Background

The idea of guidelines for the safe development of MNT (Molecular Nanotechnology) has been discussed within the Foresight community for over a decade. It is inevitable that any guidelines put forth today will be further discussed and perhaps substantively changed; but the dialog on specific proposals must begin somewhere. This latest version of the Foresight Guidelines represents another step in an ongoing discussion.

In spite of the diversity of briefing materials and views represented at the Monterey workshop in February of 1999, the participants managed to discuss the technical and policy issues with both intensity and civility. While any one participant might have preferred more or less emphasis on a particular issue, the group was able to converge on a common set of draft guidelines for the development of MNT.

The group agreed to review the Guidelines among themselves, discuss them in wider Foresight meetings during 1999, and then release them on the WWW for review by the larger community. The goal was to get the Guidelines endorsed and adopted by organizations sponsoring MNT research and development projects, and to inspire effective self-regulation wherever necessary and possible.

Another goal of the Workshop members was to educate MNT researchers about the potential benefits and risks of the technology. The long-term goal was to eventually produce a dialog and set of Guidelines that would be useful to policy makers, the public, and the MNT research and development community.

The Foresight Guidelines are intended as a living document, subject to modification and revision. Early drafts have been reviewed and revised several times since the Monterey workshop, including during Foresight / IMM sponsored discussions led by Neil Jacobstein in May and November of 1999. They were also provided in the attachments to Ralph Merkle's June 1999 Congressional testimony on MNT, and referenced in Neil Jacobstein's presentation on Nanotechnology and Molecular Manufacturing: Opportunities and Risks at Stanford University's Colloquium for Doug Engelbart in January of 2000. The Workshop participants debated whether the Guidelines were sufficiently developed for widespread publication, when Bill Joy's article: "Why the Future Doesn't Need Us" was published in the April 2000 issue of *Wired* Magazine. This article raised public awareness of the potential dangers of self-replicating technologies, including nanotechnology.

Since that time, the Guidelines were reviewed critically by Robert Freitas, and revised by Ralph Merkle and Neil Jacobstein. Version 3.6 of the Guidelines was discussed in a May 2000 Foresight workshop session led by Neil Jacobstein. Bill Joy was invited to participate in this discussion. He made several constructive suggestions, including one that outlined a guideline on closing the economic incentives loop via an insurance policy requirement for developers. Jacobstein incorporated the feedback from this and subsequent discussions into version 3.7 of the Guidelines, and they were then published for open review on the web.

Neil Jacobstein rewrote the Guidelines as a form of self assessment scorecards for version 4.0, based on his observation that this kind of self assessment is becoming a standard part of quality and six sigma programs in industry and government. He combined and added some new guidelines, including a guideline based on a paper by Eric Drexler and Chris Phoenix in the Journal of Nanotechnology on "Safe Exponential Manufacturing". This paper made the case for MNT and molecular manufacturing using a hierarchy of machine tools, without the need for general purpose self-replicating assemblers. Glenn Reynolds edited the draft and provided an analysis by his law students on current treaties and the fact that weaponized MNT might not be covered by them. Eric Drexler also reviewed the draft and made additional editorial suggestions.

Version 4.0 of the Guidelines is available at the Foresight web URL: http://www.foresight.org/guidelines/. We encourage your ideas and constructive criticism about how to improve the Guidelines.

Eventually, the Guidelines need to become sufficiently specific that they can form the basis for a legally enforceable framework within which MNT development can be safely pursued. Future versions of the MNT Guidelines might eventually be enforced via a variety of means, possibly including lab certifications, randomized open inspections, professional society guidelines and peer pressure, insurance requirements and policies, stiff legal and economic penalties for violations, and other sanctions. Enforcement will be inherently imperfect, but the deterrent effect of unpredictable inspection, combined with predictable and swift consequences for violations, may prove preferable to the available alternatives.

Care must be taken that future revisions of the Guidelines do not become so restrictive that they simply drive MNT research and development underground. This could expose compliant countries to the increased risks associated with decreased technical, economic, and military capabilities. It would also sacrifice the many significant economic, environmental, and medical benefits of MNT that counteract serious and certain risks that society now faces in industrialized countries and particularly in the developing world.

