

Also by Nicholas Carr

Does IT Matter?

The
Big
Switch

Rewiring the World, From Edison to Google

Nicholas Carr



W. W. NORTON & COMPANY

New York London

Excerpt from "All Watched Over by Machines of Loving Grace" from *The Pill Versus the Springhill Mine Disaster* by Richard Brautigan. Copyright © 1968 by Richard Brautigan. Reprinted by permission of Houghton Mifflin Company.
All rights reserved.

Copyright © 2009, 2008 by Nicholas Carr

All rights reserved
Printed in the United States of America
First published as a Norton paperback 2009

For information about permission to reproduce
selections from this book, write to Permissions,
W. W. Norton & Company, Inc., 500 Fifth Avenue, New York, NY 10110

For information about special discounts for bulk purchases, please contact
W. W. Norton Special Sales at specialsales@wwnorton.com or 800-233-4830

Manufacturing by Courier Westford
Book design by Chris Welch
Production manager: Anna Oler

Library of Congress Cataloging-in-Publication Data

Carr, Nicholas G., 1959—
The big switch : rewiring the world, from Edison to Google / Nicholas
Carr. — 1st ed.
p. cm.

Includes bibliographical references and index.

ISBN 978-0-393-06228-1 (hardcover)

1. Computers and civilization. 2. Information technology—
Social aspects. 3. Technological innovations. 4. Internet. I. Title.

QA76.9.C66C38 2008

305.48'34—dc22

2007038084

ISBN 978-0-393-33394-7 pbk.

W. W. Norton & Company, Inc.
500 Fifth Avenue, New York, N.Y. 10110
www.wwnorton.com

W. W. Norton & Company Ltd.
Castle House, 75/76 Wells Street, London W1T 3QT

3 4 5 6 7 8 9 0

One lingered long among the dynamos,
for they were new, and they gave to history a new phase.

—Henry Adams

iGod

IN THE SUMMER of 2004, Google's founders, Larry Page and Sergey Brin, nearly sabotaged their own company. They sat down for a long interview with *Playboy*, and the magazine published the transcript in early August, just days before Google's scheduled debut on the NASDAQ stock exchange. The appearance of the interview roiled Wall Street, as it seemed to violate the Securities and Exchange Commission's prohibition on unauthorized disclosures of information during the "quiet period" leading up to an IPO. Investors feared that the SEC might force the company to cancel its stock offering. But after Google hurriedly distributed a revised prospectus, including the entire text of the *Playboy* interview as an appendix, the SEC cleared the stock sale, and on August 19 Google became a public company.

Lost in the hubbub was the interview itself, which provided a fascinating look into the thoughts and motivations of a pair of brilliant young mathematicians who were about to join the ranks of the world's wealthiest and most powerful businessmen. Toward the end of the interview, Page and Brin gave voice to their deepest ambition. They weren't just interested in perfecting their search engine, they said. What they really looked forward to was melding

their technology with the human brain itself. "You want access to as much [information] as possible so you can discern what is most relevant and correct," explained Brin. "The solution isn't to limit the information you receive. Ultimately you want to have the entire world's knowledge connected directly to your mind."

The interviewer was taken aback. "Is that what we have to look forward to?" he asked.

"I hope so," said Brin. "At least a version of that. We probably won't be looking up everything on a computer."

The interviewer probed again: "Is your goal to have the entire world's knowledge connected directly to our minds?"

"To get closer to that—as close as possible," replied Brin. "The smarter we can make the search engine, the better. Where will it lead? Who knows? But it's credible to imagine a leap as great as that from hunting through library stacks to a Google session, when we leap from today's search engines to having the entirety of the world's information as just one of our thoughts."

It wasn't the first time that Brin and Page had talked about their desire to tinker with the human brain—and it wouldn't be the last. In fact, the creation of an artificial intelligence that extends or even replaces the mind is a theme they return to again and again. "Every time I talk about Google's future with Larry Page," reports Steve Jurvetson, a prominent Silicon Valley venture capitalist, "he argues that it will become an artificial intelligence." During a question-and-answer session after a presentation at his alma mater, Stanford University, in May 2002, Page said that Google would fulfill its mission only when its search engine was "AI-complete." "You guys know what that means?" he quizzed the audience of students. "That's artificial intelligence."

