PUBLISHED OCTOBER 2003, REVISED NOVEMBER 2003

Three Systems of Action:

A Proposed Application for Effective

Administration of Molecular Nanotechnology

by Chris Phoenix and Mike Treder

Introduction

Within the next few decades, and perhaps sooner, a new type of manufacturing will be made possible by molecular nanotechnology (MNT). Considering its enormous potential for profound social, environmental, economic, and military impacts, MNT has received insufficient attention in ethical and policy discussions. The first section of this paper provides a brief introduction to MNT, in order to establish the need for increased policy attention. The second section describes three different approaches to policymaking, each based on a different system of action, or set of principles, used for solving various kinds of problems. The third section demonstrates that MNT, as a flexible *general purpose technology*, will require a flexible approach to policymaking that encompasses all three systems of action. The fourth section presents specific recommendations and possibilities for accomplishing this difficult balance between incompatible policy styles.

MOLECULAR NANOTECHNOLOGY

“The principles of physics, as far as I can see, do not speak against the possibility of maneuvering things atom by atom. It is not an attempt to violate any laws; it is something, in principle, that can be done; but in practice, it has not been done because we are too big.”[1] With these words, in 1959, Nobel Prize winning physicist Richard Feynman introduced the concepts that would eventually be known as nanotechnology.

Nanotechnology has several meanings and covers many diverse fields. The U.S. National Science Foundation defines it as: “Research and technology development at the atomic, molecular or macromolecular levels, in the length scale of approximately 1 - 100 nanometer range, to provide a fundamental understanding of phenomena and materials at the nanoscale and to create and use structures, devices and systems that have novel properties and functions because of their small and/or intermediate size."[2] By this definition, some kinds of nanotechnology exist already. Computer manufacturers are exploring new types of circuits built with individual molecules, and other technologies too numerous to list are being developed or already in use. Current products of nanoscale technology include powders, films, and chemicals. Computer parts and sensors may be only a few years away. But this type of nanotechnology simply produces specialized materials and components, one innovation at a time. Nanoscale technology is not general purpose.

Molecular nanotechnology (MNT) is different. The field was initially defined by popular and even fictional portrayals. Often misunderstood or overlooked in discussions of nanotechnology, it is nevertheless a field of active and serious study. [See APPENDIX] Starting with the publication of *Nanosystems* by Eric Drexler in 1992, a coherent and reliable picture has been built over the last decade of a plausible—and possibly imminent—manufacturing system: a *limited* molecular nanotechnology (LMNT) based on programmable, mechanically guided carbon-lattice chemistry workstations. Such a manufacturing system could make a wide variety of functional molecular shapes, large enough to have stable properties but small enough to make dinosaurs out of today's cutting-edge products.

From the beginning, MNT was claimed to be dangerously transformative. It was forewarned that once a general purpose nanoscale manufacturing capability was developed, a flood of diverse and powerful products would appear almost overnight. Distributed, cheap, clean manufacturing would upset many aspects of the economy. Manufacturing systems would be able to duplicate themselves speedily, leading to extremely low capital costs. Such claims seemed fantastic, and this perception has perhaps contributed to the poor reputation of the field. However, work over the past decade has made it clear that many of the initial warnings were justified.

Advances in conventional manufacturing are happening continually: parts are getting smaller and more precise, robots are doing more of the fabrication and assembly, inventories are shrinking while speed increases, and industrial processes are producing less waste. The natural destination of these trends is a completely automated manufacturing system that can rapidly build products with atomic precision directly from raw materials, while making use of every molecule. In the normal course of innovation, this might take many decades or even centuries to achieve. However, MNT appears to provide a shortcut: by mixing digital control with chemistry, a mechanochemical system could do for manufacturing what computers have done for information processing.

In studying the social impacts of MNT, it is important to know when it might be developed. If it were to be delayed by many decades, its capabilities would not be revolutionary. There is good reason to believe, however, that it could be developed much sooner—possibly within a single decade.[3] This would be cause for grave concern.

Certainly, improved manufacturing is desirable for a wide variety of reasons. But it's also true that MNT—even in its limited form—would be able to create products so advanced as to destabilize many social and political institutions. Given the immense power and resources that an LMNT capability would provide, we cannot afford to merely hope that no nation will realize this for the next thirty years; we must assume that somewhere in the world an LMNT development project will be started as soon as the cost falls within a large nation's budget.

The barriers to an LMNT project may already be surprisingly low. Mechanochemistry has been demonstrated on a silicon lattice. New nanoscale sensing and manipulation technologies are being developed at a rapid pace. The main technical barriers are the lack of a detailed design, and the need for a better understanding of the laboratory techniques required to build the first molecular fabricator. Both of these problems should be solvable with sufficient resources. In the bigger picture, the main institutional barrier is a widespread belief that MNT is not worth pursuing at this time. Even if this belief does not shift soon in the United States or Europe, other nations may be quicker to recognize the potential and more politically able to fund an LMNT program.

Recent work by the Center for Responsible Nanotechnology has demonstrated that development of even a primitive mechanochemical fabrication ability could be followed rapidly by development of an integrated, automated, self-contained, human-scale factory.[4] Such a sudden increase in advanced manufacturing capability could have substantial military as well as economic advantages for the nation or bloc that controls it. This greatly increases the incentive for rapid and early development of MNT.

A full exploration of the possible societal impacts of nanotechnology is beyond the scope of this paper. But regardless of the form they take, it seems certain that these impacts will be extensive. Our purpose here is to examine the types of institutions that will best be suited for administering MNT, to avoid the dangers and maximize the benefits.

