Pervasive Computing

An Overview of the Concept and exploration of the Public Policy Implications

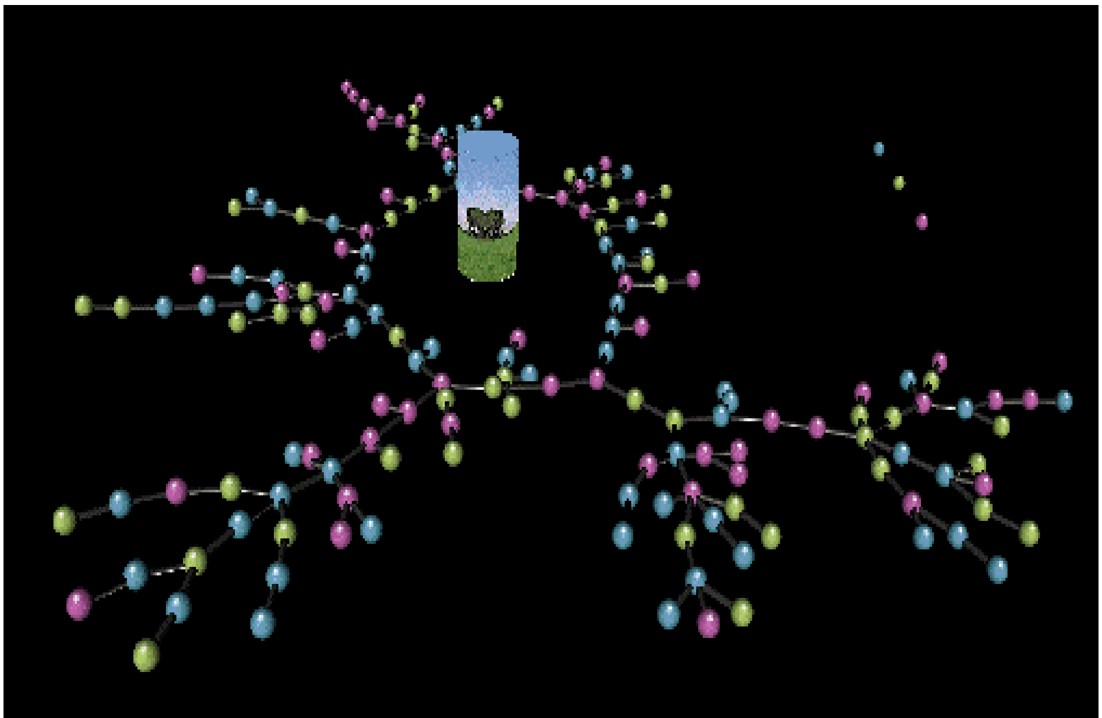

PUBLICATION 2003-3

Dr. M. Mitchell Waldrop

Foresight and Governance Project

Woodrow Wilson International Center for Scholars

http://wwics.si.edu/foresight

ABOUT THE **AUTHOR:**

Dr. M. Mitchell Waldrop, an independent writer living in Washington, DC, is the author of Man-Made Minds (1987), a book about artificial intelligence; Complexity (1992), a book about the Santa Fe Institute and the new sciences of complexity; and The Dream Machine (2001), a book about the history of computing. He earned his Ph.D. in elementary particle physics at the University of Wisconsin, and from 1980 to 1991 was a senior writer at Science magazine.

This paper was developed as part of a project on **The Future of Computing** at the Woodrow Wilson International Center of Scholars. The series of seminars and white papers is being implemented by the Center's Foresight and Governance Project with support from IBM. You can learn more at the web site: www.thefutureofcomputing.org
Cover illustration from *The Atlas of Cyberspace* at:
www.cybergeography.org/atlas/topology.html

I. The Convergence

Think of it as technology's Perfect Storm: the convergence of three massive waves of innovation, each driving the others on to new heights. Somewhere in the collision zone we're witnessing the birth of "pervasive computing," a.k.a. "ubiquitous computing," or "ambient computing." The three waves:

- A Proliferation of Devices
- Wireless Networking
- Mobile Software

1. A Proliferation of Devices

Back in the 1960s, the word "computer" still meant a big, expensive mainframe sitting off in an air-conditioned room somewhere, with lots of users lining up to submit their jobs on punch cards. Or, if the users were lucky and the mainframe was actually a time-sharing machine, they might be interacting with it in real time via remote terminals. Either way, the 1960s was the era of *many-to-one* computing: many users sharing one machine. Yet even then, just as Intel co-founder Gordon Moore had first predicted in 1965, engineers were cramming more and more processing power into smaller and smaller packages; indeed, Moore noted, the power that was available for any given price was doubling every 18 months or so.[1] By the late 1970s, this evolution had progressed to where vendors could put an entire small computer on the user's desktop—and build it cheaply enough that that user could actually afford to buy it. This brought us into the era of *one-to-one* computing: one person, one personal computer. The ensuing PC revolution, together with the Internet revolution of the 1990s, has transformed the industry, and has dominated our thinking about the technology ever since.

But now Moore's Law has taken us so far that a typical user not only owns a desktop or laptop PC—or two, or three—but may find herself carrying around a cell phone, a PDA (Personal Digital Assistant), a pager, a smart card, a GPS (Global Positioning Satellite) receiver, an MP3 player, and more. At her office she may likewise find herself surrounded by printers, copiers, fax machines, scanners, servers, and routers, all of which have some form of on-board processing power. And that's not even counting all the microprocessors and sensors that may be operating invisibly inside her microwave, her refrigerator, her car, or even her wristwatch.

As we enter the new millennium, in short, the era of *one-to-one* computing is fast giving way to the *one-to-many* era. Each of us has more and more processors at our fingertips, know it or not, and the devices they power are becoming increasingly mobile.

As we enter the new millennium… the era of one-to-one computing is fast giving way to the one-to-many era.

2. Wireless networking

Not only are the processors themselves proliferating, but more and more often they're connected. In fact, this is one of the hottest areas of innovation in high tech today, with several wireless technologies already competing in the marketplace. (A group at the University of Michigan offers a useful online tutorial.[2]) Wi-Fi,[3] for example, is a wireless analog of Ethernet, meaning that it's designed to create a "local area network" linking computers and other devices within an office or building. First introduced in 1997, it is rapidly growing in popularity for home and office use, and has recently begun to pop up in airports, coffee shops, and other such public spaces. (The idea is to provide Wi-Fi "hotspots" where anyone with a properly equipped laptop or PDA—and a subscription to the correct account—can access the Internet at will.) The current Wi-Fi standard, known formally as 802.11b, has an indoor range of roughly 150 feet. It operates at the same 2.4 gigahertz frequency band used by high-end cordless telephones, and can deliver data at up to 11 million bits per second, or just a bit higher than standard Ethernet. Products based on a new 802.11g standard, approved in early 2003, promise data rates of up to 54 million bits per second.

