

M I C R O P R O C E S S O R

www.MPRonline.com

THE INSIDER'S GUIDE TO MICROPROCESSOR HARDWARE

INTEL MAPS WIRELESS FUTURE

Goals Include Reconfigurable Radio Chips, Digital Radio on CMOS

By Tom R. Halfhill {6/9/03-01}

As the semiconductor slump lurches into its third year, more and more companies are seizing upon wireless communication as their savior. Just about the only other thing growing fast enough to justify a marketing department's hockey-stick graph is spam.

Relatively speaking, Intel isn't a recent convert to the wireless religion. The company most famous for its PC processors began investing heavily in communications in the mid-1990s by acquiring startups, absorbing competitors, and launching new product lines. (See *MPR 9/13/99-01*, "Intel Network Processor Targets Routers.") Intel's Centrino platform, delivered with the new Pentium M mobile processor in March, is an aggressive attempt to make wireless LANs pervasive in PCs. (See *MPR 3/31/03-01*, "Pentium M Hits the Street.") Now Intel is mapping an even more ambitious strategy to virtually eliminate the hardware cost of wireless integration by making digital radios inexpensive enough to build into almost any chip.

Two thrusts of the so-called Radio Free Intel initiative are a new microprocessor architecture and better radio integration with mainstream fabrication technology. The first goal is to create multiband communications processors that can automatically reconfigure themselves on the fly for different wireless standards—a necessity, says Intel, for a near-future world that will be saturated with multiple, rapidly evolving radio protocols. The other goal is to integrate a wireless-baseband processor and analog front end on a single CMOS chip—without the extra costs and complications of external components, exotic semiconductors, or additional processing steps during fabrication. Ultimately, Intel wants to enable designers to integrate radios on their chips as freely as they integrate UARTs today.

Radio Free Intel is a long-term strategy, not a next-quarter cure for a becalmed semiconductor industry. It's a

series of research and development projects managed by Intel's Corporate Technology Group, which reports to chief technology officer Pat Gelsinger. All told, the projects involve 50–60 people in Arizona, California, Oregon, England, Japan, and Russia. Eventually, they will hand the technology over to product groups for implementation. Although the engineers and executives leading these projects have laid some important groundwork, they admit they may not achieve some of their longest-term goals before they retire.

Wireless Without Strings Attached

Even today, users and wireless service providers are wrestling with multiple standards that overlap to some degree. For wireless LANs, 802.11b Wi-Fi has surged to an early lead, but 802.11a offers more bandwidth at the expense of backward compatibility, whereas 802.11g promises better bandwidth and compatibility at the expense of a late start. Even more standards in the 802.11 series are in the works, as are new standards, such as 802.16a, also known as WiMAX. All these wireless LAN protocols may coexist with wide-area wireless networks, such as 3G-cellular and W-CDMA, and with shorter-range wireless networks, such as Bluetooth and ultra-wideband. The ether is thickening with data traffic.

Users will expect their laptops, PDAs, cell phones, and other wireless clients to work in this environment—seamlessly and transparently. When they move from their living room to a Starbucks, they won't want to consult an FCC manual to figure out which wireless standard the coffee

shop's base station supports so they can reconfigure their network card and protocol stack. They'll expect the convenience of wireless Internet access with no strings attached.

That's the reason Intel is working on multiprotocol baseband processors that can automatically detect and identify multiple wireless standards and reconfigure their internal processing units to match whatever is within earshot. What's surprising is how far Intel has quietly proceeded along this path; what's not surprising is how closely its prototype processor architecture resembles similar designs from other companies.

Intel has already defined a reconfigurable communications architecture and has created a functional software model of a multiband processor. The architecture is based on a mesh of heterogeneous processing elements (PE) that is scalable to smaller or larger processors, depending on the number and types of wireless protocols it must support. One example design supports 802.11a/b/g. Intel has designed some ALUs and control units, has taped out the designs, and is expecting silicon later this year—but only for parts of the mesh, not for the whole processor. The architects say they need more time to understand the routing among the various function units and PEs. Figure 1 is a block diagram of their early design.

Intel claims its new architecture will offer a better balance of power and flexibility than other solutions do. It will probably consume more power than a hard-wired ASIC, but it adapts more easily to new wireless standards. It may be less flexible than a general-purpose processor, DSP, or FPGA, but it's supposed to consume less power in systems that support multiple protocols. Intel thinks the breakeven point for power consumption will be three protocols. Dedicated hardware will maintain an advantage in systems that support fewer protocols than that, but systems supporting four or more will do much better with a reconfigurable baseband.

