CMP

ELECTRONIC ENGINEERING

EE TIMES

Issue 1124 www.eet.com The industry newspaper for engineers and technical management Monday July 31, 2000

Vision Quest

BY JAYME GANEY

Chicago-area surgeons, below, implanted silicon retinas in the eyes of three patients (detail, left).

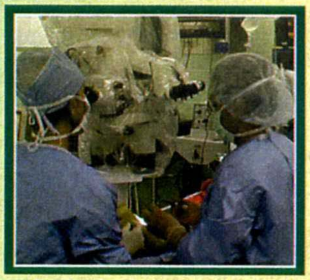

A company launched by an electrical engineer and his ophthalmologist brother is about to see results from the first implants in human patients of its solar-powered silicon retina. In late June, test chips from Optobionics Corp. (Wheaton, Ill.) were surgically placed in the eyes of three blind patients. Whether the Artificial Silicon Retina (ASR) helps these victims of retinitis pigmentosa regain at least some of their vision will be known in roughly 10 weeks, when the patients' sight is formally evaluated.

Ophthalmologist Alan Chow, who is Optobionics' president and chief execu- ▸▸CONTINUED ON PAGE 123

Duo faces off in mobile data

By Junko Yoshida

SAN MATEO, CALIF. – The World Wide Web Consortium (W3C) and the Wireless Access Protocol Forum are working to craft the next-generation wireless data delivery specification, WAP 2.0. The delicate work could forge a merger of two technologies battling it out in today's wireless data services market: NTT Docomo Inc.'s successful i-mode and the existing Wireless Application Protocol.

At the heart of the effort is a drive to leverage the groundwork laid by W3C in hopes of converging HTML subsets and variants into a common subset of XHTML. The W3C has produced a working draft of XHTML Basic, edited by three companies, two of which have been in direct competition: Phone.com (Redwood City, Calif.), which pioneered WAP, and Access Co. (Tokyo), which designed Compact HTML, the underlying markup language for i-mode. The third player, Panasonic, is joining the collaboration to seek a common ▸CONTINUED ON PAGE 136

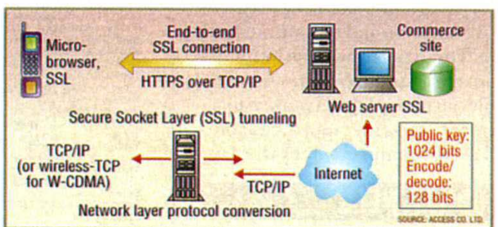

Parts of Japan's i-mode service could be built into the WAP 2.0 specification that's being developed.

Xilinx, Altera tap PowerPC in system-on-chip bids

FPGAs muscle in on ASICs' embedded turf

By Craig Matsumoto and Rick Merritt

SAN JOSE, CALIF. – The two programmable-logic titans, Altera and Xilinx, are lining up heavyweight CPU and intellectual-property partners in an attempt to propel high-density PLDs into communications and embedded designs, areas previously ruled by ASIC vendors.

Xilinx's Roelandts: 'Doubling opportunity.'

Last week, Xilinx Inc. and IBM Corp. announced a deal to meld Xilinx's Virtex-II FPGAs with IBM's PowerPC cores. Xilinx (San Jose) will license the rights to the cores and the CoreConnect bus, in a relationship that may expand to other "hard" intellectual-property (IP) cores.

Also, Xilinx will enlist IBM as a foundry, initially manufacturing Virtex-II FPGAs with the PowerPC cores at 130-nanometer (0.13-micron) design rules. However, the open-ended agreement is expected to include manufacturing of a wider number of Xilinx parts later on.

Meanwhile, Altera Corp. is in negotiations for its own PowerPC license, from Motorola Inc. "We would expect to close that discussion in three months," said Cliff Tong, vice president of corporate marketing at Altera. That company, also based in San Jose, recently unwrapped its Excalibur program, which aims to blend ARM, MIPS and Altera's own Nios processor cores into Altera's Apex programmable FPGAs (see June 12, page 53).

▸CONTINUED ON PAGE 138

Plea to alter DTV spec goes to Congress

By George Leopold

WASHINGTON – A rift that has rattled the deployment plan for digital TV was explored at a hearing before the House Commerce Telecommunications Subcommittee last week, as critics and backers of the U.S. DTV standard sparred over whether to reopen the spec to include an alternative modulation scheme or press on with a delayed rollout.

