



FAB

The Coming Revolution on
Your Desktop—from Personal
Computers to Personal Fabrication

Neil Gershenfeld

 **BASIC**
BOOKS
A MEMBER OF THE PERSEUS BOOKS GROUP
NEW YORK

Making Sense	161
Instrumentation	173
Net Work	181
Communication	197
Art and Artillery	207
Interaction	219
The Future	227
Joy	245
The Details	255
Index	265



Mainframe computers were expensive machines with limited markets, used by skilled operators working in specialized rooms to perform repetitive industrial operations. We can laugh in retrospect at the small size of the early sales forecasts for computers; when the packaging of computation made it accessible to ordinary people in the form of personal computers, the result was a unprecedented outpouring of new ways to work and play.

However, the machines that make computers (and most everything else) remain expensive tools with limited markets, used by skilled operators working in specialized rooms to perform repetitive industrial operations. Like the earlier transition from mainframes to PCs, the capabilities of machine tools will become accessible to ordinary people in the form of personal fabricators (PFs). This time around, though, the implications are likely to be even greater because what's being personalized is our physical world of atoms rather than the computer's digital world of bits.

A PF is a machine that makes machines; it's like a printer that can print *things* rather than images. By personal fabrication, I mean not only the creation of three-dimensional structures but also the integration

of logic, sensing, actuation, and display—everything that's needed to make a complete functioning system. With a PF, instead of shopping for and ordering a product, you could download or develop its description, supplying the fabricator with designs and raw materials.

Programmable personal fabricators are not just a prediction, they're a reality. The world of tomorrow can be glimpsed in tools available today. *Fab* tells the stories of these remarkable tools and their equally remarkable users around the world. It explains what can be made, and why, and how.

I first encountered the possibility of personal fabrication through the unexpectedly enthusiastic student response to a class that I teach at MIT, modestly titled "How To Make (almost) Anything." At MIT I direct the Center for Bits and Atoms. CBA comprises fifteen or so faculty from across campus: physicists, chemists, biologists, mathematicians, and mechanical and electrical engineers. They all, like me, never fit into the artificial separation of computer science from physical science.

The universe is literally as well as metaphorically a computer. Atoms, molecules, bacteria, and billiard balls can all store and transform information. Using the discrete language of computation rather than the continuous equations of calculus to describe the behavior of physical systems is not only leading to the practical development of new and more powerful kinds of information technologies, such as quantum computers, it's also leading to new kinds of insights into the nature of the universe itself, such as the long-term behavior of black holes. If the world is a computer, then the science of computing is really the science of science.

At the intersection of physical science and computer science, programs can process atoms as well as bits, digitizing fabrication in the same way that communications and computation were earlier digitized. Ultimately, this means that a programmable personal fabricator will be able to make anything, including itself, by assembling atoms. It will be a self-reproducing machine. That idea has been a long-standing science fiction staple for better or, sometimes, much worse.

In *Star Trek: The Next Generation*, the replicator is an essential plot element that is capable of making whatever is needed for each episode. It looks like an overgrown drinks dispenser, but it has the useful feature of being able to dispense anything. In theory, it does this by following stored instructions to put together subatomic particles to make atoms, atoms to make molecules, and molecules to make whatever you want. For Captain Picard, that was frequently a steaming mug of his preferred tea, obtained from the replicator with the command "Tea, Earl Grey, hot."

The less fortunate Arthur Dent in the *Hitchhiker's Guide to the Galaxy* had to contend with the infamous Nutri-Matic machine to obtain his cup of tea. Rather than storing the molecular specification in advance, the Nutri-Matic attempted to personalize Arthur's beverage by performing a spectroscopic analysis of his metabolism, and then probing the taste centers in his brain. As with Captain Picard's tea, Arthur's drink is synthesized by assembling its molecular constituents. However, in the Nutri-Matic's case the inevitable result was a plastic cup filled with a liquid that was almost, but not quite, entirely unlike tea.

None of this violates any physical laws, and in fact such atomic-scale programmable assembly is already possible in the lab today (as long as your tastes don't run to anything much larger than a few atoms).

To develop real working personal fabricators that can operate on a larger scale, my colleagues at MIT and I assembled an array of machines to make the machines that make machines. These tools used supersonic jets of water, or powerful lasers, or microscopic beams of atoms to make—well, almost anything. The problem we quickly ran into was that it would take a lifetime of classes for students to master all of the tools, and even then the students would get little practical experience in combining these tools to create complete working systems. So, we thought, why not offer a single-semester course that would provide a hands-on introduction to all the machines?