Bluetooth,[4] meanwhile, is designed to be a wireless replacement for that rat's nest of cables going into the back of your PC. Originally devised by Stockholm-based cell-phone giant Ericsson, Bluetooth was named after the tenth-century Danish king who united the unruly tribes of Scandinavia. It could be used in the obvious way—linking computers to peripheral devices such as printers and scanners, for example—but could also allow for less familiar applications, such as using your cell phone as an office cordless phone. It operates in the same 2.4 gigahertz band as 802.11b and g, but has a much shorter range and data rate (no more than 1 megabit per second.) On the plus side, Bluetooth uses considerably less power, and is supposedly more resistant to interference.

Finally, there are the many technologies that allow for data transmission over cellular telephone networks. Data rates tend to be comparatively low—the "second-generation" digital cell phones that most of now use can't do much better than a 14.4 kilobits per second dial-up modem—but the range is effectively nationwide, if not worldwide. Think of the cell system as a wireless wide-area network; certainly there are already millions of people, especially in Europe and Asia, who are happily accessing the Internet this way. And work is proceeding apace on a cluster of "third-generation," or 3G[5] technologies that promise 144 kilobits per second for highly mobile receivers moving at automobile speeds; 384 kilobits per second roughly equivalent to a cable modem) for pedestrians moving at walking speed; and up to 2 megabits per second for stationary receivers.

It's anyone's guess how all of this will shake out in the marketplace. The Wi-Fi, Bluetooth, and 3G technologies may each find a niche of its own, and co-exist. Or one may end up displacing the others. (Many observers, for example, think that Wi-Fi has too much of a head start for 3G to overcome.) Or in the long run they may all be deposed by some new technology, such as the proposed Ultra Wide Band.[6] But whatever happens, we can expect mobile devices in the future to be in almost constant communication with one another.

3. Mobile Software

Much of the Internet's magic and power stems from the fact that it's an information commons. Indeed, that's been a central design goal ever since the network first took life as the Arpanet, back in the 1960s: you could tap into it from anywhere there was a portal, and get access to any information it contained (or at any rate, any information that wasn't fenced off behind some kind of firewall.) This magic was democratized in the 1990s by the World Wide Web, which not only gave the Internet a graphical interface, but hyperlinks. Suddenly, we could go browsing through online documents by the billions, almost as if the documents had all taken up residence on our local hard disk. And now, in the new millennium, a cluster of rapidly emerging technologies is allowing us to work much the same magic with other computational resources—everything from databases, simulation packages, and visualization tools, to the number-crunching power of the computers themselves. Known variously as web services,[7] grid protocols,[8] and peer-to-peer computing,[9] these new technologies emerged in the 1990s quite independently, but are so similar in spirit and purpose that they are now in the process of coalescing. Whatever the name, they promise to give our multiple connected devices the power to reach out into cyberspace, find computational resources wherever they may be, and then assemble them on the fly into whatever applications we need—without our ever having to know or care where the resources are, what computers they are running on, or how any of this occurs.

II. Pervasive Computing Is…

The trick, of course, is to figure out what will actually emerge from this Perfect Storm. It's conceivable that the maelstrom of mobile devices, wireless connectivity, and mobile software will just continue on in chaos, leaving in its wake nothing more than a bigger pile of gadgets. But to an increasing number of people in the computer and telecommunications field, it seems much more likely that the collision zone will produce a discontinuity: a change in our relationship with technology that will be at least as radical as the PC revolution.

Among the earliest and most influential advocates of this view was the late Mark Weiser,[10] chief technologist at PARC, the Xerox Palo Alto Research Center. In 1988, Weiser articulated a vision of "ubiquitous computing," which he later described this way:

> *For thirty years most interface design, and most computer design, has been headed down the path of the "dramatic" machine. Its highest ideal is to make a computer so exciting, so wonderful, so interesting, that we never want to be without it. A less-traveled path I call the "invisible"; its highest ideal is to make a computer so imbedded, so fitting, so natural, that we use it without even thinking about it.*[11]
>
> *…Our preliminary approach: Activate the world. Provide hundreds of wireless computing devices per person per office, of all scales (from 1" displays to wall sized). This has required new work in operating systems, user interfaces, networks, wireless, displays, and many other areas. We call our work "ubiquitous computing"…*[12]

As Weiser often pointed out, this was the exact inverse of that much-hyped notion, "virtual reality": instead of asking us to enter into the computer's world, ubiquitous computing would bring computers into *our* world. By the early 1990s, he and his colleagues had an experimental environment operating within PARC, using a wireless network and three

…instead of asking us to enter into the computer's world, ubiquitous computing would bring computers into our world

types of computers: “tabs,” which were essentially smart tags that could be attached to a book, say, or used as a personal ID badge; “pads,” which were handheld machines that could be written on with a stylus, much like today’s tablet computers; and “boards,” which could be mounted on a wall and used like a conventional whiteboard—except that the “ink” could be stored and processed electronically. (The latter was later marketed as the Xerox LiveBoard.) In the process, they began to work out many of the ideas discussed below.

Sadly, Weiser died in April 1999 after a brief battle with cancer. By that point, however, his ideas had already achieved widespread currency among other researchers—albeit often under the alternative name “pervasive computing.” (Many people found the word “ubiquitous” to be a mouthful, and didn’t like the frequently used contraction “ubicomp” much better.) With the three technological trends mentioned above becoming abundantly apparent, industry heavyweights such as IBM and Hewlett-Packard were already beginning to make significant investments in the area, as were federal funding agencies such as the National Science Foundation and the Defense Advanced Research Projects Agency. And a wide variety of demonstration projects were getting underway—most of which are still going strong even in the face of the downturn. Among the most notable are IBM’s Pervasive Computing Lab[13] in Austin, Texas; Hewlett Packard’s Cooltown[14] in Palo Alto; MIT’s Project Oxygen[15]; the California Institute for Telecommunications and Information Technology—Cal-(IT)²—in San Diego[16]; the Georgia Tech Aware Home[17]; and Royal Philips Electronics’ Homelab[18] in Eindhoven, The Netherlands (where it’s called “ambient computing”).