THREE SYSTEMS OF ACTION

Consider this question: "How much should a corporation pay for the right to kill someone?" The question is appalling. Obviously corporations have no right to kill anyone, under any circumstances. Yet other institutions in our society sometimes do have that right. A policeman has every right to kill a criminal who is trying to kill him. And yet, the question, "How much should a policeman pay for the right to kill someone?" is also appalling. The words "pay" and "kill" simply don't belong together—except in the Mafia. An institutional system that involves the ability to kill people should not also involve money.

As the example above demonstrates, there are some groups whose conduct includes the use of force, and other groups that engage in trade—but a single group that deals in both can be very dangerous. This section explores the systems of action of these two types of groups, and the principles that make up those systems. We also introduce a third system that involves neither force nor trade, but information.

Guardian Principles: Provide Security

Think of a fortress guarding a frontier. The soldiers must always be prepared to fight, but most of the time they are training or relaxing. Strict discipline is necessary to make them a unified fighting force. One traitor, or paid spy, can get them all killed. Visiting merchants are a distraction and a security problem; too much money floating around can weaken their dedication to the task.

Guardian principles are appropriate for governments and police forces, organizations that defend laws and land. In such a group, betrayal can cause disaster; force is frequently necessary; tradition is valuable; and loyalty is more important than money (see Table 1).

If someone steals your car, they have one more car and you have one less. The thief has created no value—in fact, once the car goes to the chop shop, value is destroyed. When one nation forcibly annexes the territory of another nation, one government gains and another government loses. These are examples of zero-sum or even negative-sum transactions. Guardian ethics are well suited to organizations dealing with zero-sum situations. In *Systems of Survival*[5], Jane Jacobs has analyzed Guardian organizations (as well as Commercial organizations, described in the next section). According to Jacobs, Guardian principles (which she terms ethics) include "Deceive for the sake of the task", "Take vengeance", and "Shun trading"—all good advice for a group surrounded by enemies. However, these principles would not work well for a commercial organization.

Commercial Principles: Optimize Trade

Think of a small neighborhood shop. The employees should be ready to do business with anyone who walks in, and must maintain a reputation of honesty with both suppliers and customers. The store must continually improve, or the other stores will lure away its customers. A small business owner does not have a lot of free time and must work efficiently.

Commercial principles are appropriate for business and trade, which seek to increase value to all parties involved. Money is the lifeblood of commerce. Innovation and efficiency are more useful than tradition, and the use of force is severely frowned on.

If one person has a surplus of wheat, and another has a surplus of eggs, they will both be better off if they trade. This is a positive-sum transaction, and the creation of value is what drives commerce. Commercial principles include, "Be honest", "Be thrifty", "Compete", and "Respect contracts". The more closely a commercial organization follows these ideals, the more trading it can engage in and the richer it will become. Commercial principles also include, "Shun force". This is good advice for companies that must focus on competing in the marketplace; coercion is not a good way to build productive trading relationships.

Information Principles: Promote Abundance

Think of a programmer working at 2:00 AM to add a feature to an Open Source program he didn't write. The programmer is not paid for this work; he does it because he wants the program to be more usable and more popular; he has been working for six hours without a break. At 2:30 AM he adds his name to the list of contributors, uploads the improved program to a website for free distribution, then spends the next hour reading free articles on-line.

We have said that Guardian principles are best for dealing with zero-sum or negative-sum situations, and Commercial principles are best for dealing with positive-sum situations. The invention of computers has created unlimited-sum situations. An unlimited-sum situation is one in which the cost of an event is both very low ("too cheap to meter")[6] and unrelated to the value. Anything that exists in the form of computer data can be copied and distributed at remarkably low cost. And a wide range of things—including music, news, blueprints, books, recipes, and scientific papers—can be represented as computer data. Of course, some things are valuable only because they are rare, so too much copying would actually reduce their value. But some information becomes more valuable to its creators the more it is copied. Many hobbyists would like their creations to be widely appreciated—as long as the viewer knew who had created it. The authors of scientific papers and the programmers of Open Source[7] software want as many people as possible to use their work—as long as they get appropriate credit. The more such information is copied, the more benefits accrue both to the inventor and to the users.

The Information system has arisen to facilitate the production and copying of freely shared information. This system of action is related to the so-called "hacker ethic" and to the older system of academic endeavor.

Table 1: Comparison of Systems[8]

Information	Commercial	Guardian
Shun force	Shun force	[Rely on force]
Shun trading	[Rely on trading]	Shun trading
Use intelligence	Use initiative and enterprise	Exert prowess
Publish all information	Be honest	Deceive for the sake of the task
Be idealistic	Be optimistic	Be fatalistic
Ignore comfort	Promote comfort and convenience	Make rich use of leisure
Respect authorship; Ignore ownership	Respect contracts	[Defend your territory]
Demonstrate the superiority of your own ideal	Dissent for the sake of the task	Be obedient and disciplined
Invent and create	Be open to inventiveness and novelty	Adhere to tradition
Shun authority	[Adapt to the system]	Respect hierarchy
Collaborate easily with strangers and aliens	Collaborate easily with strangers and aliens	Be exclusive
Accept largesse	Be thrifty	Dispense largesse
Be unique; Develop a reputation	Be industrious	Be ostentatious
Be productive	Invest for productive purposes	Take vengeance

Rise of the Information System

The invention of writing allowed information to be stored for later use, and even copied verbatim. The printing press made the copying process much easier, giving many more people access to the information. Computers, with networks and word processors, have reached another level.[9] The cost of copying most information is now essentially zero. Using computers, a person can write a book, email it to friends and colleagues, and put it on a website for the whole world to see. The same can be done for recipes, music, and drawings. With training, that person could even write a better word processing program so that friends—or the whole world—will find it easier to write books. There are programs to help write that word processing program, and there are people working to improve those programs. The whole system can be improved exponentially, as long as people are willing to "give away" their work. It may seem strange at first, but many people will do just that.