In 1998 we tried teaching "How To Make (almost) Anything" for the first time. The course was aimed at the small group of advanced

students who would be using these tools in their research. Imagine our surprise, then, when a hundred or so students showed up for a class that could hold only ten. They weren't the ones we expected, either; there were as many artists and architects as engineers. And student after student said something along the lines of "All my life I've been waiting to take a class like this," or "I'll do anything to get into this class." Then they'd quietly ask, "This seems to be too useful for a place like MIT—are you really allowed to teach it here?"

Students don't usually behave that way. Something had to be wrong with this class, or with all the other classes I taught. I began to suspect the latter.

The overwhelming interest from students with relatively little technical experience (for MIT) was only the first surprise. The next was the reason why they wanted to take the class. Virtually no one was doing this for research. Instead, they were motivated by the desire to make things they'd always wanted, but that didn't exist. These ranged from practical (an alarm clock that needs to be wrestled into turning off), to fanciful (a Web browser for parrots), to profoundly quirky (a portable personal space for screaming). Their inspiration wasn't professional; it was personal. The goal was not to publish a paper, or file a patent, or market a product. Rather, their motivation was their own pleasure in making and using their inventions.

The third surprise was what these students managed to accomplish. Starting out with skills more suited to arts and crafts than advanced engineering, they routinely and single-handedly managed to design and build complete functioning systems. Doing this entailed creating both the physical form—mastering the use of computer-controlled tools that produce three-dimensional shapes by adding or removing material—and the logical function—designing and building circuits containing embedded computer chips interfaced with input and output devices. In an industrial setting these tasks are distributed over whole teams of people who conceive, design, and produce a product. No one member of such a team could do all of this, and even if they

could, they wouldn't: personal screaming technology is unlikely to emerge as a product plan from a marketing meeting (even if the participants might secretly long for it).

The final surprise was how these students learned to do what they did: the class turned out to be something of an intellectual pyramid scheme. Just as a typical working engineer would not have the design and manufacturing skills to personally produce one of these projects, no single curriculum or teacher could cover the needs of such a heterogeneous group of people and machines. Instead, the learning process was driven by the demand for, rather than supply of, knowledge. Once students mastered a new capability, such as waterjet cutting or micro-controller programming, they had a near-evangelical interest in showing others how to use it. As students needed new skills for their projects they would learn them from their peers and then in turn pass them on. Along the way, they would leave behind extensive tutorial material that they assembled as they worked. This phase might last a month or so, after which they were so busy using the tools that they couldn't be bothered to document anything, but by then others had taken their place. This process can be thought of as a "just-in-time" educational model, teaching on demand, rather than the more traditional "just-in-case" model that covers a curriculum fixed in advance in the hopes that it will include something that will later be useful.

These surprises have recurred with such certainty year after year that I began to realize that these students were doing much more than taking a class; they were inventing a new physical notion of literacy. The common understanding of "literacy" has narrowed down to reading and writing, but when the term emerged in the Renaissance it had a much broader meaning as a mastery of the available means of expression. However, physical fabrication was thrown out as an "illiberal art," pursued for mere commercial gain. These students were correcting a historical error, using millions of dollars' worth of machinery for technological expression every bit as eloquent as a sonnet or a painting.

Today there aren't many places where these kinds of tools are available for play rather than work, but their capabilities will be integrated into accessible and affordable consumer versions. Such a future really represents a return to our industrial roots, before art was separated from artisans, when production was done for individuals rather than masses. Life without the infrastructure we take for granted today required invention as a matter of survival rather than as a specialized profession. The design, production, and use of engineered artifacts—agricultural implements, housewares, weapons, and armor—all took place locally. The purpose of bringing tool-making back into the home is not to recreate the hardships of frontier living, just as it's not to run personal-scream-container production lines out of the family room. Rather, it's to put control of the creation of technology back in the hands of its users.

