Thirty Essential Studies

Chris Phoenix, Director of Research, Center for Responsible Nanotechnology

CRN believes it is urgent to understand several issues related to molecular manufacturing (MM), to prepare for its possible development sometime in the next decade. The technology will be more transformative than most people expect, and could develop too rapidly for reactive policy to succeed. MM is the result of convergence of many technologies, and will benefit from synergies between them. It will be more powerful than most people will be able to comprehend without serious study.

Molecular manufacturing, along with other technologies that it will enhance or enable, will create new problems and new opportunities that require new solutions. To date, there has not been anything approaching an adequate study of these issues. This document presents some of these issues in the form of recommended studies. CRN's preliminary answers are included to reinforce the relevance and urgency of the investigation.

The studies are organized in several sections. The first section covers the fundamental theory: insights that may be counterintuitive or unobvious and need explanation, but that can be double-checked by simple thought. The second section addresses technological capabilities of possible molecular manufacturing technologies. The third section addresses 'bootstrapping'—the development of the first self-contained molecular manufacturing system (which will then be able to produce duplicates at an exponential rate), including schedule considerations. The fourth section explores the capabilities of products, building toward the fifth section, which raises serious questions about policies and policymaking.

These studies should *not* be undertaken sequentially. Our understanding points to a crisis, and the answers may be needed ASAP. We urge a process of repeated refinement and increasing attention to all the studies in parallel. Of course, these studies can be reorganized and recombined.

To begin this iterative process, we have supplied provisional answers to each study, with supporting data where available. Several preliminary conclusions should be noted here:

■Programmable positional chemistry, with the ability to fabricate nanocomponents, can be the basis of an extremely powerful manufacturing technology. The importance of this is substantially unrecognized.

■Development of molecular manufacturing may be imminent, depending on whether any of several actors has begun investigating it already. We believe that a program started today, even outside the United States, could finish in under a decade, including development of a substantial product design capability.

■Development activity may be very difficult to detect.

■Several considerations, including economics and product sophistication, point to MM being a transformative, disruptive, destabilizing, and potentially dangerous technology.

■Although the technology may be quite dangerous, avoidance and prevention are not viable options. Simple attempts to dominate or control the capability will also be unworkable.

■MM will also have many productive uses, and policy must account for the global-scale problems it can solve as well as a possible high level of civilian demand/utilization.

■Policymaking and preparation will be complex and difficult, and will require substantial time.

Study Titles:

Product performance

Policies and policymaking

Technical/Foundational:

1. Is mechanically guided chemistry a viable basis for a manufacturing technology?

Molecular manufacturing is based on the idea of using physical manipulation to cause reliable chemical reactions, building components for products (including manufacturing systems) from precise molecular fragments. Although several flavors of this have been demonstrated (including the ribosome), there is still skepticism in some circles as to whether a self-contained manufacturing technology can be based on this.

■Is there anything wrong with the basic theory of using programmably controlled nanoscale actuators and mechanics to do chemistry?

To the best of our knowledge, there is nothing wrong with the theory, and it has been demonstrated in certain cases: semi-programmable nanoscale ribosomes do positional chemistry. Nanoscale actuators and mechanical devices exist in a variety of forms and designs. Sub-angstrom-scale precision adequate to do reliable chemistry may be achieved by any of several mechanisms. The question is what families of chemistry are possible. Quite a few have been proposed.

■Can engineered biomolecules (e.g. DNA) do solution chemistry to synthesize more biomolecules with low error rates?

It may be possible to 'cap' and 'uncap' the end of a growing DNA strand with an enzyme-like molecular system, programmable or controllable by any of several signals. By washing chemicals through in sequence, multiple strands of DNA could be grown with different programmed patterns. Note this is only one of several ways to build DNA with desired sequences.

■Can diamond robotics do scanning-probe vacuum chemistry to build diamond with low error rates? Even at room temperature?

Scanning probe microscopes have already done several kinds of covalent chemistry, with and without electric currents. Basic theory says that a stiff low-energy covalent surface should not reconstruct or deform easily, even if one or two reactive atoms are brought near it; those atoms can then be applied to a chosen spot on the surface and perform a predictable reaction.

It has not been difficult to find deposition reactions that, in simulation, can be used to build diamond. These reactions or similar ones will probably work in practice.

According to Drexler's analysis in *Nanosystems*[1], achieving the necessary precision for diamond synthesis at room temperature appears to require an overall stiffness between workpiece and probe of 10 N/m. This assumes that the required precision is on the order of a bond length, 1.5 Angstrom. Diamond nanoscale components can probably satisfy this requirement for room-temperature diamond mechanosynthesis.

Freitas and Merkle[2] have studied a dimer deposition reaction on the (110) diamond face. They found that for this particular tool tip and reaction, positional accuracy of 0.1 angstrom was required to distinguish between configurations. If this is the case in general, it may affect the temperature at which the synthesis can be carried out reliably. Note, however, that low temperatures are good because they improve the efficiency of computation.

■What other chemical methods will allow molecular machines to build molecular machine parts? (e.g. turning benzene rings into graphene)

---

[1] K. Eric Drexler, Nanosystems: Molecular Machinery, Manufacturing, and Computation, John Wiley & Sons, 1992.

[2] Freitas and Merkle have published their work in mainstream journals. From http://www.foresight.org/stage2/mechsynthbib.html, a useful bibliography of mechanosynthesis papers:

- Ralph C. Merkle, Robert A. Freitas Jr., "Theoretical analysis of a carbon-carbon dimer placement tool for diamond mechanosynthesis," J. Nanosci. Nanotechnol. 3(August 2003):319-324; http://www.rfreitas.com/Nano/JNNDimerTool.pdf or http://www.rfreitas.com/Nano/DimerTool.htm

- Jingping Peng, Robert A. Freitas Jr., Ralph C. Merkle, "Theoretical analysis of diamond mechanosynthesis. Part I. Stability of C2 mediated growth of nanocrystalline diamond C(110) surface," J. Comp. Theor. Nanosci. 1(March 2004).

- David J. Mann, Jingping Peng, Robert A. Freitas Jr., Ralph C. Merkle, "Theoretical analysis of diamond mechanosynthesis. Part II. C2 mediated growth of diamond C(110) surface via Si/Ge-triadamantane dimer placement tools," J. Comp. Theor. Nanosci. 1(March 2004).

This is an open-ended question. One possibility, as mentioned in the question, is using organic chemistry to create graphite-like (graphene or fullerene) shapes and components. The bigger question is: what simple, programmable, high-reliability, high-throughput, autoproductive methods are waiting to be invented?

■Will there be substantial difficulty in automating and scaling up fabrication chemistry or subsequent assembly of parts?

This depends on many factors: whether the actuation method can easily be controlled in parallel, whether the chemistry is reliable enough to proceed without error checking, whether the parts will be easy to grip and manipulate, whether the parts will stick easily when assembled correctly (and not before), and for scale-up, whether control and actuation can be implemented in suitable nanoscale technology. Architecture-level designs and calculations have been done for diamondoid mechanosynthesis systems,[3] and they appear to scale quite well to tabletop systems making integrated decimeter-scale products and fabricating their own mass in a few hours.

Conclusion: Any of several types of mechanically guided chemistry appear to be viable technologies for inexpensive, high-volume molecular manufacturing of complex, high-performance products.

2. To what extent is molecular manufacturing counterintuitive and underappreciated in a way that causes underestimation of its importance?

To the extent that the importance of molecular manufacturing is underestimated, it may not be adequately studied or prepared for. Several factors may combine to create substantial underestimates of MM's significance.

■Benefits are concentrated at the end of development — will projections from partial progress or spinoffs underestimate benefits?

The benefits of molecular manufacturing come from automation and autoproductivity. For example: suppose that parts and labor to build a 1-kg nanofactory cost $1000 per gram, and a million-dollar factory can make 100 kg of product in its lifetime. Then factory cost contributes $10

---

[3] See Drexler: Nanosystems, Phoenix: Nanofactory, Merkle (various); Freitas and Merkle, *Kinematic self-replicating machines* http://www.MolecularAssembler.com/KSRM.htm (this has a new design for a basic mechanosynthetic fabricator).

per gram of product cost. If the factory can make 90% of its own parts with 90% automation, then factory cost drops to about $110,000. But if the factory can make and assemble 100% of its parts with full automation, then factory cost (and product cost) drop to cost of raw materials: probably a few dollars per kilogram.

The first 90% saves one order of magnitude product cost. The last 10% saves another *three* orders of magnitude. And because molecular manufacturing builds everything using the same bottom-up processes, the last 10% will probably be the easiest to design—very different from conventional engineering.

■Product complexity and functionality is not limited by manufacturing system complexity — will projections from MM development difficulty overestimate product development difficulty?

A computer built with a $4 billion semiconductor plant, containing a billion transistors and millions of lines of software, can be programmed by a child to do simple tasks. The software is key: it translates meaningful, easy-to-learn commands into long sequences of basic operations. Likewise, once a product design methodology is worked out that translates useful, easy-to-learn CAD specifications into molecular manufacturing operations, anyone who can create a CAD specification can design a product.

That same computer can be programmed by an expert to do trillions of operations and produce a result more complex than its own physical structure, such as a design for a better computer. Again, information is key: memory is physically repetitive but can hold very complex patterns of data. Likewise, a programmable nanofactory can make products physically more complex than itself by running sufficiently complex blueprints.

■Molecular manufacturing may be overshadowed by superficially similar technologies — is there a risk that people will think they're studying MM when they're actually studying something else?

Popular concepts of nanotechnology include molecular manufacturing, and may even be identified with it, since that was the original meaning of the word as coined by Drexler. However, the loose constellation of fields called ‘nanotechnology’ covers everything from photonics to nanoparticles to molecular electronics. Most nanoscale technology research today is unrelated to molecular manufacturing. Current work in nanotechnology pursues nanoscale *products*, not nanoscale *productive systems* (which can also make large products). Policymakers who want to promote molecular manufacturing, but are unaware of the distinction, may feel a false sense of security from reports of successes in nanotechnology.

■Molecular manufacturing is opposed by special interests — is study of it likely to be stunted by political maneuvering?

Study of molecular manufacturing has already been stunted by politics. Mark Modzelewski, founder of the American NanoBusiness Alliance, has launched vituperative attacks against commentators who dare to suggest that molecular manufacturing is possible. Richard Smalley, advisor to the U.S. National Nanotechnology Initiative leadership, has called for chemists to oppose the "fuzzy-minded nightmare dream". The NNI website declares that "nanobots" are "science fiction" and refers to them as "creatures".[4]

This probably has multiple motivations. Some researchers seem to be afraid that refocusing the NNI toward molecular manufacturing would threaten their research funding. Others might fear that admitting the possibility of nanobots (while failing to distinguish simple industrial mechanisms from complex life-like systems) would increase public fear of destructive or runaway nanotechnology. Some opposition probably stems from simple incomprehension.

■Engineering benefits of nanoscale physics (near-frictionless interfaces; perfectly precise construction; scaling laws) are not widely known — would better knowledge increase research and development?

The problems of nanoscale engineering are famous, perhaps overly so: thermal noise, sticky surfaces, etc. But some alleged problems, like friction, go away when atomically precise machines can be built. And no one talks about the benefits, which are substantial.

Covalent molecules are perfectly precise in their formulation: an atom is either in the right place, or you have a different molecule. This means that fabrication can benefit from absolute precision: there's no need to specify or account for a manufacturing tolerance.

Sliding interfaces that are atomically precise can be almost completely frictionless. This quality, called 'superlubricity', was analyzed by Drexler in connection with nanosystems and has

[4] http://www.nano.gov/html/facts/faqs.html

recently been observed. Experience from high-friction MEMS is misleading, since MEMS surfaces are quite imprecise and rough.

Unfamiliar nanoscale effects, including thermal noise and springiness of molecules, are generally seen as problems; their engineering benefits are substantial but not generally appreciated. For example, thermal noise reduces friction and can allow jammed machines to unjam themselves. Springy molecules allow less exacting mechanical design.

Things are inherently more efficient at smaller scales. For example, a meter-scale robot arm may handle (produce) 1 kg/s with 100 W of friction. Eight half-meter arms (the same mass) could handle 2 kg/s with 200 W of friction at the same speed (twice the operating frequency). But throughput scales linearly with speed, while friction in sliding interfaces scales roughly as the square of the speed. So handling 1 kg/s should require only 50 W. If this is scaled to 100-nm arms, then 10,000 kg/s can be handled with 1000 W of friction.[5]

■The operations of programmable, automated manufacturing may be easier at the nanoscale — will projections from conventional engineering overestimate difficulty?

Macro-scale engineering uses many different parts built many different ways, usually with top-down processes that must be re-engineered for each product and involve many idiosyncratic operations. Programmable manufacturing is therefore difficult and must be specially designed for each part and process. By contrast, bottom-up manufacturing uses very few operations in programmable sequence. It should be relatively easy to generate the sequence algorithmically to produce the desired shapes and structures.