In another presentation at Stanford a few months later, Page reiterated the goal: "The ultimate search engine is something as smart

as people—or smarter. . . . For us, working on search is a way to work on artificial intelligence." Around the same time, in an interview on public television's *NewsHour*, Brin explained that the "ultimate search engine" would resemble the talking supercomputer HAL in the movie *2001: A Space Odyssey*. "Now, hopefully," said Brin, "it would never have a bug like HAL did where he killed the occupants of the spaceship. But that's what we're striving for, and I think we've made it a part of the way there."

In July 2003, during a talk at a technology conference, Brin and Page went into more detail about their aspiration to use artificial intelligence to make us smarter. Brin suggested, according to a report from a member of the audience, that "wireless brain appliances" might be used to automate the delivery of information. Page elaborated on that idea in a February 2004 interview with Reuters, saying, "On the more exciting front, you can imagine your brain being augmented by Google. For example you think about something and your cell phone could whisper the answer into your ear."

Brin also discussed Google's progress toward its ultimate goal in an interview with *Newsweek* writer Steven Levy. "I think we're pretty far along [with Internet searching] compared to ten years ago," he said. "At the same time, where can you go? Certainly if you had all the world's information directly attached to your brain, or an artificial brain that was smarter than your brain, you'd be better off. Between that and today, there's plenty of space to cover." David Vise relates a similar remark by Brin in his 2005 book *The Google Story*. "Why not improve the brain?" Brin muses at one point. "Perhaps in the future, we can attach a little version of Google that you just plug into your brain."

At a London conference in May 2006, Larry Page again spoke of Google's pursuit of artificial intelligence. "We want to create the ultimate search engine," he said. "The ultimate search engine

would understand everything in the world." A year later, in February 2007, he told a group of scientists that Google has a team of employees who are "really trying to build an artificial intelligence and to do it on a large scale." The fulfillment of their goal, he said, is "not as far off as people think."

In taking a transcendental view of information technology, seeing it as a way to overcome what they perceive to be the physical limitations of the human brain, Brin and Page are expressing a desire that has long been a hallmark of the mathematicians and computer scientists who have devoted themselves to the creation of artificial intelligence. It's a desire that, as David Noble notes in *The Religion of Technology*, can be traced all the way back to the seventeenth-century French philosopher René Descartes, who argued that "the body is always a hindrance to the mind in its thinking" and saw in mathematics a model for "pure understanding." The Cartesian ideal runs through the work of mathematicians like George Boole, Alfred North Whitehead, and Alan Turing, whose breakthroughs in algebraic logic set the stage for the modern binary computer.

In her 1979 book *Machines Who Think*, Pamela McCorduck wrote that artificial intelligence promises to provide "an extension of those human capacities we value most." She quoted MIT professor Edward Fredkin's claim that "artificial intelligence is the next step in evolution." Danny Hillis, whose pioneering work in parallel computing paved the way for Google's systems, argued in a 1992 interview that AI could provide a means of remedying man's mental shortcomings, of fixing the "bugs left over history, back from when we were animals," and lead to the creation of beings who are "better than us." In "Reinventing Humanity," a 2006 article, the acclaimed inventor and author Ray Kurzweil predicted that artificial intelligence "will vastly exceed biological intelligence by the mid-2040s," resulting in "a world where there is no distinction

between the biological and the mechanical, or between physical and virtual reality."

To most of us, the desire of the AI advocates to merge computers and people, to erase or blur the boundary between man and machine, is troubling. It's not just that we detect in their enthusiasm a disturbing misanthropy—Hillis dismisses the human body as "the monkey that walks around," while Marvin Minsky, the former director of MIT's artificial intelligence program, calls the human brain a "bloody mess of organic matter"—it's also that we naturally sense in their quest a threat to our integrity as freethinking individuals. Even Bill Gates finds the concept discomforting. In a 2005 talk in Singapore, he discussed the possibility of connecting people's bodies and brains directly to computers. One of his Microsoft colleagues, he told the audience, "always says to me, 'I'm ready, plug me in.'" But Gates said that he was wary of the idea: "I don't feel quite the same way. I'm happy to have the computer over there and I'm over here."

