

09-91 SCI AMER WEISER *** 1

Scientific American Ubicomp Paper after Sci Am editing

one more final edit from me to go

The Computer for the 21st Century

Mark Weiser

The most profound technologies are those that disappear. They weave themselves into the fabric of everyday life until they are indistinguishable from it.

Consider writing, perhaps the first information technology: The ability to capture a symbolic representation of spoken language for long-term storage freed information from the limits of individual memory. Today this technology is ubiquitous in industrialized countries. Not only do books, magazines and newspapers convey written information, but so do street signs, billboards, shop signs and even graffiti. Candy wrappers are covered in writing. The constant background presence of these products of "literacy technology" does not require active attention, but the information to be conveyed is ready for use at a glance. It is difficult to imagine modern life otherwise.

Silicon-based information technology, in contrast, is far from having become part of the environment. More than 50 million personal computers have been sold, and nonetheless the computer remains largely in a world of its own. It is approachable only through complex jargon that has nothing to do with the tasks for which which people actually use computers. The state of the art is perhaps analogous to the period when scribes had to know as much about making ink or baking clay as they did about writing.

The arcane aura that surrounds personal computers is not just a "user interface" problem. My colleagues and I at PARC think that the idea of a "personal" computer itself is misplaced, and that the vision of laptop machines, dynabooks and "knowledge navigators" is only a transitional step toward achieving the real potential of information technology. Such machines cannot truly make computing an integral, invisible part of the way people live their lives. Therefore we are trying to conceive a new way of thinking about computers in the world, one that takes into account the natural human environment and allows the computers themselves to vanish into the background.

Such a disappearance is a fundamental consequence not of technology, but of human psychology. Whenever people learn something sufficiently well, they cease to be aware of it. When you look at a street sign, for example, you absorb its information without consciously performing the act of reading. Computer scientist, economist, and Nobelist Herb Simon calls this phenomenon "compiling"; philosopher Michael Polanyi calls it the "tacit dimension"; psychologist TK Gibson calls it "visual invariants"; philosophers Georg Gadamer and Martin Heidegger call it "the horizon" and the "ready-to-hand", John Seely Brown at PARC calls it the "periphery". All say, in essence, that only when things disappear in this way are we freed to use them without thinking and so to focus beyond them on new goals.

The idea of integrating computers seamlessly into the world at large runs counter to a number of present-day trends. "Ubiquitous computing" in this context does not just mean computers that can be carried to the

beach, jungle or airport. Even the most powerful notebook computer, with access to a worldwide information network, still focuses attention on a single box. By analogy to writing, carrying a super-laptop is like owning just one very important book. Customizing this book, even writing millions of other books, does not begin to capture the real power of literacy.

Furthermore, although ubiquitous computers may employ sound and video in addition to text and graphics, that does not make them "multimedia computers." Today's multimedia machine makes the computer screen into a demanding focus of attention rather than allowing it to fade into the background.

Perhaps most diametrically opposed to our vision is the notion of "virtual reality," which attempts to make a world inside the computer. Users don special goggles that project an artificial scene on their eyes; they wear gloves or even body suits that sense their motions and gestures so that they can move about and manipulate virtual objects. Although it may have its purpose in allowing people to explore realms otherwise inaccessible -- the insides of cells, the surfaces of distant planets, the information web of complex databases -- virtual reality is only a map, not a territory. It excludes desks, offices, other people not wearing goggles and body suits, weather, grass, trees, walks, chance encounters and in general the infinite richness of the universe. Virtual reality focuses an enormous apparatus on simulating the world rather than on invisibly enhancing the world that already exists.

Indeed, the opposition between the notion of virtual reality and ubiquitous, invisible computing is so strong that some of us use the term "embodied virtuality" to refer to the process of drawing computers out of their electronic shells. The "virtuality" of computer-readable data -- all the different ways in which it can be altered, processed and analyzed -- is brought into the physical world.

How do technologies disappear into the background? The vanishing of electric motors may serve as an instructive example: At the turn of the century, a typical workshop or factory contained a single engine that drove dozens or hundreds of different machines through a system of shafts and pulleys. Cheap, small, efficient electric motors made it possible first to give each machine or tool its own source of motive force, then to put many motors into a single machine.