Are 'millions being spent for naught'?

While broadcasters and equipment makers haggled over technical issues, lawmakers and regulators worried about the impact of the standards rift on the deadline for completing the DTV transition, reassigning precious analog TV spectrum for other applications and ensuring that rural parts of the country ▸CONTINUED ON PAGE 16

What's Hot

NORTEL READIES BID FOR ALTEON

Nortel Networks late last week said it plans to acquire Alteon Web Systems Inc., which provides technology for handling network switching based on content. Nortel is offering roughly $7.8 billion in stock.

STMICROELECTRONICS BUYS WSI

STMicroelectronics last week bought Waferscale Integration Inc. (Fremont, Calif.), a supplier of programmable system devices, for $68 million. ST has been involved in designing and making WSI products for years.

Picking alternative interface for upcoming Willamette could relegate Rambus to niche status

At eleventh hour, Intel wavers on Rambus commitment

By Will Wade and Anthony Cataldo

ANALYSIS

SAN MATEO, CALIF. – Just days after Intel Corp. publicly insisted it would stick with only the GHz Rambus memory, the company officially hedged its memory-interface bet last week by opting to use standard memory for its forthcoming Willamette processor. The abrupt change in direction will create waves that will travel up and down the systems food chain and stretch well into next year.

Intel remains firm in its support of Rambus DRAMs as well, but industry watchers observed that RDRAM might be unable to emerge from its niche as a solution for high-end systems.

"We have been in conversations with our OEM customers and have decided to produce a chip set for PC 133 SDRAM," an Intel spokesman said. The chip set will be available in 2001. Intel is also investigating the possibility of a chip set that would support the double-data-rate (DDR) DRAM format, he said.

Infineon's du Preez: Bullish on Rambus.

Acer Laboratories Inc. (ALI) confirmed that it is in talks with Intel for a bus license to interface with the Willamette's 400-MHz front-side bus. The Willamette will be marketed under the brand name Pentium 4 and uses a different bus architecture from the P6 bus used in current Pentium-class chips. ALI expects that chip set to be released early next year. Via Technolo- ▸CONTINUED ON PAGE 6

THE CENTURY OF THE ENGINEER: THE NETWORKED MARKETPLACE

Applying PC lessons to embedded

BY DEREK SPRATT

The key dynamic of embedded-systems development in the new millennium will revolve around device networking and the network-based management of devices. With ubiquitous computing will come the enormous and daunting task of managing the millions of new end points of the Internet. It's no longer just about chips and hardware–networking software is now becoming paramount in helping companies develop smart devices and intelligently manage those devices once they've been deployed. The old saying that truth is stranger than fiction applies to the sea change in computing that is now before us. For the last two decades we've lived through the PC era, a period where embedded systems were more or less relegated to the purview of hard-core developers who found themselves outside the mainstream of PC-based computing. Today, the future lies beyond the PC in what to many of us is old, familiar territory: embedded systems. As we enter the new millennium our beloved space has been reinvented. It's now referred to as "pervasive computing" and the term "embedded system" has been displaced by terms such as "intelligent computing devices" or "Internet appliances." It's the hot place to be at the moment, or so we're told.

To those of us who have been developing embedded systems since the days of the 6800 and the 8080, it's a homecoming of sorts, but it's also the rebirth of an industry with a broad set of new dynamics to grapple with. In a world that has become fixated on the Internet, the old embedded model of the device-centric system is quickly being replaced with a new, powerful network-centric computing model.

This is for a very good reason: The value of the Internet lies not with the end points or devices themselves, but with the transactional data that flows through the network. The ability to efficiently develop, deploy and manage networks of collaborating smart devices will quickly become a key underlying driver.

The network is finally becoming the computer, and a bold new vision of embedded systems is being born along with it. Think in terms of real-time data flow and what groups of collaborating devices with robust access to enterprise-level systems can accomplish in almost any market segment you care to imagine. But who is going to develop, build, install, service and support all of those millions of smart, networked devices? There aren't enough educated and experienced embedded-systems developers and networking professionals in the world today to even manage the challenges involved in building out the current PC-based Internet computing model, so something has to change.

