

leading edge

What's hot
in the
design
community

Edited by
Fran Granville



Yeah, but, where are you *now*?

"We now have people who perform running commentary on the minutiae of their daily lives as they travel down a grocery line—'Honey, I'm at the avocados'—or round the corner. A favorite *New Yorker* cartoon tracks a cell-phone narrator through three panels: 'I'm boarding the train. I'm on the train. I'm leaving the train.'"

—Ellen Goodman,
The Boston Sunday Globe,
June 13, 2004

CompactPCI computer breaks \$1000 barrier

By Warren Webb

TARGETING INDUSTRIAL, transportation, and military applications, the new ICP-CM single-board computer from Inova Computers delivers midrange performance at an entry-level price.

Inova based the product on the Intel Celeron M microprocessor, and it offers low-power operation, a minimum five-year availability, and a MTBF of more than 200,000 hours in a 3U CompactPCI package. In its basic single-width configuration, the board provides one USB 2.0 interface, a pair of independent Fast Ethernet ports, and VGA graphics.

The board also includes a CompactFlash socket for a flash-memory card or a rug-



The ICP-CM single-board computer offers low-power, high-reliability performance in a 3U Compact PCI form factor.

gedized Micro-drive. The ICP-CM is available in double- and triple-width versions to extend its suite of storage and I/O capabilities. It accommodates either a 600-MHz or a 1.3-GHz Celeron M microprocessor and comes with 256 Mbytes of DDR 266 SDRAM, expandable to 1 Gbyte. The 600-MHz Celeron M version dissipates a maximum of 9W, and the 1.3-GHz version dissipates 22W. The ICP-CM can host Windows, Linux, and a range of real-time operating systems, such as VxWorks and QNX. The ICP-CM sells for \$995 (100).

► **Inova Computers Inc**, 1-508-771-4415, www.inovacomputers.com.

TDM operates over PSN without headaches

ZARLINK SEMICONDUCTOR's new ZL50111 family of TDM-over-IP packet processors meets the protocol requirements in the recently announced ITU-T (International Telecommunications Union Telecommunications) TDM-MPLS (Multiprotocol Label Switching) networking spec. The specification, called User Plane Interworking, Recommendation Y.1413, outlines the TDM timing, signaling, service quality, and alarm integrity necessary to carry legacy voice, data, and multimedia traffic over an MPLS network at up to DS3 or E3 rates.

Service providers hoping to consolidate networks and reduce operating costs can use MPLS switches to carry network traffic, such as voice-over-packet, over a converged PSN (packet-switched network) and still cost-effectively transport TDM-based traffic with guaranteed quality of service.

The ZL50111 family complies with both the structure-aware and structure-agnostic modes of operation that the recommendation outlines and uses CESOP (circuit-emulation-services-over-packet) technology to tunnel many forms of Layer 2 TDM traffic over the PSN. The TDM-over-IP packet processors maintain critical timing and service-quality requirements by guaranteeing less-than-1-msec latency through the devices with accurate clock recovery and synchronization. On-chip queuing mechanisms minimize the effect of network latency by prioritizing time-sensitive TDM packets over data packets in processing queues. Currently in production, the ZL50111 sells for \$276.47 (1000).

—by Nicholas Cravotta

► **Zarlink Semiconductor**, 1-613-592-0200, www.zarlink.com.

Tool provides hierarchical verification

DESIGNERS KNOW that using a hierarchical approach to circuit design improves the quality of results and often decreases development time.

Verification engineers often separately verify blocks

of a design before integrating them into circuits to perform validation tests on the integrated system. Formal-verification methods are still evolving, and engineers have little tool support in using hierarchical methods for formal verification. In many cases, formal verification tools cannot handle large designs, so engineers have to manually

keep track of verified blocks and invent their own integration-verification procedures.

Jasper Design Automation has addressed these problems with the JasperGold 3.0 tool, which supports “provably correct design” methodology. This approach to verification advocates early application of formal methods before designers check in RTL blocks.