Numerous organizations have developed to create, promote, distribute, and use information that is freely copyable. A computer operating system called Linux[10], a competitor to Microsoft Windows, is a good example of this. Linux is perhaps the most famous product of the Open Source software movement. It has no owner in the traditional sense: anyone who wants to can obtain a copy for free and use it on as many computers as they like.[11] The creators of Linux—thousands of programmers worldwide—are quite happy with this state of affairs. They do not want to sell the software; they simply want their name to be included in the credits. As more people use it, the authors gain bragging rights among their fellow programmers, and they also know that they have made the world a better place by saving money for each user. Microsoft operating systems cost hundreds of dollars. The Linux operating system is absolutely free. Thousands of Linux programmers work without pay to make this possible.

If this mindset sounds implausible, consider this story of how it actually worked. A hacker (originally the word meant simply a good programmer) named Peter Samson wrote a program to make primitive computers play music—and then gave it away. "Samson proudly presented the music compiler to DEC to distribute to anyone who wanted it. He was proud that other people would be using his program. The team that worked on the new assembler felt likewise. .... They felt honored when DEC asked for the program so it could offer it to other PDP-1 owners. .... As for royalties, wasn't software more like a gift to the world, something that was reward in itself? The idea was to make a computer more usable, to make it more exciting to users... When you wrote a fine program you were building a community, not churning out a product."[12]

In a system where information can be copied perfectly at low cost, it is tempting to treat all information that way. Steven Levy asserts that the experience of working with an early computer at MIT led to the "hacker ethic", one tenet of which was, "All information should be free." More recently, this has mutated into the slogan, "Information wants to be free." In a sense, this is true: some kinds of information are designed to entice us to copy them. Songwriters try to make their tunes "catchy". Programmers try to make their programs useful. The most successful information almost seems to "want" people to copy it.[13]

A backlash has been building against the free copying of information, with Digital Rights Management,[14] the Digital Millennium Copyright Act[15] (in the U.S.), and network snooping and lawsuits aiming to curtail most illegal copying—and some legal copying as well. Although there are some partnerships between Commercial and Information entities (such as between the Red Hat software company and the Linux project), there are many problems yet to be resolved before Information ethics can find its proper place in society. Molecular nanotechnology will exacerbate these problems, by making Information actions relevant not only to information but to physical products. Additionally, the sudden advances in manufacturing and product capability will create substantial security issues, possibly equivalent to or even greater than nuclear weapons, and economic changes comparable to the Industrial Revolution but far more rapid. Better methods of inter-group collaboration will need to be developed, and quickly.

Interactions Between the Three Systems of Action

Trust is important in commerce, even between competing companies; by contrast, competing Guardian organizations frequently are enemies and cannot afford to trust each other. It is obvious that the principles of these organizations are different: buying products is appropriate, while buying pardons is not. What is less obvious at first is that these systems are actually incompatible. Many centuries of development have created distinct traditions and expectations for each kind of task. There are strong reasons why sets of principles should not be mixed, and why each system should be applied only to the tasks it is suited for. History has shown what happens when this advice is ignored: as discussed below, the failure of the Soviet economy is one example.

Guardian and Commercial systems have learned to coexist, and even to benefit each other. A healthy flow of commerce needs Guardian organizations to minimize the problems of theft, fraud, and piracy. Guardian organizations don't actually need Commercial organizations, but without commerce the system reverts to something like feudalism: warlords fighting to maintain and extend their land, and peasants engaged in heavily-taxed zero-sum farming when they're not being drafted for cannon fodder. Without organized force, trade and wealth are impossible; without commerce, society stagnates. However, it is very important that the two systems of action be embodied in different organizations.

As Jacobs points out, when government takes control of commerce, the problems of Soviet Communism may arise: severe lack of competition, innovation, and incentive.

When commerce takes control of government, the result may reflect the worst abuses of notorious “company towns”, where workers can become indistinguishable from slaves. It is also likely that consumers will suffer from higher prices, unchecked deception, and lack of competition. If the systems are mixed, you might get something like the Mafia: engaged in both commerce and force, willing to destroy in order to advance its goals, with no financial checks and balances on its financial activities and few legal checks on its forcible (criminal) activities. Another common result of mixing principles is a government where everything is for sale—you can literally get away with murder if you know whom to bribe.

This understanding provides a broad foundation for public policy. To solve problems related to the minimizing of harm (theft, invasion), an organization following the Guardian system is best. To solve problems related to maximizing wealth (trade, invention), an organization with Commercial principles is preferable. And an organization that mixes the two systems—for example, one that makes laws but is subject to financial influence—is likely to be dangerous, or at least counterproductive. It's also worth noting that a large Commercial monopoly may take on Guardian characteristics—using deceptive or unfair business practices, threatening legal action to harm competitors, and buying favorable laws. If the government is willing to be bought, consumers and citizens will have no protection and unethical conditions will multiply. “Free market” must not be implemented to the extent that Commercial entities are allowed to engage in Guardian behavior.

