1

THE DNA OF INFORMATION

BITS AND ATOMS

The best way to appreciate the merits and consequences of *being digital* is to reflect on the difference between bits and atoms. While we are undoubtedly in an information age, most information is delivered to us in the form of atoms: newspapers, magazines, and books (like this one). Our economy may be moving toward an information economy, but we measure trade and we write our balance sheets with atoms in mind. GATT is about atoms.

I recently visited the headquarters of one of America's top five integrated circuit manufacturers. I was asked to sign in and, in the process, was asked whether I had a laptop computer with me. Of course I did. The receptionist asked for the model and serial number and for its value. "Roughly, between one and two million dollars," I said. "Oh, that cannot be, sir," she replied. "What do you mean? Let me see it." I showed her my old PowerBook and she estimated its value at $2,000. She wrote down that amount and I was allowed to enter the premises. The point is that while the atoms were not worth that much, the bits were almost priceless.

Not long ago I attended a management retreat for senior executives of PolyGram in Vancouver, British Columbia. The purpose was to enhance communications among senior management and to give everybody an overview of the year to come, including many samples of soon-to-be-released music, movies, games, and rock videos. These samples were to be shipped by FedEx to the meeting in the form of CDs, videocassettes, and CD-ROMs, physical material in real packages that have weight and size. By misfortune, some of the material was held up in customs. That same day, I had been in my hotel room shipping bits back and forth over the Internet, to and from MIT and elsewhere in the world. My bits, unlike PolyGram's atoms, were not caught in customs.

The information superhighway is about the global movement of weightless bits at the speed of light. As one industry after another looks at itself in the mirror and asks about its future in a digital world, that future is driven almost 100 percent by the ability of that company's product or services to be rendered in digital form. If you make cashmere sweaters or Chinese food, it will be a long time before we can convert them to bits. "Beam me up, Scotty" is a wonderful dream, but not likely

to come true for several centuries. Until then you will have to rely on FedEx, bicycles, and sneakers to get your atoms from one place to another. This is not to say that digital technologies will be of no help in design, manufacturing, marketing, and management of atom-based businesses. I am only saying that the core business won't change and your product won't have bits standing in for atoms.

In the information and entertainment industries, bits and atoms often are confused. Is the publisher of a book in the information delivery business (bits) or in the manufacturing business (atoms)? The historical answer is both, but that will change rapidly as information appliances become more ubiquitous and user-friendly. Right now it is hard, but not impossible, to compete with the qualities of a printed book.

A book has a high-contrast display, is lightweight, easy to "thumb" through, and not very expensive. But getting it to you includes shipping and inventory. In the case of textbooks, 45 percent of the cost is inventory, shipping, and returns. Worse, a book can go out of print. Digital books never go out of print. They are always there.

Other media has even more immediate risk and opportunity. The first entertainment atoms to be displaced and become bits will be those of videocassettes in the rental business, where consumers have the added inconvenience of having to return the atoms and being fined if they are forgotten under a couch ($3 billion of the $12 billion of the U.S. video rental business is said to be late fines). Other media will become digitally driven by the combined forces of convenience, economic imperative, and deregulation. And it will happen fast.

WHAT IS A BIT, ANYHOW?

A bit has no color, size, or weight, and it can travel at the speed of light. It is the smallest atomic element in the DNA of information. It is a state of being: on or off, true or false, up or down, in or out, black or white. For practical purposes we consider a bit to be a 1 or a 0. The meaning of the 1 or the 0 is a separate matter. In the early days of computing, a string of bits most commonly represented numerical information.

Try counting, but skip all the numbers that have anything other than a 1 and a 0 in them. You end up with the following: 1, 10, 11, 100, 101, 110, 111, etc. Those are the respective binary representations for the numbers 1, 2, 3, 4, 5, 6, 7, etc.

Bits have always been the underlying particle of digital computing, but over the past twenty-five years we have greatly expanded our binary vocabulary to include much more than just numbers. We have been able to digitize more and more types of information, like audio and video, rendering them into a similar reduction of 1s and 0s.

Digitizing a signal is to take samples of it, which, if closely spaced, can be used to play back a seemingly perfect replica. In an audio CD, for example, the sound has been sampled 44.1 thousand times a second. The audio waveform (sound pressure level measured as voltage) is recorded as discrete numbers (themselves turned into bits). Those bit strings, when played back 44.1 thousand times a second, provide a continuous-sounding rendition of the original music. The successive and discrete measures are so closely spaced in time that we cannot hear them as a staircase of separate sounds, but experience them as a continuous tone.

The same can be true for a black-and-white photograph. Imagine an electronic camera as laying a fine grid over an image and then recording the level of gray it sees in each cell. If we set the value of black to be 0 and the value of white to be 255, then any gray is somewhere between the two. Conveniently, a string of 8 bits (called a byte) has 256 permutations of 1s and 0s, starting with 00000000 and ending with 11111111. With such fine gradations and with a fine grid, you can perfectly reconstruct the picture for the human eye. As soon as you use a coarser grid or an insufficient number of gray levels, you start to see digital artifacts, like contours and blockiness.