Users can specify a high-level requirement of the block under creation and then use JasperGold to either find a sequential combination of inputs that violates the requirement or confirm that the design meets the requirement under all possible circumstances. In case of failure, JasperGold returns the scenario that would trigger the failure and enables users to find the cause of the design error. Designers can then stitch together formally verified blocks and have only to verify that they have properly implemented the interface protocols among the various blocks.

Designers can interactively use the tool, and verification engineers can set up batch execution of the product for regression testing. Combining JasperGold with Formal Testplanner provides a complete environment for hierarchical formal verification.

Price for JasperGold 3.0, including Formal Testplanner, is \$225,000 for a one-year time-based license for the base engine and \$45,000 per user license.

—by Gabe Moretti

► **Jasper Design Automation**, 1-650-966-0200, www.jasper-da.com.

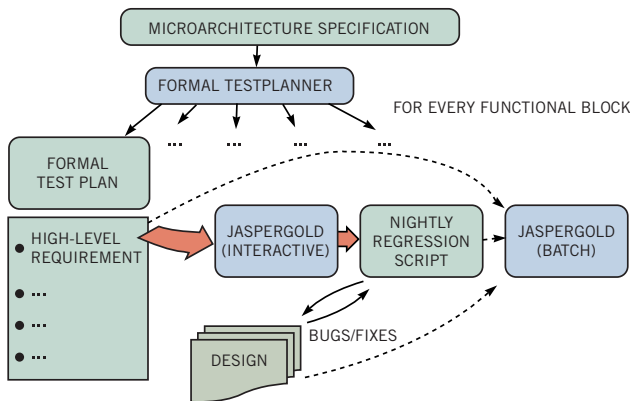
MOVE AWAY FROM FIXED HARDWARE TO SPI-4.2

Fulcrum Microsystems’ new PivotPoint SPI-4.2 switch chip targets high-performance networking and storage applications. The six-port FM1010 PivotPoint SPI-4.2 seamlessly interconnects devices such as transceivers, network processors, traffic managers, coprocessors, and custom ASICs and FPGAs. By employing an SPI-4.2 switch, developers can eliminate the use of discrete daisy-chain designs and system buses, moving from a fixed hardware configuration to a dynamic fabric of computing and packet-processing resources.

Based on Fulcrum’s Nexus technology, a nonblocking terabit crossbar circuit, the FM1010 offers data rates as high as 14.4 Gbps per SPI-4.2 interface with a ball-to-ball latency of less than 250 nsec. Available now, the FM1010 comes in a 1036-ball BGA package and sells for \$225 (1000). A hardware-development kit, including a modular board, reference software, and a design guide, simplifies evaluation of the chip with other SPI-4.2 devices.


—by Nicholas Cravotta

► **Fulcrum Microsystems Inc**, 1-818-871-8100, www.fulcrum-micro.com.



JasperGold 3.0 supports “provably correct design,” which advocates early application of formal methods before designers check in RTL blocks.


DILBERT By Scott Adams



► With more than 260 million mobile subscribers and more than 4 million new subscribers every month in 2003, China is the largest cellular market in the world and a potential hotbed of 3G activities, according to InStat/MDR.

ZigBee now comes in a box

ATMEL HAS introduced a controller, radio, and protocol-stack implementation of ZigBee, a low-data-rate, low-power networked-wireless standard for control and monitoring applications. Atmel based ZigBee on the IEEE 802.15.4 standard, which stipulates a 20-, 40-, or 250-kbps data rate and that devices using it should run for at least two years without a battery change. Target applications include environmental control, security, industrial sensors, and medical monitoring systems that typically transfer small amounts of data with a low duty cycle. Other applica-



Atmel's ZigBee offering includes an 802.15.4 radio, a ZigBee-specific controller, and both 802.15.4 and ZigBee protocol stacks.

tions include home and office lighting, although ZigBee currently finds use only in high-end systems because of price.

The 868-MHz/902- to 928-MHz AT86RF210 Z-Link 802.15.4 radio supports 10 40-kbps channels in the 915-MHz ISM (industrial, scientific, and medical) band and one 20-kbps channel in the 868-MHz European band. The radio uses a DSSS (direct-sequence-spread-spectrum) interface, including spreading and despreading, using BPSK (binary phase-shift keying). The PHY (physical-layer) fea-

tures include receiver-energy detection, link-quality indication, and clear-channel assessment, and the MAC (media-access controller) supports an optional superframe structure with beacons for time synchronization and a guaranteed-time-slot mechanism for high-priority communications.