Of course, because the research is still in its blind-men-and-the-elephant phase, each of these projects (and each of the people involved in the projects) has a different view of precisely what pervasive computing will be. Speaking generally, however, most of the properties they describe can be grouped under three headings: **nomadic**, **embedded**, and **invisible**.

- **Nomadic**

The idea here is to let users set up their tents anywhere they like, so to speak, instead of being tied to a specific place or specific machine. Or to put it another way, the idea is to let users’ software and data be just as mobile as they are.

To some extent, this is starting to happen already. In 2002, for example, the Toronto Police Service began its deployment of eCOPS[19]: an integrated database system in which (among other things) officers on patrol can use wireless-equipped laptops to get an immediate listing of any prior occurrences at a particular address, any outstanding warrants, the records and mug shots of any individuals they've placed under arrest, and other such critical information. In the Washington, DC, area, meanwhile, Maryland, Virginia, District, and federal authorities are pushing ahead with CapWIN, the Capital Wireless Integrated Network.[20] Inspired in part by 9/11, CapWIN is intended to provide fully integrated communication and data access among the region's multitudinous federal, state, and local law enforcement organizations, fire and emergency medical services, transportation, and other public safety agencies.

Similar efforts are underway in the business world, where there's widespread interest in taking full advantage of all those smart phones, wireless PDAs, and in-vehicle information systems. "Given the current economic environment, deployment is not as robust as could be," admits Rod Adkins, IBM's vice president for pervasive computing. "But in some areas, the advantages are very compelling. For example, keeping your sales force mobile in the field." With wireless connectivity, he says, "they can handle more customers and make more service calls, while still getting the data the need at point of engagement." Indeed, says Adkins, IBM—which is lead contractor on both the eCOPS and CapWIN efforts—sees potential markets for pervasive computing in areas that range from financial services, to product distribution, to travel planning. In healthcare, to take a major example, physicians who carry PDAs and/or tablet computers on their rounds could get immediate access to critical lab results, as well as a menu-based prescription system that minimizes the chances of an error—a major source of adverse reactions and even death in hospitals.

As useful as such applications may be, however, they're just the beginning of real data mobility. Say you need to catch up on some work while you're on the road. In the pervasive world, say the visionaries, you'll be able to walk up to any free PC and beam a network address at it from your favorite handheld device. That PC will then pull down your entire desktop environment, along with all your files and applications, so that you can work on them exactly as you would at home. (All that personal information will be stored back to the network, of course. And once you're finished, security being as vital as it is, it will vanish from the host machine without a trace.) Or say you've been invited to give a seminar. You'll be able to walk into the conference room, beam another network address at the projector, and have it instantly bring up your PowerPoint slides. At the same time, with the projector feeding instructions back to your handheld device telling it how to become a controller, the device's display will suddenly show touchsensitive buttons for *forward*, *back*, and all the rest, and you'll use it to launch right into your presentation.

Applications such as these are currently under development at HP Cooltown, Project Oxygen, and many other pervasive computing labs. However, the technical challenges are large. Take that handheld device, for example. The proliferation of such devices is out of control already, which is why there's strong market pressure for integration; no one, with the possible exception of the geekiest computer nerds, wants to walk around wearing a utility belt that's full of PDAs, cell phones, MP3 players, and whatnot. Thus we're seeing PDAs that are also cell phones, cell phones that can browse the Web, and so on. But in a truly nomadic world, where we'll be using the things constantly, the challenge to integration will be extreme. The actual shape of our favorite gadget will no doubt be a matter of personal preference; one model may look a lot like a current-generation tablet PC, another like a PDA, and so on. But with one or two of the things, at most, we'll want to have the functionality of a cordless phone at home, a cell phone on the road, a tablet PC, a PDA, a GPS receiver, a projector controller, and much more. We'll also want our devices to automatically upgrade themselves whenever some new functionality comes along, so that we don't have to run out and buy a new gadget every time.

In short, we'll want what some researchers have dubbed the universal information appliance.[21] What makes this hard is that designers can't just keep on cramming in more and more specialized chips, which is how integration is done now; not only does that approach make for a tough upgrade path, but even with Moore's Law, there are limits. One possible alternative is the RAW chip architecture[22] developed for Project Oxygen by MIT's Anant Agarwal. Think of it as "soft hardware." Instead of making just one very complicated microprocessor per chip, as Intel does when it fabricates, say, a Pentium, Agarwal's design calls for laying down a few zillion very small and very simple processors in a regular array, each of them identical. By routing and rerouting the way data flows through this array—in effect, rewiring the chip on the fly—the software can turn the RAW chip into whatever processor it needs to be for the task at hand.

And then there's the knotty issue of connectivity. The world of pervasive computing will be much like the world is today: full of people, cars, printers, handhelds, and a myriad other entities, all of which are constantly moving (or being moved) from place to place. The difference is that every one of those entities will need at least intermittent connections to the network, on demand, often *while* they are moving. This means that each device will have to reach out when necessary, and automatically hook itself into the cellular phone system, or a conference room's local Wi-Fi network, or a desktop computer's Bluetooth connection, or whatever else gives it the best data rate and most reliable connection at that particular moment. If it's moving around, it will also have to shift to new connections as needed—and do so in mid-conversation, without the user's being aware of anything except a possible speeding up or slowing down of the data rate. As it enters each new location, moreover, each device will have to reach out and identify any other devices and software resources that are there, while simultaneously introducing itself to *them*; soon enough, after all, they may be given a new task that they will have to carry out as a team. (E.g., "Send this document to the nearest printer.") In short, as the Project Oxygen web site puts it, the networks of the pervasive world will become "*ad hoc* collaborating communities of people and computing devices, [which] configure and reconfigure themselves automatically, as nodes appear, migrate, and disappear."[23]

Unfortunately, this is *not* how networks operate today. The basic "middleware" that controls the Internet builds in the tacit assumption that computers are stationary, meaning that they can be served by a more or less static set of data pipes that feed a more or less static set of data outlets—each of which has an absolutely static, fixed-in-stone "IP address." (That's the numeric address that your web browser looks up when you type in, say, *www.ibm.com*.) Moreover, that address exists only in a kind of abstract network space; there's no way to direct something to "the nearest printer" because the system itself has no concept of physical location.