Assembling parts into products may also be easier to automate. Improved precision, material properties, and feature size will make simple assembly techniques (e.g. snap-fit) applicable to a wide variety of products.

■Nanotechnology has been the domain of scientists. Engineers have a much faster approach to development. How will this affect progress?

We have known that the nanoscale existed since atoms and molecules were discovered. But only recently has it become a realm where we can engineer, rather than merely investigate.

[5] This example is from recent talks by Drexler.

Investigation requires science, slow and careful experiment punctuated by unpredictable insight. Engineering uses known rules to achieve predictable results.

We now know enough of the nanoscale to predict, with the help of modeling software, what a particular molecule or system will do. This knowledge is imperfect, but sufficient to guide design. We also know some basic rule sets that appear sufficient to design systems for a desired purpose. A novel protein fold has been designed and tested. Many engineered shapes have been made with DNA. Although we don't know nearly all there is to know about the nanoscale, we can design shapes and interactions in a few key domains.

Scientists focus on what we don't know. Engineers focus on what we do know, and what can be done with it. Nanoscale engineering, now that we know enough to do it, will go much faster than scientists would estimate.

Conclusion: The importance of molecular manufacturing is likely to be substantially underestimated by any particular body. However, it is not hard to realize its importance, and the relevant information and theory have been available for many years. If one group comprehends the implications of the theory while others ignore it, then that group may go ahead and develop the technology while others are not even looking. This could lead to unpleasant surprises.

Capabilities of molecular manufacturing technologies:

Molecular manufacturing is the use of programmable chemistry to make programmable products, including duplicate manufacturing systems. Programmability implies automation, and duplication implies low capital cost. MM may drastically reduce the cost of both products and manufacturing capacity. In addition, precise control of chemistry should produce very strong structure and very compact functionality. High performance products imply high performance manufacturing. Quantifying these advantages is necessary to understand the impact and desirability of MM.

3. What is the performance and potential of diamondoid machine-phase chemical manufacturing and products?

Diamondoid molecular manufacturing systems were described and analyzed in some detail in *Nanosystems*[6]. They would do scanning-probe chemistry in vacuum to build diamondoid machine parts, including bearings, motors, cams, and scanning probe systems.

■Can a simple set of chemical cycles be developed to process simple feedstock molecules into renewable chemical 'tool tips' suitable for deposition fabrication?

Refer to Merkle's study on Hydrocarbon Metabolism.[7] Preliminary investigation says the answer is: probably.

■Can a simple set of deposition reactions be developed to build programmable diamondoid parts with the 'tool tips'?

Freitas & Merkle report that they have found one, and think that six to ten are necessary; see their Foresight proposal.[8] Experience based on computational chemistry investigation says the answer is: yes.

■Can diamondoid parts be combined into machines that can manipulate 'tool tips' with the required precision, as well as supplying components for other types of products?

---

[6] Ibid.

[7] http://www.zyvex.com/nanotech/hydroCarbonMetabolism.html

[8] http://www.foresight.org/stage2/project1A.html

Based on Drexler's *Nanosystems*, it appears the answer is: yes, diamondoid (3D carbon-based solid) is a great material for nanoscale machines, is stiff enough to achieve sub-angstrom precision at room temperature (with careful design), and also makes great bearings, motors, etc.

■What would be the performance of nanostructured, atomically precise diamond machines, including strength, power handling, and digital logic?

According to *Nanosystems*: 100 times as strong as steel, 10^{15} W/m^3 electromechanical power conversion (10^8 increase in power density?), 10^{16} instructions/sec/W (10^6 increase in computer power?), 10^4 sec to double manufacturing capital.

■Can nanoscale fabricators be combined into an efficient scalable manufacturing system to build large products?

Based on Phoenix's "Design of a Primitive Nanofactory",[9] it should be straightforward to build an integrated tabletop manufacturing system producing kg-scale products (not just kg's of mg-scale products) at kg/hour rates. The basic architecture should scale quite a bit larger than that without sacrificing much efficiency. This work builds on *Nanosystems* and Merkle's work, and shows that a much simpler design should come within an order of magnitude of Drexler's performance numbers (though Drexler's numbers may themselves be a substantial underestimate).

■How difficult will product design be?

Once basic 5-50 nm molecular components are designed and characterized, they can be combined to make a vast range of products without further molecular design. Software engineering methods will help, including modular design and levels of abstraction. Reliability will be an issue but should be solvable by simple redundancy. Above the molecular scale, products should not be much harder to design than familiar products of similar complexity. (Note that complexity of large products can often be reduced substantially by duplication of simple designs.) One factor that should make design easier is the ability to build cheap prototypes rapidly.

Conclusion: Diamond machine-phase manufacturing has the potential to be an extremely powerful technology.

---

[9] http://www.jetpress.org/volume13/Nanofactory.htm

4. What is the performance and potential of biological programmable manufacturing and products?

Biology has been making complex molecules and structures for billions of years, and self-replicators already exist and produce cheap valuable products. Can this be harnessed to produce engineered products?

■Can the rules of protein folding and self-assembly be accessed to design novel proteins, structures, and machines?

Progress is preliminary, but encouraging. A new protein fold has been designed and tested. Drexler[10] pointed out that protein engineering should be much easier than solving the protein folding problem for natural proteins.

■Can intracellular transport mechanisms be adapted to increase the programmability of part assembly?

Biological motors have been extracted from cells and made to run. Programmability would depend on whether some way other than diffusing chemicals could be found to power them.

■How efficiently can new genetic specifications be synthesized and transferred into cells?

Progress is being made... Study the cost per nucleotide vs. time. Also look at plasmid and artificial chromosome development.

■Can the rules of multicellular structure formation (analogous to ontology or cellular specialization) be accessed to design larger products?

Good question. MIT work on amorphous computing may be relevant.

■What would be the performance of engineered systems based on biological materials, with or without augmented biochemistry?

Strength: perhaps comparable to modern polymers. Computation: with augmented chemistry, could include molecular electronics. This depends largely on covalent bond density.

---

[10] Molecular engineering: An approach to the development of general capabilities for molecular manipulation, PNAS, 78(9), Sept. 1981

■What would be the production speed of a biology-based manufacturing system?

Unknown.

■What is the smallest size (genome and physical) of a viable cell?

Unknown.

■Can extracellular protein synthesis systems improve any of these answers?

Unknown.

Conclusion: More research will be needed to tell whether this technology can be revolutionary, but it looks promising so far.

5. What is the performance and potential of nucleic acid manufacturing and products?

Nucleic acids fold and self-assemble into predictable three-dimensional shapes. Motors and truss-like structures have already been built. Several families of nucleic acid polymer are being investigated, including DNA, RNA, and PNA (peptide nucleic acid). Modifications including polyamide (nylon-like) backbones have been demonstrated for increased strength. A robotic system based on this might go beyond self-assembly to active templating or programmable assembly. This might form the basis for a programmable manufacturing system capable of building complex products from simple parts.

■Can required nucleic acid sequences be calculated directly from the desired shape of the resultant parts?

This has already been done, with a bit of human post-processing, for the recent single-strand octahedron.[11]

■Can a mechanically actuated system be built to allow for programmable assembly of simple sequences (reducing the complexity and number of input sequences)?

Almost certainly. The required precision appears feasible.

---

[11] http://www.scripps.edu/news/press/021104.html

■What would be the speed and accuracy of such a manufacturing system?

Compare with current accuracy for DNA synthesis and binding in sensors.

■What would be the performance of machines built of nucleic acids, including strength, power handling, and digital logic?

DNA is fairly weak; PNA is stronger but has less chemistry developed to handle it; DNA with polyamide backbone has already been demonstrated. The system will also be limited by packing/conjugation strength unless a crosslinking chemistry is used. DNA-conjugation actuators are likely to be weak, but other actuators could probably be integrated.

Conclusion: More research will be needed to tell whether this technology can be revolutionary, but it looks promising so far.

6. What other chemistries and options should be studied?

This is a grab bag of questions intended to suggest possibilities that may have been overlooked.

■What other chemistries may be suitable for atomically precise programmable assembly?

Merkle has suggested small cubical molecules with boron and nitrogen. Or perhaps precise metal nanoparticles could be fused. Other possibilities no doubt will be suggested.

■What is the potential of top-down technologies using imprecise chemistry, in terms of self-manufacture and device performance? (e.g. extrusion, DPN, metal-over-buckytube, MEMS, inkjet, stereolithography, masked or hologram-switched optical surface activation)

Some of these appear to have fairly high throughput. Many are flexible in the materials they can deposit. More work will be needed to determine what kinds of devices, especially bearing surfaces, can be made with these imprecise technologies.

■What about atom holograms and atom lasers?

Unknown. Atom holograms, a way of programmably redirecting a beam of atoms into complex deposition patterns, were demonstrated in Japan several years ago,[12] and have not made a lot of news since. Atom laser is a confusingly similar name for a very different technology: a way to reduce a cloud of atoms to a single quantum state, making them extremely controllable. The technologies may be synergistic.

■Are there synergies between any of the considered technologies, making problems easier to solve or improving performance of a technology?

Almost certainly.

Conclusion: Molecular manufacturing may be easier than we realize. Many possibly helpful technologies have not yet been assessed. There's no way to know without studying multiple alternatives.

Development of molecular manufacturing technologies:

Molecular manufacturing does not exist today. This section explores the requirements of developing a molecular manufacturing technology.

7. What applicable sensing, manipulation, and fabrication tools exist?

Development efforts will be aided by the ability to directly interact with the nanoscale, to manipulate nanoscale objects and to sense nanoscale structures. In particular, a combination of sensing and manipulation in the same platform will be very helpful.

■What nanometer or angstrom-level sensing modalities exist or can be developed for off-the-shelf use? In particular, can sub-wavelength nanometer-scale optical non-proximal video imaging be developed?

Sensing at the nanoscale has been difficult, because traditional optics can't ‘see’ smaller than a few hundred nanometers. However, a variety of sub-wavelength technologies do exist. For example, two-part fluorescent systems can detect nanometer displacements. Electron microscopes can image down to angstrom levels. Scanning proximal near-field technologies can bypass the

---

[12]http://web.archive.org/web/20021214175826/http://www.academicpress.com/inscight/03252002/graphb.htm

diffraction limit. Other scanning technologies can reach atomic resolution (AFM, STM, even MFM). Most interestingly, it appears that near-field effects can be extracted and detected, allowing parallel (video-like) 3D non-proximal imaging of nanometer-scale features. AngstroVision claims to have developed a system that can detect 12x12x4 nm at 1-3 frames per second.[13] NASA has also published theoretical work leading to a sub-wavelength non-proximal imaging system using incoherent light.[14]

■What manipulation technologies exist or can be developed for off-the-shelf use?

For positioning: piezo-driven probes; optical tweezers. For gripping: antibodies; recent work on engineering RNA to grip arbitrary shape; perhaps EBD-fabricated tweezers.

■What combinations of sensing and manipulation can be integrated?

Piezo probes have been placed inside a SEM and integrated with EBD in Denmark. AngstroVision claims their system will work in a shirtsleeve environment; possibly in conjunction with optical tweezers.

■What environments can be supported by the various techniques and combinations? High temperature, room temperature, low/cryogenic temperature? High vacuum? Solvated? Micro environments (e.g. droplets)?

Detailed engineering studies needed here.

■What nano-fabrication technologies exist or can be developed for off-the-shelf use? Special attention should be given to technologies that produce rapid and low-cost results.

Direct-write lithography: laser, e-beam, dip-pen nanolithography (DPN). Gel deposition, possibly with glass precursor coating/baking for further miniaturization. 3D inkjet? Chemistry plus self-assembly: a very large field with lots of possibilities. Nanotube welding. Electron beam deposition (EBD). Et cetera.

---

[13] http://www.parc.xerox.com/cms/get_article.php?id=223

[14] http://www.nasatech.com/Briefs/Sept00/NPO20687.html

■What are compatible combinations of nano-fabrication and real-time sensing? What nano-fabrication technologies are well enough modeled for reliable CAD-to-product workflow?

EBD and nanotube welding with SEM. Chemistry with fluorescence and maybe with non-proximal near-field imaging as described above. DPN with scanning tactile probe. Unknown what technologies are compatible with CAD-to-product; to some extent this depends on the required product characteristics. But note that a major DPN manufacturer is now selling text-writing software.

■What handling technologies exist for moving samples between environments and/or locations efficiently?

Unknown.

■Which of these technologies is compatible with automation and/or high throughput?

Unknown. Probably most are compatible with automation. Chemistry plus self-assembly is generally compatible with high throughput.

Conclusion: Many relevant fabrication and sensing tools exist off-the-shelf. Single-nanometer optical open-air video imaging is a strong possibility. Chemistry and lithography (bottom-up and top-down) have already met in the middle.