As the copying of information becomes cheaper, Information principles and actions become increasingly significant. When thousands of copies can be made at negligible cost, the concept of ownership becomes fuzzy and leads to ethical problems. Almost everyone would agree that the purchaser of a music CD has every right to copy the tunes to her computer and listen at her desk (note that U.S. copyright law contains a concept of “space-shifting”[16]). But should that person be able to copy “her” tunes to her friend's computer and let him listen as well? If so, then what's wrong with putting the tunes in a file-sharing network like Napster so that everyone can enjoy them? At some point, the principle of free copying collides with the principle of commercial rights. Napster was shut down after a lengthy legal battle. Despite the wishes of many computer programmers, hobbyists, and users of information, the new Information system cannot take over the world. But it can, and should, lead to cheaper or even free access to some forms of information. The ongoing struggle between the Recording Industry Association of America (RIAA) and music sharers is evidence that policy issues have not been satisfactorily resolved.[17]

The Danger of Monstrous Hybrids

One of the most important concepts in Jacobs' theory is that of “Monstrous moral hybrids”, which are created when an organization tries to adopt inappropriate principles, or systems of action. A government that took bribes would certainly qualify, as would a government that tried to regulate all commerce by force and central planning. Both problems arise from a Guardian entity involving itself too much with money. Even a subtle mixture can have distinctly bad effects. Jacobs gives the example of a police department that tried to make itself more efficient by giving bonuses to police officers for making arrests. Note that "Be efficient" is a Commercial and not a Guardian principle, and Guardians are supposed to "Shun trading." The result was predictable in hindsight: many false arrests were made in order to get the bonuses. The Mafia engages in trade and force both, using loyalty, greed, and coercion as motivators. Most people outside of the Mafia would agree that it is monstrous.

Applying a consistent system of action to the wrong situation can be as bad as mixing and matching principles for convenience. Commerce stagnates if it is centrally regulated; conversely, mercenaries do not make trustworthy or effective soldiers. Information principles, when applied to other people's private property, result in actions that are indistinguishable from theft. Organizations should not try to extend their influence to situations that fall outside their ability to address appropriately. Table 2 gives some indication of which kinds of issues are best addressed by which systems.

Table 2: Characteristics of Systems

	Information	Commercial	Guardian
Transaction Benefit	Unlimited Sum	Positive Sum	Zero Sum
Planning Horizon	Months	Years	Decades
Openness	Very Open	Semi-open	Secretive
Cautiousness	Low	Medium	High
Use of Force	Unthinkable	Deprecated	Common
Goals	Freedom	Profit	Safety
Personal Identity	Actions, Skills	Possessions	Affiliations
Organizational Structure	Collaboration	Loose Hierarchy	Strict Hierarchy

A corollary of the dictum that systems of action must not be mixed is that no single organization should be expected to solve all problems. We should not expect a government to give non-citizens the same privileges as citizens. A corporation should not be expected to do business with people who have no money. Information creators should not be expected to decide whether or how to restrict their work or the information they produce. This is probably unwelcome news; it is strongly tempting to make any powerful organization responsible for anything it touches. However, to do this would create unhealthy, inefficient, or even tragic situations: e.g., a state forced to violate its own security; a corporation forced to waste money; a creator prevented from creating. To force an organization to adopt alien principles (or to solve alien problems) is to force it to act unethically. Unfortunately, this means that many organizations will create problems that they are not equipped to solve, and almost all organizations will confront problems that they cannot address.

Since any single organization can only deal with a fraction of the possible problems, the solution is to have organizations of each type working together, keeping each other in check, and letting solutions to problems emerge from their interaction. Governments and commercial entities have had many centuries to learn how to work together. Information principles are somewhat newer, since they only became widespread with the availability of cheap computers and the Internet.

The Internet investing bubble of the 1990's, typified by the phrase "Information Economy", is a clear result of attempting to mix two systems of action. Companies that tried to make money via the Internet in the Information Economy were doomed from the start. One after another, they found that they could get a large number of users as long as the "customers" were not paying them any money; the customer base grew exponentially, as might be expected in an unlimited-sum system. Yet when the companies began charging fees, most of the users went elsewhere. We may suspect that many of these users felt that charging fees for what had been a free service was unethical. In terms of Information principles, it is in fact improper to charge fees. The companies that succeeded, such as Amazon.com, were the ones that offered a traditional Commercial service, using the vast potential of the Internet simply as a communication channel.

The Open Source movement illustrates the potential for all three systems to work together productively. Open Source was inspired by the Free Software movement, which is militantly dedicated to purely Information principles; for example, the Free Software Foundation does not approve of distributing free software together with non-free software.[18] Open Source, by contrast, welcomes the collaboration of Commercial (non-free) software manufacturers. Some companies, such as Red Hat, make money by selling and supporting Open Source software. The Open Source licenses, which keep the software from being hijacked by commercial interests, make use of copyright law, so that Open Source and Free Software rely on our Guardian legal system for protection just as much as any Commercial entity does.

Molecular nanotechnology will make it possible to manufacture or copy physical objects almost as easily and cheaply as data is produced and copied now. This will provide a broader domain for the advantages of unlimited-sum transactions, but will also pose additional problems and tangible dangers. An understanding of all three systems of action will help in policy-making, administering the new technologies to maximize the benefits while minimizing the problems.

POLICY ISSUES OF MOLECULAR NANOTECHNOLOGY

Advanced nanotechnology will present a wide range of problems and opportunities: not just diverse issues, but different kinds of issues. Many of these issues have arisen already, with older technologies and institutions. The anticipated impacts of nanotechnology are so extensive that some analysts[19] classify it as a general purpose technology (GPT). A GPT offers transformative applications across a wide spectrum of industries, and can have massive economic and social implications. Previous examples include steam engines, railways, internal combustion engines, electricity, telecommunications, and computers.[20]

Some of the issues raised by molecular nanotechnology are new and unfamiliar, and even the old issues take on new urgency when they occur in new combinations. The promise of MNT is material abundance and rapid improvement of technology at low cost and high convenience.[21] The threat of MNT is the potential of developing and fabricating dangerous weapons and other undesirables covertly or in large quantity.[22] To minimize the threat while maximizing the benefit will require the cooperation of many organizations of several distinct types.