Of course, it's certainly possible to transcend those limitations. For example, devices can establish their unique identity through radio frequency ID tags, or RFIDs,[24] which have been commercially available since the 1980s. They can monitor their location and context through GPS signals. (Or, if they are indoors where GPS is unreliable, they can triangulate the signals from ceiling-mounted radio beacons, as in MIT's Cricket system.[25]) And they can learn about other devices in their vicinity—"What's in this room? How do I access it?"—via the kind of resource discovery mechanisms that are already built into web services, grid protocols, and peer-to-peer networking.

Nonetheless, it's still a cutting-edge challenge to get all these capabilities working together as a robust and seamless system. Indeed, it's a challenge that extends beyond networking middleware to encompass software in general—and arguably even computer science in general. After all, a standard shrink-wrapped application like Microsoft Office builds in some equally strong assumptions about the kind of system it's going to be running on. ("Pentium 166 or higher processor required...") But in a pervasive computing environment, there will be no way to know in advance what kind of processors will be present, not to mention what kind of memory resources, network connections, and all the rest. The “applications” that carry out any given task will have to be put together on the fly by the system itself, with the available devices collectively making use of the available resources as best they can.

Call it the challenge of Computing in Uncertain Environments. “It’s changing teaching and research all through computer science,” declares Victor Zue, director of MIT’s Laboratory for Computer Science. And not just in academia. In 2001, to take a notable example, IBM announced that it was reorienting its advanced research program around the notion of autonomic computing:[26] building “computer systems that regulate themselves much in the same way our autonomic nervous system regulates and protects our bodies.” The IBM vision was that these autonomic systems would be *self-configuring*, meaning that each new device would automatically fit itself into the existing community of devices in the way described above (think Plug-and-Play on steroids); *self-optimizing*, meaning that the system would automatically distribute tasks and subtasks to the various devices in the most efficient way possible; *self-protecting*, meaning that the systems would detect and guard itself against damage from accidents, equipment failure, or outside attacks; and *self-healing*, meaning that the systems could detect, diagnose, and recover from any damage that did occur.

This is an ambitious and long-term research agenda, to put it mildly. But in the meantime, the nervous-system metaphor brings us to the second characteristic of pervasive computing—

- **Embedded**

Consider the modern refrigerator, which has a built-in electronics package that beeps at you if the door is slightly ajar, and warns you if there’s been a power failure that might have let food spoil. Or consider the modern automobile, which uses dozens of embedded microprocessors and micro-sensors to adjust the fuel mix, trigger warning lights if an air bag needs attention, and generally keep a constant watch on the vehicle’s health and performance—and then, if need be, provide diagnostic readouts for the technicians at the garage.

The idea behind “embedded” pervasive computing is to extend that kind of electronic nervous system to virtually everything else in the world.

A prime example is the Ecological Observatory[27] now being demonstrated by Cal(IT)^2: the California Institute for Telecommunications and Information Technology. (The institute, which has emerged a leading testbed for the embedded approach, is one of the four California Institutes for Science and Innovation first proposed in December 2000 by Governor Gray Davis. It has 40 industrial participants, including such heavyweights as IBM, Hewlett-Packard, and Microsoft, and encompasses some 200 faculty members and students at UC San Diego and UC Irvine.) The Ecological Observatory is located in the 4000-acre Santa Margarita Ecological Reserve, in northeastern San Diego County. The goal is to monitor the environment there by deploying the widest possible range of sensing devices: a "sensornet" whose elements will communicate with headquarters (and each other) using the cell network, satellite links, and a variety of other wireless connections. According to the observatory web site, the sensornet will keep watch on "pollutants in the Santa Margarita River; seismic events; the moisture 'load' of the vegetation (which affects the degree of fire danger); air quality; the presence and spread of invasive botanical species; and the movement, behavior, and numbers of animal populations, such as mountain lions and bats." The Cal(IT)^2 observatory will also serve as a prototype for the National Science Foundation's much broader-scale National Ecological Observatory Network.[28]

Another example from Cal(IT)^2 is Autonet,[29] which seeks to create a central nervous system for urban transportation. Researchers on the project are still developing the basic technology. But their long-term vision includes 1). a dense network of roadside stations that will collect second-by-second traffic flow data, while simultaneously broadcasting that data to passing vehicles; 2). on-board navigation devices in each vehicle that can augment the real-time data with advanced traffic-prediction algorithms, presumably downloaded from commercial route-guidance services; and 3). wireless vehicle-to-vehicle communication that will allow the on-board devices to share traffic and other data. (As Cal(IT)^2 founding director Larry Smarr asks in mock horror, "How can it be that all the cars today don't know where all the other cars are, and what speed they're going?") The hoped-for result: a much more efficient traffic flow—a prospect that has considerable appeal in a state where traffic congestion is chronic, and one of the biggest obstacles to continued economic growth.

Meanwhile, it's easy to imagine similar electronic nervous systems monitoring the roadways and bridges themselves, as well as other elements of the civil infrastructure. After an earthquake, for example, embedded sensors could instantly alert repair crews (and drivers) that such-and-such a bridge was no longer safe. Likewise, in the much-discussed “smart home” and “smart office,” embedded electronic nervous systems could monitor and manage each building for energy efficiency, safety, and security. Say you’re planning a vacation, says Bill Bodin, director of IBM’s Pervasive Computing Lab in Austin: a properly wired house could not only tell you whether it was safe to leave—Did you really turn that iron off?—but whether it was safe to come back: Has a window been broken? Are there warm human bodies waiting inside a house that’s supposed to be empty?

Indeed, the notion of an electronic nervous system could even be extended to individuals. In principle, at least, an array of unobtrusive and non-invasive sensors built into your watch, jewelry, and/or clothing could monitor our risk of, say, a heart attack—and then automatically alert the emergency rescue services via wireless if your warning signs suddenly spiked into the red zone. Indeed, declares Smarr, “to me, it’s clear that everybody should be having their body read out on the Internet.” True, he adds, “people tend to have a horrified reaction to that idea”—not least because of the privacy implications. But if strong privacy safeguards were built in from the outset, he says, such a system could not only safeguard personal health, but could provide an incomparable trove of statistics for public health. For example, within a few years, given the way the technology is progressing, getting your own, personal genome sequenced could become as standard as getting a regular physical. “So you could take that genomic sequence, and combine it with the metabolic readout as you interacted with environment,” says Smarr. “And now if you had that kind of information for millions of people, you could use very sophisticated data mining techniques to understand how different populations of people will react to a new drug.”