8. What will be required to develop diamondoid machine-phase chemical manufacturing and products?

This explores the various steps needed to develop a complete manufacturing system based on diamondoid vacuum mechanosynthesis.

■How much computer time and human creativity would it take to invent, then simulate and verify a set of diamondoid-building (and/or graphene-building) reactions?

Robert Freitas has proposed a $5 million, five-year project to do just that; the project would also simulate the construction of nanodevices using these reactions.[15]

[15] http://www.foresight.org/stage2/project1A.html

■What will be involved in developing a non-diamondoid manipulation system that can carry out the required manipulations to build the first system?

Unknown, but it should be noted that we can now lithographically fabricate features that are smaller than the molecules we can engineer. In other words, we can build pretty much any shape at any size scale.

■How reliably can the operation of diamondoid machine parts be simulated? What would be the cost and development time of a CAD/simulation system capable of extracting mechanical characterization from molecular dynamics simulation of such parts?

Unknown, but this is a much easier problem than characterizing proteins: the parts involved are much stiffer, and energetic computations can afford to be much less accurate. Hydrocarbon MM packages have been around for years (e.g. Brenner) and are now appearing in open source software (e.g. NanoHive).

■How many parts and surfaces would be needed to constitute a complete set of low-level structural and functional components? How much human effort would be required to develop them?

Unknown. Low-level components include rotational, helical, and flat bearings; conductive and insulating components; molecular interfaces between different surfaces and crystal orientations. Note that Freitas expects to design at least some working components as part of his $5 million proposal.

■What would be the cost and development time of a CAD/simulation/tracking system that could support the design of machines and systems from low-level components?

Unknown. Probably comparable to high-end software design tools, or semiconductor design tools circa 1990. It wouldn't have to handle a lot of different parts or physics, at least in early versions where performance can be sacrificed to reduce undesired interactions between parts.

■What would be the cost of developing a design for an integrated, hierarchical manufacturing system to build large products?

An architecture for such a design has been worked out.[16] The molecular fabrication in that design is based on a simple robotic-chemistry design by Ralph Merkle. Many fabricators make parts in parallel, and the parts are then combined via convergent assembly. Merkle's design requires perhaps 100 moving parts and half a billion atoms (most of which don't have to be individually specified). Convergent assembly appears to require only simple robotics at several scales. Assembly and fabrication appear to require only simple control software. Much of the engineering, even at nanometer scales, will be more or less familiar to mechanical engineers. Overall engineering difficulty might be comparable to an aerospace project.

■How many of these steps could be accomplished concurrently in a crash program?

All of these steps could be started concurrently, with successive refinement. This may not happen due to caution on the part of the funders. However, a funding organization that was willing to fund a crash program could probably do all these steps in parallel.

■How precisely can costs and schedules be estimated?

Due to lack of study, very little information is available. For the sub-projects that we can estimate, the cost is consistently under $1 billion, and several appear to cost just a few million. Also, all of them (with the exception of software engineering, which should not be a major fraction of the total cost) appear to be getting easier rapidly. We can't rule out the possibility that the whole thing might cost less than $1 billion; in fact, that appears likely to us, though we don't say it loudly because it sounds too implausible. A project starting five or ten years from now very likely would find the cost greatly reduced. (However, other studies indicate that this is not a sufficient reason to delay; it's simply evidence that if we do delay, a rapidly increasing set of organizations will be able to do it.)

About schedules, again, very little information is available. The argument parallels the cost discussion. The project can be divided cleanly into sub-projects. In the areas where we can make estimates for the sub-projects, the estimates are surprisingly short. We don't see any sub-project that needs to take more than five years. Doing all sub-projects in parallel would require excellent management, visionary funding, and good communication to ensure smooth integration. But this appears feasible, and implies that the whole thing might be done in five years with sufficient effort and skill. (But government bureaucracy is not well suited to do this.)

---

[16] http://www.jetpress.org/volume13/Nanofactory.htm

Conclusion: At a guess, the difficulty and schedule of developing a tabletop kg-scale manufacturing system producing kg-scale nano-featured products may be comparable to the Apollo Program. Or it may be quite a bit easier; we can't know without more engineering investigation. At this point, we can't rule out the possibility that it could be done in five years for less than $1 billion. Note also that work on this may have already started somewhere, and may be quite close to completion.

9. What will be required to develop biological programmable manufacturing and products?

This study would explore the various steps involved in harnessing biology to produce engineered products.

[Answers in italics in this topic are provided by Robert Bradbury.]

■How much time and effort would be required to develop the ability to design predictable protein folding, possibly by introducing novel amino acids?

Unknown, but a novel protein fold has been successfully designed and tested. Increasing computer power will make this rapidly easier.

It would not be difficult to integrate novel amino acids using a standard protein synthesis robot. It is more difficult to integrate to integrate them into bacteria (as my "Protein Based Assembler" paper discusses), but it has been done.

■How difficult would it be to automate all steps of new-protein synthesis? How long would a fully automated system need to produce and characterize a new protein?

New protein synthesis is already automated (its a volume/cost issue that can be a hang-up -- which is why bacteria are used to produce things like insulin, antibiotics, etc.). The NSF is pushing rapidly on the automation of the characterization problem (everything from X-ray crystallography to computers figuring out the structure). I've read that they are trying to push it to 30,000 structures per year. Though I'm not sure if I can believe that number -- if you look at the growth of the contents of PDB it may be a reality in the near future.

There are *structures that are difficult to characterize -- these are usually proteins that normally reside in cell membranes of one form or another. So it's a limited subset -- perhaps 20-*

30% of all proteins. Some novel techniques have been reported for dealing with this but this is ultimately just going to require a lot of work and clever ideas.

■What software support must be developed to allow design and testing of novel protein-based machines?

Tough question -- we already have the software to design proteins (and the machines to manufacture at least the smaller ones). Testing isn't really a problem. The problem is the creation of a 'novel' machine design.

■How much time and/or research will be required before we know how cell signaling/differentiation/gene expression works?

We know how gene expression works reasonably well (something like 3 classes of transcription factors, the structures of which tend to be very standardized, etc.). We also know a lot *about signaling and differentiation. We've got hundreds of extracellular molecules and receptors pinned down at this point. The problem is the molecules involved within the cell from the membrane to the nucleus. These are very complex. There is a company in Germany that has worked out much of this in yeast and the #1 priority on the NIH Nanomedicine goal list is to extend this to determine all of the protein complexes in humans.*

■How can cell toxicity or metabolic interference from novel chemicals be predicted and avoided?

This is relevant because one method of protein synthesis involves using gene-spliced cells to synthesize the protein. However, there are ways of manufacturing proteins that do not require cells.

The simple answer is knowledge of the structures of most if not all of the enzymes, receptors, etc. in the body, knowledge of the structure of the novel chemicals and a heck of a lot of computer power to see when/how the structures can interfere. A more complex answer would involve actual toxicity tests at a MEMS scale level to determine when chemicals interfere with the functioning of a protein. (This isn't too different from the work that has been done to synthesize large chemical/drug libraries -- but requires that one understand the metabolic pathways involved and devise individual tests to see when there is interference.)

Conclusion: This deserves further investigation.

10. What will be required to develop nucleic acid manufacturing and products?

This study would explore the development of nucleic acid manufacturing.

■What is required (research and software) to automate the design, production, and characterization of nucleic acid molecules directly from specification of shape and properties?

We are close to this today; see the single-strand octahedron announcement.[17]

■What actuation techniques (chemical, electrical, other method?) are available? How fast, reliably, forcefully can they operate?

DNA-conjugation actuation is fairly slow but very programmable. Actuation by redox sliding rings (catenane, rotaxane) is faster and allows either chemical or electrical actuation. This can provide significant (~nN?) force; see the "elevator".[18] Several bio-based motors are being investigated. These are switched by simple chemicals and may be hard to select or control.

■What chemistry (steric mechanism) could be used to allow programmable fabrication? How small could the selectable units be? (Atoms? Nucleic acid monomers? Short chains?) Can the selected fabrication chemistry produce the required mechanism?

Good questions...

■How much additional design would be required to scale up/duplicate a fabrication system for large-scale production?

The system might be attached to beads for large surface area. This might be more, or less, difficult than scaling up other surface-catalyzed chemical synthesis processes.

■How much additional design would be required for a scaled-up system to produce monolithic heterogeneous products?

This might require nanoscale computation to control local actuators, and better attachment, localization, and control of the individual production systems. Biomimetic (e.g. amorphous

---

[17] http://www.scripps.edu/news/press/021104.html

[18] http://nanotechweb.org/articles/news/3/3/14/1

computing) and mechanistic approaches should both be investigated; very little work has been done to date.

Conclusion: This deserves further investigation.

11. How rapidly will the cost of development decrease?

How long will it be before development of molecular manufacturing becomes attractive to large corporations? How long before it can be done in a garage or a developing nation? How long before it falls off the radar of any reasonable detection effort? It is crucial that we learn the answers to these questions.

- ■How rapidly is the cost of computer time falling? How much additional advantage could be gained by innovative computation (distributed computing, special-purpose logic, etc)?

In general, computer costs fall according to Moore's Law. Additionally, new ways of using existing resources such as distributed computing (SETI@Home) and massive clusters of cheap computers (Google) may reduce the cost for big projects. Special-purpose hardware may improve price/performance by multiple orders of magnitude.

- ■What software is being developed (commercial as well as Open Source) for physics simulation, chemical simulation, and CAD?

Lots.

- ■How quickly are sub-nanometer or even sub-angstrom sensing and manipulation technologies becoming cheap, simple, readily available, and well understood?

A group in Russia has developed an SPM with angstrom resolution that sells for US$30,000.[19] SPMs have been available for over a decade and are not hard to use. New tools are generally computer-controlled, making it possible to design intuitive interfaces. The Russian system deserves special attention because it combines several capabilities that appear targeted at atomically precise mechanosynthesis: gas flow-through (for deposition); STM (for imaging and surface modification); and equipment for rapid sample changing.

---

[19] http://www.nanotech.ru/cn/e/cs/tech7.pdf

http://www.nanotech.ru/cn/e/tech6.php

■How rapidly is the cost of top-down nanofabrication falling, the resolution shrinking, and the lag time decreasing?

We need more numbers on cost. Resolution is down to ~50 nm or better for optical litho, ~20 nm for e-beam and two-photon polymerization, ~15 nm for DPN. Some litho technologies have lag time of hours.

■Can the increasing size, functionality, and programmability of molecules be plotted or projected? (E.g. dendrimers, precise polymers, nucleic acids)

Good question. Metrics can be developed for assessing recent trends.

■How rapidly are these techniques and capabilities filtering down to postdocs and other readily available workers?

Our impression is that postdocs can easily learn these technologies.

■How rapidly is the cost of mechanical design, including CAD software, decreasing?

To some extent, this depends on computer power. To some extent, on writing new software, which will probably remain the same—but probably won't be a significant expense. To some extent, on creativity, which is very hard to quantify. But it should be noted how much has been accomplished by just a few unfunded researchers over the past decade.

Conclusion: Computer and lab resources are becoming rapidly less expensive. The speed will surprise anyone not familiar with the computer industry. Although it's hard to quantify, our current estimate (based also on tracking previous difficulty estimates) is that the cost will decrease exponentially more or less like the cost of computers: falling by half every two years or so.

12. How could an effective development program be structured?

We need to understand the factors that will affect the success of a targeted or crash program.

■How can the scientists and engineers be engaged in the project?

A lesson from the computer industry may be relevant here: Hire people who are too young to know what's impossible. Once feasibility is established (or assumed, for a crash project), skeptical scientists should not be put in charge of research. In fact, the people in charge should probably be engineers, not scientists.

■How could it be funded?

An incremental project, funded by spinoff developments and near-term goals, would take too long. A crash project will probably be funded by a military budget or by politics of national pride. Since the biggest results will come at the end, funding will have to be based on long-term thinking. This may be hard to do in either U.S. business or political system, but might be more achievable in other systems including U.S. military and foreign top-down planning systems.

■How could bureaucratic friction be minimized?

As with funding, a minimum of interference from outside once the project is started will be a big help. Organizational design and culture will be important to minimize internal politics. Trust in team leaders will be crucial to minimize the need for detailed oversight.

■How could innovation be maximized?

Don't let the most cautious/skeptical people control the funding. Make sure that the goal is to weed out approaches that can't work rather than to fund only projects that are sure to work.

■How can the shortest path be rapidly invented and identified?

A contest would be a good way to generate lots of suggestions. If a short path is not obvious, then investigate in parallel with a goal of rapidly establishing feasibility of each path.

■How should the overall project be structured?

This depends on how much of a hurry you're in, and how early a development pathway can be identified. If you're in a big hurry, start work in parallel on CAD software, mechanosynthesis, nanomachine design, and nanofactory design.

■Under what (corporate or governmental) cultures could an effective program take place?

Given the likely intense competition, an effective program would have to be fast. Silicon Valley is probably a good place to look for inspiration.

■How can development time be minimized?