The emergence of continuity from individual pixels is analogous to a similar phenomenon on a much finer scale in the familiar world of matter. Matter is made of atoms. If you could look at a smoothly polished metal surface at a subatomic scale, you would see mostly holes. It appears smooth and solid because the discrete pieces are so small. Likewise digital output.

But the world, as we experience it, is a very analog place. From a macroscopic point of view, it is not digital at all but continuous. Nothing goes suddenly on or off, turns from black to white, or changes from one state to another without going through a transition. This may not be true at microscopic level, where things that we interact with (electrons in a wire or photons in our eye) are discrete. But there are so many of them that we approximate them as continuous. There are, after all, roughly 1,000,000,000,000,000,000,000,000 atoms in this book (a very analog medium).

There are many merits to digitization. Some obvious ones include data compression and error correction, which is important in the delivery of information through a costly or noisy channel. Broadcasters, for example, can save money, and viewers can see or hear studio-quality picture and sound. But we are discovering that the consequences of being digital are far more important than those.

When using bits to describe sound and picture, there is a natural advantage to using as few bits as possible. This is akin to energy conservation. However, the number of bits you devote per second or per square inch relates directly to the fidelity of the music or image. Typically, one has an interest in digitizing at a very high resolution and then using a less-resolved version of the sound or picture for one application or another. For example, a color image might be digitized at very high resolution for final print copy but used at a lower resolution for a computer-aided page layout system. The economy of bits is driven in part by the constraints of the medium on which it is stored or through which it is delivered.

The number of bits that can be transmitted per second through a given channel (like copper wire, radio spectrum, or optical fiber) is the bandwidth of that channel. It is a measure of how many bits can get down a given pipe. That number or capacity needs to be matched carefully with the number of bits needed to render a given type of data (voice, music, video): 64,000 bits per second is more than ample for high-quality voice, 1.2 million bits per second is more than sufficient for high-fidelity music, and 45 million bits per second is terrific for rendering video.

Over the past fifteen years, however, we have learned how to compress the raw digital form of sound and picture by looking at the bits over time, space, or both and removing the intrinsic redundancies and repetitions. In fact, one of the reasons that all media has become digital so quickly is that we achieved very high levels of compression much sooner than most people predicted. In fact, as recently as 1993, some Europeans were arguing that digital video would not be a reality until the next millennium.

Five years ago, most people did not believe you could reduce the 45 million bits per second of raw digital video to 1.2 million bits per second. Yet in 1995 we can compress and decompress, encode and decode, video at that rate, inexpensively and with high quality. It is as if we suddenly have been able to make freeze-dried cappuccino, which is so good that by adding water, it comes back to us as rich and aromatic as any freshly brewed in an Italian café.

WHEN ALL MEDIA ARE BITS

Being digital allows you to deliver a signal with information added to correct errors such as telephone static, radio hiss, or television snow. These artifacts can be removed from the digital signal using a few extra bits and increasingly sophisticated error-correction techniques that are applied to one form of noise or another, in one medium or another. On your audio CD, one-third of the bits are used for error correction. Similar techniques can be applied to existing television so that each home receives studio-quality broadcast—so much clearer than what you get today that you could mistake it for so-called high definition.

Error correction and data compression are the two obvious reasons for digital television. You can put four studio-quality digital TV signals into the same bandwidth that previously accommodated one noisy analog television transmission. You deliver a better picture, and, using the same channel, you potentially get four times the audience share and advertising revenue.

Better and more efficient delivery of what already exists is what most media executives think and talk about in the context of being digital. But like the Trojan horse, the consequence of this gift will be surprising. Wholly new content will emerge from being digital, as will new players, new economic models, and a likely cottage industry of information and entertainment providers.

When all media is digital—because bits are bits—two fundamental and immediate results will be observed.

First, bits commingle effortlessly. They start to get mixed up and can be used and reused together or separately. The mixing of audio, video, and data is called *multimedia*; it sounds complicated, but is nothing more than commingled bits.

Second, a new kind of bit is born—a bit that tells you about the other bits. These new bits are typically "headers," which are well known to newspaper reporters who file "slugs" (which we never see) to identify a story. Such headers are also familiar to scientific authors who are asked to provide key words with their journal articles. These header bits can be a table of contents or a description of the data that follow. On your CD today, you have simple headers that allow you to skip from song to song and, in some cases, get more data about the music. These bits are not visible or audible but tell you, your computer, or a special-purpose entertainment appliance about the signal.