That’s a ways in the future, of course. But in the meantime, it’s clear enough that much of the embedded infrastructure will actually be mobile, and that nomadic devices will actually spend much of their time communicating with embedded ones. These two characteristics of pervasive computing are completely complementary and synergistic—which leaves us with the third characteristic…

- **Invisible**

Go back to the modern automobile again and consider the anti-lock braking system: the cluster of embedded sensors and microprocessors that makes the brakes pulse on and off during an emergency stop, so that the car doesn’t skid. This a classic example of what Mark Weiser meant by “invisible” computing: pulsing the brakes isn’t something you do *on* the computer. It's something that just happens in the background—automatically—while you stay focused on the primary task of stopping the car.

The goal is to make this kind of invisibility the norm in pervasive computing. After all, goes the argument, computation is hardly a scarce resource anymore. What *is* in chronically short supply is human attention—a resource that our current-generation cell phones, pagers, Web browsers, email clients, TVs, etc. tend to squander. And when you multiply that by the overwhelmingly larger number of processors that could surround us in the pervasive world, a certain invisibility seems essential. Better to save our attention for what really matters, instead of frittering it away on tools that are forever asking us to scroll through menus, or find the right button to click. In the mid-1990s, in fact, Weiser and his then-boss, PARC director John Seely Brown, suggested that we should be working toward a whole new generation of "calm technology": stuff that would sooth the users, and actually help them stay focused.[30]

…our pervasive computational environments will have to interact and respond to us in much more humanly intuitive and natural way.

In practice, this means that our pervasive computational environments will have to interact and respond to us in much more humanly intuitive and natural way. Specifically, they will have to become a). *multimodal*, responding appropriately to speech, gesture, and perhaps even facial expressions; b). *proactive*, anticipating our needs and requests in context, and fulfilling them before we have to ask explicitly; and c). *semantic*, understanding and doing what we mean, without insisting on a precise definition—or violating common sense. ("The nearest printer," for example, probably doesn't mean "the printer that's directly below this office three floors down," even if it actually is the closest in a straight line.)

Of course, none of this is going to happen easily, or soon; artificial intelligence researchers learned a long time ago that "common sense" is a lot slipperier and more subtle than it looks, to take just one example, and they've been struggling to get decent speech understanding for at least fifty years now. True invisibility is by far the toughest challenge in pervasive computing. Nonetheless, any number of people are trying.

At MIT's Project Oxygen, for example, Howard Shrobe and his co-workers are developing technology for an "Intelligent Room." What's been achieved so far is admittedly pretty mundane, says Shrobe. When one or more people walk into such a space—the group has installed prototypes in several offices and conference rooms—microphone and camera arrays embedded in the walls and ceiling allow the room to respond in certain ways. (*Human*: "Computer, what's the weather tomorrow?" *Room*: THE FORECAST FOR

WEDNESDAY IS PARTLY CLOUDY, WITH A HIGH OF 67° FAHRENHEIT.) On request, the room can lower the screens for a presentation and dim the lights, or videotape a meeting. But hopefully, says Shrobe, this is just the walk-before-you-run stage, preparing for a much more sophisticated kind of human-environment interaction. In the future, according to the group's web site, a meeting might go something like this: "[the conference room environment] knows the interests, organizational roles, and skills of all team members, and it understands the application domain within which the team functions. For example, it tracks action items within the group and dependencies with other groups, retrieving relevant information and bringing it to the attention of the most appropriate individuals. The collaboration system plays the role of an active participant, noticing tasks that need to be undertaken, noticing when information required for those tasks has been developed, and making conclusions when appropriate."

Meanwhile, the MIT researchers hope that one of the most useful furnishings in any Intelligent Room of the future will be the "smart surface," which will be able to take freehand sketches and turn them into working simulations. MIT's Christine Alvarado and Randall Davis have devised a working prototype of such a system called Assist,[31] a.k.a. "Magic Paper." Sketch on the wall with a special wireless-equipped "pen," and the system will project a trail of digital "ink." It's almost as if you were drawing on an ordinary whiteboard—except that your rough circles automatically become perfect circles, your shaky lines become straight, and so on. Now sketch a little cart with wheels, sitting at the top of an inclined plane, and tap the button labeled "Run": the cart rolls down the plane and falls off the end. In a similar way, if you draw an electric circuit, the system will recognize the various capacitors, resistors, and so on, pretty them up, and then simulate the flow of current through the circuit.

It's a beginning, says Davis. But think of the possibilities down the road for, say, education: "Part of the dream is that in the textbook of the future, you'll never see a static diagram. You'll always be able to ask it, "I wonder what would happen if...?"

III. Policy Challenges

Indeed, the techno-visionaries have given us a beautiful dream of what pervasive computing could be. But how much of that dream will ever become real?

As we've seen, the technical challenges are daunting. Pervasive computing in 2003 is roughly where personal computers were circa 1976, when the machines existed only as hobbyist kits that you had to put together yourself. But then, for an optimist, that's good news—a strong indication from history that pervasive computing's growth curve from here on out will be at least as explosive as PCs in the 1980s. And indeed, there's evidence to support that view. In New York's ultra-trendy Soho district, for example, each item of clothing in Prada's new Epicenter store has an RFID tag attached. When you find, say, a suit that you want to try on, and hang it in the dressing room's "smart closet," antennas in the walls scan the tag and activate the room's LCD touch screen—whereupon you suddenly find yourself paging through a display of the same item in different colors, or suggestions for accessories that would go with the item's look. "That's a very interesting first experiment in creating a brand new experience in the physical places people go," says HP Cooltown director Gene Becker. "And there are plenty of other companies experimenting with this technology, as well. So when economy picks up, you can expect to see people deploying it quite rapidly."

Be that as it may, however, history is a decidedly imperfect guide to the future. There's also good reason to think that we *won't* see a repeat of the wide-open, wild and wooly days of the early PC revolution (or for that matter, the early Internet revolution.) After all, the PC pioneers were not only operating in a consumer electronics market that was largely unregulated, so that competition had free reign, but they were offering up a technology that seemed to have no serious downside. You may or may not have *liked* PCs, but there was no strong reason to be afraid of them.