The analogy between the personalization of fabrication and computation is close, and instructive. Remember that mainframes were behemoths; in 1949 *Popular Mechanics* famously forecast that "Computers in the future may weigh no more than 1.5 tons." The Digital Equipment Corporation (DEC) pioneered computers that were the size of desks rather than rooms. They called these "Programmed Data Processors" (PDPs) rather than computers because the market for computers was seen as being too small to be viable. But over time this class of devices came to be called *minicomputers*. A 1964 DEC ad proudly proclaimed, "Now you can own the PDP-5 computer for what a core memory alone used to cost: \$27,000." That's a lot more than a PC costs today, but a lot less than the price of a mainframe. Minicomputers were thus accessible to small groups of users rather than just large corporations, and consequently their uses migrated from meeting the needs of corporations to satisfying individuals. PDPs did eventually shrink down so that they could fit on a desk, but they didn't end up there, because the engineers developing them didn't see why nonengineers would want them. DEC's president, Ken Olsen, said in 1977 that "there is no reason for any individual to have a computer in their home." PCs are now in the home; DEC is now defunct.

The adoption of PCs was driven by "killer apps," applications that were so compelling they motivated people to buy the systems to run them. The classic killer app was the original spreadsheet, VisiCalc, which in 1979 turned the Apple II from a hobbyist's toy to a serious business tool, and helped propel IBM into the PC business. VisiCalc's successor, Lotus 1-2-3, did the same in 1983 for IBM's PCs. The machine tools that the students taking "How To Make (almost) Anything" use today are much like mainframes, filling rooms and costing hundreds of thousands of dollars. But despite their size and expense, they've been adequate to show, year in and year out, that the killer app for personal fabrication is fulfilling individual desires rather than merely meeting mass-market needs. For one student this meant putting a parrot online; for another an alarm clock that got her up in the morning. None of the students needed to convince anyone else of the value of their ideas; they just created them themselves.

The invention that made PCs possible was integrated circuits, leading up to the development of microprocessors that put the heart of a computer onto a single silicon chip. The invention that promises to put the capabilities of a roomful of machine tools onto a desktop is the printing of functional materials. A printer's ink-jet cartridge contains reservoirs of cyan, magenta, yellow, and black inks; by precisely placing tiny drops it's possible to mix these to produce what appear to be perfect reproductions of any image. In the research lab today, there are similar inks that can be used to print insulators, conductors, and semiconductors to make circuits, as well as structural materials that can be deposited to make three-dimensional shapes. The integration of these functional materials into a printing cartridge will make possible faithful reproductions of arbitrary objects as well as images. At MIT we have a joke that we now take seriously: a student working on this project can graduate when their thesis can walk out of the printer. In other words, along with printing the text of the document, the printer must produce the means for the thesis to move.

Ultimately, rather than relying on a printer to place droplets of material, the logic for assembling an object will be built into the materials themselves. This is exactly how our bodies are made; a molecular machine called the ribosome translates instructions from genes into the series of steps required to assemble all of the proteins in our bodies out of the twenty amino acids. The discovery of building with logic is actually a few billion years old; it's fundamental to the emergence of life. Current research is now seeking to do the same with functional materials, creating a fundamentally digital fabrication process based on programming the assembly of microscopic building blocks. This mechanism will be embodied in personal fabricators fed by such structured materials. Much as a machine today might need supplies of air, water, and electricity, a digital personal fabricator will use as raw feedstocks streams of conductors, semiconductors, and insulators.

Unlike machines of today, though, but just like a child's building blocks, personal fabricators will also be able to disassemble something and sort its constituents, because the assembled objects are constructed from a fixed set of parts. The inverse of digital fabrication is digital recycling. An object built with digital materials can contain enough information to describe its construction, and hence its deconstruction, so that an assembler can run in reverse to take it apart and reuse its raw materials.

We're now on the threshold of a digital revolution in fabrication. The earlier revolutions in digitizing communications and computation allowed equipment made from unreliable components to reliably send messages and perform computations; the digitization of fabrication will allow perfect macroscopic objects to be made out of imperfect microscopic components, by correcting errors in the assembly of their constituents.

Return now to the mainframe analogy. The essential step between mainframes and PCs was minicomputers, and a similar sequence is happening along the way to personal fabrication. It's possible to approximate the end point of that evolution today with a few thousand

dollars of equipment on a desktop, because engineering in space and time has become cheap.

First the space part. An inexpensive CD player places its read head with a resolution of a millionth of a meter, a micron. Combining this metrology with a desktop computer-controlled milling machine enables it to move a cutting tool in three dimensions with that same resolution, patterning shapes by removing material with tolerances approaching the limits of human perception. For example, it can cut out circuit boards with features as fine as the tiniest components.

Then the time part. A one-dollar embedded computer chip can operate faster than a millionth of a second, a microsecond. This is fast enough to use software to perform functions that traditionally required custom hardware, such as generating communications signals and controlling displays. It is possible to program one such chip to take on the functions of many different kinds of circuits.