Reducing the Risks of Nanotechnology: Guardian Principles

In 2002, researchers used DNA purchased over the Internet to build fully infectious poliovirus.[23] MNT-based rapid-prototyping or manufacturing capability, whether available in the home or by mail order, would be able to build things far more intricate and functional than a simple virus. Nanometer-scale computer circuitry will probably be one of the first products; this enables all sorts of computing, communication, and surveillance devices. If the technology is based on biochemistry, then medicines, drugs, and poisons may be available; indeed, a wide range of custom-designed chemicals. If the technology is based on rigid machine parts, then a wide range of shapes and manipulations will be possible; it has already been demonstrated that a cavity in plastic can act as a "binding site" to trap chemicals,[24] so even a purely mechanical nanotechnology should be able to interact with biochemical systems to some extent.

Today's supercomputers can be used for tasks of military significance, such as simulating nuclear explosions and cracking codes. In fact, sufficiently powerful computers are considered a munition (armament), and their export is controlled. Certain software, including some common encryption software, is also considered a munition. But the supercomputers of today are the desktop computers of tomorrow and the palmtops of the day after, and nanotechnology will certainly allow the building of computers that are immensely powerful by today's standards. These computers may be integrated with devices of varying degrees of sophistication; for example, a near-magical surveillance technology could be packaged into a device too small to see. Even the mundane types of nanotechnology may need to be controlled.

The more exotic suggestions, such as “gray goo”,[25] can get quite scary. Gray goo is a kind of nanodevice that takes in biomass and turns it into copies of the gray goo device. In theory, if such a device were not countered, it could “eat” the biosphere. Fortunately, the design of such a device would be quite difficult, and devices of the gray goo class would have no commercial or even military use, since more specialized non-replicating devices would be far more efficient. It is thus highly unlikely that anyone would build a gray goo, or device that could run amok and become gray goo, by accident, and military or commercial organizations would have little interest in building such a thing on purpose. However, the prevalence of computer worms and viruses indicates that some people do build things like this for fun.

If approached with pessimism, MNT appears far too dangerous to be allowed to develop to anywhere near its full potential. However, a naive approach to limiting R&D, such as “relinquishment”,[26] is flawed for at least two reasons. First, it will almost certainly be impossible to prevent the development of MNT somewhere in the world. China, Japan, and other Asian nations have thriving nanotechnology programs, and the rapid advance of “enabling technologies” such as biotechnology, MEMS, and scanning-probe microscopy (SPM) ensures that R&D efforts will be far easier in the near future than they are today. Second, MNT will provide benefits that are simply too good to pass up, including environmental repair; clean, cheap, and efficient manufacturing; medical breakthroughs; immensely powerful computers; and easier access to space.

If the spread of molecular manufacturing ability cannot be prevented, then the resulting problems must be dealt with piecemeal. Some sort of watch will have to be kept, on an ongoing basis, to reduce the number of criminals, terrorists, and hobbyists building dangerous nanodevices, and to clean up problems when they do occur. A deployment of nasty stuff such as “time bomb dust” or “gray goo” might have to be dealt with quickly, forcefully, and invasively. A nanotech-focused police organization would need broad powers to exert force on random people or property. It is crucial that such an organization be incorruptible. Several of the Guardian principles, such as “Shun trading” and “Be exclusive”, work to minimize corruption. Others, such as “Exert prowess”, “Take vengeance”, and “Deceive for the sake of the task”, are uncomfortable to many people but probably necessary.

One might think that a Guardian organization would need to adapt and change swiftly when facing a new and rapidly improving technology. However, there is a good reason why Guardian principles include “Adhere to tradition” and “Be fatalistic.” An organization that allows itself to change too quickly may lose focus. In addition, if the Guardians are too quick to try new things, they may create cures worse than the diseases. This is one of several reasons why a Guardian-only solution cannot work. Instead, the Guardians must be willing to allow a broad range of innovation, carried out by more than one type of organization, and then adapt the most suitable technologies to do their own work.

Nanotechnology for Profit: Commercial Principles

“Be thrifty. Be optimistic. Be efficient. Be industrious. Be honest.” These are a few of the Commercial principles (see Table 1). Commercial organizations will work very hard to give people what they want—or what they will pay for, which often is the same thing. If a product needs to be developed to satisfy a market, some company somewhere is probably working on it. Money can be a great incentive.

Commercial organizations must compete, but they are not allowed to use force—that is reserved for Guardian organizations. So they try to make their products better, and sell them to more people. They are willing to invest in developing products, and making them easier to use, and making consumers aware of them. They are accustomed to collaborating and innovating, and to making and keeping contracts. Molecular nanotechnology will enable the creation of many new technologies and products, many of them quite specialized and quite useful. The development of all of this potential—not to mention the development of basic LMNT capabilities—requires an incentive, and Commercial organizations will use the incentive of money to bring the benefits of molecular manufacturing to a wide swath of the population.

In fact, the advantages of capitalism are so well established that it would be redundant to spend more time on them here. We will, however, mention some limitations. Motivated by money, corporations do not tend to consider factors that cannot be quantified. Some industries are notorious for creating environmental or health problems, simply because it is more profitable to do so than to pursue alternatives. In addition, money tends to flow where it can achieve a quick return on investment; venture capitalists won't invest in a project that takes longer than a few years to reach fruition. Long-term benefits, like hidden costs, are largely ignored by the Commercial mindset.

This is not bad; it simply means that Commercial entities cannot be the sole decision-makers. Some opportunities would be missed due to long development cycles or unwillingness to do “pure research”, and some problems would be created that, even if they were acknowledged, would not be avoided due to institutionalized short-sightedness.