Foresight Institute is a non-profit 501(C)3 organization
Foresight Institute • P.O. Box 61058 • Palo Alto, CA 94306 • USA, tel. +1(650)917-1122 • foresight@foresight.org

FORESIGHT INSTITUTE®
Preparing for nanotechnology

Foresight Guidelines Version 4.0:
Self Assessment Scorecards for Safer Development of Nanotechnology

by Neil Jacobstein and Glenn Harlan Reynolds

Nanotechnology, and the future molecular manufacturing it enables, will alter our relationship with molecules and matter as profoundly as the computer changed our relationship with bits and information. Molecular manufacturing will enable us to program matter precisely, and inexpensively. The Foresight Guidelines are designed to address the potential positive and negative consequences of this new technology base in an open and scientifically accurate matter. The objective is to provide a basis for informed policy decisions by citizens and governments, and specific guidelines for the responsible development of nanotechnology-based molecular manufacturing by practitioners.

Version 4 of the Guidelines is presented in the active format of a set of self-assessment scorecards for nanotechnology professionals, industry organizations, and regulatory agencies. Industry organizations for example can now assess and score their own degree of compliance with the Guidelines. This allows the dialog to move from loose recommendations to self assessment of compliance. Precise scoring is not necessary to move the dialog towards self-assessment of compliance with a more operational set of nanotechnology development guidelines. As the dialog progresses, more precise scoring guidelines are likely to evolve.

This version also includes consideration of near and long term forms of nanotechnology, as well as distinctions between specialized and non-self-replicating molecular manufacturing machinery, and general purpose assemblers. It takes the position that molecular manufacturing can be accomplished without self replication or general purpose assemblers, but the probable development of assembler technology is still an important and prudent area for consideration. There are also several new and reworded guidelines. As always, feedback is welcome for the next version.

Preamble

The term “Molecular Nanotechnology” (MNT) refers to the ability to program matter with molecular precision, and scale it to three-dimensional products of arbitrary size. This developing technology presents an unprecedented new set of technical and economic opportunities. The opportunities include: the development of inexpensive and abundant diamondoid building materials with a strength-to-weight ratio 50 times greater than titanium, the possibility of widespread material abundance for all the Earth’s people, the development of revolutionary new techniques in medicine, and the opening of the space frontier for development. Along with these new capabilities come new risks, and new responsibilities. The acceptance of these responsibilities is not optional. The future capabilities of MNT also raise an unprecedented set of military, security and environmental issues. Dealing with these issues proactively, neither amplifying nor downplaying potential risks, will be critical to the positive development of the field.

In all cases, “Molecular Nanotechnology” that provides massively scalable and precise control of molecules needs to be distinguished from less advanced, present day nanoscale technologies such as nanoelectronics or the nanoparticles used in sunblocks and coatings. Further, self-replicating nanomachines or assemblers are unlikely to be the norm for molecular manufacturing. Molecular machine systems can be completely non-biological, and self-replicating assemblers are not necessary to achieve molecular manufacturing capabilities. As Drexler and Phoenix have shown in their Safe Exponential Manufacturing paper (2004 *Nanotechnology* 15 869-872), developing manufacturing systems that use self-replicating assemblers would be needlessly inefficient and complicated. The simpler, more efficient, and safer approach is to make nanoscale tools and put them together in factories big enough to make what is needed. People use tools to make more and better tools, from blacksmiths’ tools to automated machinery. One schema that develops this idea is based on the convergent assembly architecture developed by Ralph Merkle (1997 *Nanotechnology* 8 18-22), where small parts are put together to form larger parts, starting with nanoscale blocks and progressing up the hierarchy to macroscopic systems. The machines in this would work like the conveyor belts and assembly robots in a factory, doing similar jobs. If you pulled one of these machines out of the system, it would pose no risk, and be as inert as a light bulb pulled from its socket.

The first version of the Foresight Guidelines was developed during and after a workshop on Molecular Nanotechnology (MNT) Research Policy Guidelines sponsored by the Foresight Institute and the Institute for Molecular Manufacturing (IMM). The workshop was conducted over the February 19-21, 1999, weekend in Monterey, California. Participants included: James Bennett, Greg Burch, K. Eric Drexler, Neil Jacobstein, Tanya Jones, Ralph Merkle, Mark Miller, Ed Niehaus, Pat Parker, Christine Peterson, Glenn Reynolds, and Philippe Van Nedervelde. The Guidelines have been revised periodically in the intervening years. The resulting Foresight Guidelines (“the Guidelines”) include assumptions, principles, and some specific recommendations intended to provide a basis for responsible development of molecular nanotechnology.