In addition to finding the prospect of being turned into computer-enhanced cyborgs unsettling, we also tend to be skeptical of the idea. It seems far-fetched, even ludicrous—like something out of a particularly fanciful piece of science fiction. Here, though, we part company with Gates. In that same speech, he made it clear that he believes the blending of computers and people is inevitable, that we will, in the foreseeable future, come to be augmented by digital processors and software. "We will have those capabilities," he declared. And evidence suggests that Microsoft, like Google, aims to be a pioneer in creating human-computer interfaces for commercial gain. In 2004, the company was granted a patent for a "method and apparatus for transmitting power and data using the human body." In its filing, Microsoft described how it is developing technology that will turn skin into a new kind of electrical conduit,

or “bus,” that can be used to connect “a network of devices coupled to a single body.” It also noted that “the network can be extended by connecting multiple bodies through physical contact [such as] a handshake. When two or more bodies are connected physically, the linked bodies form one large bus over which power and/or communications signals can be transmitted.”

Microsoft’s patent is just one manifestation of the many corporate and academic research programs that are aimed at merging computers and people and, in particular, at incorporating human beings more fully into the Internet’s computing web. A 2006 study sponsored by the British government’s Office of Science and Innovation surveyed some of the most promising of these initiatives. In addition to confirming that our bodies are fated to become data-transmission buses—leading to the rise of “computing on the human platform”—the study’s authors document the rapid advances taking place in the melding of the real and virtual worlds. New “ambient displays,” they write, promise to make computing “ubiquitous,” surrounding us with data and software everywhere we go: “In ubiquitous computing, the physical location of data and processing power is not apparent to the user. Rather, information is made available to the user in a transparent and contextually relevant manner.” Within ten years, we won’t even have to use keystrokes and mouse clicks to tell computers what we want them to do. There will be “new ways of interacting with computers in which delegated systems perform tasks proactively on users’ behalf, tuned precisely to the momentary requirements of time and place.”

The researchers also predict that the Google founders’ dream of a direct link between the brain and the Internet should become a reality by 2020. That’s when we’re likely to see “the first physical neural interface,” providing “a direct connection between a human

or animal brain and nervous system and a computer or computer network.” At that point, we’ll be able “to interact directly with computers by merely thinking.” Such a neural interface promises to be a blessing to many people afflicted with severe disabilities. It could help the blind to see and the paralyzed to move. But its applications go well beyond medicine, the researchers note. It also offers the “potential for outside control of human behavior through digital media.” We will become programmable, too.

THE INTERNET DOESN’T just connect information-processing machines. It connects people. It connects us with each other, and it connects us with the machines. Our intelligence is as much a part of the power of the World Wide Computer as the intelligence embedded in software code or microchips. When we go online, we become nodes on the Internet. That’s not just a metaphor. It’s a reflection of the hyperlinked structure that has from the beginning defined the Web and our use of it. The Internet, and all the devices connected to it, is not simply a passive machine that responds to our commands. It’s a thinking machine, if as yet a rudimentary one, that actively collects and analyzes our thoughts and desires as we express them through the choices we make while online—what we do, where we go, whom we talk to, what we upload, what we download, which links we click on, which links we ignore. By assembling and storing billions upon billions of tiny bits of intelligence, the Web forms what the writer John Battelle calls “a database of human intentions.” As we spend more time and transact more of our commercial and social business online, that database will grow ever wider and deeper. Figuring out new ways for people—and machines—to tap into the storehouse of intelligence is likely to be the central enterprise of the future.

On November 2, 2005, we got a glimpse of what lies ahead for

the World Wide Computer when Amazon.com began testing a new service with a strange name: Mechanical Turk. The name, it turned out, was borrowed from an infamous chess-playing “automaton” that was built in 1770 by a Hungarian baron named Wolfgang von Kempelen. The wooden machine, fashioned to look like a Turkish sorcerer sitting in front of a large cabinet, appeared to play chess automatically, using an elaborate system of gears and levers to move the pieces. In its debut, at the Schönbrunn Palace in Vienna, the Mechanical Turk quickly dispatched its first opponent, a Count Cobenzl, to the delight of the assembled courtiers. News of the remarkably intelligent robot spread rapidly, and von Kempelen took the Turk on a tour of Europe, where it defeated a series of famous challengers, including Napoleon Bonaparte and Benjamin Franklin. It was not until years later, after von Kempelen’s death, that the hoax was revealed. Hidden inside the cabinet had been a chess master, who used a system of magnets to follow opponents’ moves and make his own. The player had been simulating an artificial intelligence.