Another limitation of the Commercial system is that there is no obligation toward people who can't afford to pay for a product or service that they need. Again, this is not bad, as long as there are non-commercial organizations that can take up the slack. A purely philanthropic, altruistic, for-profit corporation is a contradiction in terms—it would become ineffective through making unprofitable choices. What this means is that commercial organizations should do what they do best, developing and selling products with a relatively short planning horizon, but they should not claim ownership of everything. A society based entirely on Commercial principles would have many people literally starving in the streets, and most of us would not want to live there. The solution is not to nationalize commercial entities. Societies that tried that saw their economies stagnate or implode. The solution is to maintain thriving commercial entities that care only about the bottom line, in balance with other entities using other principles.

Fortunately, molecular nanotechnology—even in its early, limited form—will allow the creation of vast amounts of non-commercial, widely distributable wealth.

Unlimited Benefits of Nanotechnology: Information Principles

It may be only a matter of time until the manufacture of products becomes as cheap as the copying of files. Molecular nanotechnology will help this process along, because the first practical self-duplicating factory will almost certainly be designed on the nanometer scale. A tabletop model might weigh a kilogram. The amount of raw materials required to produce a new factory might cost only a few dollars or Euros, and a well-designed factory could process that much material in an hour or so. Once one such factory exists, it and its copies can be used to make an unlimited number of tabletop factories, cheap enough to give away. Building a new product would be as simple as emailing its blueprint to the factory—which might be sitting beside your computer.

If nanofactories were ubiquitous, then their products would be readily available. The only limits would be raw materials—which would be completely renewable—and licensing fees for the products. If a product design were created and given away, as Linux is given away, anyone who wanted or needed one could have it. Any product that could alleviate poverty or suffering might be instantly available to everyone. As soon as a need was recognized, designers and programmers would be motivated by the desire to gain reputation through filling the need and by the knowledge that their work could improve the lives of millions of people.

This would only be possible, however, if nanofactories were not restricted to prevent the making of free products. Commercial entities, of course, would have a strong interest in preventing competition from products that people didn't pay for. And Guardian entities would be sweating over the malicious ways that an unrestricted factory might be used. A completely unrestricted factory looks like a bad idea for several reasons, including intellectual property violations and dangerous products. A tightly controlled factory is a bad idea for at least three reasons: heavy restrictions would prevent the alleviation of vast amounts of human suffering, would hinder the creation of an undreamed-of level of prosperity, and would also make a black market inevitable. Some sort of compromise must be reached.

There are two ways to exert control over nanofactory manufacturing. One way is to maintain tight limits on distribution of nanofactories and/or access to usage, which is inadvisable for several reasons, as we have seen. Another option is tight built-in technical restrictions on individual nanofactories.[27] Preventing a nanofactory from building unapproved products likely could be done using technologies already in use today. In fact, it appears that the nanofactory control structure could be made virtually unbreakable. This does not mean that implementing technical controls on production necessarily will be easy. A great deal of work must be done to develop an effective strategy for product certification.

Table 3 illustrates four separate options of combining controlled or open distribution with restricted or unrestricted nanofactory design. Use of a restricted nanofactory design combined with widespread deployment and ample access appears to be the best solution.

Table 3: Matrix of Nanofactory Control Options

		Distribution & Access	
		TIGHT	**LOOSE**
Technical Restrictions	**LOOSE**	· Intellectual property violations · Manufacture of dangerous products · Makes black market inevitable	· Intellectual property violations · Manufacture of dangerous products · Removes incentive for black market
	TIGHT	· Prevents alleviation of human suffering · Hinders creation of prosperity · Makes black market inevitable	· Allows alleviation of human suffering · Fosters creation of unprecedented prosperity · Reduces incentive for black market

A SPECIFIC PROPOSAL

Nanofactories as Infrastructure

The Internet was originally created by a governmental agency.[28] As it grew, it was supported by funding and technology from many sources, including Guardian (government), Commercial (corporate web sites and private users paying ISPs), and Information (hobbyist programmers, responsible among other things for creating the Internet Protocol). The Internet is an infrastructure, usable by any group. A comparable *molecular manufacturing infrastructure* could provide a project to which all three types of groups could contribute. Guardians could regulate usage, Commerce could charge tolls, and Information groups could enhance both the infrastructure and the products available through it.

As explained above, a tabletop nanofactory appears to be quite feasible even with limited molecular nanotechnology. Such a factory could form the core of a molecular manufacturing infrastructure. Depending on the cost of the factory, it could either be available in service bureaus or in individual homes. It would be able to produce an immense range of incredibly useful products at very low cost. The benefits to society would be almost incalculable: the financial and environmental costs of manufacturing and transportation would be greatly reduced, and new products would be available far more quickly, customized for each user.

Once nanofactories can be built, people will demand access to them. If legitimate access is not provided, some of the "have-nots" will obtain black market devices of comparable functionality. Such devices would presumably be uncontrolled, thwarting any attempt to regulate, tax, or charge royalties on products they produce. Since a small nanofactory can make a bigger one, and a large one can make thousands of duplicates, smuggling would be impossible to prevent.

To minimize the black market, it is in the interests of both Guardian and Commercial organizations to supply nanofactories, as capable and flexible as possible, to the entire global population. This flexibility must include the ability to build certain products with minimal royalties or taxes—preferably zero added cost, because anything else would only encourage illicit factories. Of course, the factories cannot be completely unrestricted. Certain weapons, substances, and dangerous nanobots should be prohibited or restricted, and all commercial intellectual property should be controlled according to the wishes of the owner. However, aside from these limitations, Information system workers should be given free rein to design and give away any product. This will greatly reduce the pressure for illicit factories.