These two phenomena, commingled bits and bits-about-bits, change the media landscape so thoroughly that concepts like video-on-demand and shipping electronic games down your local cable are just trivial applications—the tip of a much more profound iceberg. Think about the consequences of a broadcast television show as data which includes a computer-readable description of itself. You could record based on content, not time of day or channel. Or, what about a single digital description that can generate a program in audio, video, or textual form at the receiver? And if moving these bits around is so effortless, what advantage would the large media companies have over you and me?

Being digital begs such questions. It creates the potential for new content to originate from a whole new combination of sources.

WHERE INTELLIGENCE LIVES

Broadcast television is an example of a medium in which all the intelligence is at the point of origin. The transmitter determines everything and the receiver just takes what it gets. In fact, per cubic inch, your current TV set is perhaps the dumbest appliance in your home (and I'm not even talking about the programs). If you have a microwave oven, it likely has more microprocessors than your TV. Instead of thinking of the next evolutionary step of television as increased resolution, better color, or more programs, think of it as a change in the distribution of intelligence—or, more precisely, the movement of some intelligence, from the transmitter to the receiver.

A newspaper is also produced with all the intelligence in the transmitter. But the medium of large-format paper provides some relief to the "sameness" of information, as it can be consumed differently, by different people, at different times. We browse and flip through pages, guided by headlines and pictures, each of us treating very differently the identical bits delivered to hundreds of thousands of people. The bits are the same, but the reading experience is different.

One way to look at the future of being digital is to ask if the quality of one medium can be transposed to another. Can the television experience be more like the newspaper experience? Many people think of newspapers as having more depth than television news. Must that be so? Similarly, television is considered a richer sensory experience than what newspapers can deliver. Must that be so?

The answer lies in creating computers to filter, sort, prioritize, and manage multimedia on our behalf—computers that read newspapers and look at television for us, and act as editors when we ask them to do so. This kind of intelligence can live in two different places.

It can live at the transmitter and behave as if you had your own staff writers—as if *The New York Times* were publishing a single newspaper tailored to your interests. In this first example, a small subset of bits has been selected especially for you. The bits are filtered, prepared, and delivered to you, perhaps to be printed in the home, perhaps to be viewed more interactively with an electronic display.

The second example is one in which your news-editing system lives in the receiver and *The New York Times* broadcasts a very large number of bits, perhaps five thousand different stories, from which your appliance grabs a select few, depending on your interests, habits, or plans for that day. In this instance, the intelligence is in the receiver, and the dumb transmitter is indiscriminately sending all the bits to everybody.

The future will not be one or the other, but both.

2

DEBUNKING BANDWIDTH

FROM A TRICKLE TO A DOWNPOUR

In the late 1960s, when I was an assistant professor of computer graphics, nobody knew what that was. Computers were totally outside everyday life. Today, I hear sixty-five-year-old tycoons boasting about how many bytes of memory they have in their Wizards or the capacity of their hard disks. Some people talk half-knowingly about the speed of their computer (thanks to the brilliant campaign of "Intel Inside") and affectionately (or not) about the flavor of their operating systems. I recently met one socialite, a wealthy and charming woman, who knew so much about Microsoft's operating system that she started a small business that provided consulting services for her "less-wired" peers. Her business card read, "I do Windows."

7

WHERE PEOPLE AND BITS MEET

FATAL REACTION

I spend a minimum of three hours a day in front of a computer and have done so for many years, and I still find it very frustrating at times. Understanding computers is about as easy as understanding a bank statement. Why do computers (and bank statements) have to be so needlessly complicated? Why is "being digital" so hard?

They don't, and it need not be. The evolution of computing has been so fast that we've only recently had enough low-cost computing power to spend it freely on improving the ease of interaction between you and your computer. It used to be considered wasteful and frivolous to devote time and money to the user interface, because computer cycles were so precious and had to be expended on the problem, not the person.

Scientists would justify stoic interfaces in many ways. In the early 1970s, for example, a handful of "scholarly" papers were published on why black-and-white displays were "better" than color. Color is not bad. The research community just wanted to vindicate its inability to deliver a good interface at reasonable cost or, to be a bit more cynical, at the expense of some imagination.

Those of us working on the human-computer interface in the late 1960s and throughout the 1970s were considered computer sissies and eyed with outright contempt. Our work was not the right stuff, even though the field was gaining acceptance. To remind yourself of how important sensing, effecting, and feedback can be, think of the last time you pushed an elevator button and the light did not come on (presumably because the bulb was burned out). The frustration is enormous: did it hear me? Interface design and function are very important.

In 1972 there were only 150,000 computers in the world, whereas five years from now, the integrated-circuit manufacturer Intel alone expects to be shipping 100 million each year (and I think they are vastly underestimating). Using a computer thirty years ago, like piloting a moon lander, was the realm of a precious few schooled in the hocus-pocus needed to drive these machines, sometimes with primitive languages or none at all (just toggle switches and blinking lights). In my opinion, there was a subconscious effort to keep it mysterious, like the monopoly of the monks or some bizarre religious rite in the Dark Ages.