Nanoscale lab techniques are developing rapidly; so is ability to test mechanosynthesis in simulation. And nanomachine design may turn out to be not all that complicated—as long as you have good software. Software is likely to take the longest to develop, since it involves an industrial-strength CAD/simulation system covering multiple length scales, several different kinds of simulation packages, and lots of physics bookkeeping. But starting software even before the preliminary science results come back would be hard to justify in terms of traditional product planning.

■What cost and time overruns should be expected?

These can't really be estimated until the project is started.

Conclusion: An effective development program would probably include several features not easily implemented in Western corporate or government-funded programs, with the possible exception of a few crash military projects. A central-planning approach to obtaining plentiful funding (probably multiple billions of US$) combined with a semi-autonomous approach to design work is probably the fastest approach.

Product performance

This section suggests metrics for manufacturing and product capability. The following studies should be run for each plausible molecular manufacturing technology.

These questions will be answered for diamondoid systems based on the Phoenix nanofactory design.[20]

13. What is the probable capability of the manufacturing system?

How much product per hour? How many features per hour? How much input, and what kind? How much waste?

■Does the system require human supervision or intervention while operating?

No. The (calculated) extremely high reliability of mechanosynthesis should allow completely autonomous operation; see Drexler, *Nanosystems*.[21] Convergent assembly can use very simple robotics[22]. With a reasonably low error rate in each fabrication unit permitting a reasonably low degree of unit-level redundancy, the nanofactory can take units offline permanently at any failure, so would not need repair.

■How many features per second (complexity) will the system produce?

Each fabrication unit might produce 1,000 to 10,000 features per second : 10 to 100 atoms per feature, 100,000 atoms placed per unit per second. A less primitive design might place a million or more atoms per second. Each unit would be independently addressable with any of several thousand or million program streams. Basically, the product complexity is limited by the information that can be downloaded into the factory over a fast network in the few-hour fabrication time. This could easily amount to several terabytes—far more complexity than would be needed for most products. (For comparison, human DNA is several gigabytes.)

■What error rate will be built into the product components?

---

[20] http://www.jetpress.org/volume13/Nanofactory.htm

[21] K. Eric Drexler, Nanosystems: Molecular Machinery, Manufacturing, and Computation, John Wiley & Sons, 1992.

[22] http://www.jetpress.org/volume13/Nanofactory.htm

With primitive mechanochemical hardware, fewer than 1 in 10^8 atoms should be out of place. Better designs should be able to achieve 1 in 10^{15}. At this point, damage from environmental radiation becomes a bigger concern.

■How many grams per hour will the system produce?

A small-scale manufacturing system with no redundancy and external computer control might fabricate its mass in several hours. Scaled to tabletop size, it could take the better part of a day, but might be much quicker with more advanced designs. A single box massing a few kg could produce ~1 kg/hr in the reference design.

■What raw materials will the system require?

Some small carbon-rich molecule, not yet specified.

■What waste will it produce?

Not yet specified. Ideally it would produce harmless or useful molecules such as water and hydrocarbons. The reference design also uses ~250 kWh/kg energy.

Conclusion: The reference design would be easy and cheap to use, producing its mass in probably less than a day. Its products could be quite complex—limited by design capabilities rather than limitations inherent in the nanofactory architecture.

14. How capable will the products be?

The more they can do, the more widely they can be used.

■What materials will the products be built of?

3D carbon lattice: basically, diamond.

■Does the product functionality include: Digital logic? Analog signal processing? Energy storage, transmission, and transformation? Linear and rotational actuators? Structure, at multiple scales? Kinematics, at multiple scales? Displays? Sensors? Biocompatibility?

Digital: yes (see *Nanosystems*). Analog: probably (physical systems—cams, springs, etc). Energy storage: atomically precise springs can store energy at near-chemical density. Energy transmission: mechanical looks quite efficient. Energy transformation: yes, electrical <-> mechanical with very high efficiency and power density. Actuators: yes, both rotational and solenoid-like. Structure: from nanometer feature size (1 nm^3 = ~176 diamond atoms) (and even individual atoms in certain components, e.g. gear teeth) to macroscale (with convergent assembly). Kinematics: yes, including near-frictionless rotational and linear bearings. Displays: yes, mechanical semaphores, maybe semiconductors also. Sensors: yes, lots.[23] Biocompatibility: looks good so far.[24]

■What will be the efficiency of the various product functionalities?

Excellent; see *Nanosystems*. Nanoscale bearings: 10^{-16} W. Logic operations: less than kT per (reversible) gate at 1 GHz.

■How much post-processing does the output need?

Probably none. Carbon is a very flexible element and the product can include a variety of structure and appearance. See "Nanofactory" paper section 7.

■Can the system produce complete products, or only components?

Complete products.

■What components of itself can the system produce ('autoproduction')?

All components.

■What new capabilities can the products implement? (Machine-phase chemistry? Plasmonic logic?)

Machine-phase chemistry: yes. Molecular electronics: Buckytube transistors have been demonstrated. Optics and plasmonics: seems likely. Building biomolecules (medicine, food): not without additional R&D.

■What subset of desirable products can known design methodologies access?

---

[23] http://www.nanomedicine.com/NMI/4.1.htm

[24] Nanomedicine Vol. IIA; http://www.nanomedicine.com/NMIIA.htm

The nanofactory is well-suited for levels of abstraction (similar to software design). A single 'nanoblock' can contain hundreds or thousands of parts, enough to implement general-purpose behavior (motor, computer, etc). The combination of these into systems, 'smart materials', and products appears to encompass most conceivable functionality at all scales above 100 nm. Smaller functions such as molecular manipulation would have to be individually designed, though this may be straightforward for many tasks.

Conclusion: The output of the nanofactory would be fully finished and highly advanced products.

15. What will the products cost?

How many dollars per feature? Per kilogram? Note: If the system can duplicate itself completely, the cost may drop by orders of magnitude.

■How much will environmental maintenance cost? Labor? Raw materials? Energy? Waste disposal?

The nanofactory is designed to operate in a shirtsleeve environment, with access to less than a megawatt of energy and comparable cooling capacity. Labor is negligible. Raw materials are likely to be cheap chemicals, though purification may add somewhat to the cost. (Some filtration/molecular sorting is inherent in the chemical uptake mechanism.) Energy (in a very primitive, inefficient design, the Phoenix nanofactory) is perhaps $20/kg at today's rates (note that one early product is very cheap solar cells). The waste should be highly pure, small organic molecules, at the worst requiring incineration.

■How much will post-processing cost?

Zero.

■How much will product design cost?

This depends largely on the functionality of the product. As a first estimate, the cost of most products will be dominated by the cost of software engineering to implement the product's functions.

■How much will the non-autoproduced components of the system cost (amortized)?

All components can be autoproduced.

■How much will the autoproduced components of the system cost (amortized)?

Nanofactories will probably be limited by policy rather than utility, so the degree of use can't be estimated. But they should be good for at least several trillion US$ worth of product per year, and the development cost probably won't go above $20 billion (and could be much less), so development cost should contribute pennies on the dollar of value.

■What will be the total product cost, per feature and per kilogram?

A primitive design may cost $10-100 per kg, based on costs for energy (as estimated in Phoenix nanofactory paper) and highly pure chemicals. However, the Phoenix design is deliberately crude: a lower bound, not a best-guess estimate. With the use of more efficient mill-type mechanosynthesis, and the use of nano-constructed filters/purifiers, cost may drop to pennies per kg.

Per feature: Since fabrication is automated and bottom-up, details don't cost any extra. One kg of product can include 10^20 features; cost per feature is negligible. Note that the superior material properties of diamond should allow products to be orders of magnitude lighter than metal, plastic, or even carbon-fiber versions; most large human-scale products will be inflatable and will require tiny fractions of a gram per cubic centimeter to maintain their shape.

Conclusion: Product cost will be highly competitive with current high-tech products: not just semiconductors, but entire telephones, computer monitors, and aerospace hardware. Present calculations indicate it will even be competitive with cheap materials in structural applications ($/strength though perhaps not $/mass).

16. How rapidly could products be designed?

What skills and time are required to design a new product?

■To what extent can components be re-used between products?

As noted in *Nanosystems* and explored in “Nanofactory,” a convergent-assembly system combining relatively large (e.g. 200-nm) functional blocks should allow a few basic types of blocks to be built into many different products. Most product designers will not have to worry about chemistry or special nanoscale physics.

■To what extent can low-level design be automated?

Levels of abstraction should allow design on the level of volume-filling specification of nanoblocks. All lower levels can be computed, right down to the mechanosynthesis.

■How quickly and cheaply can prototypes be built?

As quickly and cheaply as any finished product. The manufacturing steps can be computed from the CAD specification of the product. There's no distinction between prototype production and mass production. This also implies immediate rollout/deployment once a product design is finished—no retooling, retraining, or design-for-manufacture.

■How directly applicable are current engineering methods?

Once a set of designs is developed to emulate familiar macro-scale structural and functional components, crude products could be developed directly with current engineering methods (with some advantages such as effectively infinite tolerance and ‘smart’ materials). More sophisticated products requiring micro- or nano-scale design may require new methods, though even here the designer's job will be made easier by careful choice of lower-level components.

■What new engineering methods (e.g. fault tolerance, emergent architecture) need to be invented to use this technology?

Fault tolerance will be a requirement. However, the extreme compactness and efficiency of actuation and computation will allow massive overdesign and redundancy. For example, a single computer may fail, but the incremental cost of three—or even 100—parallel voting computers will be negligible in most applications.

Emergent architecture and complicated software architectures will not be necessary for products comparable to today's in functionality.

Mass-saving structures will be desirable, especially in aerospace applications. Fractal trusses and inflatable compression members are two simple possibilities.

Conclusion: Design of products comparable to today's cutting edge may be even easier than today's design methods.

Policies and policymaking

This section assumes the existence of a general-purpose molecular manufacturing system. It suggests problems and opportunities raised by molecular manufacturing, and hints at the difficulties of making policy to deal with them.

The answers in this section, as in the previous section, assume a diamondoid nanofactory technology.

17. Which of today's products will the system make more accessible or cheaper?

For each suggested product, determine if the cost, compactness, or functionality could be enhanced by an order of magnitude or more, compared to present alternatives.

■Computers (logic)?

More efficient by *six* orders of magnitude. Smaller by perhaps four (vs. transistor) or seven (vs. packaged chip) orders of magnitude.

■Physical structure?

Maybe two orders lighter for tension, more for compression. Due to more efficient use of material, the cost of finished products may be substantially less than today's raw materials cost for a comparable product.

■Actuators?

Eight orders of magnitude smaller vs. today's electric motors.

■Avionics?

Perhaps three or four orders of magnitude lighter.

■Medical devices?

Molecular sensors may be sub-micron; actuators likewise; whole new classes of device will become possible. These new classes will show improvements of 10-1000 fold over natural biological systems (a technically defensible claim, based on Freitas's device design papers, *Nanomedicine*, etc.)

■Sensors?

Many sensors will be many orders of magnitude smaller and cheaper. More precise for nearly all sensors, due to more precise manufacturing and accessibility of higher-tech detection and amplification.

■Integrated systems (e.g. robotics)?

Similar to avionics. Orders of magnitude more integrated computer power will allow greater functionality.

■Compact systems (e.g. surveillance, medical)?

Yes.

■Energy systems (e.g. solar collection, storage, transport/transmission)?

Several kinds of solar collector should be buildable with a few grams per square meter/kilowatt. Several kinds of efficient energy storage are possible.

■Large systems (e.g. infrastructure, civil engineering)?

Cheap, fast manufacturing of strong materials should allow large projects to be undertaken. Fast design of special-purpose robotics should reduce labor costs of installation, including for projects that must be fabricated in pieces.

Conclusion: Diamondoid nanofactory molecular manufacturing will be revolutionary and highly disruptive in many areas of high-tech as well as low-tech manufacturing, including aerospace, energy, and medical technologies.

18. What new products will the system make accessible?

For each suggested field, determine if a molecular manufacturing system would allow significant advances compared with what can be built by alternative systems.

■Aerospace?

Yes. Figure an airplane (spacecraft, missile) might weigh 1% of its current dry weight, with essentially unlimited onboard computer power. Also, smaller actuators will make shape-changing and active skin feasible and even cheap. Continuously inverting skin might allow more efficient and higher-speed designs.

■Computational systems: data mining, strong artificial intelligence?

For comparison, the NEC Earth Simulator could be built in a cubic millimeter and draw 2 watts.

■Medical, including human enhancement?

Yes. A basic computer/sensor package could be smaller than a neuron. This could easily allow direct high-bandwidth brain-to-computer communication; already, primitive brain-implanted electrode arrays have detected a rat's position with respect to its cage, a monkey's intention to move its arm, and the image from a cat's optical nerve.

■Weapons (a very broad category)?

