World Technology Evaluation Center (WTEC)

WTEC Workshop Report on

R&D Status and Trends in Nanoparticles, Nanostructured Materials, and Nanodevices in the United States

Proceedings of the May 8-9, 1997 Workshop

Richard W. Siegel, WTEC Panel Chair
Evelyn Hu, Panel Co-Chair
M.C. Roco, NSF Coordinator

JANUARY 1998

International Technology Research Institute
R.D. Shelton, Director
Geoffrey M. Holdridge, WTEC Director and Series Editor

4501 North Charles Street
Baltimore, Maryland 21210-2699

WHITHER NANOTECHNOLOGY?[1]

Ralph C. Merkle
Xerox PARC
3333 Coyote Hill Road
Palo Alto, CA 94304

Introduction

A new manufacturing technology looms on the horizon: molecular nanotechnology (http://nano.xerox.com/nano). Its roots date back to a 1959 talk by Richard Feynman (http://nano.xerox.com/nanotech/feynman.html) in which he said, "The principles of physics, as far as I can see, do not speak against the possibility of maneuvering things atom by atom. It is not an attempt to violate any laws; it is something, in principle, that can be done; but in practice, it has not been done because we are too big."

In the last few years the idea that we should be able to economically arrange atoms in most of the ways permitted by physical law has gained fairly general acceptance. This can be viewed as simply the culmination of a centuries-old trend: the basic objectives of manufacturing are lower cost, greater precision, and greater flexibility in what can be manufactured: as the decades have gone by, we've gotten better and better at it. The limit of low cost is set by the cost of the raw materials and energy involved in manufacture, the limit of precision is the ability to get every atom where we want it, and the limit of flexibility is the ability to arrange atoms in whatever patterns are permitted by physical law. While it seems unlikely that we will ever completely reach these limits, the objective of molecular nanotechnology is to approach them. Manufacturing costs should be low — a dollar a pound or less — almost regardless of what is being manufactured. Almost every atom should be in the right place — while background radiation limits this, error rates of a single atom out of place among many tens of billions seem feasible in properly designed structures under "normal" conditions. And finally, we should be able to make most of the stable structures that are consistent with physical law. As structures become less stable they become more difficult and arguably impossible to make, but this still leaves a vast space of possible structures that are beyond the reach of current methods. In addition, some structures might be stable if only we could make them, but all intermediate states would be unstable. Drexler, for example, has argued that the molecular equivalent of a stone arch (http://www.asiapac.com/EnginesOfCreation/EOC_References.html#0025) would be unstable unless all its pieces were in place. The final result would be stable, but all synthetic pathways leading to this result would have to pass through an unstable state, making synthesis impossible.

While the broad objective has gained acceptance, as a community we have still not agreed on how best to proceed, nor on what this future technology will look like, nor on how long it will take to develop. The purpose of this paper is not primarily to focus on specific technical approaches, but to ask, "What do we need to do, as a community, to speed the development of this new technology?"

The Goal

Before going further we need to make sure we are in broad agreement about the goal. Molecular nanotechnology should, by definition, permit us to manufacture (among other things) molecular computers with mole quantities of switches, connected in the intricate patterns required by today's complex computers, at a cost of perhaps a dollar a pound (or less). Today's computer — if we weigh the thin layer on top of a computer chip — costs tens of millions of dollars per pound. This thin layer, only a few microns thick, contains almost all the complexity of the modern computer. The rest of the wafer is mere sand, dragged along as a convenient mounting platform for the active region but doing little else. Viewed in this light, lithography falls woefully short of the cost goal.

Molecular nanotechnology should let us extend this very thin and complex layer into three dimensions while greatly shrinking the size of the switches. It should let us position dopant atoms at specific lattice sites

[1] This rough draft was written over a few days for the May 8-9 WTEC workshop on nanotechnology.

(chosen by design to optimize device function) while simultaneously keeping the manufacturing costs as low as the manufacturing costs of a piece of wood.

Besides computers, molecular nanotechnology should let us make inexpensive materials with a strength-to-weight ratio similar to that of diamond. These would have wide ranging applications in structural and load bearing applications. Manufactured with precisely the desired shape and structured at the molecular scale to optimize material properties, we should be able to make a jet, a rocket, a car or even a chair that would, by today's standards, be remarkably light, strong, and cheap.