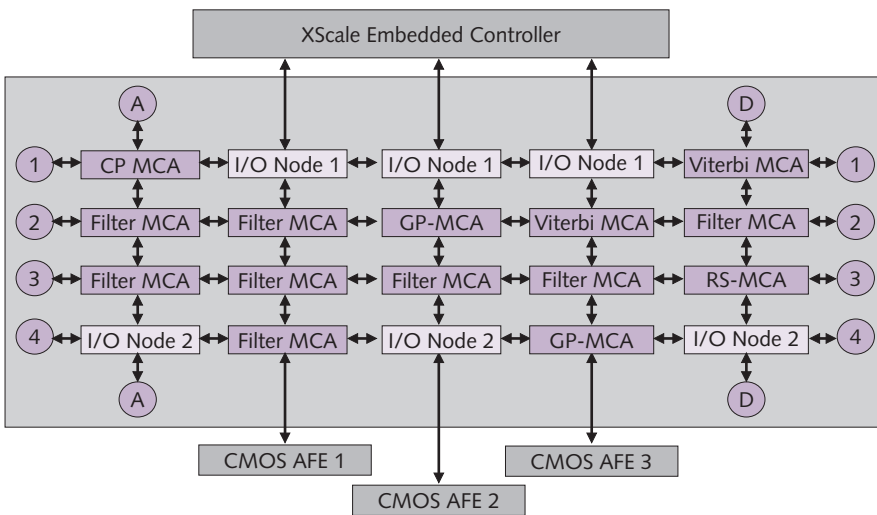


Figure 1. Intel's reconfigurable communications architecture uses a mesh of processing elements that can handle the baseband processing for different wireless protocols. This example design targets the three most popular wireless LAN standards in the 802.11 series and couples to analog front ends (AFE) implemented in CMOS.

Other companies have also designed reconfigurable communications processors with meshlike architectures. The idea of weaving together a fabric of PEs, ALUs, or small microprocessor cores is attractive for its flexibility, although penetrating the market remains a challenge.

In 2001, Chameleon Systems introduced its CS2112 reconfigurable communications processor, which has a fabric of 84 ALUs and 24 integer multipliers distributed in four slices, each with three tiles that contain their own local memory and control logic. (See *MPR 6/12/2000-01*, "Chameleon Crosses CPU, FPGA.") In 1999, Cradle Technologies announced the first chip in its Universal Microsystem architecture, which uses parallel arrays of microprocessor and signal-processing cores, programmable logic elements, and protocol engines. (See *MPR 10/6/99-05*, "Cradle Chip Does Anything.") And at **Embedded Processor Forum** on June 18, picoChip will reveal the technology behind its picoArray PC101, a massively parallel array processor that integrates more than 400 heterogeneous 16-bit processors on a single die.

Intel's marketing muscle gives the company an advantage over smaller competitors, but until we see a final implementation, it will be difficult to compare Intel's new architecture with competing solutions. One possible difference is that Intel seems to be targeting mobile clients, at least in part, whereas the other processors mentioned are more suitable for fixed infrastructure applications, such as wireless access points. Of course, all the companies in this arena emphasize the scalability of their architectures, so it's likely that different implementations will target a variety of wireless applications. The ability to support multiple protocols is as valuable in a wireless base station as it is in a wireless client.

Many other questions also will remain unanswered until Intel's new architecture emerges from the research lab and reaches the eager hands of the product groups. For instance, if a multiband-capable system detects overlapping wireless LANs, to which one will it connect? Will it choose the WLAN with the strongest signal or the one with the most bandwidth? Or will it prompt the user to choose between them? Could a multiband communications processor reconfigure itself quickly enough to maintain multiple wireless connections simultaneously—such as Wi-Fi and Bluetooth, which use different protocols but share the same radio spectrum? It may be years before Intel resolves these and other implementation details.

Low-Cost Radios for the Masses

Inexpensive, palm-sized transistor radios made broadcast news and entertainment available everywhere in the late twentieth century, and Intel wants to do

the same thing for two-way wireless communications in the early twenty-first. This time, the driving force will be inexpensive baseband processors integrated on the same CMOS chip with the analog front end—the antenna subsystem that captures and digitizes radio signals for the baseband.

Today, integrating the analog front end of a radio on the same digital CMOS substrate with the baseband processor creates problems with cross interference. The noisy digital circuits can desensitize the radio components, reducing their range and data throughput. Alternative semiconductors—such as gallium arsenide and silicon germanium—are more resistant to substrate coupling, but they're impractical for large digital circuits at commodity prices, so they're unsuitable for the kind of analog/digital integration Intel envisions. In the short term, Intel plans to add a new process step (a deep *n*-well layer) to improve the isolation characteristics of regular CMOS and will have the option of using silicon germanium in its new 90nm CMOS process, which begins production later this year.