The radio has a receiving current of 14.5 mA; a transmitting current of 4 mW at 6 dBm and 1.8V; a sleep mode of 1 μ A; and the ability to support as many as 65,000 nodes in star, cluster, or mesh configurations at a range as long as 100m.

The full-function AT86Z-13201 Z-Link ZigBee-specific controller routes networks and coordinates links. It meets the memory, security, and peripheral requirements of ZigBee to both decrease costs and optimize power consumption. On-chip memory includes 8 kbytes of ROM preloaded with the stable 802.15.4 protocol stack, 32 kbytes of flash for storing the still evolving ZigBee networking and application-protocol stacks with the appropriate profiles and application code, and 8 kbytes of SRAM, providing headroom for any increases in ZigBee data structures or to support larger networks.

Using ROM for the fixed code 802.15.4 stack reduces the nonvolatile memory portion of the die by 15%, helping to reduce costs. The controller has SPI and buffered transmitter/receiver serial points to interface to 802.15.4 radios. Atmel based the device on the 8-bit AVR RISC processor, so developers can use standard AVR-based code-development, debug-

ging, and emulation tools for application development.

The Z-Link controller also integrates 128-bit symmetric key AES (Advanced Encryption Standard) encryption hardware with an enhanced counter-cipher block-chaining message-authentication-code mode and hardware pseudorandom-number generator. Implementing encryption in hardware reduces overall latency to less than 30 msec and reduces power consumption compared with

software-based encryption.

Protocol support includes both the 802.15.4 and the ZigBee protocol stacks in a verified implementation, enabling developers to focus design efforts on application development. Currently available for sampling, the radio, full-feature-device controller, and both 802.15.4 and ZigBee protocol stacks sell for \$6.74 (100,000).

—by Nicholas Cravotta

► **Atmel Corp.**, www.atmel.com.

MEMS-DESIGN SOFTWARE PROVIDES EDA APPROACH

The integration of electronic and mechanical parts of a system has always been a challenge for designers. Normally, teams using different tools design the two portions of the system, and the integration often occurs only after the vendor fabricates the devices. Developing a computer model of a system requires using approximations that can lead to inaccurate simulation results when dealing with nanometer processes. Yet, the drive to increase the integration and miniaturizations of systems in products such as automobiles and cameras is increasing the need for system simulation before the mechanical and electronic components go to manufacturing.

To address these challenges, Coventor, a provider of tools for the development of MEMS (microelectromechanical systems), microfluidics, and semiconductor-process applications, has released the CoventorWare suite. The tools allow developers to combine MEMS models with electronic models they designed in either Cadence (www.cadence.com) Virtuoso or Synopsys (www.synopsys.com) Saber environments. CoventorWare comprises products to create, model, analyze, and integrate MEMS devices.

The Architect, Designer, Analyzer, and Integrator tools support the development of sensor, actuator, RF, and microfluidic MEMS. The Architect library contains new models, including-beam actuator models that cover thermally or piezoelectrically actuated devices, such as mirrors, and membrane models for modeling pressure sensors. Designer includes enhanced layout-editor features, including design-rule checks for small edges and file-transfers options.

Analyzer has new capabilities, including postprocessing options for designing piezoelectric RF resonators and enhanced electro wetting on dielectric and microheater analysis for microdroplet applications. You use Integrator to export reduced-order macromodels of MEMS to the Virtuoso or the Saber environment.

Prices for CoventorWare range from \$25,000 to \$100,000, depending on the configuration. The tools operate on both Windows and Solaris operating systems.—by Gabe Moretti

► **Coventor**, 1-919-657-8114, www.coventor.com.