Amazon’s Mechanical Turk accomplishes a similar feat. It “hides” people inside a software program, using them to carry out tasks that computers aren’t yet very good at. Say, for example, that a programmer is writing an application that includes, as one of its steps, the identification of buildings in digital photographs—a job that baffles today’s computers but is easy for people to do. Using the Mechanical Turk service, the programmer can write a few simple lines of code to tap into the required intelligence. At the designated point in the running of the program, a request to carry out the “human task” automatically gets posted on Amazon’s Turk site, where people compete to perform it for a fee set by the programmer.

As Amazon explains on its Web site, Mechanical Turk stands the usual relationship between computers and people on its head:

“When we think of interfaces between human beings and computers, we usually assume that the human being is the one requesting that a task be completed, and the computer is completing the task and providing the results. What if this process were reversed and a computer program could ask a human being to perform a task and return the results?” That’s exactly what Mechanical Turk does. It turns people’s actions and judgments into functions in a software program. Rather than the machine working for us, we work for the machine.

We play a similar role, without even realizing it, in the operation of Google’s search engine. At the heart of that engine is the Page-Rank algorithm that Brin and Page wrote while they were graduate students at Stanford in the 1990s. They saw that every time a person with a Web site links to another site, he is expressing a judgment. He is declaring that he considers the other site important. They further realized that while every link on the Web contains a little bit of human intelligence, all the links combined contain a great deal of intelligence—far more, in fact, than any individual mind could possibly possess. Google’s search engine mines that intelligence, link by link, and uses it to determine the importance of all the pages on the Web. The greater the number of links that lead to a site, the greater its value. As John Markoff puts it, Google’s software “systematically exploits human knowledge and decisions about what is significant.” Every time we write a link, or even click on one, we are feeding our intelligence into Google’s system. We are making the machine a little smarter—and Brin, Page, and all of Google’s shareholders a little richer.

In Mechanical Turk and the Google search engine, we begin to see the human mind merging into the artificial mind of the World Wide Computer. In both services, people become subservient to the machine. With Mechanical Turk, we’re incorporated into a soft-

ware program, carrying out a small function without being aware of the greater purpose—just as manual laborers became cogs in long assembly lines. In Google's search engine, our contributions are made unconsciously. Brin and Page have programmed their machine to gather the crumbs of intelligence that we leave behind on the Web as we go about our everyday business.

As the computing cloud grows, as it becomes ubiquitous, we will feed ever more intelligence into it. Using global positioning satellites and tiny radio transmitters, it will track our movements through the physical world as meticulously as it today tracks our clicks through the virtual world. And as the types of commercial and social transactions performed through the Internet proliferate, many more kinds of data will be collected, stored, analyzed, and made available to software programs. The World Wide Computer will become immeasurably smarter. The transfer of our intelligence into the machine will happen, in other words, whether or not we allow chips or sockets to be embedded in our skulls.

Computer scientists are now in the process of creating a new language for the Internet that promises to make it a far more sophisticated medium for expressing and exchanging intelligence. In creating Web pages today, programmers have limited options for using codes, or tags, to describe text, images, and other content. The Web's traditional hypertext markup language, or HTML, concentrates on simple formatting commands—on instructing, for instance, a Web browser to put a line of text into italics or to center it on a page. The new language will allow programmers to go much further. They'll be able to use tags to describe the meaning of objects like words and pictures as well as the associations between different objects. A person's name, for instance, could carry with it information about the person's address and job, likes and dislikes, and relationships to other people. A product's name could have tags

describing its price, availability, manufacturer, and compatibility with other products.

This new language, software engineers believe, will pave the way for much more intelligent "conversations" between computers on the Internet. It will turn the Web of information into a Web of meaning—a "Semantic Web," as it's usually called. HTML's inventor, Tim Berners-Lee, is also spearheading the development of its replacement. In a speech before the 2006 International World Wide Web Conference in Scotland, he said that "the Web is only going to get more revolutionary" and that "twenty years from now, we'll look back and say this was the embryonic period." He foresees a day when the "mechanisms of trade, bureaucracy and our daily lives will be handled by machines talking to machines."