Neither is the case with pervasive computing. The concept is utterly dependent on wireless networking, after all—and wireless is part of the (partially) regulated telecom industry. Moreover, the very pervasiveness of the technology raises obvious concerns about privacy, security, and trust. Some of the issues that policy-makers and business planners may want to consider:

…pervasive computing…is utterly dependent on wireless networking,…and wireless is part of the (partially) regulated telecom industry.

- **Spectrum allocation.** Nature gave us only so much space in the radio spectrum, and there's no way to make more of it. Thus the FCC, and the perennial headache of trying to satisfy more and more users with smaller and smaller slices of a fixed resource—a problem that could be greatly exacerbated if the market for pervasive-style devices really does take off. On the other hand, thanks to the basic physics of channel capacity, which can be quite subtle and counter-intuitive, it's theoretically possible for distributed networks of devices to make collective use of the available bandwidth in ways that individual devices cannot. So maybe technology can keep ahead of the game. And maybe the FCC will need to do some fundamental rethinking of it's whole approach to allocation.

- **Interoperability, Open Standards, and the Business Model.** "This is one of the big issues that [pervasive computing] has to face as an industry, " says HP Cooltown's Becker. "The last thing we want a bunch of competing, completely incompatible infrastructures, so that, when I'm at home, I have a company A experience, but when I go to the office, none of my home devices work; I can only have a company B experience. And then when I go to my dentist, I can only have a company C experience." Indeed, that kind of fragmentation could cripple pervasive computing before it really gets started; much of the dream' appeal comes from its vision of utterly seamless connectivity, with each device having open access to cell networks, Wi-Fi links, and all the rest, as needed.

  For better or worse, however, we live in a world in which fragmentation is the norm. "'Open' is *not* the way the wireless world works," emphasizes MIT's David Clark, one of the architects of the modern Internet. "If I change cell providers, I have to throw away my phone! The market as vertical as they can make it." Likewise the rapidly expanding network of Wi-Fi hotspots: unless you've got a paid-up subscription with the particular service that operates the particular hotspot you're in, don't bother trying to log on.

  The challenge, says Clark, is to find a way for vendors to provide universal service—and still make money.

  That's trickier than it sounds, he says. Witness the Internet, which is the only universal standard that's ever been a real success. "The reason the Internet looks the way it does is because it was funded by the government," says Clark, "which meant that the academics who built it had nothing to cut off their idealized bias towards open access and open standards." And that was a good thing, too, since the very openness of the Internet was what allowed it to function as an information commons, thus paving the way for the World Wide Web, e-commerce, and all the rest. "But the open architecture of the Internet also means that the guys at the bottom, the ones deploying the infrastructure, have no way to capture the money being made at the top," says Clark. Quite the opposite: the Internet's open architecture has given rise to open competition at the infrastructure level, which has tended to drive the profit margins toward zero.

Thus the wireless providers' desire to keep their own markets as segmented as they can, as a way to get a healthy return on their initial investment. And thus the need to find a better way. After all, says Clark, "history suggests that what's good for an individual player isn't necessarily good for the structure of the industry." One top priority is the search for alternative business models for pervasive computing—models that would allow for an open architecture *and* profits. (Clark is looking at several possibilities in collaboration with researchers from the MIT's Sloan School of Business.) Another is a top-to-bottom rethinking of the regulatory environment, starting with the Telecommunications Act of 1996.

But the one thing we should *not* do, says Clark, is assume that the pervasive infrastructure will just spring to life of its own accord. "There is no Moore's Law for deployment," he says. "If we get this wrong in one way, we'll get an infrastructure that's completely vertical. And if we get it wrong in the other way, we won't have any infrastructure at all."

- **Privacy.** This is the elephant in the room—the one issue that everybody involved with pervasive computing lists as a key concern. The word "pervasive" can have quite a sinister connotation, after all. And world full of electronic devices that are always on and always connected can sound very much like a world that will monitor exactly where you are and exactly what you're doing—every minute of every day. What could your friendly neighborhood hacker do with such a system? What would J. Edgar Hoover have done?

  The word "pervasive" can have quite a sinister connotation.

  "None of us wants to be the guy who invented Big Brother," says MIT's Shrobe. Victor Zue, director of MIT's Laboratory for Computer Science, couldn't agree more: "All of us are committed to privacy and security as being not the last thing we add on, but something that's built into the system from the start."

  Indeed, there is a lot that technology could do. As mentioned earlier, for example, pervasive technology could build in a certain anonymity: once you've finished using a projector, or a guest PC, or whatever, that particular device would retain no information whatsoever indicating that you've been there. Likewise, once your personal, handheld device has finished asking the surrounding intelligent environment about its location, that environment would retain no information whatsoever about the device, the time, the location, or you.

Still, you don't want to pin all your hopes on well-behaved technology, not when that technology is the creation of fallible (or duplicitous) human beings. And in any case, technology alone can only *safeguard* privacy; figuring out what the safeguards ought to be, and where our zone of privacy actually lies, is a matter of policy, law, and ultimately social norms.

To put the issue in perspective, says Rodney Brooks, director of the MIT Artificial Intelligence Laboratory, "How would you like to have a live microphone in your bedroom 24 hours a day? Well, many of us do: it's called a telephone. Or how would you like to carry around a computer that always knows where you are to within a few hundred meters? Again, many of us do: it's called a cell phone." We've come to accept such "Big Brother-ish" technologies without worrying too much about them, mainly because we trust that the phone companies won't use the technology against us. And that, says Brooks, is a social compact as much as anything.

Of course, new technology does have a way of confronting us with new social dilemmas; the norms for the pervasive world will have to be worked out as we go along. For example, says Cooltown's Becker, "there should be consistent ways for people to know what's going on—say, that a meeting they're in will be recorded by cameras and microphones hidden in the ceiling." Or, as MIT's Shrobe puts it, "What rights do people have *not* to be on camera?"

Eventually, of course, we'll have to address the question of law: to what extent should these evolving expectations about privacy in the pervasive world be legislated? Today, for example, you have a legal right to know if you're being taped. So in the future, will offices be required to put up signs, "Caution—you are entering a camera zone"? In that context, perhaps it's worth pointing out that the concern about privacy is ultimately part of the larger concern about control: being able to have some say in what information is collected about you, and how that information is going to be used. "One possible way to handle the privacy issue is via the debate going on now about who owns your data," says Larry Smarr. "For example, who owns your DNA sample? So maybe that's part of it: you have a right to your data stream."

But however we handle the issue, notes Becker, "the one thing we can't afford to do is stick our heads in the sand. Technology becoming pervasive whether we want it to or not."