Dueling graphics juggernauts trade announcement, improvement jabs

WITH NVIDIA'S latest graphics flagship now disrobed and in feisty fighting shape, you had to know that ATI Technologies' counterpunch wouldn't be too long in arriving (see "Graphics advancements span PCs to cell phones," www.reed-electronics.com/edn/mag/article/CA412448). ATI's mainstream X800 Pro and high-end X800 XT Platinum Edition devices might strike you, at first glance, as nearly identical to their Nvidia competitors, the GeForce 6800 and 6800 Ultra.

Both vendors' mainstream parts have six vertex-shader units and 12 pixel pipelines in four-pipeline-per-group granularity with two pixel-shader units per pixel pipeline. Both high-end parts bump the pixel pipeline count to 16 and interface to a 550-MHz GDDR-3 SDRAM over a 256-bit memory bus with 1.1-Mbps-per-pin bandwidth driven by a four-way crossbar memory controller.

ATI and Nvidia claim that all their devices hardware-accelerate MPEG and Windows Media Video decoding along with MPEG encoding, although ATI accomplishes the tasks in multifunction shaders, whereas Nvidia employs a dedicated video processor. Peer closer, though, and some significant differences emerge, differences that,

at least in the near term, will likely play in ATI's favor.

ATI has consciously chosen to not yet support DirectX 9.0c's Shader Model 3 instructions and their 32-bit precision, instead continuing to rely on the claimed "more-than-good-enough-for-now," 24-bit precision that the company pioneered in the previous generation Radeon 9700 family and Microsoft supported in DirectX 9.0b and Shader Model 2. This decision has enabled ATI to construct its X800 devices from "only" 160 million transistors versus 222 million, or 39% more, for Nvidia. Couple that fact with the ATI devices' low-K, 0.13-micron fabrication process, and the result is that the X800 XT Platinum Edition runs at a 520-MHz core frequency—30% faster than the Nvidia GeForce 6800 Ultra—and burns only an estimated 65W of power when running Fu-

turemark's (www.futuremark.com) 3DMark03 benchmark suite—60% of the 6800 Ultra's estimated power draw.

Citing the high cost of 4 Mbit×32-bit memories required to construct 128-Mbyte frame buffers with a 256-bit aggregate-bus interface, ATI plans to offer no X800 board variants that employ them. The X800 Pro, running at a 475-MHz core frequency and with 256 Mbytes of 450-MHz DDR SDRAM, costs \$399, and the company is now shipping it. The premium X800 XT Platinum Edition costs \$499 and hit store shelves in May.

At the 11th hour, Nvidia pulled two additional GeForce 6800 stock-keeping units out of its hat in a competitive response. The \$399 6800 GT runs at 350-MHz core and 500-MHz memory speeds with a 256-Mbyte frame buffer. (Recall that the \$299, conventional 6800 board offers only a 128-Mbyte frame buffer and reportedly runs at 375-MHz core and 350-MHz memory speeds.) Nvidia also announced the 6800 Ultra Extreme, with unspecified higher-than-6800 Ultra pricing and clock rates.

Nvidia's first wave of GeForce 6xxx desktop-graphics introductions is now out of the way, and the company has turned its attention to preparing the industry for upcoming mobile

variants of the architecture. At the same time, Nvidia is responding to those of you who want to upgrade your notebook PCs' graphics subsystems but without locking yourself into the proprietary-graphics upgrade modules from companies such as Alienware (www.alienware.com) and Dell (www.dell.com). Nvidia's proposed MXM (Mobile PCI Express Module) standard, which it hopes that other graphics vendors will also embrace, comes in three module sizes commensurate with graphic-chip dimensions and frame-buffer densities. It builds on the proven, cost-reduced SO-DIMM socket, bumping up the connector's contact count to 230 pins and beefing up its sturdiness and current-carrying capability to reflect the larger, heavier, and hungrier graphics modules.

What does Nvidia plan to do with all those contacts? Among other things, they'll provide sufficient power-supply flexibility to feed numerous graphics- and memory-chip subsystems with varying voltage and tolerance needs. Also, MXM handles a mind-boggling array of simultaneous displays: dual-link, 24-bit LVDS, VGA, S-Video, composite video, component video, and two independent DVI ports.

—by Brian Dipert

► **ATI Technologies**, 1-905-882-2600, www.ati.com.

► **Nvidia**, 1-408-486-2000, www.nvidia.com.