The increasing accessibility of space and time means that a relatively modest facility (on the scale of the MIT class) can be used to create physical forms as fine as microns and program logical functions as fast as microseconds. Such a lab needs more complex consumables than the ink required by a printer, including copper-clad boards to make circuits and computer chips to embed into projects. But, as the students found at MIT, these capabilities can be combined to create complete functioning systems. The end result is much the same as what an integrated PF will be able to do, allowing technological invention by end users.

Minicomputers showed how PCs would ultimately be used, from word processing to e-mail to the Internet, long before technological development caught up to make them cheap enough and simple enough for wider adoption. Struck by this analogy, I wondered if it would be possible to deploy proto-personal fabricators in order to learn now about how they'll be used instead of waiting for all of the research to be completed.

This thought led to the launch of a project to create field "fab labs" for exploring the implications and applications of personal fabrication

in those parts of the planet that don't get to go to MIT. As you wish, "fab lab" can mean a lab for fabrication, or simply a fabulous laboratory. Just as a minicomputer combined components—the processor, the tape drive, the keypunch, and so forth—that were originally housed in separate cabinets, a fab lab is a collection of commercially available machines and parts linked by software and processes we developed for making things. The first fab labs have a laser cutter to cut out two-dimensional shapes that can be assembled into three-dimensional structures, a sign cutter that uses a computer-controlled knife to plot flexible electrical connections and antennas, a milling machine that moves a rotating cutting tool in three dimensions to make circuit boards and precision parts, and the tools for programming tiny high-speed microcontrollers to embed logic. A bit like the original PDP minicomputers, all of this could be called a Programmed Materials Processor. This is not a static configuration; the intention over time is to replace parts of the fab lab with parts made in the fab lab, until eventually the labs themselves are self-reproducing.

The National Science Foundation (NSF) provided the seed funding for fab labs through its support of the Center for Bits and Atoms (CBA). NSF expects research activities that it funds on the scale of CBA to have an educational outreach component, which all too often is limited to teaching some classes at a local school, or creating a Web site describing the research. Instead, my CBA colleagues and our NSF counterparts agreed to try equipping ordinary people to actually *do* what we're studying at MIT instead of just talking about it. It's possible in the fab labs to do work that not too long ago required the resources of a place like MIT (an observation that's not lost on me on days when the administrative overhead of working at a place like MIT is particularly onerous).

Starting in 2002, the first fab labs went to rural India, Costa Rica, northern Norway, inner-city Boston, and Ghana. The equipment and supplies for each site initially cost about twenty thousand dollars. Knowing that that cost will come down as the technology progresses, the first fab labs weren't meant to be economically self-sustaining. One

of the first surprises from the field was the demand for duplicating the labs even at that cost.

In keeping with the fab lab project's goal of discovering which tools and processes would be most useful in the field, we started setting up these labs long before we knew how best to do it. The response in the field was as immediate as it had been at MIT. We ended up working in so many far-flung locations because we found a demand for these capabilities around the world that was every bit as strong as that around campus. In the village of Pabal in western India, there was interest in using the lab to develop measurement devices for applications ranging from milk safety to agricultural engine efficiency. In Bithoor, on the bank of the Ganges, local women wanted to do three-dimensional scanning and printing of the carved wooden blocks used for chikan, a local kind of embroidery. Sami herders in the Lyngen Alps of northern Norway wanted wireless networks and animal tags so that their data could be as nomadic as their animals. People in Ghana wanted to create machines directly powered from their abundant sunlight instead of scarce electricity. Children in inner-city Boston used their fab lab to turn scrap material into sellable jewelry.

For all the attention to the "digital divide" in access to computers between developed and developing countries, these recurring examples suggest that there is an even more significant divide in access to tools for fabrication and instrumentation. Desktop computers are of little use in places that don't have desks; all too often they sit idle in isolated rooms built by aid agencies to house them. Appropriate computing requires the means to make, measure, and modify the physical world of atoms as well as the virtual world of bits. And instead of bringing information technology (IT) to the masses, fab labs show that it's possible to bring the tools for IT development, in order to develop and produce local technological solutions to local problems.