At the University of Washington's Turing Center, a leading artificial intelligence laboratory, researchers have already succeeded in creating a software program that can, at a very basic level, "read" sentences on Web pages and extract meaning from them—without requiring any tags from programmers. The software, called TextRunner, scans sentences and identifies the relationships between words or phrases. In reading the sentence "Thoreau wrote *Walden* after leaving his cabin in the woods," for instance, TextRunner would recognize that the verb "wrote" describes a relationship between "Thoreau" and "*Walden*." As it scans more pages and sees hundreds or thousands of similar constructions, it would be able to hypothesize that Thoreau is a writer and *Walden* is a book. Because TextRunner is able to read at an extraordinary rate—in one test, it extracted a billion textual relationships from 90 million Web pages—it can learn quickly. Its developers see it as a promising prototype of "machine reading," which they define as "the automatic, unsupervised understanding of text" by computers.

Scientists are also teaching machines how to see. Google has

been working with researchers at the University of California at San Diego to perfect a system for training computers to interpret photographs and other images. The system combines textual tags describing an image's contents with a statistical analysis of the image. A computer is first trained to recognize an object—a tree, say—by being shown many images containing the object that have been tagged with the description “tree” by people. The computer learns to make an association between the tag and a mathematical analysis of the shapes appearing in the images. It learns, in effect, to spot a tree, regardless of where the tree happens to appear in a given picture. Having been seeded with the human intelligence, the computer can then begin to interpret images on its own, supplying its own tags with ever increasing accuracy. Eventually, it becomes so adept at “seeing” that it can dispense with the trainers altogether. It thinks for itself.

In 1945, the Princeton physicist John von Neumann sketched out the first plan for building an electronic computer that could store in its memory the instructions for its use. His plan became the blueprint for all modern digital computers. The immediate application of von Neumann's revolutionary machine was military—designing nuclear bombs and other weapons—but the scientist knew from the start that he had created a general purpose technology, one that would come to be used in ways that could not be foretold. “I am sure that the projected device, or rather the species of devices of which it is to be the first representative, is so radically new that many of its uses will become clear only after it has been put into operation,” he wrote to Lewis Strauss, the future chairman of the Atomic Energy Commission, on October 24, 1945. “Uses which are likely to be the most important are by definition those which we do not recognize at present because they are farthest removed from our present sphere.”

We are today at a similar point in the history of the World Wide Computer. We have built it and are beginning to program it, but we are a long way from knowing all the ways it will come to be used. We can anticipate, however, that unlike von Neumann's machine, the World Wide Computer will not just follow our instructions. It will learn from us and, eventually, it will write its own instructions.

GEORGE DYSON, A historian of technology and the son of another renowned Princeton physicist, Freeman Dyson, was invited to Google's headquarters in Mountain View, California, in October 2005 to give a speech at a party celebrating the sixtieth anniversary of von Neumann's invention. “Despite the whimsical furniture and other toys,” Dyson would later recall of his visit, “I felt I was entering a 14th-century cathedral — not in the 14th century but in the 12th century, while it was being built. Everyone was busy carving one stone here and another stone there, with some invisible architect getting everything to fit. The mood was playful, yet there was a palpable reverence in the air.” After his talk, Dyson found himself chatting with a Google engineer about the company's controversial plan to scan the contents of the world's libraries into its database. “We are not scanning all those books to be read by people,” the engineer told him. “We are scanning them to be read by an AI.”

The visit inspired Dyson to write an essay for the online journal *Edge*, in which he argues that we've reached a turning point in the history of computing. The computer we use today, von Neumann's computer, uses a physical matrix as its memory. Each bit of data is stored in a precise location on that matrix, with a unique address, and software consists of a set of instructions for finding bits of data at specified addresses and doing something with them. It's a process that, as Dyson explains, “translates informally into ‘DO THIS with what you find HERE and go THERE with the result.’ Everything

depends not only on precise instructions, but on **HERE**, **THERE**, and **WHEN** being exactly defined.”

As we know today—and as von Neumann foresaw in 1945—this machine can be programmed to perform an amazing variety of tasks. But it has a fundamental limitation: it can only do what it’s told. It depends entirely on the instructions provided by the programmer, and hence it can only perform tasks that a programmer can conceive of and write instructions for. As Dyson writes, “Computers have been getting better and better at providing answers—but only to questions that programmers are able to ask.”