We are still paying the price today.

When people talk about the look and feel of computers, they are referring to the graphical user interface, which "professionals" call a GUI. The GUI improved enormously starting around 1971 with work at Xerox and, shortly after, at MIT and a few other places, and it culminated in a real product a decade later when Steve Jobs had the wisdom and perseverance to introduce the Macintosh. The Mac was a major step forward in the marketplace and, by comparison, almost nothing has happened since. It took all the other computer companies more than five years to copy Apple and, in some cases, they have done so with inferior results, even today.

The history of human endeavors to make machines more usable is almost exclusively devoted to enhancing the sensory points of contact and evolving better physical designs. The interface was treated largely as a traditional industrial design problem. The designers of teapots and rakes might consider the handle in terms of shape, heat transfer, and the prevention of blisters.

Cockpit design is a daunting challenge, not only because there are so many switches, knobs, dials, and gauges, but also because two or three sensory inputs of a similar kind can interfere with one another. In 1972 an Eastern Airlines L1011 crashed because its landing gear was not down. The voice of the air-traffic controller and the beeping of the on-board computer resulted in the crew not hearing the warning message. Deadly interface design.

At home I used to have a very intelligent VCR with near-perfect voice recognition and knowledge of me. I could ask it to record programs by name and, in some cases, even assume it would do so automatically, without my asking. Then, all of a sudden, my son went to college.

I have not recorded a TV program in more than six years. Not because I can't. It is because the value is too low for the effort. It is needlessly hard. More important, VCR usage and remote-control units in general have been treated as a button-pushing problem. Likewise, the general interface with personal computers has been treated as a physical design problem. But interface is not just about the look and feel of a computer. It is about the creation of personality, the design of intelligence, and building machines that can recognize human expression.

A dog can recognize you from your gait more than one hundred yards away, whereas a computer does not even know you are there. Almost any pet can tell when you are angry, but a computer does not have a clue. Even puppies know when they have done wrong; computers don't.

The challenge for the next decade is not just to give people bigger screens, better sound quality, and easier-to-use graphical input devices. It is to make computers that know you, learn about your needs, and understand verbal and nonverbal languages. A computer should know the difference between your saying "Kissinger" and "kissing her," not because it can find the small acoustic difference, but because it can understand the meaning. That's good interface design.

The burden of interaction today has been placed totally on the shoulders of the human party. Something as banal as printing a computer file can be a debilitating exercise that resembles voodoo more than respectable human behavior. As a result, many adults are turned off and claim to be hopelessly computer illiterate.

This will change.

ODYSSEYS

In 1968 Arthur C. Clarke shared an Oscar nomination with Stanley Kubrick for the movie *2001: A Space Odyssey*. Oddly, the

movie came out before the book. Clarke was able to revise his manuscript after viewing the rushes (based upon an earlier version of the story). In a very real sense, Clarke was able to simulate his story line and refine his concepts. He was able to see and hear his ideas before committing them to print.

This may explain why HAL, the film's starring computer, was such a brilliant (if lethal) vision of a future human-computer interface. HAL (whose name does not come from the respective preceding letters of IBM) had a perfect command of speech (understanding and enunciating), excellent vision, and humor, which is the supreme test of intelligence.

Almost a quarter of a century passed before another example of interface excellence emerged: *The Knowledge Navigator*. This videotape, also a theatrical production, a so-called video prototype, was commissioned by the then CEO of Apple, John Sculley, whose own book was also called *Odyssey*. Sculley's book ended with ideas for a "knowledge navigator," later to become a video. He wanted to illustrate an interface of the future, beyond mice and menus. He did an excellent job.

The Knowledge Navigator depicts a flat book-like device open on the desk of a tweedy-looking professor. In one corner of its display is a bow-tied person who turns out to be the persona of the machine. The professor asks this agent to assist him in preparing for a lecture, delegates a handful of tasks, and on a couple of occasions is reminded of other matters. The agent can see, hear, and respond intelligently, like any human assistant.

What HAL and the Knowledge Navigator have in common is that they exhibit intelligence to such a degree that the physical interface itself almost goes away. Therein lies the secret to interface design: make it go away. When you meet somebody for the first time, you may be very conscious of their looks, speech, and gestures. But quickly the content of your communication dominates, even if it is largely expressed through tone of voice or the language of facial expressions. A good computer interface should behave similarly. The problem is less like designing a dashboard and more like designing a human.

Most interface designers, on the other hand, have been stubbornly trying to make dumb machines easier to use by smart people. They have taken their lead from the field called "human factors" in the United States and "ergonomics" in Europe, which is about how the human body uses its sensors and effectors to work with tools in its immediate surroundings.

A telephone handset is probably the most redesigned and overdesigned appliance on earth, yet remains utterly unsatisfactory. Cellular telephones make VCRs pale with their unusable interface. A Bang & Olufsen telephone is sculpture, not telephony, harder not easier to use than an old black rotary telephone.