Yes. For example, micro UAVs with sufficient functionality to be loaded with chemical poisons. Electrical power density high enough to enable new classes of projectile weapons. Cheap deployment of massive systems or networks. Expanding-microtruss fuel-air explosives. Much easier access to space.

Conclusion: Diamondoid nanofactory molecular manufacturing would allow fundamentally new products in several important and problematic areas. It is impossible to predict or make a comprehensive list of all products that could be created.

19. What impact will the system have on production and distribution?

Determine how diamondoid nanofactory molecular manufacturing will change the way products are made and delivered.

■How close can the factory be placed to place and time of product use?

The factory should be able to be placed almost anywhere and might even be suitable for use as a home appliance. Products could be built in a few hours. High functional density implies that most of a product will be inert, so basic functional nanoblocks could be pre-built and simply rearranged into the final product; this would take only a few seconds.

■How easily can new products be designed?

With CAD-to-prototype costing very little in time or money, new product development should be comparable to digital graphic arts combined with software engineering. In other words, simple products could be designed in a day. Without the need to retool factories or stock a supply chain, market testing of new products could involve much lower commercial risk.

■For what products will this out-compete traditional systems by an order of magnitude or more?

With a suitable palette of appearance and functional units, almost any manufactured product could be built with this. Manufacturing cost should be significantly lower. Transportation and storage costs should be near zero. Design costs may be higher at first but will drop rapidly.

Conclusion: This will give serious competition to, and probably displace, a large fraction of extraction, manufacturing, transport, and storage.

20. What effect will molecular manufacturing have on military and government capability and planning, considering the implications of arms races and unbalanced development?

It has been predicted that a sufficiently advanced and general-purpose molecular manufacturing (MM) technology could have a significant destabilizing effect. This must be explored.

■How quickly can new weapons be invented, designed and deployed?

Very quickly. (See the previous few studies.)

■What new theatres or contexts for conflict will be created? (Outer space, cyberspace, underground, other?)

It will become quite important to be able to detect very small devices—perhaps even sub-microscopic devices. Outer space will become much easier to reach. Millionfold increases in computer power will create new opportunities. Extremely large-scale sensor networks, backed by large-scale computers, may make some environments (such as the ocean) less opaque. Living organisms (especially humans) are high-value and perhaps high-resource targets, and may require advanced engineering to monitor and protect without excessive disruption. Data-mining from massive sensor arrays and human transaction monitoring may be crucial; this will probably be limited more by software than by hardware. The sensor networks themselves, and disrupting or hiding from them, may be a focus of conflict, but one that is likely to be won by the sensors (see David Brin, *The Transparent Society*[25]).

■To what extent will portable manufacturing allow forces to be autonomous of supply?

Manufacturing of just about anything from clothing to missiles should be feasible with only raw materials. Advances in thermal depolymerization technology may allow conversion of local plant matter into feedstock with a relatively small (man-portable) chemical plant.

■To what extent will advanced technology allow forces to be remotely or autonomously controlled?

Any algorithm that can be run on a supercomputer today will be able to run onboard even a bullet or insect-format robot. This implies rather good image recognition. Also, the ability to field as many UAV or smart dust relays as desired will allow very high-bandwidth networking. Improved robotics, displays, and sensory or even neural interfaces can greatly enhance telepresence.

■What impacts will human augmentation (including direct brain interface) have?

Unknown at this time, but probably includes significantly improved reaction time, situational awareness, telepresence, teleoperation of robots, fully immersive VR, and enhanced memory/cognition.

■What impacts will advanced data gathering and data processing have?

A full-coverage sensor network with full storage seems plausible. This would give the ability to see and hear anything from any angle at any time in the present or past (after the network was installed, of course). Image processing should allow tracking of people through time. Data mining based on image processing should allow connections to be found and highlighted (for example, full speech-to-text conversion of all conversations, followed by text searching to determine where the other end of a phone call went).

---

[25] Brin, David, The Transparent Society: Will Technology Force Us to Choose Between Privacy and Freedom?, Perseus Publishing, 1999

This could greatly surpass DARPA's TIA, and enable DARPA's LifeLog: “an electronic diary to help the individual more accurately recall and use his or her past experiences to be more effective in current or future tasks.”[26]

■To what extent will rapidly advancing technology reduce the enemy's predictability?

If a full sensor network can be installed, the enemy may be come extremely predictable. However, in the absence of direct sensing, the speed with which new products and new types of weapons can be conceptualized, developed, and deployed argues that it will be very hard to know what the enemy's capability is or will be.

■How quickly and effectively can new doctrine be invented or adapted to new capabilities on either side?

This is an institutional question. Note that a failure of human institutions will tempt the development of automated or adaptive threat detection and response, comparable to automated computer virus characterization. Note further that such automated response systems could be extremely dangerous.

■Will offense or defense be fundamentally stronger?

Since this question must be answered for each possible class of weapon, and since MM makes many new classes of weapon possible, it appears that offense will probably win. However, this analysis is shallow; and because of the crucial importance of this question, it should be studied carefully.

■How well can military targets be protected?

Military targets can be dispersed, miniaturized, hardened with advanced materials, and rebuilt quickly. The main vulnerability will be the people, which again argues for automation.

■How well can civilian targets be protected?

Billions of toxin-carrying insectoid nanobots can fit in a small packing crate. Orbital or UAV-based weapons can be deployed on a large scale. It looks like civilians and civilian property may not be defensible without major lifestyle changes. It's possible that a comprehensive shield

---

[26] http://www.darpa.mil/ipto/Programs/lifelog/

could protect against some forms of attack, possibly including nano-scale robots, but long-range high-energy weapons may require impractical amounts of shielding.

The alternative is to prevent the deployment of such weapons in the first place, but this would be quite difficult to achieve by any means. A control-freak approach would be hugely oppressive (for the protected civilians as well as non-citizens) and may not be sustainable, and an effective policy-based approach will be difficult to design.

■Is an arms race likely to be unstable?

Yes. The nuclear arms race was stable for several reasons. In virtually every way, the nano-arms race will be the opposite.

Nuclear weapons are hard to design, hard to build, require easily monitored testing, do indiscriminate and lasting damage, do not rapidly become obsolete, have almost no peaceful use, and are universally abhorred. Nano capability is easy to build (given a nanofactory), will allow easily concealable testing, will be relatively easy to control and deactivate, would become obsolete very rapidly, almost every design is dual-use, and peaceful and non-lethal (police) use will be common. Nukes are easier to stockpile than to use; nano weapons are the opposite.

Also, as Mark Gubrud pointed out, a deployed rapid-response net would be unstable.[27] (A hair-trigger complex system eventually *will* suffer a false alarm.) One reader argued that immune systems are not generally unstable, and humans should be able to do even better. We disagree on three counts. First, humans aren't close to understanding the immune system yet, and we may have to design military systems before we do understand it. Second, what's needed is not very comparable to a biological immune system, so we'll be doing a lot of new engineering that'll be hard either to test or to analyze. Third, the instability that Gubrud analyzed is not from one defensive system reacting to disorganized and localized threats—it's from two defensive systems reacting to each other. The closest analogy from immunology would be graft-vs-host disease, which is a great example of instability.

■How hard will it be to recover from a nanotech gap?

---

27 http://www.foresight.org/Conferences/MNT05/Papers/Gubrud/

At the point where a nanofactory or equivalent system is developed, even a few months difference could be unrecoverable. The more advanced side would have access to vastly better computers, and the technology would advance as rapidly as their creativity allowed. There is no obvious plateau in capability that would allow a laggard to catch up. Also, the advanced side would be in a much better position to thwart development in its opponents, with or without all-out war.

■Could a non-nano power defend itself against a nano power?

No. And even a nuclear power might not be able to deter a nano power: aerospace superiority (with rapid prototyping and cheap manufacturing) could make it much easier to build an effective missile shield.

■How could governments use molecular manufacturing in their own countries?

This deserves a whole study of its own. Abusive and oppressive governments could become far worse. Any country could modernize (and militarize) very fast, depending on how much expertise it can buy or train locally. MM could enhance national character, for example: Americans could become more independent / off-grid (which could reduce vulnerability to terrorism); others could become more socially linked through high-bandwidth connection and data-sharing; there'll be plenty of opportunity for both laziness and productivity.

Conclusion: Military practice and planning will have to change a lot. Unstable arms race looks like a definite possibility. Substantial innovation will be required to even begin to protect civilians. Development of molecular manufacturing may have a crucial impact on national strength.

21. What effect will this have on macro- and microeconomics?

It has been predicted that a sufficiently advanced and general-purpose molecular manufacturing (MM) technology could have a significant transformative and potentially disruptive effect. This must be explored.

■How quickly can new products be invented, designed and distributed?

As described in other studies, this can be extremely quick due to fast prototyping, point-of-use manufacture, and low risk.

■How will distributed manufacturing affect the supply chain?

It will eliminate the supply chain for superseded products and their components.

■How will automated manufacturing affect jobs?

It will eliminate manufacturing jobs for superseded products, as well as related transportation, storage, and extraction jobs. It may create design and installation jobs (though a lot of installation can be done robotically). Compare manufacturing jobs with the percentage of population involved in agriculture from 1900 to today: 37.5% to 0.5%, almost a two order of magnitude decrease.

■How will increased material self-sufficiency affect international and local trade?

Trade in raw materials and finished products will be reduced. Depending on policy, trade in intellectual property may be either reduced or increased.

■How will simpler material requirements affect extraction?

Extraction will probably be unnecessary to support diamondoid manufacturing, though limited quantities of fossil fuels may be useful as a carbon source.

■Will energy production, storage, and/or distribution be impacted?

The ability to collect and store solar energy cheaply will greatly reduce the need for fossil fuels and the power grid. Also, products (including houses) can become a lot more efficient, further reducing energy demand.

■How much incentive will there be to use molecular manufacturing?

Its products will be multiple orders of magnitude better on several counts. Also, they'll probably be at least one order of magnitude cheaper to produce, and completely bypass substantial parts of the current production and distribution chain.

It's said that a tenfold improvement (one order of magnitude) is sufficient for a new product or method to displace existing ones. Nanofactory-built products greatly exceed this criterion, so new companies could out-compete existing ones that are not quick to adopt it.

Conclusion: This is likely to have a large and rapid effect on economics of manufactured products. Existing businesses that don't adopt it will be out-competed by new businesses.

22. How can proliferation and use of nanofactories and their products be limited?

This study will explore the challenge of preventing black markets, independent development, etc.

■How easy will it be to detect a development program?

Probably quite difficult. Development does not require exotic materials or massive industrial activity. It may require mainly off-the-shelf technology. Researchers will be from diverse and common fields like software engineering and computational chemistry, not concentrated in one exotic field. Depending on the bootstrapping 'recipe', the design effort might be dispersed (networked/teleconferenced), and the entire physical operation might be carried out in one moderate-sized laboratory. And most of the research would not require world-class talent, though a successful program today might well require world-class leadership.

■How much easier will it be to develop a second nanofactory, compared with developing the first one?

Reverse engineering will give hints as to which path to take. The definite knowledge that it can be done at all will reduce institutional friction. General technology advances will give a second program more to work with. Any leaks of know-how or software will further reduce the difficulty. It seems likely that the second nanofactory will be an order of magnitude less costly.

■How can nanoscale products be detected?

Unknown. Nanoporous filters can trap them. Non-proximal sub-wavelength optics, if they work as claimed, may be able to scan for them at a distance—but there are lots of natural nanoparticles, so recognition is also a problem. MRI may be able to detect at a distance, though resolution is a problem and there may be a theoretical limit.

■How easy will it be to smuggle nanofactories?

A fully functional nanofactory, able (given a supply of feedstock, energy, and blueprint software) to make one twice as big (and so on) and thus recreate a full manufacturing capacity, may be just a few microns on a side--small enough to hide inside a human cell. Or any convenient size in between. We don't know of any way to detect something like that without total intrusion of the volume being searched, which probably implies destruction.

■How easy will it be to detect proliferation-related activity?

Quite difficult. Especially once the 'recipe' is known, it'll be very hard to spot a project—R&D for a nanofactory project may require only a single small lab and a few computers. (For comparison, consider Zyvex.)

■How effective will deterrence be?

To someone lacking a comparable capability, a nanofactory would be incredibly valuable. This implies that deterrence will not be successful.

Conclusion: It will be very difficult to limit proliferation of nanofactory technology and possession of bootleg nanofactories.

23. What effect will this have on policing?

Determine how difficult it would be to make and enforce laws if novel products are readily available through molecular manufacturing.

■Could a 'home appliance' version of the manufacturing technology be used to produce undesirable products?

Yes. Just download the blueprint from the Internet. It could be as easy as printing a picture from a Web browser today.

■Could medical advances lead to new and controversial pleasure devices/drugs?

Yes. Although the chemistry may not be able to make medical chemical compounds, it could make very sophisticated surgical robots. For example, 'acupuncture needle' type probes (with antibiotic surfaces) that can be used for direct brain stimulation ('wireheading') with relatively low medical risk. Or new kinds of sexual appliances.