There's been a long-standing bias that technology's role in global development mirrors its own history, progressing from low- to high-tech. The fab lab experience suggests instead that some of the least

developed parts of the world need some of the most advanced technologies. That observation has led me to spend some head-spinning days in Washington, DC, going from the World Bank to the National Academy of Sciences to Capitol Hill to the Pentagon, having essentially the same meeting at each place. Fab labs challenge assumptions that are fundamental to each of these institutions. Instead of spending vast sums to send computers around the world, it's possible to send the means to make them. Instead of trying to interest kids in science as received knowledge, it's possible to equip them to *do* science, giving them both the knowledge and the tools to discover it. Instead of building better bombs, emerging technology can help build better communities.

The problem I found on these trips was that none of these institutions knew how to pay for this kind of work; there isn't a Pentagon Office of Advanced Technologies for Avoiding Wars. Financing personal fabrication in underserved communities is too directed a goal for traditional basic research funding, and too speculative for conventional aid organizations or donors. The closest precedent is microcredit lending, which provides small loans to help support financial cooperatives, typically run by women, in developing countries. The loans are used to acquire an asset such as a cell phone that can then be used to generate income. But that model doesn't help when the financing is for an invention. What's needed are the skills of a good venture capitalist rather than a banker. That's not an oxymoron; the best venture capitalists add value to their investments by helping shepherd, prune, and protect ideas, build operational teams, and develop business models.

The historical parallel between personal computation and personal fabrication provides a guide to what those business models might look like. Commercial software was first written by and for big companies, because only they could afford the mainframe computers needed to run it. When PCs came along anyone could become a software developer, but a big company was still required to develop and distribute big programs, notably the operating systems used to run other programs.

Finally, the technical engineering of computer networks combined with the social engineering of human networks allowed distributed teams of individual developers to collaborate on the creation of the most complex software.

Programmers write source code that people can understand, which gets turned into executable code that computers can understand. Commercial companies have protected the former and distributed the latter to their customers. But individuals who share the source code that they write can collaborate in ad hoc groups, which might never meet physically, on the creation of programs that are larger than any one of the members could write alone. The Linux operating system is built out of such "open source" software. Much like the way science progresses by researchers building on one another's publications, a programmer can make available a piece of code that might then get taken up and improved by someone on the opposite end of the earth.

In a world of open-source software, ownership of neither computers nor code alone provides the basis for a proprietary business model; what's left is the value added to them by creating content and delivering services. Profitable old and new computer companies are making money from freely available software in just this way: by charging for their role in solving problems.

Similarly, possession of the means for industrial production has long been the dividing line between workers and owners. But if those means are easily acquired, and designs freely shared, then hardware is likely to follow the evolution of software. Like its software counterpart, open-source hardware is starting with simple fabrication functions, while nipping at the heels of complacent companies that don't believe that personal fabrication "toys" can do the work of their "real" machines. That boundary will recede until today's marketplace evolves into a continuum from creators to consumers, servicing markets ranging from one to one billion.

That transition has already happened in the case of two-dimensional printers, the immediate predecessor to personal fabricators. High-quality

printing, once solely a commercial service, came into the home via laser printers. The most important attribute of an industrial printing press is its throughput, which is the number of pages it can produce per minute. Laser printers started going down this technological scaling curve with the development of ever-faster printers needed for serious on-demand commercial printing. Within Hewlett-Packard, a competing group of engineers had the idea that squirting individual drops of ink could make beautiful images more cheaply than transferring toner onto sheets. Ink-jet printing would be slower than laser printing, but they reasoned that for a printer in the home, quality mattered much more than speed. This was such a heretical idea that this group decamped from HP's headquarters in Palo Alto and set up shop out of sight in Corvallis, Oregon. The rest is business history: the relative cost of producing and selling ink-jet cartridges has been about the closest a company has come to legally printing money. Ink-jet printing hasn't replaced commercial printing; it's created an entirely new—and enormous—market driven by quality and access rather than speed.

Similarly, the emerging personal fabrication tools I've been describing are intended for personal rather than mass production. Their development was originally driven in industry by the need to quickly create prototypes of products to catch errors before they became much more expensive to correct in production. Machine-tool shows relegate such rapid-prototyping machines to a sleepy corner away from the giant cutting, stamping, and molding tools that are at the top of the machine food chain. But if the market is just one person, then the prototype *is* the product. The big machines will continue to mass-produce things used in large quantities; nuts and bolts are valuable because they're identical rather than unique. But little machines will custom-make the products that depend on differences, the kinds of things being made in the fab class and fab labs.

The biggest impediment to personal fabrication is not technical; it's already possible to effectively do it. And it's not training; the just-in-time peer-to-peer project-based model works as well in the field as at

MIT. Rather, the biggest limitation is simply the lack of knowledge that this is even possible. Hence this book.

