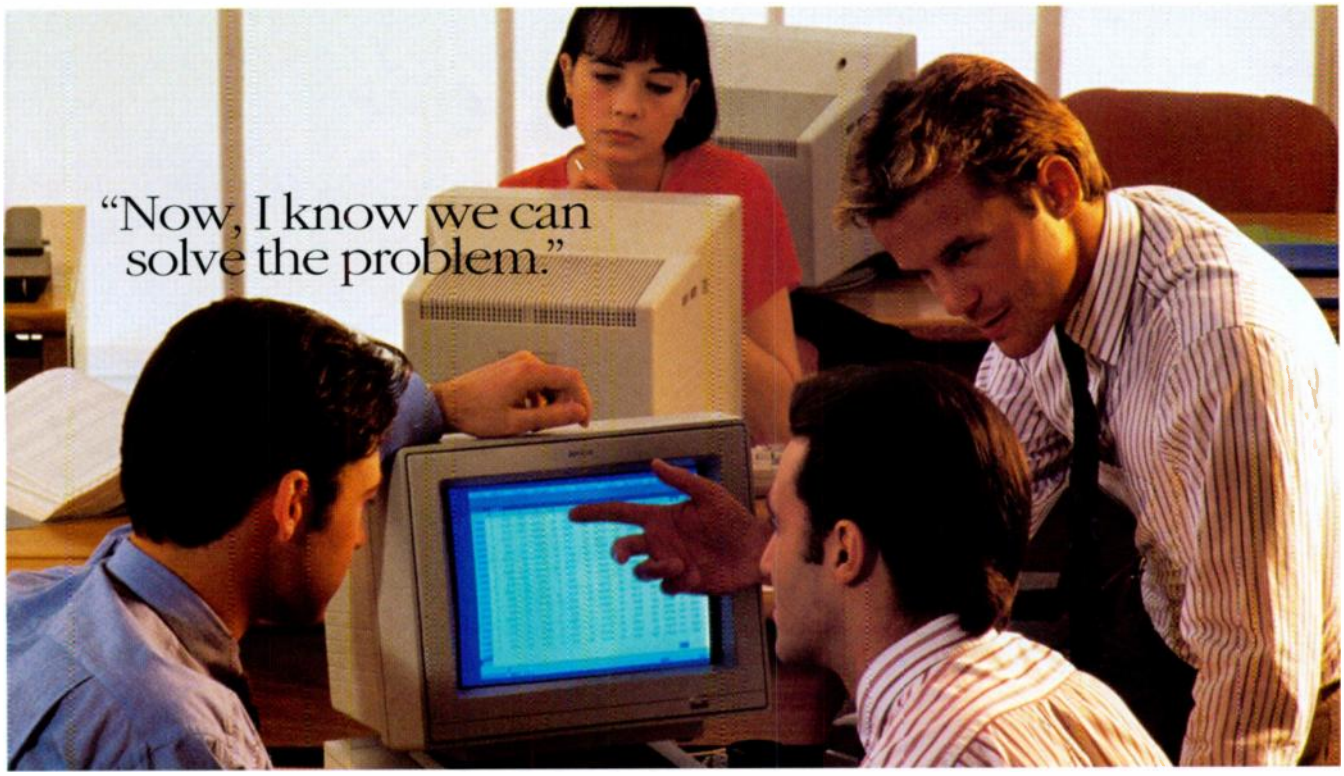

Introducing Three New Pentium® Tools

PowerPack® EA-Pentium®
High Performance.

PowerPack® SW-Pentium®
Fast Solutions.

PowerPack® ITP-Pentium®
In-target Probe for Software Development

These new Pentium® in-circuit emulators show software and hardware events that are invisible with any other tools.

Software debuggers lack hardware event triggers and full-speed trace, making it difficult to isolate and identify real-time conflicts.

NEW **SWAT™ Software Analysis Tool**

Now you can check code coverage and monitor performance without inserting instrumentation tags into your source code.

Microtek's unique SWAT™ software analysis tool offers the same transparent interactive control you get with our high performance emulators.

Code Changes

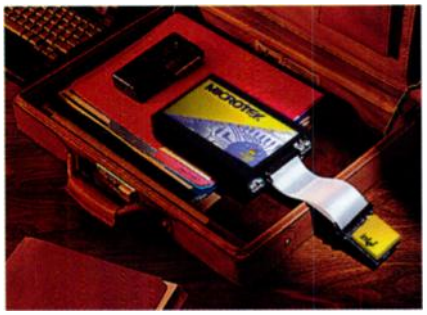

"I was surprised to find that Microtek already has three different Pentium tools that fit into my briefcase."
Gary Rains
Development Systems Director

Clock-Edge Event Triggers
Microtek PowerPack® High Performance Emulators relate hardware events back to the source code with clock-edge resolution.

160-Bit Wide Trace
Now you can see what is happening in real-time without stopping the target. With 160-bits wide by 256k frames deep trace, virtually no target activity goes unrecorded.

Shrinking Probe Head
The new PowerPack® EA-Pentium emulator is only 7.2" x 4.6". And the probe tip is only 3" x 1.9", barely larger than the target microprocessor!

Microtek Emulators available for:

Pentium® • Intel486™ • National NS486™ AMD AM486™ • Intel386™EX • 386DX 386CX/SX • 80C186

68360 • 68340 • 68F333 • 68332 68331 • 68330 • 68HC16 • 68328 • ColdFire

To find the solutions to your problems, please call:

1 (800) 886-7333

(503) 645-7333

Email: info@microtekintl.com

Fax: (503) 629-8460

MICROTEK
IN-CIRCUIT EMULATORS

Web: www.microtekintl.com

READER SERVICE 89

Banner Solutions for PowerPC on VME!

V•I Computer has a wealth of experience in PowerPC based VME boards.

Breadth of product offering and ability to customize designs for your specific requirements make V•I Computer the first place to call when you need Power and Performance on VME.

The PowerPC offers you unsurpassed VME Price/Performance capability.

More PowerPC Processors let you select performance level

V•I Computer has a PowerPC solution for every system performance need. Choose from models ranging from the economical Power•2™ (33-MHz 403 processor) to the lightning-fast Power•4B™ (166-MHz 604 processor).

PCI Mezzanine board for high bandwidth I/O

There's finally an *open standardized* mezzanine bus for VME with multiple vendor support. PMC cards offer faster throughput than any of the earlier mezzanines. All models above the Power•2 have PMC mezzanine slots. The Power•2 supports IP or M modules.

Software Support

All boards have Open Firmware (IEEE-1275 compliant). Support is available for many real time and UNIX OS's, .

Call V•I today, and get the whole story.

1-800-VME-CPUs

(800) 863-2787

Model	Processor Speed (MHz)	Max Memory (MB)	
Power•4B	166	64	Ultra SCSI
Power•4A	100	32	100 MB Ethermet
Power•4	100	256	100 MB Ethernet
Power•3e	100	256	Low cost SBC
Power•3ec	100	16	8 MB Flash
Power•2	33	16	Lowest cost PowerPC

V•I computer
A Subsidiary of Themis Computer

531 Encinitas Blvd., Bldg 114
Encinitas, CA 92024
Tel. (619) 632-5823 • Fax 619-632-5829
e-mail: sales@vicomp.com

READER SERVICE 99