Nvidia's MXM aims to deliver industry-standard graphics flexibility to notebook-PC builders and buyers.

► **The Labor Department reported in June that its consumer-price index, a common measure of inflation, posted its strongest monthly advance since January 2001, rising 0.6% in May. Through the first five months of the year, inflation is running at a 5.1% annual rate, compared to 1.9% in all of 2003.—The Boston Globe, June 16, 2004**

Microprocessors offer connectivity options

FREESCALE Semiconductor's four new ColdFire families—the MCF547x, MCF548x, MCF523x, and MCF527x 32-bit microprocessor families—comprise more than 20 32-bit microprocessor devices. The MCF547x and MCF548x families are the first to implement the V4e ColdFire core at speeds as high as 266 MHz. These devices include an MMU, an FPU, as many as two 10/100 Ethernet controllers, a 32-kbyte SRAM, 32-kbyte instruction and data caches, and a DDR/SDR-memory controller in a 388-ball PBGA package supporting industrial-temperature ranges.

The MCF548x adds support for dual CANs (controller-area networks). Other optional modules include

USB 2.0 high speed with integrated PHY (physical layer) and hardware-accelerated encryption to support SSI (synchronous-serial interface), SSH (secure shell), or IPsec (Internet Protocol security). All of the family devices are pin-compatible.

Freescall based the MCF523x family on the V2 ColdFire core, operating as fast as 150 MHz. It targets industrial- and motor-control applications that require connectivity. This family is the first to implement the eTPU (enhanced time-processing unit), a programmable I/O controller with as many as 32 channels. The eTPU has its own core and memory system that enables it to perform complex timing and I/O-management independently

of the CPU. This family of devices integrates 64 kbytes of SRAM, 8 kbytes of configurable instruction/data cache, an enhanced MAC (multiply-accumulate) unit, hardware-cryptography acceleration, a 32-bit nonmultiplexed bus with as many as eight chip selects, Ethernet, and one or two CAN interfaces. The MCF523x family can provide higher performance and networking support for Freescall MC68332 users.

Freescall based the MCF527x family on the V2 ColdFire core, operating as fast as 166 MHz. It targets cost-sensitive applications that require control processing for file management, connectivity, data buffering, user interface, and signal processing for security, imaging, networking,

and gaming applications. MCF5272 users can upgrade to these devices to gain additional performance, a second 10/100-Mbps Ethernet MAC, hardware encryption, and a 16-bit DDR-memory controller. The MCF527x family includes options for USB 2.0 full-speed-device and hardware-accelerated encryption to support SSI, SSH, or IPsec.

The MCF34710 companion power-management device supports user-adjustable voltages to enable you to reduce the power dissipation by removing or lowering voltages under software control. All of these devices are available now, and prices range from \$7 to \$27 (10,000).

—by Robert Cravotta

► **Freescall Semiconductor**, www.freescall.com.

Devices allow transparent expansion of PCI and PCI-X to PCI Express

PLX TECHNOLOGY has announced its first three PCI Express devices, the PEX 8516 and PEX 8532 switches, supporting 16 lanes with four ports and 32 lanes with eight ports, respectively, and the PEX 8104-to-PCI/PCI-X bridge. The flexible port configuration of the switches supports a mix of fan-out, fan-in, and peer-to-peer applications. Nontransparent-port capacity, compatible with the PCI Express and Advanced Switching architectures, enables multihost, load-sharing, and high-availability systems.

Key features of the switches include a fully nonblocking fabric; an integrated serializer/deserializer; peer-to-peer transfer, allowing I/Os to transfer data without host involvement;

two virtual channels to support quality-of-service features, such as priority handling and bandwidth management; advanced error reporting; symmetric and asymmetric port combinations; both link and end-to-end CRC (cyclic redundancy check); hot-plugging; and JTAG capabilities.

The four-lane 8104 bridge enables interoperability to PCI and PCI-X subsystems. Nontransparency during forward bridging enables intelligent-adaptor and multihost applications. Reverse bridging avoids fork-lift upgrades by allowing legacy subsystems to control the advanced features of PCI Express. This feature allows PCI and PCI-X motherboards to directly use native PCI Express components and adaptors and facilitates seamless

expansion of I/O subsystems via PCI Express cabling.