Worse, telephone designs have been "featured" to death. Number storing, redialing, credit card management, call waiting, call forwarding, autoanswering, number screening, and on and on are constantly being squeezed onto the real estate of a thin appliance that fits in the palm of your hand, making it virtually impossible to use.

Not only do I not want all those features; I don't want to dial the telephone at all. Why can't telephone designers understand that none of us want to dial telephones? *We want to reach people on the telephone!*

Given half a chance, we would delegate that task, which suggests to me that the problem of a telephone may not be in the design of the handset, but in the design of a robot secretary that can fit in your pocket.

BEYOND ETCH-A-SKETCH

Computer interface design started in March 1960, when J. C. R. Licklider published his paper "Man-Computer Symbiosis." Lick (as he was called) was an experimental psychologist and acoustician by training who became a convert to and a messiah for computing, leading ARPA's initial computer efforts. He was asked in the mid-1960s to write an appendix for the Carnegie Commission report on the future of television. It was in that appendix that Lick coined the term *narrowcasting*. What Lick did not know at the time was that these two contributions, man-computer symbiosis and narrowcasting, were destined to converge in the 1990s.

Very early human-computer interface research, in the beginning of the 1960s, split into two parts that were not to rejoin for twenty years. One addressed interactivity, and the other focused on sensory richness.

Interactivity was tackled by solving the problem of sharing a computer, then an expensive and monolithic resource. In the 1950s and early 1960s, a computer was so valuable that you made every effort to keep it running nonstop. It was unthinkable to connect a keyboard and have, for example, the computer type out a question and then have it sit there idly while a human read, thought about, and replied to it. The invention, called time sharing, was a method by which multiple users could share a single machine from remote locations. If you divided the resource among, say, ten people, it was not just that each person could have one-tenth of the machine, but that one person's moment of reflection could be another person's full use of the computer.

Such sharing of the digital pie worked under the condition that no single user was a hog, needing vast amounts of computation or bandwidth. Early terminals ran at 110 baud. I remember vividly when that was increased to 300 baud; it seemed so fast.

By contrast, sensory richness was addressed with very high bandwidth graphical interaction. Early computer graphics required a machine fully dedicated to providing the image. It was no different in principle than today's personal computer, but it filled a large room and cost millions of dollars. Computer graphics was born as a line-drawing medium that demanded a great deal of computing power to control directly the beam of the cathode ray tube.

Only ten years later did computer graphics start to move away from a graphics of line drawing to a graphics of shapes and images. These new displays, called raster scan displays, required a lot of memory to store the picture point by point. They are so common today that most people don't know that they were originally considered heretical. (Almost nobody believed in 1970 that computer memory would ever be cheap enough to devote so much of it to graphics.)

Time sharing and computer graphics were poor bedfellows for the next two decades. Sensory-impoverished time-sharing systems emerged as the accepted tool for business and academic computing, giving birth to electronic banking and airline reservation systems, which we take for granted today. Commercial time sharing went hand in hand with very parsimonious interface design, usually typewriter output, almost purposely wanting the system to be slow enough for any single user, so others could get their fair share.

Computer graphics, on the other hand, developed for the most part with stand-alone computing. By 1968 so-called minicomputers in the $20,000 range started to emerge, largely because factory and machine automation needed very precise and real-time controls. So did computer graphics. Coupled with display devices, these stand-alone computer graphics systems were precursors to what we know today as workstations, which are nothing more than personal computers with long pants.

MULTIMODAL INTERFACE

Redundancy is generally considered a bad symptom, implying needless verbosity or careless repetition. In early human interface design, people studied techniques for interaction and tried to select judiciously one means or the other, for one set of circumstances or the other. Was a light pen better than a data tablet? The "either/or" mentality was driven by the false belief that there was a universal "best" solution for any given situation; it is false because people are different, situations change, and the circumstances of a particular interaction may well be driven by the channel you have available. There is no best in interface design.

I recall visiting an admiral in the mid-1970s who had one of the most advanced command-and-control systems. He would bark orders to a junior seaman, who would dutifully type in the proper commands. So, in this sense, the system had a terrific interface: it had speech-recognition facilities, and patience as well. The admiral could walk around the room, talk, and gesture. He could be himself.

Nonetheless, the admiral was unprepared to plan an attack through such an indirect interface. He knew the seaman was looking at the situation through the keyhole of the computer system's small display. The admiral preferred interacting directly with a large wall map of the "theater" on which he would thumbtack little blue and red ships of appropriate shape. (At the time we always joked about how the Russians were using the same colors.)

The admiral was comfortable with the map, not because it was old-fashioned and very high resolution but because he engaged with it with his whole body. As he moved ships, his gestures and motor actions reinforced his memory. He was deeply involved with the display, right down to his neck muscles. It was not an either/or interface; it was both/and.