That’s very different from how living systems, such as our brains, process information. As we navigate our lives, our minds devote most of their time and energy to computing answers to questions that haven’t been asked, or at least haven’t been asked in precise terms. “In the real world, most of the time,” Dyson explains, “finding an answer is easier than defining the question. It’s easier to draw something that looks like a cat, for instance, than to describe what, exactly, makes something look like a cat. A child scribbles indiscriminately, and eventually something appears that resembles a cat. A solution finds the problem, not the other way around.” What makes us so smart is that our minds are constantly providing answers without knowing the questions. They’re making sense rather than performing calculations.

For a machine to demonstrate, or at least simulate, that kind of intelligence, it cannot be restricted to a set of unambiguous instructions for acting on a rigidly defined set of data. It needs to be freed from its fixed memory. It needs to lose its machine-ness and begin acting more like a biological system. That is exactly what’s becoming possible as the Internet itself becomes a computer. Suddenly, rather than having a finite set of data arrayed precisely in a matrix, we have a superfluity of data floating around in a great unbounded

cloud. We have, to switch metaphors, a primordial soup of information that demands to be made sense of. To do that, we need software that acts more like a sense-making brain than like von Neumann’s calculating machine—software with instructions that, as Dyson writes, “say simply ‘DO THIS with the next copy of THAT which comes along.’”

We see this new kind of software, in embryonic form, in Google’s search engine and in other programs designed to mine information from the Web. Google’s engineers recognize that, as Dyson puts it, “a network, whether of neurons, computers, words, or ideas, contains solutions, waiting to be discovered, to problems that need not be explicitly defined.” The algorithms of the company’s search engine already do a very good job of drawing out of the Internet answers to questions that we pose, even when we phrase our questions in ambiguous terms. We don’t always know precisely what we’re looking for when we do a Google search, but we often find it nonetheless. If the World Wide Computer is a new kind of computer, then the Google search engine is a preview of the new kind of software that will run on it.

Eric Schmidt has said that the company’s ultimate product, the one he’s “always wanted to build,” would not wait to respond to his query but would “tell me what I should be typing.” It would, in other words, provide the answer without hearing the question. The product would be an artificial intelligence. It might even be, to quote Sergey Brin again, “an artificial brain that was smarter than your brain.”

AND WHAT OF our brains? As we come to rely ever more heavily on the Internet’s vast storehouse of information as an extension of or even a substitute for our own memory, will it change the way we think? Will it alter the way we conceive of ourselves and our

relationship to the world? As we put ever more intelligence into the Web, will we, individually, become more intelligent, or less so?

In describing the future of the World Wide Computer—the “Machine,” in his terminology—Kevin Kelly writes, “What will most surprise us is how dependent we will be on what the Machine knows—about us and about what we want to know. We already find it easier to Google something a second or third time rather than remember it ourselves. The more we teach this megacomputer, the more it will assume responsibility for our knowing. It will become our memory. Then it will become our identity. In 2015 many people, when divorced from the Machine, won’t feel like themselves—as if they’d had a lobotomy.”* Kelly welcomes the prospect. He believes that the submergence of our minds and our selves into a greater intelligence will mark the fulfillment of our destiny. The human race, he says, finds itself today at a new beginning, a moment when “the strands of mind, once achingly isolated, have started to come together.”

Others are less sanguine about our prospects. In early 2005, the playwright Richard Foreman staged his surrealist drama *The Gods Are Pounding My Head* on a stage at St. Mark’s Church in Man-

* Kelly’s description of man’s growing dependency on computers carries a disquieting, if inadvertent, echo of a passage in the notorious manifesto written by Theodore Kaczynski, the Unabomber. “[As] machines become more and more intelligent,” Kaczynski wrote, “people will let machines make more of their decisions for them, simply because machine-made decisions will bring better results than man-made ones. Eventually a stage may be reached at which the decisions necessary to keep the system running will be so complex that human beings will be incapable of making them intelligently. At that stage the machines will be in effective control. People won’t be able to just turn the machines off, because they will be so dependent on them that turning them off would amount to suicide.” What was for Kaczynski a paranoia-making nightmare is for Kelly a vision of utopia.