Fab tells the stories of pioneering personal fabrication users, and the tools they're using. Because both are so striking, I've interwoven their stories in pairs of chapters that explore emerging applications and the processes that make them possible. The not-so-hidden agenda is to describe not only who is doing what but also how, providing introductions to the tools that are similar to the orientations we give at MIT and in the field. These stop just short of hands-on training; a final section gives enough detail on the products, programs, and processes used to duplicate what's shown in the book.

Throughout *Fab* I use what are known as “hello world” examples. In 1978, the instruction manual for the then new C programming language written at Bell Labs used as an example a simple program that printed out the words “hello world.” This is more exciting than it sounds, because it requires an understanding of how to write a rudimentary program, compile it into computer code, and cause the program to print the text. “Hello world” programs have since become staples for introducing new computer languages; the Association for Computing Machinery currently lists 204 examples for languages from A+ to zsh.

The difference between those “hello world” programs and the examples I give in this book is that mine arrange atoms as well as bits, moving material as well as instructions. But the principle is the same: show the minimum specification needed to get each of the tools to demonstrate its successful operation. Taken together, these examples provide a fairly complete picture of the means for almost anyone to make almost anything.

My hope is that *Fab* will inspire more people to start creating their own technological futures. We've had a digital revolution, but we don't need to keep having it. Personal fabrication will bring the programmability of the digital worlds we've invented to the physical world we inhabit. While armies of entrepreneurs, engineers, and pundits search for the next killer computer application, the biggest thing of all coming in computing lies quite literally out of the box, in making the box.

and explain the benefits with the data for the same, with concrete
evidence, numbers and sources for the numbers.

Below is my digital fabrication manifesto, but as it grew over time
my original idea for fabrication manifested itself as a series
of questions to address in a digital fabrication manifesto. The original
idea for digital fabrication of the same type. The same people who
were making smart self-organizing materials were also thinking
about how to make computers. And even smarter were the people
who were thinking about how to make computers. I know
the difference is computer science is computer science and it's not
just that, it's a way of thinking about things with computers
and by making things computationally.

Actually, the difference is a few billion years old. Going back to
when the biological evolution for protein construction first evolved.
Today, the theory of evolution that is currently in the field is
evolution theory. It's not just a theory, it's a hypothesis. It helps explain the
central role of coding and state structures in gene biology, and
today, the development of various digital fabrication projects of
interest in digital biology's future.

There are three main reasons why it is not between biology and
engineering or a similar position. Beyond the engineering field, computer
science is the philosophical and religious culture that the definition of
life, we need to understand generally a very complex nature that is
digital, biological, and complex. It has the essential nature of a
living system. With biological systems and biological systems in
other systems, the progress is that understanding is now making
up for the character of understanding with the flow. Over a structure
between a model of reproducing itself, there is no stream, different field
between a few and fabrication and the idea of a person.

Joy

**In 2000, Sun Microsystems' chief scientist Bill Joy published an
intelligent and influential essay in *Wired Magazine* titled "Why the
Future Doesn't Need Us." In it he argued that the development of self-
reproducing molecular assemblers of the kind I described in "The
Future" would be so dangerous to our survival as a species that we
should consider voluntarily abandoning the effort. In considering the
future of personal fabrication, it's indeed worth thinking about where
it's headed, and why, as well as how it will get there. The state of the
technology today does provide intriguing insight into whether personal
fabrication will improve or destroy the world of tomorrow.**

Technologists do have a terrible track record for reflecting on the
wisdom of where their work might lead. Bill Joy's sensible point was
that robotics, genetic engineering, and biotechnology are so fundamen-
tally and dangerously unlike any technology that has come before that
we can no longer afford to blindly develop them and hope that it will
all come out OK. There are two reasons to see these technologies as dif-
ferent from all the rest. First, the only barrier to entry in each of these
areas is access to knowledge, not scarcity of resources. It requires an enor-
mous and enormously expensive supply chain to produce plutonium to

make a nuclear weapon; neither the machines nor the money are easily hidden. But producing a protein just requires easily available reagents and a laboratory bench to mix them on. Ideas are so powerful precisely because they can't be controlled in the same way that scarce physical resources can.