[1] Moore, Gordon E. "Cramming More Components Onto Integrated Circuits." *Electronics* (1965); http://www.intel.com/research/silicon/mooreslaw.htm

[2] http://www-personal.umich.edu/~acarbon/Alejandro/wifi1.htm

[3] http://www.webopedia.com/TERM/W/Wi_Fi.html

[4] http://www.bluetooth.com/

[5] http://80211-planet.webopedia.com/TERM/3/3G.html

[6] http://www.webopedia.com/TERM/u/UWB.html

[7] http://www.webopedia.com/TERM/W/Web_services.html

[8] http://www.webopedia.com/TERM/g/grid_computing.html

[9] http://www.webopedia.com/TERM/p/peer_to_peer_architecture.html

[10] http://www.ubiq.com/hypertext/weiser/UbiHome.html; Weiser, Mark. "The Computer for the Twenty-First Century." *Scientific American* (1991): 94-100.

[11] Weiser, Mark. "Creating the Invisible Interface." *Symposium on User Interface Software and Technology* New York, NY: ACM Press, 1994.

[12] Weiser, Mark. "The World Is Not a Desktop." *ACM Interactions* 1, no. 1 (1994): 7-8.

[13] http://www-106.ibm.com/developerworks/wireless/library/wi-pvc/

[14] http://cooltown.hp.com/cooltownhome/index.asp

[15] http://oxygen.lcs.mit.edu

[16] http://www.calit2.net

[17] http://www.cc.gatech.edu/fce/ahri/

[18] http://www.philips.com/homelab

[19] http://www-3.ibm.com/software/success/cssdb.nsf/cs/NAVO-4U634A?OpenDocument&Site=indwireless

[20] http://www.capwinproject.com/index.html

[21] K.F. Eustice, et al. "A Universal Information Appliance." *IBM Systems Journal* 38, no. 4 (1999): 575-601; http://www.research.ibm.com/journal/sj38-4.html

[22] http://cag.lcs.mit.edu/raw/

[23] http://oxygen.lcs.mit.edu/Network.html

[24] http://www.aimglobal.org/technologies/rfid/

[25] http://nms.lcs.mit.edu/projects/cricket/

[26] http://www.research.ibm.com/autonomic/

[27] http://www.calit2.net/observatory/index.html

[28] http://www.nsf.gov/search97cgi/vtopic

[29] http://www.calit2.net/transportation/index.html

[30] http://www.ubiq.com/hypertext/weiser/acmfuture2endnote.htm; Weiser, Mark, and John Seely Brown. "The Coming Age of Calm Technology." *Beyond Calculation: The Next Fifty Years of Computing*. Editors Peter J. Denning, and Robert M. Metcalfe, 75-86. New York: Copernicus, 1997.

[31] http://www.ai.mit.edu/projects/rationale/project_assist

Learning from Prada

The Italian clothier has put RFID tags on every item in its New York Epicenter store. This is more than just a clever marketing ploy.

June 24, 2002 - On the northwest corner of Prince Street and Broadway in Manhattan's excruciatingly trendy Soho district, there's a nondescript brick building with green trim. The only sign on it is for the Guggenheim museum. The store on the ground level is not identified. That is only one of many things that make the Prada Epicenter store an unusual experiment in retailing.

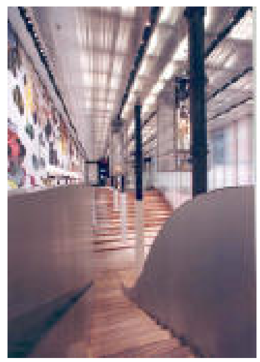

Inside the Soho store

The store opened on Dec. 15, 2001, and immediately garnered extensive media coverage for its bold redesign of the retail space and the use of technology throughout the store. But while many articles have either lauded or derided the high-tech design, few have given any real thought to the concepts being tested in the store.

The Prada store, no doubt, is a clever marketing strategy. The Italian clothier spent a reported $40 million on it. The company hired noted Dutch architect Rem Koolhaas to reinvent the retail environment in an effort to reinvigorate the Prada brand. Some architectural critics like the bold, two level store, which features a sloping wood floor, which Koolhaas calls "the wave," a large staircase for displaying shoes, and cages and shelving that slide on tracks, so the space can be changed easily.

Technology plays a big role in the store. There are video monitors that hang from racks or are embedded in tables. An alcove at the back of the basement level is lined with small video monitors playing a steady stream of random clips. A large round, glass enclosed elevator that shuttles customers between the ground floor and basement level.

And you've probably read about the dressing rooms, which feature clear glass that turns opaque when you step on a round black button on the floor. Inside are two boxes made of thick, semi-transparent Lucite. Both have thick, flat bronze ribbons embedded in them – the RFID antennas. One is small and square and is used for shoes and purses and whatnot. The other is long and narrow, for hanging clothes.

OPINION

Consumer Benefits

Contrary to what you read, RFID doesn't just benefit big companies.

FULL STORY

Pick up any pair of shoes or handbag or dress and you'll find a clear RFID tag, with the antenna and chip clearly visible. On shelves around the store are small handheld readers. Staff – mostly good-looking guys in black suits – can scan the tag on a $2,000 suit. They can then use a video monitor to show the suit on the runway, show a collection photographs and designer sketches, or providing more in-depth information about the color, cut, fabric and materials used to create it.

If you want to try on your suit, you enter the high-tech dressing room and hang it in the smart closet. The closet reads the RFID tag and displays information about the suit on a liquid crystal screen with a touch-screen overlay. You can flick through accessories or see the same item in different colors. The content displayed is all related to the item in the closet, part of the same line or "look" in Prada's parlance. The sales associate can use the screen to up-sell by showing you items that might go well with your suit.

Prada hired IconNicholson, the New York arm of IconMedialab, to do much of the integration work and to write the software for the screens that display the Prada products. A British company called KTP (now part of TrenStar) handled the RFID technology. Texas Instruments supplied the RFID tags.

Bruce Eckfeldt, Icon's engagement manager on the project, describes the Epicenter store as a cross between a concept store and a laboratory. "Prada wanted to test some concepts in the real world," he says. "They will see what works and then deploy it in other stores, as is or in some modified way. It gives them the opportunity to learn based on feedback and to build a better and more successful retailing environment in the future."