Continued research and education are needed to create a shared understanding and sufficient knowledge base on the entire set of MNT development and risk management issues that must be addressed. While discussion of guidelines can begin today, the scientific and technical community will continue to evolve its understanding of the issues. The Guidelines have already changed over time to reflect

Thus, self-replicating machines are designed to be incapable of replication in any natural environment.

8. Self-replicating machines (if any) are incapable of evolutionary change. For example, the information that specifies their construction is stored and copied in encoded form, and the encoding is such that any error in copying randomizes and thus destroys the decoded information.

Scorecard 3
Government Policy Guidelines

Self Scoring: 0-5, 0 = no compliance, 5 = high compliance

Best Score in this section = 55

1. Regulatory controls distinguish the wide variety of nanotechnologies, and recognize that their different risk profiles require different regulatory policies. Nanomaterials and non-self-replicating nanotechnologies and their end products are distinguished from potentially self-replicating technologies.
2. Regulations promulgated by researchers, industry, or government provide specific and clear guidelines, and encourage inherently safer designs for nanotechnology and molecular manufacturing.
3. Regulators have specific responsibilities and authorities, for providing efficient and fair methods for identifying different classes of hazards, providing approvals when necessary, and for carrying out inspection and enforcement. The goal is to provide the minimum effective regulatory environment to ensure the safe and secure development of various forms of nanotechnology.
4. Economic incentives are provided through discounts on insurance policies for molecular manufacturing and development organizations that certify Guidelines compliance. Willingness to provide self-regulation and open access for third party inspection that safeguards proprietary technology are a condition to utilize advanced forms of molecular nanotechnology.
5. Access to non-self-replicating special purpose molecular manufacturing systems and products is unrestricted unless the special purpose capabilities pose a specific risk.
6. The community of nations and non-governmental organizations practice an effective international means of restricting the deliberate misuse of molecular nanotechnology. Such means should not restrict the development of non-self-replicating nanoscale materials, molecular manufacturing systems, or defensive measures.
7. Accidental or willful misuse of nanotechnology is constrained by legal liability and, where appropriate, subject to criminal investigation and prosecution.
8. Eventual distribution of self-replicating molecular manufacturing development capability is restricted, whenever possible, to responsible actors that have agreed to practice these Guidelines. No such restriction need apply to special-purpose, non-self-replicating molecular machine systems, or to the end products of molecular manufacturing that satisfy the Guidelines.
9. Governments, companies, and individuals who fail to follow reasonable principles and guidelines for development and dissemination of MNT are placed at a substantial competitive disadvantage with respect to access to companies, collaborative organizations, R&D funding, plans, designs, software, hardware, and cooperative market relationships.
10. Industry and government developers collaborate on continuous improvement and use of best practices in nanotechnology and risk management, including the theory, mechanisms, and experimental designs for inherently safer molecular manufacturing, monitoring, and control systems.
11. Regulatory entities sponsor research on increasing the accuracy and fidelity of environmental models of nanotechnology and risk management, as well as the theory, mechanisms, and experimental designs for built-in safeguards and advanced nanodevice defensive or immune systems.

Scorecard 1
Nanotechnology Professional Guidelines

Self Scoring: 0-5, 0 = no compliance, 5 = high compliance

Best Score in this section = 40

1. Nanotechnology developers adopt and practice professional guidelines relevant to the responsible development of both near term and advanced nanotechnology.
2. Nanotechnologists attempt to consider proactively and systematically the environmental and health consequences of their specific technologies. They recognize that the scope and magnitude of potential problems are reduced to the extent that they consider the possibilities, and plan to minimize their effects.
3. Nanotechnology research and development is conducted with due regard to accepted principles and practices of environmental science and public health, with the understanding that significant changes in physical and physiological properties may occur when macroscale materials are developed and utilized on the nanoscale.
4. Nanotechnology products are conceived and developed using total product lifecycle analysis.
5. Molecular manufacturing system designs make no use of self-replicating machines.
6. When controversy exists concerning the theoretical feasibility or implementation timing of advanced molecular nanotechnologies, such as specialized molecular manufacturing components or assemblers, researchers address and clarify the issues rapidly, and attempt to resolve any controversy openly.
7. Any use of self-replicating systems is avoided except in approved and controlled circumstances.
8. Any developers who design or build self-replicating machines adopt systematic security measures to avoid unplanned distribution of their designs and technical capabilities. Both potential benefits and risks of alternative technologies are explored actively, in a balanced and rigorous manner.