■How easily could a black market in these technologies be maintained?

For some, more easily than today's drug market.

■How well could lawmaking keep up with newly invented products?

Whole new classes of pleasure device? It'll be hard even to decide what's socially acceptable and what's not.

■How much would new weaponry endanger police?

See the study on military implications (#20). There won't be parity between police and criminals. If criminals have access to advanced weapons, any flesh-and-blood policeman will be in the position of a civilian and police would have to depend on systemic incentives not to kill them. The next likely alternative is that police become paramilitary—SWAT team or “Robocop” —or use remote sensor nets and telepresence.

■How would the ‘arms race’ between invention and detection/defense affect crime? Terrorism?

Criminals and terrorists tend to be stupid and unimaginative, but so do bureaucracies. A smart bad guy would find a large range of new opportunities. Again, it'll be difficult to ‘harden’ civilian targets against crime as well as destructive attack.

Law enforcement expert Tom Cowper suggests that “the biggest unknown is how effective the public police can become—effectively stopping criminals while effectively preserving civil liberties. This is where concepts such as Net-Centric Policing/Government come into play.” In previous conversations with us, he's argued that a key factor is whether we or the terrorists become better at using networks, “augmented reality,” and other tech tools.

We think Tom's emphasis on police (as opposed to military) as a counter to terrorism is worth further attention. Most counter-terrorism involves interaction with civilian populations, and police will do that more sustainably than military (both at home and abroad).

■How much would new sensing and data-mining capabilities help police? How much would they help criminals?

It may be that universal sensor nets would make things better. But they also seem to offer new opportunities for planning crime and for extortion.

Conclusion: Distributed manufacturing of advanced products will pose several substantial challenges to traditional police operations.

In response to this question, Tom Cowper (Police Futurists International, etc) writes: “Let me start this by saying that the issue of molecular manufacturing (MM) mandates dramatic improvements in the way we do policing in the free world. If we are to maintain a free society in an MM world we will have to become very effective at identifying, stopping and incapacitating criminals and terrorists of the future, and do so in a way that does not violate civil liberties. Admittedly a tall order. But as you have pointed out here several times a police state is one definite possibility for the future if government and law enforcement doesn't get its act together and find ways to provide both safety and security, which includes regulating MM to some extent. We don't have to become entirely paramilitary to accomplish this but we will have to employ advanced technologies, including MM created weapons and IT capabilities like TIA. One of the things that we have to keep in mind and you have to include in this paper I think is the understanding that MM won't exist in a vacuum. The future world within which MM will exist will also be a world where MM will facilitate and be facilitated by advanced AI, macro-robots, intelligent environments, cybernetics, etc. Within that world, our notions of privacy and liberty, derived exclusively from Agricultural and Industrial Age circumstances will have to change. Brin's Transparent Society[28] is one future concept within which effective policing might be capable of providing both safety and liberty. There may be others.”

24. What beneficial or desirable effects could this have?

Explore positive factors that will promote the development and deployment of molecular manufacturing (MM).

■How much could the technology reduce illness and disability?

Simple things like water filters and fast, cheap, easy medical sensors could make a big difference. At first, rapid diagnosis of disease would allow effective quarantine. Later, the ability to rapidly develop products should accelerate medical research and speed the process of finding cures; large-scale quarantine operations may become unnecessary even for new diseases. And the ability to monitor a body in detail and in real time should reduce the risks of new therapies, streamlining research still further. Prosthetic devices, including sensory prosthetics, would be greatly improved.

---

[28] Brin, David, The Transparent Society: Will Technology Force Us to Choose Between Privacy and Freedom?, Perseus Publishing, 1999

Advanced automated treatment devices could be made very cheaply, allowing semi-skilled delivery of medical care. Think of automatic defibrillators in airports. Now project that approach into devices with wide-spectrum real-time biochemical sensors that can dispense appropriate medicines.

Surgical robots could become far smaller, more capable and automated, less invasive. Even without bloodstream robots, a catheter-based approach can be used to clean important blood vessels or repair cartilage. A smart catheter could be smaller than a hair, and used by a general practitioner in an outpatient context.

■To what extent could the technology alleviate underdevelopment?

A general-purpose self-contained factory could bootstrap a region's productivity in a matter of weeks. The main limiting factor would be the availability of designs to solve local problems. But see Gershenfeld on “fab labs”.[29]

■Could this help with food and water shortages?

Diamond-building chemistry could not directly make food. But it could make greenhouses, allowing more reliable food production with less resource usage. It could also make water filters and the required energy supply (solar), both for increasing fresh water supplies and treating runoff or wastewater.

■How much and in what ways (e.g. replacing manufacturing, infrastructure, extraction) could it alleviate environmental problems?

Most of today's components that rely on extracted materials, such as metals and rare earths, could be emulated with higher performance by nano-built systems. Carbon-based products could be disposed of by clean combustion. More automation means fewer people have to work in factories, reducing transportation requirements for both people and materials.[30] More efficient agriculture could reduce soil loss, water use, and agricultural runoff. Cleanup of existing problems would be easier with better and cheaper sensors and robotics.

---

[29] http://www.edge.org/3rd_culture/gershenfeld03/gershenfeld_index.html

[30] http://home.earthlink.net/~durable/

Some serious thinkers are concerned about a global environmental collapse in the next few decades, even apart from the Peak Oil problem. Large-scale use of MM could alleviate much environmental pressure, and actively correct many problems.

■Which natural disasters could it prevent or alleviate?

Easier access to space makes it much easier to deal with asteroids. Also, vastly cheaper construction of telescopes makes it easier to spot them. Large-scale engineering projects could defuse volcanoes and even calderas by turning them into massive geothermal energy projects. Stronger construction could resist earthquakes and hurricanes. Also, large-scale construction of automated aircraft/helicopters could suppress wildfires and aid in rapid evacuations. Better sensors would allow better prediction of weather and climate.

■How much could these benefits reduce social unrest?

Poverty, contagious and parasitic disease, and hunger could be drastically reduced at extremely low cost. To the extent that these fuel social unrest, the application of these technologies would reduce the unrest. However, new problems such as social disruption and boredom may emerge.

■How much cost savings does this represent?

Most sources of product cost would virtually disappear. Even design cost might decrease, as shown by the Open Source software movement. Indirect costs of technological activity, such as pollution, could be substantially reduced.

■How much commercial incentive is suggested by these questions?

The difference between production cost and user value of nano-built products will be astronomical. This provides a high incentive for developing the technology—and then manipulating policy so as to maintain artificial scarcity. Artificial scarcity would cancel many of these benefits.

Conclusion: Molecular manufacturing could be a major benefit to humanity, saving lives, mitigating environmental problems and hazards, and reducing misery enough to substantially reduce social unrest. However, this all depends on policy.

25. What effect could this have on civil rights and liberties?

Study the extent to which advanced technologies will allow violation or protection of civil rights.

■What effect will new surveillance capabilities have on privacy (used by government or privately)?

Extremely cheap manufacture of tiny integrated sensor/network/self-positioning packages, as well as sufficient computer power to store and integrate the information, could completely destroy privacy, unless strenuous decontamination efforts are used.

■What effects will new surveillance capabilities and/or weapons have on governments and other power wielders?

An unaided human would be completely defenseless against even primitive versions of a sensor web and telepresence robotics.

■What effects could new medical technologies have on personal autonomy and sanctity of thought?

Implanted chemical monitors could indicate emotional state. Implanted dispensers could manipulate it. We don't know how feasible or difficult it would be to read thoughts from brain electrode arrays, but we can already read intentions to move muscles (in monkeys).

■To what extent will abuses and crimes increase demand for security and control?

This is way too much power to allow criminals to have. It would send us back to a "state of nature" where no one is safe from anyone else without constant vigilance. For comparison, consider the vulnerability of most home computers to worms and viruses. Compare with the effects of 9/11 on public acceptance of government monitoring (PATRIOT Act, etc).

■To what extent will new capabilities increase demand for autonomy?

It will be much easier to live 'off grid', perhaps even off earth. There will be strong demand for health improvement, which leads naturally to human augmentation.

■To what extent can manufacturing breakthroughs alleviate poverty and misery?

This is important because poverty and misery are breeding grounds for instability and terrorism, and extreme poverty is a human rights violation according to the UN Declaration. It should be possible to eradicate poverty and misery worldwide with very little effort or cost.

Conclusion: Molecular manufacturing technology will force some very hard choices about civil rights. A nano-enabled group that does not consider human rights to be of fundamental importance will be able to violate them utterly. Even when human rights are respected, our concept of them may have to evolve to deal with new and pervasive technological capabilities.

26. What are the disaster/disruption scenarios?

Determine which of the following scenarios are plausible, and if so, whether they are survivable or preventable.

■Massive war?

Highly plausible. A nano arms race appears almost inevitable, and would probably be unstable as discussed in the military capabilities study (#20).

A nano-enabled war would probably be very lethal to civilians. As pointed out by Tom McCarthy, "Military planners will seek a target that is large enough to find and hit, and that cannot be easily replaced. The natural choice, given the circumstances, will be civilian populations."[31] Both full-scale war and unconventional/terroristic war will target civilians, who will be nearly impossible to defend without major lifestyle changes. It would be easy to deploy enough antipersonnel weapons to make the earth unsurvivable by unprotected humans.

■Economic meltdown?

It's easy to imagine a nanofactory package that allows completely self-sufficient living, off grid and without money, while retaining modern first-world comfort levels. However, a modest amount of advertising would make this unattractive to most people.

As discussed elsewhere, we can expect a large fraction of jobs in a wide range of areas related to manufacturing, extraction, and supply to disappear. This problem is already appearing with increased automation and efficiency, but could rapidly get worse.

---

[31] http://www.mccarthy.cx/WorldSystem/war.htm

The factors that lead to economic meltdown also provide increased self-sufficiency, so it ought to be survivable in the absence of oppressive policy (maintaining artificial scarcity while removing sources of income). Secondary effects from social disruption may be problematic but ought to be survivable.

Attempts to subsidize dead-end jobs will probably be harmful in the long run. Some amount of economic disruption should be expected. Social engineering to reduce the stigma of unemployment (why should unearned income be good for the rich and bad for the poor?) and policy to allow displaced workers to share in the benefits of the new technology will be helpful.

■Runaway self-replication?

Also known as the 'gray goo' scenario, this is perhaps the earliest and most famous concern related to molecular manufacturing. Contrary to early statements by Drexler, this could not happen accidentally; manufacturing systems, even early lab versions, will not remotely have the capability to become self-contained free-range self-replicators. However, the deliberate combination of a very small nanofactory, a very small chemical plant to convert organic chemicals into feedstock, and some robotics, could be a substantial nuisance or even threat. Eventually, the technology will develop to the point where it'll be easy to make a device that requires active cleanup to avoid widespread environmental damage. The prevalence of computer viruses implies that creating such devices will be attractive to certain personality types, and eventually within their capability.

So, although runaway self-replication is not a first-rank concern, eventually it will need to be studied, and some combination of prevention and cleanup capability probably will have to be implemented. In theory, this could pose an existential threat.

■Dangerous software?

An arms race (either military or corporate—in fact, conducted by any organization) could involve the development of increasingly capable AIs for the purpose of manipulating or coercing people. Note that this does not require full general intelligence. A variety of manipulative techniques (on either human psychology or other complex systems) can be imagined using only specialized data-processing.

Some theorists believe that a self-improving AI could pose an existential threat: almost any command would cause unexpected and massively disruptive side effects. We do not know whether this is plausible. But nanotech development will certainly be an enabling technology for powerful AI, though we may face this problem even before nanotech is developed. Robert Freitas cites some of these concerns going back decades in *Kinematic Self-Replicating Machines.* Already, enough infrastructure is computer-controlled to make a cyberspace attack potentially very destructive. As more products become computer-integrated, a software attack could shut down, damage, or subvert increasingly crucial functions.

The variety of possible impacts on human psychology, computer-integrated infrastructure, and other systems (e.g. the effect of computer trading on the stock market) implies that this whole area should be extensively and creatively studied.

■Moral or social meltdown?

The availability of new products and lifestyles may cause disruption in social fabric, especially in conservative societies that may actively resist change. This may inspire a backlash, possibly including force. It is likely to destroy some cultures. Broader effects are unknown.

■Environmental devastation by overproduction?

It would be easy to build enough nano-litter to cause serious pollution problems. Small nano-built devices in particular will be difficult to collect after use. It will also be easy to consume enough energy to change microclimate and even global climate.

Overpopulation is probably not a concern, even in the event of extreme life/health extension. The more people use high technology, the fewer children they seem to have.

Conclusion: Several plausible disaster scenarios appear to pose existential threats to the human race.

27. What effect could this have on geopolitics?

Explore the impact that molecular manufacturing will have on the current habit of maintaining sovereign nations.

■What would be the effects on international relations of reduced international trade, especially in oil?