A glance through the shop manual of a typical automobile, for example, reveals twenty-two motors and twenty-five more solenoids. They start the engine, clean the windshield, lock and unlock the doors, and so on. By paying careful attention it might be possible to know whenever one activated a motor, but there would be no point to it.

Most of the computers that participate in embodied virtuality will be invisible in fact as well as in metaphor. Already computers in light switches, thermostats, stereos and ovens help to activate the world. These machines and more will be interconnected in a ubiquitous network. As computer scientists, however, my colleagues and I have focused on devices that transmit and display information more directly. We have found two issues of crucial importance: location and scale. Little is more basic to human perception than physical juxtaposition, and so ubiquitous computers must know where they are. (Today's computers, in contrast, have no idea of their location and surroundings.) If a computer merely knows what room it is in, it can adapt its behavior in significant ways without requiring even a hint of artificial intelligence.

Ubiquitous computers will also come in different sizes, each suited to a particular task. My colleagues and I have built what we call tabs, pads and boards: inch-scale machines that approximate active Post-It notes, foot-scale ones that behave something like a sheet of paper (or a book or a magazine), and yard-scale displays that are the equivalent of a blackboard or bulletin board.

How many tabs, pads, and board-sized writing and display surfaces are there in a typical room? Look around you: at the inch scale include wall notes, titles on book spines, labels on controls, thermostats and clocks, as well as small pieces of paper. Depending upon the room you may see more than a hundred tabs, ten or twenty pads, and one or two boards. This leads to our goals for initially deploying the hardware of embodied virtuality: hundreds of computers per room.

Hundreds of computers in a room could seem intimidating at first, just as hundreds of volts coursing through wires in the walls did at one time. But like the wires in the walls, these hundreds of computers will come to be invisible to common awareness. People will simply use them unconsciously to accomplish everyday tasks.

Tabs are the smallest components of embodied virtuality. Because they are interconnected, tabs will expand on the usefulness of existing inch-scale computers such as the pocket calculator and the pocket organizer. Tabs will also take on functions that no computer performs today. For example, Olivetti Cambridge Research Labs pioneered active badges, and now computer scientists at PARC and other research laboratories around the world are working with these clip-on computers roughly the size of an employee ID card. These badges can identify themselves to receivers placed throughout a building, thus making it possible to keep track of the people or objects to which they are attached.

In our experimental embodied virtuality, doors open only to the right badge wearer, rooms greet people by name, telephone calls can be automatically forwarded to wherever the recipient may be, receptionists actually know where people are, computer terminals retrieve the preferences of whoever is sitting at them, and appointment diaries write themselves. No revolution in artificial intelligence is needed--just the proper imbedding of computers into the everyday world. The automatic diary shows how such a simple thing as knowing where people are can yield complex dividends: meetings, for example, consist of several people spending time in the same room, and the subject of a meeting is most likely the files called up on that room's display screen while the people are there.

My colleague Roy Want has designed a tab incorporating a small display that can serve simultaneously as an active badge, calendar and diary. It will also act as an extension to computer screens: instead of shrinking a program window down to a small icon on the screen, for example, a user will be able to shrink the window onto a tab display. This will leave the screen free for information and also let people arrange their computer-based projects in the area around their terminals, much as they now arrange paper-based projects in piles on desks and tables. Carrying a project to a different office for discussion is as simple as gathering up its tabs; the associated programs and files can be called up on any terminal.

The next step up in size is the pad, something of a cross between a sheet of paper and current laptop and palmtop computers. Bob Krivacic at PARC has built a prototype pad that uses two microprocessors, a workstation-sized display, a multi-button stylus, and a radio network that can potentially handle hundreds of devices per person per room.

Pads differ from conventional portable computers in one crucial way. Whereas portable computers go everywhere with their owners, the pad that must be carried from place to place is a failure. Pads are intended to be "scrap computers" (analogous to scrap paper) that can be grabbed and used anywhere; they have no individualized identity or importance.

One way to think of pads is as an antidote to windows. Windows were invented at PARC and popularized by Apple in the Macintosh as a way of fitting several different activities onto the small space of a computer screen at the same time. In twenty years computer screens have not grown much larger. Computer window systems are often said to be based on the desktop metaphor--but who would ever use a desk whose surface area is only 9" by 11"?