And second, genes and robots can make more genes and robots. Self-reproduction, as I described it in the last chapter, implies exponential growth. One bacterium can divide and make two bacteria, which divide and make four, then eight, then sixteen, on up to a plague. If a new technology is truly capable of self-reproduction, then it could cause a technological plague, growing without bound until it consumes all available resources on earth.

Eric Drexler, a pioneer in nanotechnology, termed this the "gray goo" scenario in his 1986 book *Engines of Creation: The Coming Era of Nanotechnology* refers to technologies working on the scale of nanometers, 10^{-9} meters, the size of atoms. Eric's colorful (colorless?) description of unchecked self-reproducing nanorobots as a gray goo follows from the likelihood that they would appear to the naked eye as a featureless mass of growing stuff.

This frightening prospect has led to opposition, sometimes violent, to unchecked nanotechnology research from unlikely bedfellows, from Prince Charles to environmental activists to survivalists. If self-reproducing machines could take over the planet, and if the only restraint on this happening is control over the spread of knowledge about how to do it, then we're doomed because ideas can't be contained. Whether out of malicious intent, benign neglect, or misplaced optimism, a single technological Typhoid Mary could destroy everyone and everything.

Rather, we will be doomed if that does happen, but there are a number of reasons for thinking that it might not. It is true that self-reproducing nanorobots are fundamentally unlike earlier technological threats to our survival because of the risk of exponential growth, and because of the difficulty of containing a technology based on ideas rather than rare raw materials. Then again, each of the other end-of-

life-on-earth risks have been fundamentally different from their predecessors. Firearms brought an immoral edge to warfare, making it possible to kill at a distance, and rendering existing armor obsolete. Then came nuclear weapons, which could erase whole cities rather than just kill individuals. More recently, the viability of the networked global economy has been challenged by computer virus writers, with lone hackers able to cause billions of dollars' worth of damage by remotely shutting down networks and computers.

We've survived so far because of coupled advances in evaluating these threats and developing countermeasures. Firearms are as deadly as ever, but police and soldiers can now wear Kevlar and ceramic body armor that provides protection, which would have been inconceivable to a blacksmith. For nuclear weapons, the attempt to build a Star Wars-type shield has remained a very expensive kind of science fiction, but the combination of a nuclear deterrent and nonproliferation regimes for monitoring and controlling nuclear materials has provided a historically unprecedented restraint on the use of such a fearsome military weapon. At the dawn of their creation, nuclear weapons represented an almost unimaginably large force; at the Trinity test of the first nuclear weapon, General Groves was annoyed by Enrico Fermi's offer to take wagers on whether the test would ignite the earth's atmosphere, or perhaps just New Mexico's. Hans Bethe was able to show that atmospheric ignition can't happen, and predicting the yield of a nuclear weapon has now become a rather exact science. It would likewise be a mistake to assume the use of twentieth-century technologies in analyzing and addressing a twenty-first-century threat.

Technologies coevolve with their design tools and countermeasures. Computer virus writers found vulnerabilities in the Internet and in operating systems; a combination of software engineering (firewalls, software patches, e-mail filters) and social engineering (financial liability, legal prosecution) has not eliminated computer viruses but has turned them into a manageable nuisance. The same is true of the existing example of a self-reproducing machine, a real virus like AIDS.

What's surprising about the AIDS epidemic is not the scale of human suffering it's already caused, horrific as that's been; it's that the epidemic hasn't been even worse. The HIV virus that causes AIDS is a fiendishly clever mechanism. HIV is a nearly pure form of information. At a size of about a hundred nanometers it's too tiny to move far on its own or to carry much in the way of energy; it obtains its resources by hijacking a host cell. HIV does this with three enzymes. The first of these, reverse transcriptase, converts HIV's RNA into DNA. The second, integrase, splices the DNA into the host cell's own DNA. When the cell translates this DNA into a protein, the third enzyme, protease, cleaves the protein to assemble a new virus. This process is so effective precisely because of the reliability of the copying done by DNA polymerase. DNA replication is so carefully controlled that the cell trusts the information in DNA and doesn't bother to check for or protect against "unauthorized" tampering. Reverse transcription runs backwards from the usual direction of transcription from DNA to RNA, and because the virus doesn't error-correct its reverse transcription this introduces a mechanism for rapid mutation that helps the virus evade our natural immune response. But it also introduces a vulnerability into the life cycle of the virus that's been a successful target for our technological immune response, through AIDS drugs that disrupt reverse transcription.