The objective of molecular nanotechnology is not simply to provide a few new products nor to greatly enhance the performance of some select high-tech devices, but to replace essentially the entire existing manufacturing base with a new, radically less expensive, radically more precise, and radically more flexible way of making products. The aim is not simply to replace today's lithographic fabrication facilities to let us make better computers, but also to replace the assembly lines for cars, televisions, telephones, books, bookcases, airplanes, tractors, etc. The objective is a pervasive change in manufacturing, a change that will leave virtually no product untouched.

It Will Take a Lot of Work

It seems likely that the development of such a capability will require (a) time and (b) resources. The development of nuclear weapons took billions of dollars and a very focused development project. The Apollo program likewise took a focused effort over many years and billions of dollars in money and vast amounts of creative talent. The development of the computer industry, while following a very different pattern (private versus governmental, incremental "pay as you go" versus large up-front funding), also involved major funding and many years.

It is too early to say exactly what pattern the development of molecular nanotechnology will follow, but it is not too early to say that it is likely to require major resources. Whoever makes the decision to commit those resources is unlikely to do so unless there is a clear picture of both the goal and how to achieve it.

Suppose a hypothetical funder came to the research community today and said, "Molecular nanotechnology has a very high payoff, and I wish to start a major new program in the area. What should I do? What should I fund?" The answer, today, would be a chorus of voices tugging in all directions.

Perhaps our hypothetical funder would fund all the different approaches. This was the basic strategy used to develop nuclear weapons. But that was a war-time effort motivated by panic and the fear of annihilation. A more likely scenario is that our hypothetical funder would say, "You are all saying different things — I won't fund a major new project until at least some substantial fraction of you have reached agreement about what to do."

What, then, is the key to developing molecular nanotechnology? Developing agreement about what it is and how to achieve it. How can we develop agreement? As a first step we must explicitly pursue research into the question, "What would a molecular manufacturing system look like?"

Self Replication and Low Cost

Take the issue of manufacturing cost. This is a primary objective of molecular nanotechnology. One way to keep manufacturing costs down would be to develop self replicating manufacturing systems (http://nano.xerox.com/nanotech/selfRep.html). The development of self replicating systems seems like a daunting task, so it is natural to ask if there are alternative ways of achieving the cost objective. To date, no alternative of similar effectiveness has been proposed. As noted earlier, lithography is perhaps seven orders of magnitude too expensive. Other approaches fall short in terms of the range of things they can make, or in terms of the precision with which they can make them. Bulk chemicals are produced today at relatively low cost, but the range of molecular structures that can be made this way is very limited. Lithography can make a great many patterns on a surface, but not with molecular precision. While self assembly is a powerful approach, the direct manufacture of (for example) a diamond rocket by self assembly seems implausible (while self assembly is likely to be important if not crucial in developing nanotechnology, it can still only make an extremely small fraction of what is possible).

It would seem that either (a) we will develop artificial self replicating systems or (b) we will not. If we do, then we can address the issue of manufacturing cost. If we do not, we must seek an alternative — and no alternative of similar effectiveness has yet been proposed.

It is worth noting that we already have self replicating systems of the biological variety. Such systems can already make desirable materials. Wood, for example, is relatively low cost and provides a reasonable strength to weight ratio. Using a programmable protein synthesizer (a.k.a. a ribosome), these self replicating cells can synthesize many compounds. Biological approaches, though, can make but an infinitesimal drop in the vast ocean of the possible. Shall we turn our backs on that ocean? Diamond semiconductors, materials that resist high temperature, structural materials with the strength to weight ratio of diamond, and a host of other examples do not seem to fall within the range of structures that biological systems can directly make.

If we pursue artificial self replicating systems, what do they look like? What are the principles on which they are based? How complex will they be? These and other questions must be systematically addressed, with a confidence and at a level of detail that lets us base major investments on the answers. (While the author has written several articles about self replicating manufacturing systems (http://nano.xerox.com/nanotech /selfRep.html) and has no doubt that they will play an essential role in future molecular manufacturing systems, the point here is that individual conclusions, regardless of how sound, aren't enough. Some substantial portion of the research community must address the issues and reach at least rough agreement about the answers).

If self replication is the right approach and we fail to pursue it, we'll make no further progress. If it's the wrong approach we must develop an alternative. No plausible alternative has been proposed which could simultaneously achieve the three objectives given above: low cost, molecular precision, and great flexibility in what can be made. Investigations to date strongly support the feasibility of programmable self replicating systems. The obvious strategy is to investigate this approach in greater depth.