Pads, in contrast, use a real desk. Spread many electronic pads around on the desk, just as you spread out papers. Have many tasks in front of you and use the pads as reminders. Go beyond the desk to drawers, shelves, coffee tables. Spread the many parts of the many tasks of the day out in front of you to fit both the task and the reach of your arms and eyes, rather than to fit the limitations of CRT glass-blowing. Someday pads may even be as small and light as actual paper, but meanwhile they can fulfill many more of paper's functions than can computer screens.

Yard-size displays (boards) serve a number of purposes: in the home, video screens and bulletin boards; in the office, bulletin boards, whiteboards or flip charts. A board might also serve as an electronic bookcase from which one might download texts to a pad or tab. For the time being, however, the ability to pull out a book and place it comfortably on one's lap remains one of the many attractions of paper. Similar objections apply to using a board as a desktop; people will have to get used to using pads and tabs on a desk as an adjunct to computer screens before taking embodied virtuality even further.

Boards built by Richard Bruce and Scott Elrod at PARC currently measure about 40 by 60 inches and display 1024x768 black-and-white pixels. To manipulate the display, users pick up a piece of wireless electronic "chalk" that can work either in contact with the surface or from a distance. Some researchers, using themselves and their colleagues as guinea pigs, can hold electronically mediated meetings or engage in other forms of collaboration around a liveboard. Others use the boards as testbeds for improved display hardware, new "chalk" and interactive software.

For both obvious and subtle reasons, the software that animates a large, shared display and its electronic chalk is not the same as that for a workstation. Switching back and forth between chalk and keyboard may involve walking several steps, and so the act is qualitatively different from using a keyboard and mouse. In addition, body size is an issue -- not everyone can reach the top of the board, so a Macintosh-style menu bar may not be a good idea.

We have built enough liveboards to permit casual use: they have been placed in ordinary conference rooms and open areas, and no one need sign up or give advance notice before using them. By building and using these boards, researchers start to experience and so understand a world in which computer interaction casually enhances every room. Liveboards can usefully be shared across rooms as well as within them. In experiments instigated by Paul Dourish of EuroPARC and Sara Bly and Frank Halasz of PARC, groups at widely separated sites gathered around boards -- each displaying the same image -- and jointly composed pictures and drawings. They have even shared two boards across the Atlantic.

Liveboards can also be used as bulletin boards. There is already too much data for people to read and comprehend all of it, and so Marvin Theimer and David Nichols at PARC have built a prototype system that attunes its public information to the people reading it. Their "scoreboard" requires little or no interaction from the user other than to look and to wear an active badge.

Prototype tabs, pads and boards are just the beginning of ubiquitous computing. The real power of the concept comes not from any one of these devices; it emerges from the interaction of all of them. The hundreds of processors and displays are not a "user interface" like a mouse and windows, just a pleasant and effective "place" to get things done.

What will be most pleasant and effective is that tabs can animate objects previously inert. They can beep to help locate mislaid papers, books or other items. File drawers can open and show the desired folder -- no searching. Tabs in library catalogs can make active maps to any book and guide searchers to it, even if it is off the shelf and on a table from the last reader.

In presentations, the size of text on overhead slides, the volume of the amplified voice, even the amount of ambient light, can be determined not by accident or guess but by the desires of the listeners in the room at that moment. Software tools for instant votes and consensus checking are already in specialized use in electronic meeting rooms of large corporations; tabs can make them widespread.

The technology required for ubiquitous computing comes in three parts: cheap, low-power computers that include equally convenient displays, a network that ties them all together, and software systems implementing ubiquitous applications. Current trends suggest that the first requirement will easily be met. Flat-panel displays containing 640x480 black-and-white pixels are now common. This is the standard size for PC's and is also about right for television. As long as laptop, palmtop and notebook computers continue to grow in popularity, display prices will fall, and resolution and quality will rise. By the end of the decade, a 1000x800-pixel high-contrast display will be a fraction of a centimeter thick and weigh perhaps 100 grams. A small battery will provide several days of continuous use.