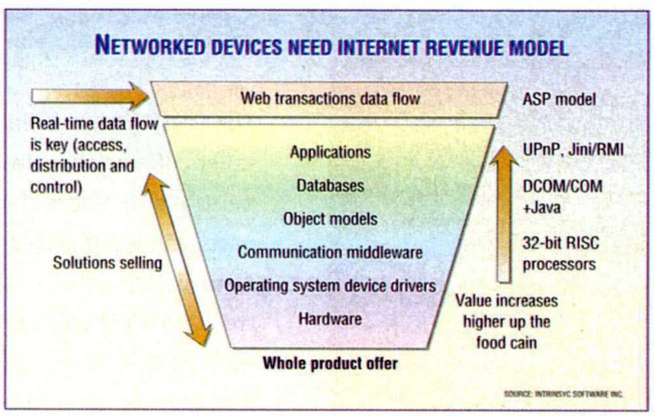

It's quite likely that all embedded-system designs in the future will be judged by the basic criteria of reducing or eliminating the need for technical resources along each step of the process: developing the device; deploying the device on the network; and maintaining and supporting the device.

In a universe with communications at its core, the software content of a networked device must be rich, deep and constantly evolving, just like the PC today. But unlike the PC, there aren't always going to be IT managers or technical end users available to install apps and configure new devices–especially when many of them won't have traditional PC user interfaces or upgrade capabilities. In a network-centric model, everything from basic software installation to configuration and maintenance must be network based and automated. And given the dictum that data flow is all important, the ability of a device to automatically install itself on the network and intelligently link into enterprise data management systems is also a looming prerequisite.

Distributed computing may be the holy grail of networked device development. Think of the Internet today as a PC-based network where only minimal collaboration occurs between clients–it's a basic client/server model. That said, the power of the Internet is in the huge number of clients and servers out there. Fast-forward to a time in the not-too-distant future when the number of embedded devices exceeds the number of PCs by 10 times or more. Where then is the power?

Embedded systems are low-cost, task-specific devices and on their own they aren't very interesting. But if they can leverage one another's capabilities, the resulting synergies can create a computing powerhouse that would dwarf the current PC Internet model. This requires a combination of auto discovery technologies (such as Microsoft's Universal Plug and Play or Sun's Jini) and distributed component object models such as Windows DCOM and COM+ protocols or various Java technologies.

Add intelligent software agents that can stream data to the right places in real-time and server-based auto profiling that will enable administrators of these new systems to automatically configure a large network of devices ➤➤ To page 93

Derek Spratt is president and chief executive officer of Intrinsyc Software Inc. (Vancouver, B.C.).

EMBEDDED PC ➤➤ From page 64

simultaneously. Suddenly, the world becomes a simpler place to live in. It's a device-to-device and device-to-enterprise networking model that in effect reverses the current client/server model–any device can be a client or server at the same time.

Embedded-systems developers have to start to think differently about approaching the design of these next-generation networked devices. The old habit of choosing the processor, then the operating system, then building the application and thinking about communication concerns is backwards. Communication is the key element in these systems and therefore must be first and foremost in the minds of developers at the inception of the product design cycle. Communication affects everything else; in many cases, developing a product with powerful distributed-computing capabilities means designing the application development from the ground up around these principles.

Distributed-system programming can be a very empowering concept for networked devices: Breaking down the barrier of whether a process is running locally or remotely on another device means having the ultimate in flexibility in both utilizing the current product design criteria and allowing the device's capabilities to be expanded and optimized in the future with very little effort or expense. Combining a distributed-computing architecture with auto discovery, intelligent agents and auto profiling creates an explosion of new opportunities.

Gone are the old days of rolling your own proprietary operating system, communication protocols, applications and remote management strategies. There's just too much going on and too much time-to-market development pressure to avoid outsourcing these critical components of the design.

There is a huge body of technology that can and should be leveraged from the success of the PC and the Internet. In a world with open standards, proprietary solutions provide protection for companies and job security for their developers–but at the risk of losing customers.

Outsourced software

When Intrinsyc was founded, it was on the premise that OEMs would outsource the software networking components that these networked devices would require. What we didn't initially anticipate was the market pull for training and supporting design services. We finally ended up broadening our product offerings to include hardware reference platforms that come loaded with our networking software for popular operating systems; the uptake on this model was quite literally an order-of-magnitude increase in sales. The common refrain heard from OEMs is that they want an integrated solution to avoid having to knit together the disparate hardware, device driver, operating system and networking layers.

It's a reflection perhaps of the steep learning curve involved in building a next-generation networked device. These are the early days in a new market.