Reduced demand for Middle East oil probably would be highly beneficial to international relations. Reduced international trade in general probably would not be beneficial, since it would reduce the interdependence of nations.

■Can a technology-driven arms race be stable?

Probably not. See study #20 for analysis of how and why a nano-weapons arms race would be more unstable on several counts than the nuclear arms race has been.

■What would be the effects of nationwide changes in lifestyle and personal resources? How quickly could those effects happen?

We might predict a lower birth rate, substantially lower death rate, and greatly increased healthspan. Access to more information could produce better democratic governance, or simply more distraction and disinterest. Other effects should be studied.

■To what extent will these technologies require worldwide policing? What problems does worldwide policing create?

An unrestricted nanofactory anywhere in the world could be used to build weapons of mass destruction with global reach. For this reason alone, it appears that either the technology or its users absolutely must be restricted/policed, unless (which we believe unlikely) it turns out that defense is superior to offense for all product technologies (see study #20).

Policing, unlike military occupation, requires that the population accept the legitimacy of the force. No legitimate worldwide policing organization exists today. Nations cannot police each other sustainably. But many nations cannot police themselves. To the extent that international policing is required, it will add to social unrest unless a new structure is developed that can coordinate and support national policing efforts while retaining national sovereignty.

■What is the possibility of preemptive strikes to prevent development in other nations?

Each nation will see only a few possibilities: 1) an arms race that will probably be unwinnable since it will develop into a disastrous war (see #20); 2) developing ahead of everyone else and establishing dominance; 3) some other nation developing earlier and establishing dominance; 4) international cooperation and trust sufficient to ensure safety; 5) a multinational organization willing and able to keep the peace.

Option 1 is undesirable; Option 3 is probably unthinkable for any of the current large powers; Option 5 is probably unacceptable to the U.S., as the world's sole superpower; Option 4 is probably unfeasible. Only one nation can succeed at Option 2. This implies that a preemptive strike option (whether military attack, or sabotage or derailment of nanotech development efforts) will appear very attractive to a number of powerful nations.

■What barriers to international cooperation could make these problems more difficult to solve?

Culture clash, lack of trust, xenophobia, religious fundamentalism, grandiose or aggressive national leadership. Increased information and reduction in poverty could reduce these factors eventually, though it could also reduce the interdependence that provides one incentive for cooperation.

Conclusion: Molecular manufacturing technology is powerful enough to require new ways of interaction between nations.

28. What policies toward development of molecular manufacturing does all this suggest?

There are several options for developing molecular manufacturing. Which ones might work as planned, and what would be their effects on post-development courses of action?

■Relinquishment: prevent development worldwide?

This is highly unlikely to work. It'll be too easy to develop, and the basic theory has been published for more than a decade.

The effect of attempted relinquishment would be to ensure that MM was developed by random outlaws. The delay would allow time for the development of more enabling technologies, probably increasing the abruptness of development and deployment.

■Asymmetric development: one nation develops in advance of the others?

This appears possible, depending on which nation. If a nation other than the U.S. tries it and does not conceal their effort successfully, the U.S. will likely be able to catch up, leading to parallel development or possibly to U.S.-led asymmetric development. A U.S. program would have to be well designed, avoiding a variety of problems common to U.S. government-funded programs.

The likely follow-up to asymmetric development would be an attempt at worldwide control. The effects of this would depend heavily on the policies adopted by the government in question.

■Parallel development: several nations develop at around the same time?

This seems quite likely, either from an arms race or from development by multinational corporations.

The result would depend heavily on policy. If an arms race can be avoided, and effective administration/policing can be implemented, it could turn out well. But an arms race looks pretty likely, and would probably be disastrous. Also, parallel development would make it harder to restrict proliferation.

■International development: explicit cooperation between nations?

Seems unlikely to be tried. If it is tried, it's likely to fail due to politics, mistrust, and inefficiency that allows a national crash/secret program to finish first.

International development would reduce the pressure for an arms race and give multiple nations a stake in setting the policy for use of MM. Paradoxically, it could reduce proliferation, since joint ownership would encourage the widespread availability of controlled versions and blunt the desire for uncontrolled versions.

Corporate development by a large, international corporation may also be an interesting possibility to study. It may even be worth working to try to make it happen that way. Corporate development is likely to be a lot more efficient and less vulnerable to politics than a project that's shared between governments. But it'll still promote the benefits listed in the previous paragraph, assuming the corporation has (and follows!) really good policy advice.

■Non-proliferation: restrict availability of the core technology?

Will probably be tried. Will probably help to some extent. Will be ineffective in the long run unless combined with two other policies: 1) reduce desire for unrestricted technology by providing easy access to restricted but useful technology; 2) develop the ability to deal with eventual proliferation.

The alternative—allowing everyone access to the unrestricted technology—appears extremely dangerous; perhaps comparable to leaving the post-Soviet nuclear infrastructure unguarded.

■Slow development: don't make special efforts?

Likely to lead to random development, rapid bootstrapping due to other nanotech advances, and lack of ability to implement policy.

If development is delayed long enough for other technologies to catch up (perhaps two or three decades) then this could give us time to adjust gradually. But that much delay appears unlikely, and we'd lose the benefits for those decades (see study #30).

■Accelerated development: put limited effort toward it?

Would likely inspire other efforts, leading to parallel development.

■Crash development: put maximum effort toward it?

Could lead to either parallel or asymmetric development. Could smooth the transition by requiring more creativity to design products.

Conclusion: Early development combined with anti-proliferation policy appears preferable, but more study is needed, and the outcome depends heavily on the actions of the developer(s).

29. What policies toward administration of molecular manufacturing does all this suggest?

There are several options for administering molecular manufacturing. Which ones might work as planned, and how desirable are they? Which classes of problem are suitable for the various options? What are the consequences if an option is tried and fails? Which options can coexist in one society, or even in one (shrinking) world?

Scope and Degree of Control

The impact of local policy may reach far beyond local borders...

- Total control: ironclad, worldwide control of all that relates to development or use of molecular manufacturing

This could work, if the controllers were sufficiently ruthless and intrusive. Obviously, unless the controllers are also saintly, it would be a human rights disaster.

- No control: let a solution emerge

The continuing problems of spam and computer viruses and intrusion indicates that this is unlikely to protect most people.

- Local control: several autonomous regions find their own solutions

Nano weapons, nanofactories, and other dangerous products can easily cross borders. Unless the regions all have an interest in keeping each other safe as well as themselves, this probably won't work.

- Coordinated or hierarchical control: a mix of local and top-down policy

This might be a good approach. Note that it would require an international organization at the top, probably with verification and enforcement capability. Note also that hierarchy is a 20th century invention, and may be outdated/surpassed by human networks. The concept of "network democracy"[32] may work better these days.

- Other structures? Implications of space access?

It's hard to control what happens several light-seconds away. This may imply a need to allow only trusted people/groups into space. (It looks like this is our unofficial global policy already.)

Approaches to Resources

[32] Garrison, Jim, America as Empire: Global Leader or Rogue Power?, Berrett-Koehler Publishers, 2004

There are several fundamentally different approaches to dealing with resource allocation and other policy issues. We have covered these in detail in “Three Systems of Action: A Proposed Application for Effective Administration of Molecular Nanotechnology”.[33]

■Security: preserve the status quo against destructive change

Prevent negative-sum transactions (e.g. theft). Deception and the use of force are acceptable. Commerce and information sharing are potential weaknesses. Loyalty, tradition, and honor are relevant values. Molecular manufacturing will raise many security issues.

■Commerce: optimize use of scarce resources; collect resources

Maximize/optimize positive-sum transactions (e.g. free market trade). Use of force is not acceptable. Efficiency, innovation, and honesty are relevant values. Several resources will still be scarce even under nearly-free manufacturing, and much work will still benefit from commercial/monetary incentives.

■Information (Non-rivalrous): maximize availability of non-scarce resources

Optimize use of unlimited information (unlimited-sum transactions: the cost is very low and is unrelated to the value). Creativity and openness are relevant values. Reputation is a major motivator. This approach may be relevant for many blueprints and nano-produced objects.

Worldviews and values

There are several cultural traditions in the world—very different, and perhaps incompatible.

■Personal freedom and opportunity, openness, free market (Western)

This has spearheaded the development of science and technology, as well as democracy. It values diversity, which makes it less destructive/oppressive. It may be unwilling or politically unable to exert sufficient force to deal with major threats to security. This will be countered to some extent by creativity in problem solving.

■Paternalism, social constraint (Tribal, Moralist)

[33] http://CRNano.org/systems.htm

Compared with Western, this has more opportunities for social engineering and less resistance/rebellion to government control. But central planning allows less creativity and diversity, and creates more oppression and limitation (at least from the Western point of view). Lack of feedback and emergence allow mistakes to persist.

■Suffering, nihilism, submission (Fundamentalist)

As far as we can see, this can only be justified by intangible values alien to Western thought (though present to some extent in America's Puritan heritage). It may also see Western tradition as dangerous and immoral; this may lead to unavoidable conflict.

Decision Making Options

There are many ways to make decisions. This is just a sample.

■Laissez-faire?

Just let things happen. This is likely to be very hard on the average person.

■Democracy?

Requires an informed electorate. Not likely—too much science and technology background required, too many counterintuitive and nonlinear effects.

■Bureaucracy?

Adds friction to the system. This is sometimes good, but unlikely to be adequately responsive for most problem solving.

■Dictatorship?

Requires a good dictator, which is not likely and perhaps not possible—there's simply too much to understand.

■Network?

Not well understood yet, but may be the best option.

Administration Options

There are lots of ways to influence or limit the use of this technology. All have limitations. These are just a sample.

■Law?

Must be backed by police force. Too much force reduces the legitimacy of the law.

■Treaty?

At best, a process for agreeing to standards and creating awareness of mutually beneficial choices. Won't work if not in the interests of all signatories, though may serve to formalize and focus the use of other incentives such as military threat (see study #27).

■Social engineering and public perception?

Will only work on some people.

■Intellectual property?

An odd convention, probably over-used in modern economies; not a good match for non-rivalrous goods.

■Commercial self-regulation?

Companies will sometimes modify their own behavior to prevent more onerous regulation. But this probably requires a substantial threat of government regulation.

■Surveillance?

Surveillance will be extremely useful and effective when sensors get cheap enough and computers get powerful enough to watch everyone full-time and highlight anomalous behavior. There are no obvious inherent limits on the use of surveillance, and several obvious benefits. This poses a severe threat to modern Western concepts of privacy. It also creates practical problems, including strong pressure for full-time behavioral conformity (since any unusual action will be scrutinized, only exhibitionists will be comfortable risking any unusual behavior) and lack of ability to oppose unjust government.

■Human modification?

Even more intrusive than surveillance: With compact technology, cheap manufacturing, and accelerated medical research, an implantable device could be developed to monitor and possibly change people's psychiatric/neurochemical profile. This threatens our concepts of autonomy and even selfhood. However, it may be the most effective way to solve the most difficult security problems, making it dangerously attractive. Efforts to design administration must take this possibility into account, either rejecting it or limiting it, with strong safeguards in either case. Conversely, the use of such technology within the administration could serve to limit the impact of destructive people and improve the effectiveness and reliability of the administration.

Extremes That Won't Work

It will be very tempting to choose simplistic extremes of policy, especially if events seem to be leading toward loss of control. But this virtually guarantees failure. Furthermore, extreme policy disasters can't be corrected by further extreme policy; in general, the bad effects will add, not cancel.

■Crash program vs. delay

As discussed in study #30, a crash program without substantial policy planning will lead to a powerful technology we don't know how to handle. But a delay, especially if it's implemented by denying the feasibility of the technology, will also lead to lack of preparedness—and reduced ability to control or predict when someone finally does develop MM.

■Restriction vs. freedom

A policy that is too restrictive will inspire attempts to circumvent it, from within the administration (idealism, high-stakes blackmail or subversion) and from without (cracking restrictions, independent development). This will require intolerable and unsustainable restrictions, and will eventually fuel a black market where one leak spreads unconstrained nanotech beyond hope of containment.

A policy that is too lax will lead to a situation that can't be controlled, a “state of nature” in which anyone can strike at anyone else unless eternal vigilance is kept. This will create a public outcry for control as well as government insecurity, leading to overly restrictive policy.

It looks like the best approach is wide availability of nanofactories with built-in technical restrictions. The more benefits are freely/widely available, the less pressure for independent development. The widespread use of 'approved' hardware allows all sorts of less-intrusive controls. See our paper on "Safe Utilization of Advanced Nanotechnology".[34]

■Global empire vs. independent states

A declared global empire will be resented, hated, and feared, no matter who is emperor. Preparation for it is likely to tempt preemptive strikes.

Independent states will not be able to coordinate the cross-border policing necessary to prevent cross-border crime and terrorism. Some states will not be able to police themselves adequately. Any state that maintains an uncontrolled nanotech capability will threaten the entire world.