Larger displays are a somewhat different issue. If an interactive computer screen is to match a whiteboard in usefulness, it must be viewable from arm's length as well as from across a room. For close viewing the density of picture elements should be no worse than on a standard computer screen, about 80 per inch. Maintaining a density of 80 pixels per inch over an area several feet on a side implies displaying tens of millions of pixels. The biggest computer screen made today has only about one fourth this capacity. Such large displays will probably be expensive, but they should certainly be available.

Central-processing unit speeds, meanwhile, reached a million instructions per second in 1986 and continue to double each year. Some industry observers believe that this exponential growth in raw chip speed may begin to level off about 1994, but that other measures of performance, including power consumption and auxiliary functions, will still improve. The 100-gram flat-panel display, then, might be driven by a single microprocessor chip that executes a billion operations per second and contains 16 megabytes of onboard memory along with sound, video and network interfaces. Such a processor would draw, on average, a few percent of the power required by the display.

Auxiliary storage devices will augment the memory capacity. Conservative extrapolation of current technology suggests that match-book size removable hard disks (or the equivalent nonvolatile memory chips) will store about 60 megabytes each. Larger disks containing several gigabytes of information will be standard, and terabyte storage -- roughly the capacity of the Library of Congress -- will be common. Such enormous stores will not necessarily be filled to capacity with usable information. Abundant space will, however, allow radically different strategies of information management. A terabyte of space makes deleting old files virtually unnecessary, for example.

Although processors and displays should be capable of offering ubiquitous computing by the end of the decade, trends in software and network technology are more problematic. Software systems today barely take any advantage of the computer network. Trends in "distributed computing" are to make networks appear like disks, memory, or other non-networked devices, rather than to exploit the unique capabilities of physical dispersion. The challenges show up in the design of operating systems and window systems.

Today's operating systems, like DOS and Unix, assume a relatively fixed configuration of hardware and software at their core. This makes sense for both mainframes and personal computers, because hardware or operating system software cannot reasonably be added without shutting down the machine. But in an embodied virtuality, local devices come and go, and depend upon the room and the people in it. New software for new devices may be needed at any time, and you'll never be able to shut off everything in the room at once. Experimental "micro-kernel" operating systems, such as those developed by Rick Rashid at Carnegie-Mellon University and Andy Tanenbaum at Vrije University in Amsterdam, offer one solution. Future operating systems based around tiny kernels of functionality may automatically shrink and grow to fit the dynamically changing needs of ubiquitous computing.

Today's window systems, like Windows 3.0 and the X Window System, assume a fixed base computer on which information will be displayed. Although they can handle multiple screens, they do not do well with applications that start out in one place (screen, computer, or room) and then move to another. For higher performance they assume a fixed screen and input mode and use the local computer to store information about the application--if any of these change, the window system stops working for that application. Even window systems like X that were designed for use over networks have this problem--X still assumes that an application once started stays put. The solutions to this problem are in their infancy. Systems for shared windows, such as those from Brown University and Hewlett-Packard Corporation, help with windows, but have problems of performance, and do not work for all applications. There are no systems that do well with the diversity of inputs to be found in an embodied virtuality. A more general solution will require changing the kinds of protocols by which application programs and windows interact.

The network connecting these computers has its own challenges. On the one hand, data transmission rates for both wired and wireless networks are increasing rapidly. Access to gigabit-per-second wired nets is already possible, although expensive, and will become progressively cheaper. (Gigabit networks will seldom devote all of their bandwidth to a single data stream; instead, they will allow enormous numbers of lower-speed transmissions to proceed simultaneously.) Small wireless networks, based on digital cellular telephone principles, currently offer data rates between two and 10 megabits per second over a range of a few hundred meters. Low-power wireless networks transmitting 250,000 bits per second to each station will eventually be available commercially.

On the other hand, the transparent linking of wired and wireless networks is an unsolved problem. Although some stop-gap methods have been developed, engineers must develop new communication protocols that explicitly recognize the concept of machines that move in physical space. Furthermore the number of channels envisioned in most wireless network schemes is still very small, and the range large (50-100 meters), so that the total number of mobile devices is severely limited. The ability of such a system to support hundreds of machines in every room is out of the question. Single-room networks based on infrared or newer electromagnetic technologies have enough channel capacity for ubiquitous computers, but they can only work indoors.