Scorecard 2
Nanotechnology Industry Guidelines

Self Scoring: 0-5, 0 = no compliance, 5 = high compliance

Best Score in this section = 40

1. Industry self-regulation is practiced proactively, and tailored to the specific risk profile of the nanotechnology under development. For example, carbon nanotubes should be developed with specialized industrial hygiene controls for particle inhalation or absorption risk. Toxicology studies relating to nanomaterials should be advanced as rapidly as is feasible.
2. Self-replicating machines are distinguished from non-self-replicating manufacturing systems and end products.
3. When molecular manufacturing systems are designed or implemented, they use no self-replicating machines.
4. Any molecular manufacturing device designs specifically limit proliferation and provide traceability and audit trails.
5. Encrypted molecular manufacturing device instruction sets are utilized to discourage irresponsible proliferation and piracy.
6. Use of self-replicating systems is avoided except in approved and controlled circumstances.
7. Self-replicating machines (if any) have absolute requirements (e.g., for externally supplied information, interventions, environmental conditions, materials, components, or exotic energy sources) that are available only where deliberately provided to enable operation of the machine.

that dynamic understanding and input by a wider community (see Background section).

Future discussions of this subject should include detailed consideration of the economic and environmental benefits of MNT, as well as the potential problems. In particular, the need for some controls should not prevent the responsible development of the technology. Rather than have reflexive, or poorly informed controls imposed upon the MNT R&D process, the developing MNT R&D community and industry should adopt appropriate self-imposed controls, formulated in light of current knowledge and the evolving state of the art. The possibility of the necessity for additional controls remains an open question, and its resolution may depend to some extent on the success of voluntary controls.

The NIH Guidelines on Recombinant DNA technology are an example of self-regulation taken by the biotechnology community almost 25 years ago. While the kind of artificial molecular machines of primary interest for nanotechnology are expected to be very different from the kind of biological systems covered by the NIH Guidelines (just as a 747 is very different from a sparrow, even though both fly), the NIH Guidelines illustrate that advance preparations are possible and can be effective. Those guidelines were so well accepted that the privately funded research community has continued to submit research protocols for juried review, in spite of the fact that it was optional for them to do so. In addition, although the NIH Guidelines have been progressively relaxed since they were first released, they did achieve their intended effect.

Experimenters and industry should have the maximum safe opportunities to develop and commercialize the molecular manufacturing industry. In addition, MNT should be developed in ways that make it possible to distribute the benefits of the technology to the four fifths of humanity currently desperate to achieve material wealth at any environmental or security cost. Providing technical abundance alone cannot make a people wealthy and secure. This also requires education, and social arrangements described as a high-trust, civil society. However, technological abundance can alleviate many of the conflicts that stem primarily from rivalry over resources. Reducing this specific cause of conflict via molecular manufacturing could make the world more secure than it is today. In addition, the release from bare economic subsistence could enable billions of people to take advantage of the emerging global classroom enabled by the World Wide Web. This education effect could compound the positive environmental and security benefits of MNT.

Relevant ecological and public health principles must be utilized in conducting MNT R&D. Nanoparticles or components may be inhaled if the proper industrial hygiene precautions are not utilized. Manufactured diamondoid products may not break down easily in the natural environment. Furthermore, consumers may not at first have means readily available to recycle them. Thus, total “product lifecycle” considerations should be taken into consideration as industry develops molecular manufacturing techniques.

Effective means of restricting the misuse of MNT in the international arena will need to be developed. The best present analysis suggests that weaponized MNT would not fall under existing arms-control treaties. Adding MNT per se to the list of technologies covered in Chemical, Biological and Nuclear Weapons treaties would be inappropriate because MNT is not a weapon, but a productive technology with broad applications. Thus, it is more similar to chemical technology than to chemical weapons, and more similar to biotechnology than to biological weapons.

Adding particular applications of MNT to the list of technologies covered in Chemical, Biological and Nuclear Weapons treaties may be appropriate. It should be remembered, however, that the productive capabilities of MNT will be extensions of general manufacturing technology. The military applications of MNT will include the manufacture of high performance aerospace vehicles and precision munitions at low cost.

Overly restrictive treaties or regulatory regimes applied to core MNT technologies could lead to the unintended consequence that only the U.S. and other rule-following nations would be at a competitive disadvantage technologically, economically, and militarily. While most nations are likely to adhere to reasonable safety restrictions, guidelines that are viewed as too restrictive will simply be ignored, paradoxically increasing risk. While a 100% effective ban could, in theory, avoid the potential risks of certain forms of molecular nanotechnology, a 99.99% effective ban could result in development and deployment by the 0.01% that evaded and ignored the ban. There are reasonable arguments on both sides of the treaty question. However, at this time, the Guideline participants as a group have not endorsed any specific initiative to address MNT safety and security concerns through treaty arrangements.