So what's working? Unfortunately, its still too early to tell. And Prada declined to be interviewed for this article. But even without hard numbers, it's clear that the Prada store is pioneering some retailing concepts that will likely become common in the years ahead. Click over to the next page and we'll explain the thinking behind the technology and tell you what Prada is doing right.

| < 1 > | 2 | 3 | Next page >>

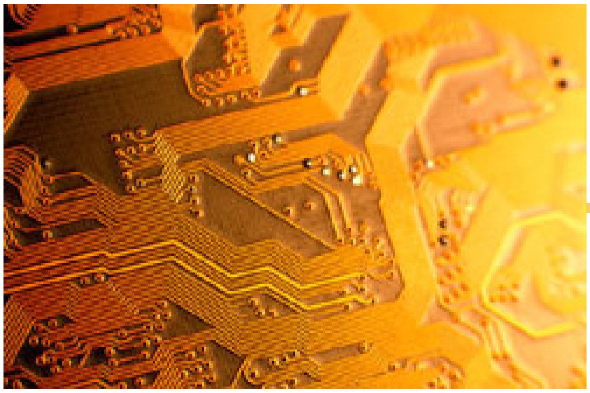

THE FUTURE OF COMPUTING

PROJECT MODULES REFERENCES

Research at Intel ▸ Research Areas ▸ Silicon ▸ Moore's Law

Silicon

Moore's Law

Overview

Gordon Moore made his famous observation in 1965, just four years after the first planar integrated circuit was discovered. The press called it "Moore's Law" and the name has stuck. In his original paper, Moore observed an exponential growth in the number of transistors per integrated circuit and predicted that this trend would continue. Through Intel's relentless technology advances, Moore's Law, the doubling of transistors every couple of years, has been maintained, and still holds true today. Intel expects that it will continue at least through the end of this decade. The mission of Intel's technology development team is to continue to break down barriers to Moore's Law.

Focus Research

- Record Setting Processes
- Lithography
- 300 mm Wafers
- Flash & Other Non-Volatile Memory Technologies
- Packaging

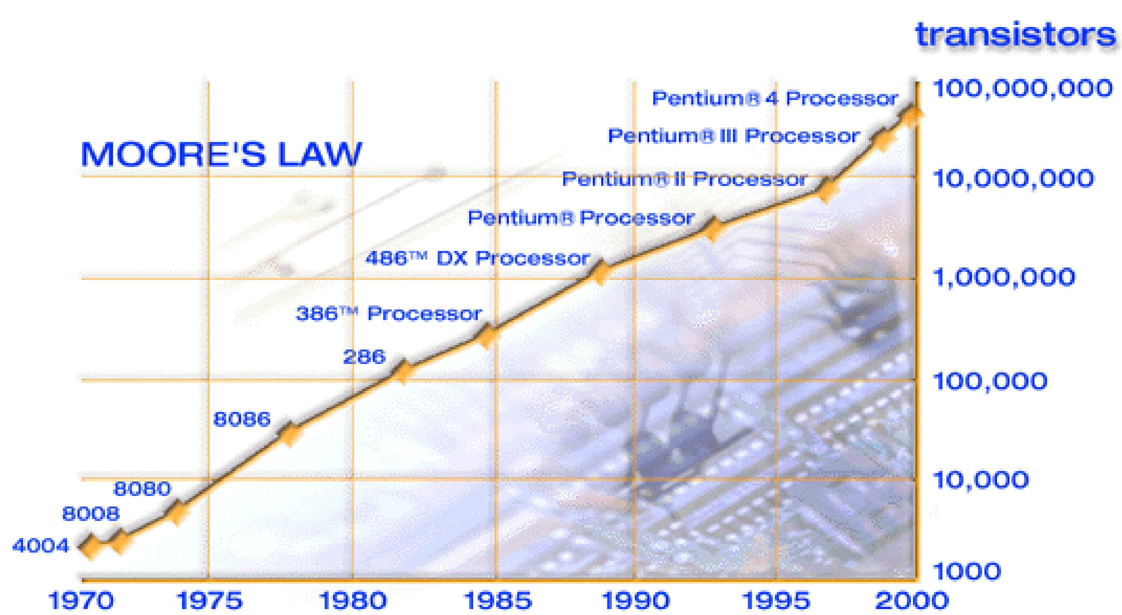

	Year of introduction	Transistors
4004	1971	2,250
8008	1972	2,500
8080	1974	5,000
8086	1978	29,000
286	1982	120,000
386™ processor	1985	275,000
486™ DX processor	1989	1,180,000

Pentium® processor	1993	3,100,000
Pentium II processor	1997	7,500,000
Pentium III processor	1999	24,000,000
Pentium 4 processor	2000	42,000,000

back to top

Gordon Moore made his famous observation in 1965, just four years after the first planar integrated circuit was discovered. The press called it "Moore's Law" and the name has stuck. In his original paper, Moore observed an exponential growth in the number of transistors per integrated circuit and predicted that this trend would continue. Through Intel's relentless technology advances, Moore's Law, the doubling of transistors every couple of years, has been maintained, and still holds true today. Intel expects that it will continue at least through the end of this decade. The mission of Intel's technology development team is to continue to break down barriers to Moore's Law.

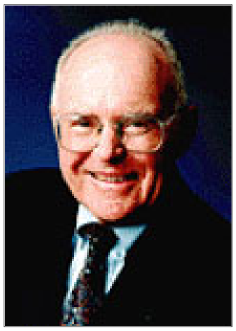

back to top

Publications

- "Cramming More Components Onto Integrated Circuits"
  (Acrobat PDF file, 167 KB)
  Author: Gordon E. Moore
  Publication: Electronics, April 19, 1965

- "Microprocessors Circa 2000"
  (Acrobat PDF file, 543 KB)
  Authors: Patrick Gelsinger, Paolo Gargini, Gerhard Parker, Albert Yu
  Publication: IEEE Spectrum, October 1989

Presentations

- No Exponential is Forever & but We Can Delay Forever
  (Acrobat PDF file, 2005 KB)
  Presenter: Gordon Moore
  Event: International Solid State Circuits Conference (ISSCC) **
  Date: February 10, 2003

back to top

**This link will take you off of the Intel Web site. Intel does not control the content of these linked Web sites.

Focus Research

- Record Setting Processes
- Lithography
- 300 mm Wafers
- Flash & Other Non-Volatile Memory Technologies
- Packaging

Related Links

- Moore's Law
- MEMS
- Lead-Free Solutions
- People
- International Technology Roadmap
- External Research Programs