Present technologies would require a mobile device to have three different network connections: tiny range wireless, long range wireless, and very high speed wired. A single kind of network connection that can somehow serve all three functions has yet to be invented.

Neither an explication of the principles of ubiquitous computing nor a list of the technologies involved really gives a sense of what it would be like to live in a world full of invisible widgets. To extrapolate from today's rudimentary fragments of embodied virtuality resembles an attempt to predict the publication of Finnegans Wake after just having invented writing on clay tablets. Nevertheless the effort is probably worthwhile:

Sal awakens: she smells coffee. A few minutes ago her alarm clock, alerted by her restless rolling before waking, had quietly asked "coffee?", and she had mumbled "yes." "Yes" and "no" are the only words it knows.

Sal looks out her windows at her neighborhood. Sunlight and a fence are visible through one, but through others she sees electronic trails that have been kept for her of neighbors coming and going during the early morning. Privacy conventions and practical data rates prevent displaying video footage, but time markers and electronic tracks on the neighborhood map let Sal feel cozy in her street.

Glancing at the windows to her kids' rooms she can see that they got up 15 and 20 minutes ago and are already in the kitchen. Noticing that she is up, they start making more noise.

At breakfast Sal reads the news. She still prefers the paper form, as do most people. She spots an interesting quote from a columnist in the business section. She wipes her pen over the newspaper's name, date, section, and page number and then circles the quote. The pen sends a message to the paper, which transmits the quote to her office.

Electronic mail arrives from the company that made her garage door opener. She lost the instruction manual, and asked them for help. They have sent her a new manual, and also something unexpected -- a way to find the old one. According to the note, she can press a code into the opener and the missing manual will find itself. In the garage, she tracks a beeping noise to where the oil-stained manual had fallen behind some boxes. Sure enough, there is the tiny tab the manufacturer had affixed in the cover to try to avoid E-mail requests like her own.

On the way to work Sal glances in the foreview mirror to check the traffic. She spots a slowdown ahead, and also notices on a side street the telltale green in the foreview of a food shop, and a new one at that. She decides to take the next exit and get a cup of coffee while avoiding the jam.

Once Sal arrives at work, the foreview helps her to quickly find a parking spot. As she walks into the building the machines in her office prepare to log her in, but don't complete the sequence until she actually enters her office. On her way, she stops by the offices of four or five colleagues to exchange greetings and news.

Sal glances out her windows: a grey day in silicon valley, 75 percent humidity and 40 percent chance of afternoon showers; meanwhile, it has been a quiet morning at the East Coast office. Usually the activity indicator shows at least one spontaneous urgent meeting by now. She chooses not to shift the window on

the home office back three hours -- too much chance of being caught by surprise. But she knows others who do, usually people who never get a call from the East but just want to feel involved.

The telltale by the door that Sal programmed her first day on the job is blinking: fresh coffee. She heads for the coffee machine.

Coming back to her office, Sal picks up a tab and "waves" it to her friend Joe in the design group, with whom she is sharing a virtual office for a few weeks. They have a joint assignment on her latest project. Virtual office sharing can take many forms--in this case the two have given each other access to their location detectors and to each other's screen contents and location. Sal chooses to keep miniature versions of all Joe's tabs and pads in view and 3-dimensionally correct in a little suite of tabs in the back corner of her desk. She can't see what anything says, but she feels more in touch with his work when noticing the displays change out of the corner of her eye, and she can easily enlarge anything if necessary.

A blank tab on Sal's desk beeps, and displays the word "Joe" on it. She picks it up and gestures with it towards her liveboard. Joe wants to discuss a document with her, and now it shows up on the wall as she hears Joe's voice:

"I've been wrestling with this third paragraph all morning and it still has the wrong tone. Would you mind reading it?"

"No problem."

Sitting back and reading the paragraph, Sal wants to point to a word. She gestures again with the "Joe" tab onto a nearby pad, and then uses the stylus to circle the word she wants:

"I think it's this term 'ubiquitous'. Its just not in common enough use, and makes the whole thing sound a little formal. Can we rephrase the sentence to get rid of it?"