The 32-lane 8532 switch consumes 6.5W, comes in a 35×35-mm HSBGA package, and costs less than \$100. The 16-lane 8516 switch consumes 3.5W, comes in a 27×27-mm HSBGA package, and costs less than \$100. Samples are available, and volume production is slated for the fourth quarter of this year. The four-lane 8104 bridge consumes 2W, comes in a 17×17-mm PBGA package, and costs less than \$30. It will become available for sampling in August, and volume production is slated for the fourth quarter.—by Nicholas Cravotta

► **PLX Technology**, 1-408-774-9060, www.plxtech.com.

► **Keynote Systems Inc**, which monitors performance of the Internet's 40 most popular sites, found that average availability of the top sites dropped from 100 to 80% during an attack against Akamai Technologies, the leading provider of Web-content-delivery services.

—*The Boston Globe*, June 16, 2004

VME board hosts 30 million gates of configurable resources

by *Graham Prophet*

RECONFIGURABLE-PROCESSING specialist Nallatech has broadened the range of hardware platforms it offers based on the latest Xilinx (www.xilinx.com) Virtex II Pro FPGAs.

You can equip VME systems with processing power into the Teraflop range and multigigabit bandwidth, using the BenNuey-VME card. The base card contains one Xilinx 2VP50 chip, plus interfaces, including Gigabit Ethernet, InfiniBand, Fibre Channel, and a bridge to PCI. A back-panel VXS port speeds VME-to-VME interfacing.

You can add as many as three of Nallatech's Dime-II modules, each of which can host a Virtex 2VP100 FPGA, on the base card. With its oth-

er resources, a single card could therefore host 30 million gates of configurable logic, 750 Mbytes of DDR SDRAM, 50 Mbytes of ZBT (zero-bus-turnaround) SRAM, and more than 40 Gbps of off-card bandwidth. System-level clocking allows you to partition application tasks across this resource and across the multiple FPGAs.

You can map custom-processing structures built on specific tasks and algorithms onto this platform, using configuration tools, such as The Math Works' ([\[www.mathworks.com\]\(http://www.mathworks.com\)\) Matlab and Simulink, Java, C++, Handel C, Xilinx system generators, and standard HDLs. Nallatech provides runtime support via its Dimetalk and Fuse software. At the implementation level, it also supports standard synthesis tools and Xilinx's ISE.](http://www.math</p></div><div data-bbox=)

The latest daughtercard is the BenData-II, carrying a Virtex FPGA with SRAM and DRAM. Eight Rocket I/O connections add to the off-card bandwidth. You can use the architecture in any situation having high processing loads that focus on specific algorithms, such as medical-image processing and software radio.

► **Nallatech**, +44 1236 789500, www.nallatech.com.

Recycle power for I/O efficiency

IN CUSTOM-IC DEVELOPMENT, controlling power demand is an ever-more-important priority in handheld products to maximise battery life. I/O drivers expend a large proportion of that power,

especially if the chip is I/O-intensive and must send high-data-rate transactions to external memory or peripherals. IP (intellectual property) from Adiabatic Logic addresses that problem with a “power-recycling” I/O cell that can reduce the overall power requirement for driving external lines by as much as 75%. “Adiabatic” implies a process that takes place without the loss or the addition of energy.

Adiabatic Logic’s IOD (intelligent output driver) provides a means of dealing with the characteristics of a transmission line; normal practice is to terminate the line with its characteristic impedance—a resistor—to absorb energy and avoid reflections. To send a rising-edge signal, you switch the sending end of the line to the rail voltage, which launches a step of the same amplitude onto the line. If you do not terminate the line at the far end, it reflects back

and, after due time, arrives back and gives 2V at the sending end (for a line voltage of “V”), a situation that you want to avoid. The IOD’s output is a three-terminal switch; it begins sending a rising edge by connecting the line to a capacitor, which has been precharged to 0.5V—that is half the line voltage—launching a rising edge of that amplitude. The line is unterminated; the reflected wave arrives back at the transmitter and adds to give twice the voltage launched into the line. At that instant, the sending switch connects to the line voltage. Charging the capacitance of the line and load and driving the load impedance itself are the only actions that expend energy. At the falling edge, the reverse switching takes place, first to 0.5V on the capacitor and then to ground at the appropriate instant. As the current sense reverses, energy returns to the capacitor.