Free availability of freely designed products will not eliminate commercial value; "Promote comfort and convenience" is a Commercial principle, but not an Information principle. We see this in practice—it is well known that Linux is more difficult for the average user than commercial operating systems. Products of the Information school are likely to be highly functional, but not especially easy to use or stylish. Still, the activities of Information-producing groups will fill needs that the Commercial groups will fail to fill. They will also serve as a source of innovation that will be usable by both Commercial and Guardian groups, which will certainly have their hands full trying to keep up with the rapidly advancing technology.

No matter who designs the products, whether it's Commercial or Information workers, a certain amount of regulation by Guardian entities will be necessary. Some classes of products, of course, will not require much scrutiny. But even the simplest medical and hygiene products will benefit greatly from MNT and will need to be studied carefully to ensure their safety. Government regulatory bodies such as the U.S. Food and Drug Administration (FDA) and private organizations such as Underwriters Laboratories (UL) have been performing this function for conventional products.

The creation of a worldwide network of restricted nanofactories appears to maximize the benefits of molecular nanotechnology while providing opportunities to minimize the risks. Guardian/regulatory, Commercial, and Information/creative organizations all will be able to pursue their goals through this infrastructure, which will provide substantial benefits to local, national, and global societies while short-circuiting many illicit uses of molecular manufacturing.

Naturally, there will be significant conflicts between groups embodying the three systems of action. People from the Information tradition will be unhappy with regulation ("Shun authority" is one of their principles), and may become frustrated with necessary limitations and try to obtain (or create) unauthorized factories. The Commercial tradition will want to maximize income-generating ability, so will try to resist or even subvert both regulation and free products. The Guardians will want to exert control over potentially dangerous technologies, even to the extent of crippling commerce and innovation. In the normal course of things, these conflicts could take decades to resolve naturally. The rapid progress of MNT, however, may not allow that much leeway. We believe that a carefully designed decision-making structure will be necessary in advance.

There is some basis for hope. Researchers running experiments on human subjects in the U.S. have been required for several decades to submit their experimental plans to an "Institutional Review Board",[29] and researchers in other sensitive fields such as genetically modified bacteria face similar restrictions. Designing a functional product to be made by molecular manufacturing may require significant training, and the training period perhaps could be used to create a culture of responsibility. Commercial and Guardian systems have a history of cooperation and balance which may stabilize their conflicts as they address a technology with such large military and commercial potential. Now is the time to begin designing the procedures, organizations, and technologies that will be required to make this transformative infrastructure a success.

MNT Administration: A Balance of Power

As we have seen, molecular nanotechnology will present a large range of new problems and new opportunities—and no single organization will be capable of addressing all of these issues.

Commercial institutions cannot adequately address either the security or the abundance issues of cheap distributed molecular manufacturing. If no institution takes responsibility for forcibly preventing the worst abuses of the new technology, commercial liability will act to reduce the risks of any given design, and this may prevent disaster for a while. But as MNT becomes more accessible, more pervasive, and more powerful, the possible dangers will grow beyond any ability to financially underwrite the risk, and businesses will be institutionally unable to address them appropriately. Likewise, if all information is owned as "intellectual property" and no institution is allowed to distribute MNT-related information and products freely, clandestine and foreign institutions will spring up to meet the demand, creating a massive exchange of "pirate" designs and a security liability.

Neither of the other systems can do the job alone either. A Guardian institution cannot administer an adequately efficient market, and a black market is the inevitable result. An Information system, dedicated to the free spread of resources, will be completely unable to plan and implement security.

In summary, Commercial and Information groups cannot be trusted to take appropriate precautions in every case, so a Guardian approach is sometimes necessary. Likewise, because Guardian and Information ethics do not create money, Commercial organizations must be involved to pay for large parts of the development and deployment of the technology. Finally, although Guardian principles include “dispense largesse”, neither Guardian nor Commercial organizations can be expected to create and distribute the almost limitless benefits that will become possible from vastly improved materials and manufacturing, so Information groups also must have a role to play.

Ideally, each organization involved in MNT would be aware of its own principles and the principles of the organizations it interacts with, and would make good decisions about which problems to tackle and which problems to leave for someone else. In practice, of course, organizations are usually not so self-aware, and even when they are, shortsighted self-interest may tempt them to expand into areas where they have no competence. In the end, an organization that overreaches itself will find that its plans don't work; it will make ineffective and antisocial decisions, and it will be out-competed by its fellow organizations and attacked by those it has encroached on. However, such a process may take much time and cause much destruction; consider the long-delayed fall of the Soviet Union, due in large part to its application of Guardian principles to commerce and information. MNT will develop far too quickly for such slow adjustments.

None of the three systems of action contains the principle, “Be aware of the limitations of your principles.” Most organizations that attempt to deal with MNT will not know why they act as they do or why they cannot successfully address certain problems. If not constrained externally, they will try and fail, perhaps creating significant waste or even tragedy. Yet the organization that is normally responsible for constraint—the government—is itself limited in its principles and its understanding of them.

Development and application of MNT policy cannot be reactive. The problems, individually and collectively, could spiral out of control before today's institutions have time to react. Prior to the advent of MNT, a collaborative *international administrative council* of some kind will have to be designed and created. However, at worldwide levels, where things move slowly, this might take as long as twenty years. If advanced nanotechnology could arrive within ten or fifteen years, urgent action is called for now.

This administrative council will need global scope, and will require careful and innovative planning to balance the requirements of Guardian, Commercial, and

Information approaches to problem solving. The proper design of such a system will be extremely difficult to accomplish, and getting it implemented will be even harder. The authors of this paper have been unable to identify any existing institution capable of performing this. Those of us who study the societal implications of nanotechnology and other advanced technologies must find ways to address this critical lack, and prepare for the radical changes that may soon be upon us.

END NOTES

[1] Richard P. Feynman, "There's Plenty of Room at the Bottom". A transcript can be found at http://www.zyvex.com/nanotech/feynman.html.