The best solution is probably an international organization, both to administer the molecular manufacturing that has been developed and to prevent possession of dangerously unrestricted versions by illicit actors. This might be modeled on the IAEA, the WHO, or UN peacekeeping forces. Unfortunately, international cooperation is not at its best right now (in mid-2004); such an organization would take time to develop, and some nations (especially, perhaps, the U.S.) may try to sabotage it and go it alone.

Both nanotech problems and nanotech solutions are international. If MM goes wrong, some of its problems may be global in scope. Gray goo and military nanorobots will not respect national borders. Economic collapse of any large nation will shake all the rest. Likewise, MM risk prevention must also be global. Programs and policies for reducing poverty must be international. Administration to detect and prevent rogue MM programs must have global jurisdiction. An accretion of national programs may be able to mitigate some problems and risks, but cannot address all of them. International policies, and international bodies, must be designed and created before molecular manufacturing arrives.[35]

[34] http://crnano.org/safe.htm

[35] http://crnano.org/int_control.htm

We'll mention again "network democracy"[36] as a possible approach. Small groups with specific focus may be both more responsive and less threatening. However, there still has to be some way to apply their recommendations.

■Guardian vs. Commercial vs. Information

As explored above, negative-sum, positive-sum, and unlimited-sum situations require very different approaches. Any single approach will be inadequate, and will not only fail but will be destructive in situations that demand a different approach. (See *Systems of Survival*[37] on "monstrous moral hybrids.") Effective administration will require application of all these approaches, chosen appropriately to address the various kinds of problems, and probably implemented by distinct but coordinated organizations.

■Capitalism vs. socialism

The goal of socialism is to make sure that everyone is provided for adequately by redistributing wealth. Molecular manufacturing will certainly produce enough wealth to make everyone (worldwide) rich by today's standards, and will probably exacerbate imbalances and inequities; this will tempt socialist policy. Socialism is great in theory, but in practice it cripples the main incentives for productivity, innovation, and trade.

The goal of capitalism is to accumulate resources and use them to generate wealth. However, it can lead to destructive imbalances of power such as monopolies. When the cost of production becomes a miniscule fraction of the value to the user, and when manufacturing capital and labor alike lose their value, capitalistic wealth accumulation may cease to provide its customary spinoff of value to the economy and to society.

The best solution may be one inspired by software development. Software is another area where the cost of duplication is far lower than the value of the product. For several decades, commercial software has coexisted with free software; each has benefited from the other, and neither has out-competed the other. Commercial software tends to be more polished, adding value; free software (and its recent cousin, Open Source software) has been an important source of innovation, and is available to people with no money to spend.

---

[36] Garrison, Jim, America as Empire: Global Leader or Rogue Power?, Berrett-Koehler Publishers, 2004

[37] Jacobs, Jane, Systems of Survival: A Dialogue on the Moral Foundations of Commerce and Politics, Vintage, 1994

Patents or other artificial scarcity applied to the nanofactory could restrict trillions of dollars of economic benefit and comparable social benefit. Since a single general-purpose manufacturing system can make millions of different kinds of products, there is plenty of opportunity for corporations to make money by designing and licensing products, and paying part of that fee to the nanofactory inventors. At the same time, vast benefits could be delivered both to poor users and to the common pool of information by designers who wish to make their designs available for free—but only if nanofactory use for producing free designs is not encumbered by heavy licensing fees. This would allow a single fundamental invention, the nanofactory, to be used in both a commercial context and a non-rivalrous, unlimited-sum context.

The difference between socialism and free sharing of non-rivalrous goods should be carefully noted. Socialism is about redistribution: something must be taken away from its owner in order to give it to someone else. By contrast, increasing the distribution of non-rivalrous goods does not require denying them to anyone. Intellectual property (both patent and copyright) is a legal construct, a right invented and maintained by society and granted for the purpose of benefiting society by stimulating innovation while maximizing distribution. Failing to maintain this artificial scarcity does not take away an inventor's intellectual property, because that property does not exist unless and until society bestows it. Under the current proposal, the inventor of a nanofactory would still become astonishingly rich by extracting whatever licensing fee the market would bear from commercial users. Thus the incentive to innovate would be preserved, while distribution would be better than if the IP were completely commercialized. (See e.g. Lawrence Lessig on upstream vs. downstream patents.)

Post-Molecular-Manufacturing "To Do list"

After molecular manufacturing is developed, the job is just beginning. This list should be expanded in consultation with various future studies groups and think tanks.

■Active shield? (Global sensor grid to detect, and possibly respond to, nanorobot activity)

If the administration fails to prevent the development of small undesirable nanorobots, it may be very important to have a system in place to rapidly detect their activity. For example, Robert Freitas has calculated that a well-dispersed airborne selfreplicator of advanced design might produce sufficient copies to block all sunlight in as little as two days.[38] If this development is possible, it obviously must be prevented with multiple levels of safeguards. Research must be done well ahead of time to determine whether such a thing may become possible; unless it can be conclusively ruled out (better than billion-to-one certainty), then deploying an early-warning sensor net and pre-positioning countermeasures would seem to be a minimal precaution.

■Artificial intelligence?

Computers will be one of the easiest things to build with molecular manufacturing. A sudden increase in available computer power by many orders of magnitude will surely make various forms of artificial intelligence more powerful, and enable new forms that are not practical with current hardware. Even if runaway AI doesn't introduce inherent danger, misused AI could be extremely powerful. Conversely, AI of various sorts—even something as straightforward as advanced data-mining—could solve several problems that currently have us stumped. It may be worth pre-planning to launch an AI research program as soon as the computer power becomes available.

■Space program?

Access to space will become cheaper by at least several orders of magnitude. This should be planned for. Space may be useful for resources, for quarantine, and for science.

Conclusion: Many options need to be considered and synthesized. Hastily chosen or simplistic policy is extremely unlikely to be wise or effective.

30. How can appropriate policy be made and implemented?

What options are still available to choose the course of molecular manufacturing and its effects, and how rapidly are the options disappearing?

■In the absence of concerted government action, when will molecular manufacturing be developed?

Several technology trends point to molecular manufacturing, or equivalent capability, being developed around the 2020-2030 time frame. However, the cost of development is falling rapidly,

---

[38] http://www.foresight.org/NanoRev/Ecophagy.html#Sec8-2_GrayDust

the paradigm is becoming well-known and plausible, and the incentives to develop it sooner appear quite significant—the impact on a single industry could easily be greater than $1 billion. This indicates that individual corporations may begin development when the cost falls below $1 billion and the time below five years.

The cost and time may already be small enough to allow development within five years and $1 billion. Several of the major sub-projects appear to cost less than $10 million apiece. We don't have a cost estimate for the lab work to build the first fabricator, and any estimate may be revised downward by invention of an easier technique. But the cost of a capability to build 3D structures with 20 nm feature sizes and thousands of features may already be less than $1 million, and limited provision of smaller feature sizes and even atomically precise features may be feasible with off-the-shelf chemistry.

■If a crash program were implemented (or has already been) in the United States or elsewhere, how soon could MM be developed, including a general design capability?

If a well-funded crash program had been started, say, in 1992 when *Nanosystems*[39] was published, it could succeed literally at any time. A program started today might be limited by software and nanosystem design, requiring maybe five years of intense Manhattan-project-level effort to develop a CAD program, a few basic molecular machines, and several designers capable of making products from them.

It should be noted that software (including computational chemistry software) is relatively cheap, and easy to work on in secret; one strategy for asymmetric development would involve developing a full set of software and nanomechanical designs in the absence of experimental verification, then waiting for lab techniques to advance to the point where the lab work could be done in just a year or two. If such program were not discovered until the lab work started, it would be extremely hard to catch up.

■How quickly could the studies listed here be completed?

A detailed analysis of all these points might easily take a year or more, even if they were approached in parallel.

---

[39] K. Eric Drexler, Nanosystems: Molecular Machinery, Manufacturing, and Computation, John Wiley & Sons, 1992.

■How high are the stakes?

It appears that development of molecular manufacturing by a hostile nation would seriously threaten the ability of any nation, including the U.S., to defend itself or to respond effectively to an attack.

Further, it appears that an arms race focused on this technology would probably end in a devastating war with extremely high civilian casualties. Several other independent disaster scenarios might cause unacceptable loss of life in the absence of effective policy administered by an acceptable controller.

■If a national crash program is necessary, how quickly could it succeed?

As implied above, a crash program started today might easily take five years. Since a major gating factor may be the development of novel software, this time may not shrink much in the future, though cost can be expected to decrease rapidly.

It must be emphasized that simply implementing a crash program is *not* an adequate strategy to avoid disaster. *In the absence of proactive policy work and implementation of the policy for effective administration, the existence of the technology very likely will lead to disaster.* However, as explained below, this is not an adequate argument for postponing the development.

■If international cooperation is necessary, how effective could it be, and how long would it take to establish?

If past experience is any guide, international cooperation could take years to establish, and would at best delay the problems.

■How detailed a plan must be worked out in advance? Who must buy into the plan?

It must start with the design of quasi-governmental administration (effectively, a constitution as well as procedures). It must also address the practical steps necessary to create that administration. Everyone who will have access to the technology (including the ability to develop it independently or acquire it through a black market) will have to buy into the plan or else be forcibly subjected to it. Note that widespread and prolonged use of force leads inevitably to an unsustainable conflict and/or a human rights disaster.

An alternate view is that it's better not to have a central over-arching administration at all: that such an administration would be too likely to abuse its power, while simultaneously suppressing the development of technologies (e.g. active shields) to mitigate bad consequences. We believe that in the absence of central administration, too much power will ‘trickle down’ to bad people and groups, then concentrate and be used for destructive purposes, creating tragedy and probably disaster.

However, central administration may not be adequate either. An effective solution may require the invention of new forms of administration/government, taking advantage of rapidly organized networks and high information flow.

■How long will it take to set up administrative structures?

Probably several years.

■What effects will public perception of ‘nanotechnology’ have?

It depends on the country. In a democracy, too much fear can remove a lot of support; conversely, realistic education about the benefits and challenges/problems can create a massive and productive effort to solve the problems. In other places, e.g. China, public perception probably doesn't matter as much.

■What could be done to delay molecular manufacturing?

Scientists such as Smalley, Whitesides, and Ratner have done a very effective job of delaying investigation and development in the U.S., but this may be about to change. CRN has heard increasing frustration and skepticism among young scientists against the position that it's impossible. Still, it would probably be possible to postpone U.S. attention for another few years if key pro-MM spokespeople could be convinced to announce that they had shifted position and now believed it was impossible to achieve.

However, now that a group in Russia appears to be working on MM, delay there may not be possible. At least one publication from Iran has announced MM as a goal of that government. So it will probably be developed somewhere in the world no matter what the U.S. does. If the U.S. doesn't work on it, it might take between 5 and 10 years; if the U.S. actively opposes development and/or sabotages programs it's aware of, it might be stretched to 10-15 years, though this doesn't appear at all certain.

■What would be the effects of delaying molecular manufacturing?

If it could be delayed three decades, its impact may already be largely eclipsed by other powerful technologies. However, this long a delay is unlikely.

If delayed by one to two decades, general nanotechnology progress combined with continuing theoretical and hobby work could make it much faster and more widely proliferated once it happens: the recipe could spread widely and quickly, and could be easily applied. The sources and timing of development would become increasingly hard to predict.

Also, realization of all the benefits (reduction in poverty, improvement in health, increased abundance providing for aging populations, avoidance of severe ecological collapse) would be delayed. This could account for tens of millions of deaths per year. Anyone who deliberately delays molecular manufacturing by even a few years could go down in history beside Stalin for mass murder by deprivation.

Robert Bradbury points out: *Leaving aside the quality of life issues -- this isn't hard to estimate. Demographers (and insurance people) typically look at this in terms of 'Years of Potential Life Lost'. If one simply takes vasculoid*[40] *(the design of which is* not *simple or cheap at least for now), one cuts out heart disease, cancer, stroke and septicemia (as well as other infectious diseases) as causes of death. This probably buys you something like 7-15 years of additional lifespan according to various papers I've seen. Given ~50 million deaths per year a good number might be 500 million years of potential life lost. I think this is a per year of delay number --* but *I'm going to have to think about it a bit to consider what the impact of a gradual phase-in of nanotechnology might involve.*

Now if *robust nanotechnology allows us to figure out the precise causes of aging, lifespan gets extended anywhere from 2000-7000 years (limited by the external hazard function). In that case one is talking something like 250 billion years of potential life lost (again subject to some discounting).*

Conclusion: The situation is extremely urgent. The stakes are unprecedented, and the world is unprepared. The basic findings of these studies should be verified as rapidly as

[40] "Vasculoid: A Personal Nanomedical Appliance to Replace Human Blood", Freitas and Phoenix, www.transhumanist.com/volume11/vasculoid.html

possible (months, not years). Policy preparation and planning for implementation, likely including a crash development program, should begin immediately.