Both/and led to a breakthrough in thinking, which, simply stated, said that redundancy is good. In fact, the best interface would have many different and concurrent channels of communication, through which a user could express and cull meaning from a number of different sensory devices (the user's and the machine's). Or, equally important, one channel of communication might provide the information missing in the other.

For example, if we are in a room with a handful of people, and I say to somebody, "What is your name?" the question really does not have any meaning unless you can see where I am looking while I speak. Namely, the adjective *your* gets its meaning from the direction of my gaze.

This was beautifully illustrated in a program called Put-That-There, developed at MIT by Dick Bolt and Chris Schmandt. The first embodiment of the program, in 1980, allowed you to speak and gesture at a wall-size display and move simple objects (later ships) around a blank screen (later the Caribbean). In a filmed demonstration of Put-That-There, it misrecognized a command. Schmandt's spontaneous "Oh, shit" was engraved in film stock to remind many future audiences of how much work was still to be done.

The idea is simple: talking, pointing, and looking should work together as part of a multimodal interface that is less about messaging back and forth (the basis of time sharing) and more like face-to-face, human-to-human conversation.

At the time, this and other early attempts at an alternate "both/and" approach to the interface design looked like sloppy science. I have little respect for testing and evaluation in interface research. My argument, perhaps arrogant, is that if you have to test something carefully to see the difference it makes, then it is not making enough of a difference in the first place.

THE NOTICEABLE DIFFERENCE

When I was a little boy, my mother had a linen closet, the back of which had a "secret wall." The secret was no big deal: a collection of pencil lines that we periodically and carefully made to mark my height. All the pencil lines were dutifully dated, and some were close together due to frequency and others were spread apart because, for example, we had been away for the summer. Using two closets did not seem to make sense.

This scale was a private matter, and I guess in some way measured my intake of milk, spinach, and other good things.

By contrast, growth has a more dramatic face. A rarely seen uncle might comment, "How much you have grown, Nicky?" (since he had seen me two years prior, we suppose). But I could not really comprehend the change. All I could see were the little lines in the linen closet.

The "just-noticeable difference," or JND, is a unit of measure in psychophysics. Its name alone has influenced human interface design. You have to ask yourself, if it is a JND, why bother? If you have to measure carefully to see any difference at all, maybe we are not working on things that matter enough.

For example, scholarly studies have suggested that speech and natural language are not appropriate channels of communication between people and computers in most applications. These technical reports are filled with tables, control groups, and the like, proving that natural language is confusing for human-computer communication.

While I certainly don't expect the pilot of a 747 to taxi and take off by singing "Up, Up and Away," I still can't fathom any reason not to use the richness of speech and gesture, even in the cockpit. Wherever the computer may be, the most effective interface design results from combining the forces of sensory richness and machine intelligence.

When this happens, we'll see a noticeable difference. We'll see what my uncle saw, instead of the little lines in the closet.

INTELLIGENT INTERFACES

My dream for the interface is that computers will be more like people. This idea is vulnerable to criticism for being too romantic, vague, or unrealizable. If anything, I would criticize it for shooting too low. There may be many exotic channels of communications of which we may not even be aware today. (As somebody married to an identical twin and with identical twin younger brothers, I am fully prepared to believe from observation that extrasensory communication is not out of the question.)

In the mid-1960s, I set my goals by trying to emulate face-to-face communication, with its languages of gesture and facial expressions and the motor involvement of our body and limbs. I used the admiral as my model.

In a landmark project called the Spatial Data Management System (circa 1976), the goal was to provide a human interface that would "bring computers directly to generals, presidents of companies, and six-year-old children." The system was designed to be learnable in thirty seconds. Familiarity with desktops and bookshelves was the tool used to browse and manipulate complex audio, video, and data.

That was radical for the late 1970s, but it still missed the more meaningful consequence of fashioning our communication after the conversation between the admiral and the seaman. Future human-computer interface will be rooted in delegation, not the vernacular of direct manipulation—pull down, pop up, click—and mouse interfaces. "Ease of use" has been such a compelling goal that we sometimes forget that many people don't want to use the machine at all. They want to get something done.

What we today call "agent-based interfaces" will emerge as the dominant means by which computers and people talk with one another. There will be specific points in space and time where bits get converted into atoms and the reverse. Whether that is the transmission of a liquid crystal or the reverberation of a speech generator, the interface will need size, shape, color, tone of voice, and all the other sensory paraphernalia.

8

GRAPHICAL PERSONA

GRAPHICS' BIG BANG

At MIT in 1963, the Ph.D. thesis of Ivan Sutherland, called "Sketchpad," exploded upon the world the idea of interactive computer graphics. Sketchpad was a real-time line-drawing system that allowed the user to interact directly with the computer screen by means of a "light pen." The achievement was of such magnitude and breadth that it took some of us a decade to understand and appreciate all of its contributions. Sketchpad introduced many new concepts: dynamic graphics, visual simulation, constraint resolution, pen tracking, and a virtually infinite coordinate system, just to name a few. Sketchpad was the big bang of computer graphics.