Molecular Modeling

If we wish to accomplish that which is new, we must at some point discuss what we have not (yet) done. If what must be done is relatively complex (a self replicating system, for example), then we must be prepared to spend substantial time and effort discussing things that have not been made and will not be made for many years.

At the same time, we must take steps to insure that our discussions of what hasn't been done remain focused and do not drift into abstract errors and vague generalities.

Fortunately, we have a tool at hand for dealing with this: molecular modeling (http://nano.xerox.com /nanotech/compNano.html). We know the laws of physics, and we do not expect them to be substantially in error as we apply them to molecular systems under "reasonable" conditions. The applicability of Schrodinger's equation to molecular machines is unlikely to change in the next several decades. We do not need, nor do we expect, any major revolutions in physical law. Our goal and our desire is to develop molecular machines that are feasible with respect to known and well understood physical law. While physical experiments let us explore a tiny fraction of what is possible, they cannot let us investigate what we do not yet know how to make.

Molecular modeling can be used to probe systems that have already been built (allowing us to check the accuracy of the models), systems that might soon be built (letting us inexpensively explore alternatives) or systems that won't be built for many years (again letting us inexpensively explore alternatives, but on a longer time horizon).

If the key to progress is developing a shared understanding of the approach or approaches which are worth pursuing, as well as some shared vision of the goal; and if the goal cannot be achieved without many years of work, then we must adopt a disciplined method of analyzing the alternative ways of achieving the objective. Molecular modeling is a major component of that discipline.

Modeling an Assembler

To sharpen the focus on this idea of modeling future molecular machines with present molecular modeling methods, let us consider the design of an assembler (http://nano.xerox.com/nanotech/nano4/merklePaper.html). Such a device is able to make copies of itself — hence achieving low cost — and can be programmed to build a wide range of useful structures. The term "assembler" actually encompasses a rather large family of possible designs. For our purposes, we wish to consider the simplest assembler able to achieve certain core objectives: make a copy of itself, and make a wide range of hydrocarbon structures (including diamond and graphite) under program control.

If we are to build an assembler, than at some point we must completely specify it: we must specify the location and element type of every atom. Interestingly enough, it should be possible to design and model such an assembler using computational chemistry software and computing hardware that are either presently available or could reasonably be developed in the near future (a few years). Molecular mechanics and dynamics models would be used to analyze the behavior of the mechanical components, while *ab initio* quantum chemistry models would be used to analyze the reactions involved (e.g., the making and breaking of chemical bonds). Some potential energy functions (such as Brenner's potential) are able to model bond formation and bond breaking. They would be used to do molecular dynamics on the chemical reactions where they were applicable.

Some of the reactions that will likely be involved in the synthesis of diamond have already been modeled. One example is the hydrogen abstraction tool (http://nano.xerox.com/nanotech/Habs/Habs.html), which has been modeled by several groups using both *ab initio* and molecular dynamics methods. Other components have been proposed, discussed, and modeled in varying levels of detail. This process can clearly be extended.

It is useful to emphasize that a design for an assembler is not the same as having an assembler. An assembler can build another assembler, but this presupposes the prior existence of an assembler. We must still build the first one using existing technology. This presents a separate design challenge — but a design challenge that can also be addressed by molecular modeling.

The Alternatives

If we are to develop molecular nanotechnology, it would seem that one of the significant tasks is to systematically investigate the various ways of achieving its basic objectives. Is this a reasonable course of action? Again, using self replication as an example, we need to ask: what are the available alternatives? To date, the only proposals for molecular manufacturing systems involve self replication. The obvious approach is to analyze in greater depth the proposals that have been advanced. If we hesitate to pursue this approach then we should explicitly seek alternatives and then analyze them to see if they are as effective at achieving the desired objectives.

If molecular nanotechnology is feasible within the existing framework of physical law — and that seems to be the predominant opinion — then unless (a) we expect physical law will change or (b) we expect molecular manufacturing systems will be easily developed without great effort, then the obvious strategy is to (c) begin the patient task of exploring the space of possibilities, winnowing out the approaches that either don't work or fail to achieve one (or more) of the objectives, and focus on the approaches that look like they should work. And when we've explored the possibilities, studied the alternatives, determined what is possible and rejected what is impossible — when we can see a clear path from where we are today to where we wish to be in the future — then we can begin in earnest.