Intel's long-term goal is to implement a complete digital radio—except for the analog front end, which is necessarily analog—on a single chip fabricated in a mainstream digital CMOS process, without special materials or manufacturing steps. In other words, Intel wants to adapt radio to its standard production methods instead of modifying its fabs for radio. Doing that could slash the cost of radio integration while allowing Intel to preserve its “copy exactly” model, which improves chip yields by outfitting every new fab with identical fabrication technology.

Another technical challenge is the integration of multiple analog front ends, because a multiband radio may require different antenna subsystems. For instance, 802.11b operates in the 2.4GHz radio-frequency band, while 802.11a operates at 5.2GHz. To go along with its future reconfigurable communications chips, Intel is working on reconfigurable “smart” antennas using microelectromechanical systems (MEMS). The tiny machines would be capable of emulating different analog front ends on demand. MEMS technology doesn't require leading-edge lithography, so Intel is building experimental devices at an old 0.25-micron fab in Israel.

Interfacing With Bureaucrats

Other hurdles to low-cost wireless integration are regulatory, which Intel admits is a big change for a company accustomed to the unregulated PC environment. For starters, there is the sheer number of regulatory agencies to satisfy: every nation has its own, and even the European Union isn't united. Some nations—such as Japan, China, and a few European countries—require a separate certification for each new radio-capable product, even if multiple products use the same components operating in the same band with the same protocols.

Individual certifications used to be the rule in the United States as well, but in 2000 Intel helped convince the FCC to adopt a streamlined approval process for radios implemented as modular components. A blanket certification now allows

For More Information

To learn more about the Radio Free Intel initiative, visit www.intel.com/labs/features/cn05031.htm.

companies to integrate a single module into any number of different products. Intel recently submitted another proposal to the FCC that would allow companies to obtain a single certification for a partitioned radio module (consisting of an antenna, an analog front end, and supporting firmware) that designers could integrate into different products in different ways. The FCC is expected to rule on the proposal this year.

Another regulatory issue is the use of wireless clients on commercial aircraft. Although it's not uncommon today for passengers with 802.11-equipped laptops to quietly establish an ad hoc wireless LAN—multiplayer gaming is reputed to be a great way to pass time on a long flight—it's actually illegal to use mobile wireless-communication devices on airliners. The most pressing question is whether the wireless data traffic could interfere with the plane's avionics. Surprisingly, there have been no official studies of this issue.

The last official U.S. study of personal electronic devices on aircraft was performed by an aviation trade organization (the RTCA) in 1996, before the advent of Wi-Fi. It focused on unintentional radio emissions from laptops and other devices, not on intentional radio transmissions. In 2000, the Civil Aviation Authority (Great Britain's counterpart to the U.S. Federal Aviation Administration) conducted a study of cell-phone emissions on aircraft, but not emissions from other kinds of devices. Since 1998, Intel and other companies have been lobbying the FAA for new studies and regulations.

A few years ago, when the industry was more optimistic about the adoption rate of Bluetooth, Intel conducted a private study on intentional radio transmissions using that technology. Intel tested Bluetooth on Boeing 727 and 747 airliners and a Gulfstream G5 corporate jet. The researchers conducted all tests with the engines running, although the planes never left the ground. (Airplanes recover from crashes less gracefully than do computers.) Even with four high-power (4W) transmitters and ten Bluetooth networks distributed throughout the cabin, the researchers detected no interference with any aircraft systems. That result should make renegade Wi-Fi users on airliners (and their fellow passengers) breathe a little easier, because Bluetooth operates in the same 2.4GHz band as 802.11b. However, nobody has studied airborne radio interference in the 5.2GHz band, where 802.11a operates. Clearly, this subject deserves immediate FAA attention.

For computing to become truly pervasive, it almost has to be wireless. The ability to economically integrate wired network connectivity is already revolutionizing the design of embedded systems, from industrial machinery to point-of-sale terminals to smart appliances. When the same connectivity is inexpensively available without any strings attached, the

number of networked clients will expand exponentially. Although the technical and regulatory obstacles are greater, so are the potential benefits. The Radio Free Intel initiative is a step in the right direction by a company that has the motivation

and wherewithal to help make it happen. Because wireless communication extends beyond the confines of a wireless client, however, Intel will have to work more closely with other companies and government agencies than ever before. ♦

To subscribe to Microprocessor Report, phone 480.609.4551 or visit www.MDRonline.com