13

THE POST-INFORMATION AGE

BEYOND DEMOGRAPHICS

The transition from an industrial age to a post-industrial or information age has been discussed so much and for so long that we may not have noticed that we are passing into a post-information age. The industrial age, very much an age of atoms, gave us the concept of mass production, with the economies that come from manufacturing with uniform and repetitious methods in any one given space and time. The information age, the age of computers, showed us the same economies of scale, but with less regard for space and time. The manufacturing of bits could happen anywhere, at any time, and, for example, move among the stock markets of New York, London, and Tokyo as if they were three adjacent machine tools.

In the information age, mass media got bigger and smaller at the same time. New forms of broadcast like CNN and *USA Today* reached larger audiences and made broadcast broader. Niche magazines, videocassette sales, and cable services were examples of narrowcasting, catering to small demographic groups. Mass media got bigger and smaller at the same time.

In the post-information age, we often have an audience the size of one. Everything is made to order, and information is extremely personalized. A widely held assumption is that individualization is the extrapolation of narrowcasting—you go from a large to a small to a smaller group, ultimately to the individual. By the time you have my address, my marital status, my age, my income, my car brand, my purchases, my drinking habits, and my taxes, you have me—a demographic unit of one.

This line of reasoning completely misses the fundamental difference between narrowcasting and being digital. In being digital I am *me*, not a statistical subset. *Me* includes information and events that have no demographic or statistical meaning. Where my mother-in-law lives, whom I had dinner with last night, and what time my flight departs for Richmond this afternoon have absolutely no correlation or statistical basis from which to derive suitable narrowcast services.

But that unique information about me determines news services I might want to receive about a small obscure town, a not so famous person, and (for today) the anticipated weather conditions in Virginia. Classic demographics do not scale down to the digital individual. Thinking of the post-information age as infinitesimal demographics or ultrafocused narrowcasting is about as personalized as Burger King's "Have It Your Way."

True personalization is now upon us. It's not just a matter of selecting relish over mustard once. The post-information age

is about acquaintance over time: machines' understanding individuals with the same degree of subtlety (or more than) we can expect from other human beings, including idiosyncrasies (like always wearing a blue-striped shirt) and totally random events, good and bad, in the unfolding narrative of our lives.

For example, having heard from the liquor store's agent, a machine could call to your attention a sale on a particular Chardonnay or beer that it knows the guests you have coming to dinner tomorrow night liked last time. It could remind you to drop the car off at a garage near where you are going, because the car told it that it needs new tires. It could clip a review of a new restaurant because you are going to that city in ten days, and in the past you seemed to agree with that reviewer. All of these are based on a model of you as an individual, not as part of a group who might buy a certain brand of soapsuds or toothpaste.

PLACE WITHOUT SPACE

In the same ways that hypertext removes the limitations of the printed page, the post-information age will remove the limitations of geography. Digital living will include less and less dependence upon being in a specific place at a specific time, and the transmission of place itself will start to become possible.

If I really could look out the electronic window of my living room in Boston and see the Alps, hear the cowbells, and smell the (digital) manure in summer, in a way I am very much in Switzerland. If instead of going to work by driving my atoms into town, I log into my office and do my work electronically, exactly where is my workplace?

In the future, we will have the telecommunications and virtual reality technologies for a doctor in Houston to perform a delicate operation on a patient in Alaska. In the nearer term, however, a brain surgeon will need to be in the same operating theater at the same time as the brain; many activities, like those of so-called knowledge workers, are not as dependent on time and place and will be decoupled from geography much sooner.

Today, writers and money managers find it practicable and far more appealing to be in the Caribbean or South Pacific while preparing their manuscripts or managing their funds. However, in some countries, like Japan, it will take longer to move away from space and time dependence, because the native culture fights the trend. (For example: one of the main reasons that Japan does not move to daylight savings time in the summer is because going home "after dark" is considered necessary, and workers try not to arrive after or go home before their bosses.)

In the post-information age, since you may live and work at one or many locations, the concept of an "address" now takes on new meaning.

When you have an account with America Online, CompuServe, or Prodigy, you know your own e-mail address, but you do not know where it physically exists. In the case of America Online, your Internet address is your ID followed by @aol.com—usable anywhere in the world. Not only do you not know where @aol.com is, whosoever sends a message to that address has no idea of where either it or you might be. The address becomes much more like a Social Security number than a street coordinate. It is a virtual address.

In my case, I happen to know where my address, @hq.media.mit.edu, is physically located. It is a ten-year-old HP-Unix machine in a closet near my office. But when people

send me messages they are sending them to me, not to that closet. They might infer I am in Boston (which is usually not the case). In fact, I am usually in a different time zone, so not only space but time is shifted as well.