HIV embodies our most advanced insights from biology, chemistry, physics, and computer science. It's hard to believe that we could accidentally or intentionally come up with a self-reproducing machine even more effective than HIV, and even if we did it would have to compete for natural resources with all of the other viruses and bacteria on turf they've spent a few billions years evolving to optimize their designs around. Molecular machines will confront the same challenges in obtaining energy, eliminating waste, and defending against threats that limit the growth of their biological predecessors.

Eric Drexler came to regret coining "gray goo" as a distraction from the far more likely development of nanotechnology based on pro-

grammed molecular manufacturing. Nanotechnology itself is a curious sort of technology. It's the subject of billion-dollar research strategies and breathless press coverage, but very few of the serious scientists studying small things call their work nanotechnology. Molecular biologists are developing minimal cells, either from scratch or by simplifying existing cells, to create controllable bioreactors for producing drugs and chemicals. Chemists are studying new materials that self-assemble from molecular building blocks. Device physicists are developing more powerful information technologies based on manipulating individual atoms, electrons, and photons. The arrival of atomic-scale digital fabrication will not be a discontinuous event; it's already here in these rudimentary forms, which will continue to emerge and converge at the intersection of these disciplines.


Planetary annihilation from the unchecked growth of self-reproducing molecular machines may be an unlikely doomsday scenario, but if molecular fabricators that can truly make anything are indeed coming, then a far more practical fear is who might make what. When I first briefed a roomful of army generals on personal fabrication and digital assembly, that question triggered a lively discussion. One contingent was appalled that we were intentionally fostering access to personal fabrication through fab labs, putting these tools in the hands of potential attackers, and they muttered about the need to control this technology. Another contingent argued that it wasn't a controllable technology, and in any case the bad guys would get their hands on it regardless of any attempted limits, and that the impact could be far more beneficial in helping address the root causes for conflict. Guess which side I was on.

Bad guys are already impressively effective at acquiring the best available technology for the destruction of their enemies; fab labs are likely to have a far greater impact on the stability of the planet by helping everyone else acquire the technology they need for their survival. Debating the wisdom of widespread access to personal fabrication tools reflects a divergence in the discussion of the appropriateness of




technology between people whose needs are and are not well met, a kind of intellectual version of not-in-my-backyard.

I first heard an early echo of the chorus of voices in this book in the response to my earlier book *When Things Start to Think*, which was about moving computing out of traditional computers and into everyday objects. On one hand, I was happily swamped with out-of-the-blue letters along the lines of "Thank heavens, finally this is information technology that's relevant for me, I need this for. . . ." They would then go on to describe their ideas for applications of embedded computing in their lives. On the extreme other hand, I was startled to read in a review that the book was irrelevant because of the well-known impossi-



bility of computers thinking. Aside from rehashing an unwinnable old philosophical argument about defining intelligence, my learned reviewer had entirely missed the point. Embedding computer chips was not an exercise in epistemology; it was a way to solve problems. As the tools for personal fabrication now make accessible not just the use but also the creation of technology, users rather than pundits can decide which problems need solving.

It's common to bemoan how kids no longer have hardware hobbies, and, as they get older, how so few are interested anymore in technical careers. These may have been temporary casualties of the digital revolution, artifacts of an overemphasis of bits over atoms as young and old sat before their computer screens. Kids' toys emulate grown-up tools: a child's play with an erector set was not too far off from the work of an




engineer. More recently, other than writing software it hasn't been feasible in the home to do anything remotely close to the frontiers of science and engineering—the required tools became too expensive and specialized. Now, though, that pendulum is swinging back so far that a child playing with personal fabrication tools can have access to capabilities beyond those available to an individual engineer working as a cog in the gears of an organization based around central scarce technical resources.

And play the kids do. In this book, I've described a deep idea: the digitization of fabrication. I've explored its implication in personal fabrication. I've presented the observation that we're already in the equivalent of the minicomputer era of personal fabrication. But the most unexpected part of this story has been the response, as the nearly universal pleasure in getting access to tools for technological develop-



ment appears to cut across ages and incomes, from tribal chiefs to MIT students to children coming in off the street. There's a shared childlike delight in invention that's shown best of all by a child.

The liberal arts were originally meant to liberate, empowering through mastery of the means of expression. The new means of



expression for personal fabrication offer technological empowerment, providing economic, intellectual, and even a kind of spiritual liberation. I believe that the best place to read about the future of personal fabrication is in the faces of a new generation getting access to prototype versions of these capabilities in fab labs. Coming from many different parts of the planet, they end up in the same place: joy.

The Details