A delay-lock loop, based on an up/down counter, sets the timing of the commutation of the three-terminal switch. The counter monitors the time of flight of the reflected wave. The silicon process—and, hence, the device speed—you use relates to the minimum line length you can match. Adiabatic officials say that, for 90- to 130-nm-silicon geometry, this length equates to about 1 to 2 cm of pc-board track.

Because the technique accurately mimics the effect of a properly terminated line, you also receive the benefit of reduced RFI. As an example of potential power savings, the company measured one handheld product and found that its interface from processor to SDRAM alone accounted for 100 mW, representing a possible reduction of 75 mW. Because the IOD has small overhead, this reduction can be cost-effective, Adiabatic asserts. The company is marketing the IOD as IP, and IMEC (www.imec.be) has incorporated it into its aerospace-process library.

—by Graham Prophet

► **Adiabatic Logic**, +44 1954 211244, www.adiabaticlogic.com.

Custom-silicon programme targets low-volume projects

THE DESIGN OF complex ASICs is becoming increasingly difficult for small and mid-sized enterprises, especially in which the final volumes of chips are also small to mid-sized. With its “high-way-to-silicon” programme, design house Accent aims to offer an affordable route to achieving that objective.

Accent echoes the input of other IC-design service companies in its assertion that complex-chip design is becoming almost an art—one in which the average design team performs too few projects to gain full mastery of the process. Therefore, Accent offers a “disaggregated” design service that offers the customer full IP (intellectual-property) protection and ownership; the customer can then transfer or migrate designs to different foundries or

processes if necessary. The company adds a contract-manufacturing package, managing production of chips through delivery of samples, volume-production units, or both. The package takes advantage of the multichip-wafer service that various foundries offer and can economically deliver high volumes at the prototyping level. It encompasses design service; wafer-fabrication outsourcing; and test, assembly, and packaging outsourcing.

Accent, which has its roots in a joint venture of STMicroelectronics (www.st.com) and Cadence (www.cadence.com), designs in a range of logic and mixed-signal technologies, including RF. It has also built ICs that incorporate on-chip FPGAs and extends the design coverage to system-in-package or multichip modules and

to system and software design if necessary.

Part of the offering is a business unit that can act as a source of IP from any provider; the company has access to the full range of STMicroelectronics’ in-house IP. The programme is flexible enough, the company claims, to handle both low-volume business for sectors such as industrial and aerospace and, at the other extreme, to design the most complex multimillion-gate chips for the catalogues of fabless semiconductor houses. The attractions of the capabilities of today’s silicon processes remain compelling, Accent says, “but all the inherent risks need managing,” according to a company spokesman.

—by Graham Prophet

►Accent, +39 039 6290126, www.accent.it.

FPGA soft cores reach 200 MIPS

WITH ITS SECOND-GENERATION soft-core-embedded Nios II processor, Altera has opted for a series of three implementations each targeting specific usage models. You can select versions that optimize for maximum performance at 200 Dhrystone MIPS, minimum logic usage, or a blend of the two. You use approximately 1800, 550, or 1100 logic elements, respectively, for the cores; all remain fully code-compatible with each other. You need to recompile code written for Nios I. Altera calls them, respectively, fast, economy, and standard. This small range of basic operations represents a compromise on the configurability of the core, which Altera intends to reduce the complexity of testing and verification versus the previous version.

Altera says its customers are using Nios for a range of functions, from simple housekeeping to datapath and multicore, parallel processing, as well as for flexible hardware/software partitioning. The design continues to be a 32-bit, pipelined-RISC core, configurable with a library of peripherals. You get features such as DSP-hardware functions in the fast and standard versions. You can also define custom instructions, which the system-level-design aspects can import and recognize, yielding, in effect, a build-your-own hardware accelerator.