[2] See http://www.nsf.gov/home/crssprgm/nano/omb_nifty50.htm.

[3] For an overview of factors supporting early development, see http://www.crnano.org/timeline.htm.

[4] See “Bootstrapping a Nanofactory” at http://crnano.org/bootstrap.htm.

[5] Jane Jacobs, *Systems of Survival: A Dialogue on the Moral Foundations of Commerce and Politics*, http://www.amazon.com/exec/obidos/ASIN/0679748164/104-2718304-3423907.

[6] Dartmouth College is providing free long distance service using VoIP technology, because billing would cost too much. The New York Times, “A New Kind of Revolution in the Dorms of Dartmouth”, Sept. 23, 2003, http://www.nytimes.com/2003/09/23/technology/23DART.html.

[7] “Open Source Initiative OSI - Welcome”, http://www.opensource.org/.

[8] Adapted from “Three Systems of Ethics for Diverse Applications”, published online at Nanotech-Now.com, http://nanotech-now.com/Chris-Phoenix/diverse-ethics.htm.

[9] Mike Treder, “Accelerating Paradigm Shifts in Information Storage and Retrieval”, http://www.lucifer.com/~exi/ideas/journal/previous/2002/09-02.html.

[10] The Linux home page is at http://www.linux.org/.

[11] Seagate, a major hard drive manufacturer, has announced that it will supply disks with a version of Linux pre-loaded at no extra charge. http://www.lindows.com/lindows_seagate.php.

[12] Steven Levy, *Hackers: Heroes of the Computer Revolution*, p. 56, http://www.amazon.com/exec/obidos/ASIN/0141000511/104-0739447-3127111.

[13] According to the theory of "memes", information evolves and copies itself in a manner similar to genes. For a further explanation, see "Meme Central - Memes, Memetics, and Mind Virus Resource" at http://www.memecentral.com/.

[14] A good resource is "Digital Rights Management and Privacy", at http://www.epic.org/privacy/drm/.

[15] For status and analysis, see http://www.arl.org/info/frn/copy/dmca.html.

[16] "EFF Fair Use FAQ" http://www.eff.org/IP/eff_fair_use_faq.html (see question 4 for a brief discussion of space-shifting).

[17] Recently the RIAA has sued individual file-sharing users, including a 65-year-old woman who does not own a computer capable of running the accused software (http://www.internet-magazine.com/news/view.asp?id=3730).

[18] "Open Source - Gnu Project - Free Software Foundation(FSF)" http://www.gnu.org/philosophy/free-software-for-freedom.html.

[19] See, for example, "Big Money in Thinking Small", authored by Michael Mauboussin and Kristen Bartholdson of Credit Suisse First Boston.

[20] For more on general purpose technologies, see http://www.sfu.ca/~rlipsey/C2.pdf.

[21] For a fuller discussion of benefits, see http://www.crnano.org/benefits.htm.

[22] Risks are explored in more detail at http://www.crnano.org/dangers.htm.

[23] Story at http://webmd.lycos.com/content/article/48/39273.htm.

[24] For example, see pubs.acs.org/cen/topstory/8030/8030notw4.html.

[25] An excellent listing of references is at http://www.foresight.org/Nanomedicine/Ecophagy.html.

[26] Bill Joy issued the most notorious call for relinquishment in his April 2000 Wired article, "Why the future doesn't need us", http://www.wired.com/wired/archive/8.04/joy.html.

[27] See "Safe Utilization of Advanced Nanotechnology", by Chris Phoenix and Mike Treder, at http://www.crnano.org/safe.htm.

[28] See "History of ARPANET" at http://www.dei.isep.ipp.pt/docs/arpa.html.

[29] See http://www.fda.gov/oc/ohrt/irbs/default.htm.

APPENDIX

Feasibility of Molecular Nanotechnology

It has been claimed repeatedly that some law of chemistry or physics will forbid molecular nanotechnology (MNT). To date, all such claims have been refuted in detail. Many of the objections, including those of chemist Richard Smalley, do not address the actual proposals. The rest are unfounded and incorrect assertions, contradicted by detailed calculations based on the relevant physical laws.

While initial descriptions of MNT included the ability to make almost any chemical substance—in Richard Feynman's words[1], "Put the atoms down where the chemist says, and so you make the substance"—recent work has focused on a much more limited chemical capability. Although the feasibility of a truly general mechanochemistry has not yet been studied in depth, the feasibility of building rigid carbon-lattice structures has been examined by careful calculation. Objections based on thermal noise, quantum uncertainty, or other "laws of physics" are invariably assertions without any investigation to back them up. All calculations to date indicate that this could be done at room temperature with high reliability by nanoscale programmable machinery.

On the chemistry front, the most notable objections have come from Smalley, who asserts that handling atoms with "fingers" would require fingers that were too "fat" and too "sticky" to be effective. These objections do not apply to any published proposal for diamond-based MNT: no such proposal has included the use of "fingers" of any kind. In fact, these proposals do not involve the handling of separate atoms at all; the atoms are always bound to larger molecules, and shift their bonds from one to another in traditional chemical fashion.

Recent objections to MNT implicitly acknowledge its theoretical feasibility, but assert that it is too complex a problem to be solved any time soon. However, these objections have also been addressed in the literature. Much mechanochemical complexity vanishes with the shift from MNT to LMNT—from general chemistry to limited diamondoid chemistry. The complexity of large MNT-based manufacturing systems turns out to be surprisingly low. A peer-reviewed technical paper by CRN's Director of Research has examined the issues involved in combining quadrillions of nanoscale fabricators into a tabletop factory, including power and heat, reliability, control, and convergent assembly of parts into products. The paper concludes that all of these issues are addressable within the scope of current engineering practice.