BEING ASYNCHRONOUS

A face-to-face or telephone conversation is real time and synchronous. Telephone tag is a game played to find the opportunity to be synchronous. Ironically, this is often done for exchanges, which themselves require no synchrony whatsoever, and could just as well be handled by non-real-time message passing. Historically, asynchronous communication, like letter writing, has tended to be more formal and less off-the-cuff exchanges. This is changing with voice mail and answering machines.

I have met people who claim they cannot understand how they (and we all) lived without answering machines at home and voice mail at the office. The advantage is less about voice and more about off-line processing and time shifting. It is about leaving messages versus engaging somebody needlessly in online discussion. In fact, answering machines are designed slightly backward. They should not only activate when you are not there or don't want to be there, but they should *always* answer the telephone and give the caller the opportunity to simply leave a message.

One of the enormous attractions of e-mail is that it is not interruptive like a telephone. You can process it at your leisure, and for this reason you may reply to messages that would not stand a chance in hell of getting through the secretarial defenses of corporate, telephonic life.

E-mail is exploding in popularity because it is *both* an asynchronous and a computer-readable medium. The latter is particularly important, because interface agents will use those bits to prioritize and deliver messages differently. Who sent the message and what it is about could determine the order in which you see it—no different from the current secretarial screening that allows a call from your six-year-old daughter to go right through, while the CEO of the XYZ Corporation is put on hold. Even on a busy workday, personal e-mail messages might drift to the top of the heap.

Not nearly as much of our communications need to be contemporaneous or in real time. We are constantly interrupted or forced into being punctual for things that truly do not merit such immediacy or promptness. We are forced into regular rhythms, not because we finished eating at 8:59 p.m., but because the TV program is about to start in one minute. Our great-grandchildren will understand our going to the theater at a given hour to benefit from the collective presence of human actors, but they will not understand the synchronous experiencing of television signals in the privacy of our home—until they look at the bizarre economic model behind it.

DEMANDING ON DEMAND

Digital life will include very little real-time broadcast. As broadcast becomes digital, the bits are not only easily time-shiftable but

need not be received in the same order or at the same rate as they will be consumed. For example, it will be possible to deliver one hour of video over fiber in a fraction of a second (some experiments today show that the time needed to deliver one hour of VHS-quality video can be as small as one-hundredth of a second). Alternately, over a thin wire or narrow radio frequency, you might use six hours of broadcast time overnight to transmit a ten-minute (personalized) video news program. The former is blasting the bits into your computer and the latter is trickle-charging it.

With the possible exception of sports and elections, technology suggests that TV and radio of the future will be delivered asynchronously. This will happen either on demand or using "broadcatching," a term coined in 1987 by Stewart Brand in his book about the Media Lab. Broadcatching is the radiation of a bit stream, most likely one with vast amounts of information pushed into the ether or down a fiber. At the receiving end, a computer catches the bits, examines them, and discards all but the few it thinks you want to consume later.

On-demand information will dominate digital life. We will ask explicitly and implicitly for what we want, when we want it. This will require a radical rethinking of advertiser-supported programming.

In 1983, when we started the Media Lab, people felt that the word *media* was pejorative, representing a one-way path to the lowest common denominator in American culture. Media, with a capital *M*, almost exclusively meant "mass media." A large audience would bring in large advertising bucks, which in turn would underwrite large production budgets. Advertising was further justified in "over-the-air" mass media on the basis that information and entertainment should be "free" to the viewer, since spectrum is public property.

Magazines, on the other hand, use a private network of distribution and share the cost between advertiser and reader. Magazines, a notably asynchronous medium, offer a much wider range of economic and demographic models and may in fact be a bellwether for the future of television. The proliferation into niche markets has not necessarily ruptured content, but it has shifted some of the burden of cost to the subscriber. In some specialty magazines there is no advertising at all.

In future digital media there will be more pay-per-view, not just on an all-or-nothing basis, but more like newspapers and magazines, where you share the cost with advertisers. In some cases, the consumer may have an option to receive material without advertising but at a higher cost. In other cases, the advertising will be so personalized that it is indistinguishable from news. It is news.

The economic models of media today are based almost exclusively on "pushing" the information and entertainment out into the public. Tomorrow's will have as much or more to do with "pulling," where you and I reach into the network and check out something the way we do in a library or video-rental store today. This can happen explicitly by asking or implicitly by an agent asking on your behalf.

The on-demand model without advertising will make content production more like theatrical Hollywood movies, with much higher risk but much higher reward as well. There will be big busts and wild successes. Make it, and they will come. If they do, great; if they don't, too bad, but Procter & Gamble is not necessarily underwriting the risk. In this sense, media

companies will be throwing bigger dice tomorrow than they do today. But there will also be smaller players, throwing smaller dice, getting part of the audience share.

The "prime" of prime time will be its quality in our eyes, not those of some collective and demographic mass of potential buyers of a new luxury car or dishwashing detergent.