A full Eclipse-based development environment, with programming and debugging via a JTAG interface, supports the core. Altera says that it paid extra attention to support for software development, reflecting the more sophisticated uses it expects for the processor. A new JTAG UART interface complements a Compact Flash interface for mass storage. You also get an RTOS from Micrium, for which a point-of-shipping licence is negotiable. Nios itself comes with a royalty-free, perpetual licence.—by Graham Prophet

►Altera, +44 1494 602000, www.altera.com.

VIRTEX-4 MAKES ITS ENTRANCE

Xilinx has formally announced the widely trailed Virtex-4 series of FPGAs, building on the ASMBL (application-specific-modular-block) architecture it introduced this year. The architecture employs “stripes” of function-specific circuitry that Xilinx then assembles into a chip that can meet the needs of a particular sector. The contents of a stripe can be I/O drivers, programmable logic cells, mixed-signal functions, or whatever else is necessary.

Accordingly, Virtex-4 initially comes in three flavours: LX for logic, SX for signal processing, and FX for embedded processing and serial connectivity. Although the fundamental FPGA logic is largely the same as that of Virtex-II-Pro, Xilinx says that density is twice and speed is as much as twice those of the previous family.

The LX family focuses on logic and memory, with proportionally fewer resources for clock structures and DSP-function blocks, whereas the SX signal-processing platform is the opposite. The FX offers a balanced mix of features, adding blocks for signal transceivers and processors. The Power PC cores that appeared in the previous generation appear here, as well, although there are now just two of them. Most critical-path logic, memory, DSP, and clock functions operate as fast as 500 MHz, and RocketIO-serial and SelectIO-parallel drivers handle fast I/O. The latest versions of Xilinx’s software support the new family, and silicon samples are expected in mid-2004.—by Graham Prophet

►Xilinx, 44 870 7350 600, www.xilinx.com.

►Sony was the third-largest handheld-system maker in 2003, shipping 1.4 million Clies into a worldwide market of about 10.7 million units, according to IDC. The Japanese market accounted for only 484,000 Clie shipments in 2003.

Nanoscale MEMS challenge flash memory

IF YOU THINK the concept of MEMS (microelectromechanical systems) is tricky, Cavendish Kinetics presents a new challenge: It takes mechanical structures to the nanodimensional region and the proposed use is as a nonvolatile memory. Cavendish has developed IP (intellectual property) in the form of a latching, cantilevered structure, which you can build in the same cell size as a conventional EEPROM and which you might consider using as nonvolatile, small to midsized memory in a range of embedded-system applications. Electrostatic forces using capacitive operation actuate the device, and surface-state atomic-scale forces latch it. (Could you therefore mechanically disrupt it by shock? No,

says Cavendish; the relative masses and attachment forces involved mean that huge g-forces would be necessary to do so.)

Why use a mechanical switch as a memory? The technology is faster by around two orders of magnitude, the power in read and write operations is three orders of magnitude lower than a conventional EEPROM or flash cell, and the switch has at least a comparable lifetime to EEPROM. You could therefore use it for extremely low-power applications, such as smart cards, especially of the noncontact variety. As a mechanical device, it is also radiation-hardened, and it runs at the native voltage of the process without boosted write or erase voltages. Cavendish built the

switch in the metal layers of a conventional-silicon process, although it does need its own silicon real estate, so you can retrofit it to an existing process. Cavendish believes that approximately 10% of today's embedded-system applications might use the technology as an economically viable alternative to adding a flash module to a microcontroller for just a small amount of memory.

The company plans to sell IP for both the memory and the process module, introducing the technology in a series of steps. In the course of its research, it has learned to manage aspects of small metal elements, and its first offering is a fuse that is approximately 15 times smaller than a laser-programmed fuse. You can use

these "e-fuses" at high density and blow them with logic voltages; they program in about 1 μ sec with 1V and 1 mA. Next step will be an ultrafast, ultra-low-power e-OTP (one-time-programming) structure, culminating in a fully rewritable e-MTP (multiple-time-programmable) device.

—by Graham Prophet

► **Cavendish Kinetics**, +31 73 62 49 120, www.cavendish-kinetics.com.