Although MNT does not require any self-replicating machines, either in development or in application, self-replicating machines are theoretically feasible, and they have played a central role in concerns about potential risks of MNT. This speaks to the need for guidelines for the development of molecular nanotechnology. It also underscores the importance of being clear about the difference between the ability to manufacture many copies of a specific product, and the ability of the manufacturing infrastructure to replicate itself.

A self-replicating machine (or replicator) is a specific kind of device that both (1) contains a set of materials processing and fabrication mechanisms sufficient to perform the operations necessary build devices like itself and (2) contains a set of instructions and instruction-interpreting mechanisms sufficient to direct the operations necessary to build a device like itself. All other machines, lacking these exceptional properties, are non-self-replicating.

The safe development and use of MNT depends, in part, on the good judgment of the researchers carrying out this work. The more clearly this is recognized, the more effective researchers are likely to be in avoiding and actively preventing unsafe designs or uses of MNT and in insuring that commercial systems have built-in safeguards. The natural and responsible path for the development of molecular manufacturing makes no use of self-replicating components. However, defense against potential rogue elements who might seek to abuse replicators is a problem not unlike the challenge of controlling the developers of viruses on the Internet. In both cases, a combination of moral and technical education, inherently safer system designs, legal frameworks, and R&D on monitoring and defensive systems may be the best solution available.

Eventually, MNT policy will have to balance potential risks with known benefits, and distinguish between different classes of risks. Molecular Manufacturing and nanotechnology are not one thing, but rather a spectrum of technologies, with radically different risk profiles. A substantial R&D program is needed to clarify the nature, magnitude and likelihood of the potential risks, as well as the options available for dealing with them effectively. In particular, toxicology analyses relating to nanomaterials should be an early priority. Nanomaterial safety is a matter that is distinct from Molecular

Nanotechnology but that may have some relevance with respect to industrial hygiene practices. As with molecular nanotechnology and manufacturing, nanomaterials should be evaluated on the basis of their own risks and benefits, and not treated differently simply because they are “nano” in nature.

There are significant risks associated with failing to address ongoing economic and environmental problems that the development of MNT could help resolve. The Guidelines were not intended to cover every risk or potential abuse of the technology. People still abuse automobiles, and society has responded both by making cars safer to operate, holding drivers accountable for their actions through laws that are enforced, and requiring drivers to pay for automobile insurance. Likewise, industry and governments are held responsible for their use of technologies that have widespread impact.

The Guidelines are intended to cover most of the risks associated with normal development and use of the technology, and to mitigate, as much as possible, the risks associated with potential abuse of the technology. Informed MNT policy could accelerate the safe development of peaceful and environmentally responsible uses of molecular nanotechnology. This includes capturing the opportunity to develop powerful new approaches to medicine, as well as energy efficient and zero emission manufacturing and transportation technologies.

The self assessment scorecards are based on the notion that the people, organizations, and governments that work in the MNT field should develop and utilize professional guidelines and practices that are grounded in science and technology principles, and knowledge of the environmental, security, ethical, and economic issues relevant to the development of MNT. This is based on the notion that “soft law” and cultural norms regarding good practice may be, in some instances, as effective as “hard law” in preventing unsafe practices, and in helping to ensure that unsafe practices are noticed and acted upon. This is not meant to suggest that “hard laws” for safety and health are not useful, and at times appropriate.

MNT research and development should be conducted with due regard to existing principles of ecological and public health. MNT products should be promoted which incorporate systems for minimizing negative ecological and public health impact, and maximizing benefits. It is important to recognize the molecular nanotechnology will eventually provide the best technologies available for remedying existing environmental damage resulting from our current, distinctly suboptimal, technology base.

The global community of nations and non-governmental organizations need to develop effective means of restricting – and responding to – the misuse of MNT. Such means should not restrict the development of peaceful applications of the technology or defensive measures by responsible members of the international community. Further research in this area is encouraged.

Any specific regulation adopted by researchers, industry or government should provide specific, clear guidelines. Regulators should have specific and clear mandates, providing efficient and fair methods for identifying different classes of hazards and for carrying out inspection and enforcement. There is great value in seeking the least-restrictive necessary legal environment to ensure the safe and secure development of each specific type of nanotechnology.