"I'll try that. Say, by the way Sal, did you ever hear from Mary Hausdorf?"

"No. Who's that?"

"You remember, she was at the meeting last week. She told me she was going to get in touch with you."

Sal doesn't remember Mary, but she does vaguely remember the meeting. She quickly starts a search for meetings in the past two weeks with more than 6 people not previously in meetings with her, and finds the one. The attendees' names pop up, and she sees Mary. As is common in meetings, Mary made some biographical information about herself available to the other attendees, and Sal sees some common background. She'll just send Mary a note and see what's up. Sal is glad Mary did not make the biography available only during the time of the meeting, as many people do...

In addition to showing some of the ways that computers can find their way invisibly into people's lives, this speculation points up some of the social issues that embodied virtuality will engender. Perhaps key among them is privacy: hundreds of computers in every room, all capable of sensing people near them and linked by high-speed networks, have the potential to make totalitarianism up to now seem like sheerest anarchy. Just as a workstation on a local-area network can be programmed to intercept messages meant for others, a single rogue tab in a room could potentially record everything that happened there.

Even today, although active badges and self-writing appointment diaries offer all kinds of convenience, in the wrong hands their information could be stifling. Not only corporate superiors or underlings, but overzealous government officials and even marketing firms could make unpleasant use of the same information that makes invisible computers so convenient.

Fortunately, cryptographic techniques already exist to secure messages from one ubiquitous computer to another and to safeguard private information stored in networked systems. If designed into systems from the outset, these techniques can ensure that private data does not become public. A well-implemented version of ubiquitous computing could even afford better privacy protection than exists today. For example, schemes based on "digital pseudonyms" could eliminate the need to give out items of personal information that are routinely entrusted to the wires today, such as credit card number, social security number and address.

Jim Morris of Carnegie-Mellon University has proposed an appealing general method for approaching these issues: build computer systems to have the same privacy safeguards as the real world, but no more, so that ethical conventions will apply regardless of setting. In the physical world, for example, burglars can break through a locked door, but they leave evidence in doing so. Computers built according to Morris's rule would not attempt to be utterly proof against cracker, but they would be impossible to enter without leaving the digital equivalent of fingerprints.

By pushing computers into the background, embodied virtuality will make individuals more aware of the people on the other ends of their computer links. This development carries the potential to reverse the unhealthy centripetal forces that conventional personal computers have introduced into life and the workplace. Even today, people holed up in windowless offices before glowing computer screens may not see their fellows for the better part of each day. And in virtual reality, the outside world and all its inhabitant effectively ceases to exist. Ubiquitous computers, in contrast, reside in the human world and pose no barrier to personal interactions. If anything, the transparent connections that they offer between different locations and times may tend to bring communities closer together.

My colleagues and I at PARC believe that what we call ubiquitous computing will gradually emerge as the dominant mode of computer access over the next twenty years. Like the personal computer, ubiquitous computing will enable nothing fundamentally new, but by making everything faster and easier to do, with less strain and mental gymnastics, it will transform what is apparently possible. Desktop publishing, for example, is fundamentally not different from computer typesetting, which dates back to the mid 1960's at least. But ease of use makes an enormous difference.

When almost every object either contains a computer or can have a tab attached to it, obtaining information will be trivial: "Who made that dress? Are there any more in the store? What was the name of the designer of that suit I liked last week?" The computing environment knows the suit you looked at for a long time last week because it knows both of your locations, and, it can retroactively find the designer's name even if it did not interest you at the time.

Sociologically, ubiquitous computing may mean the decline of the computer addict. In the 1910's and 1920's many people "hacked" on crystal sets to take advantage of the new high tech world of radio. Now crystal-and-cat's whisker receivers are rare, because radios are ubiquitous. In addition, embodied virtuality will bring computers to the presidents of industries and countries for nearly the first time. Computer access will penetrate all groups in society.

Most important, ubiquitous computers will help overcome the problem of information overload. There is more information available at our fingertips during a walk in the woods than in any computer system, yet people find a walk among trees relaxing and computers frustrating. Machines that fit the human environment, instead of forcing humans to enter theirs, will make using a computer as refreshing as taking a walk in the woods.