hattan. It was a bleak work, featuring two exhausted lumberjacks wandering through a wasteland of cultural detritus—a “paper-thin world,” as one character puts it—and muttering broken, incoherent sentences. In a note to the audience, Foreman described the inspiration for his “elegiac play.” “I come from a tradition of Western culture,” he wrote, “in which the ideal (my ideal) was the complex, dense and ‘cathedral-like’ structure of the highly educated and articulate personality—a man or woman who carried inside themselves a personally constructed and unique version of the entire heritage of the West.” He feared, however, that this tradition is fading, that it is being erased as we come to draw more of our sense of the world not from the stores of our memory but from the databases of the Internet: “I see within us all (myself included) the replacement of complex inner density with a new kind of self—evolving under the pressure of information overload and the technology of the ‘instantly available.’” As we are emptied of our “inner repertory of dense cultural inheritance,” Foreman concluded, we seem to be turning into “pancake people—spread wide and thin as we connect with that vast network of information accessed by the mere touch of a button.”

It will be years before there are any definitive studies of the effect of extensive Internet use on our memories and thought processes. But anyone who has spent a lot of time online will likely feel at least a little kinship with Foreman. The common term “surfing the Web” perfectly captures the essential superficiality of our relationship with the information we find in such great quantities on the Internet. The English biologist J. Z. Young, in his Reith Lectures of 1950, collected in the book *Doubt and Certainty in Science*, eloquently described the subtle ways our perceptions, ideas, and language change whenever we begin using a new tool. Our technologies, he explained, make us as surely as we make our tech-

nologies. That's been true of the tools we use to process matter and energy, but it's been particularly true of the tools we use to process information, from the map to the clock to the computer.

The medium is not only the message. The medium is the mind. It shapes what we see and how we see it. The printed page, the dominant information medium of the past 500 years, molded our thinking through, to quote Neil Postman, "its emphasis on logic, sequence, history, exposition, objectivity, detachment, and discipline." The emphasis of the Internet, our new universal medium, is altogether different. It stresses immediacy, simultaneity, contingency, subjectivity, disposability, and, above all, speed. The Net provides no incentive to stop and think deeply about anything, to construct in our memory that "dense repository" of knowledge that Foreman cherishes. It's easier, as Kelly says, "to Google something a second or third time rather than remember it ourselves." On the Internet, we seem impelled to glide across the slick surface of data as we make our rushed passage from link to link.

And this is precisely the behavior that the Internet, as a commercial system, is designed to promote. We are the Web's neurons, and the more links we click, pages we view, and transactions we make—the faster we fire—the more intelligence the Web collects, the more economic value it gains, and the more profit it throws off. We feel like "pancake people" on the Web because that's the role we are assigned to play. The World Wide Computer and those who program it have little interest in our exhibiting what Foreman calls "the thick and multi-textured density of deeply evolved personality." They want us to act as hyperefficient data processors, as cogs in an intellectual machine whose workings and ends are beyond us. The most revolutionary consequence of the expansion of the Internet's power, scope, and usefulness may not be that computers will start to think like us but that we will come to think like

computers. Our consciousness will thin out, flatten, as our minds are trained, link by link, to "DO THIS with what you find HERE and go THERE with the result." The artificial intelligence we're creating may turn out to be our own.

A HUNDRED YEARS ago, the utility executives and electrical engineers who joined the Jovian Society saw themselves as the architects of a new and more perfect world. To them, God was "the Great Electrician," animating the universe with an invisible but all-powerful spirit. In pursuing their work they were doing His work as well; His designs were their designs. "The idea of electricity," the Jovians announced, is "binding the world together in a body of brotherhood."

Many of the computer scientists and software engineers who are building the great computing grid of the twenty-first century share a similar sense of the importance—and the beneficence—of their work. It's only the metaphor that has changed. God is no longer the Great Electrician. He has become the Great Programmer. The universe is not the emanation of a mysterious spirit. It is the logical output of a computer. "As soon as the universe began, it began computing," writes MIT professor Seth Lloyd in his 2006 book *Programming the Universe*. "Life, language, human beings, society, culture—all owe their existence to the intrinsic ability of matter and energy to process information." "All living creatures are information-processing machines at some level," argues Charles Seife in another, similarly titled 2006 book, *Decoding the Universe*. "In a sense, the universe as a whole is behaving like a giant information processor—a computer."

Our past and our destiny are inscribed in software code. And now, as all the world's computers are wired together into one machine, we have finally been given the opportunity, or at least the temptation, to perfect the code.