

Advice to Aspiring Students (from Engines of Creation 2.0)

By K. Eric Drexler

<http://www.kurzweilai.net/meme/frame.html?main=/articles/art0699.html>

It makes no practical sense to try to build a molecular assembler today. But we can build enabling technologies today, including protein engineering, general macromolecular engineering, and micromanipulation techniques, which will make it easier to build assemblers tomorrow. So, students preparing for a career in nanotech are advised to learn the fundamentals of molecular science and technology.

Originally published in Engines of Creation 2.0, WOWIO LLC, February 2007. Published with permission on KurzweilAI.net March 16, 2007.

Studying Nanotechnology

Many students have asked what they should study to prepare for careers in nanotechnology. Giving a decent answer requires outlining the different fields of research that fall under the nanotechnology umbrella and describing the background knowledge required to work in them. It also seems wise to say something about the different levels of knowledge and modes of learning that are relevant to such a broad, interdisciplinary area. The following is a personal view, based on what I have learned (and wished I had learned), and on how learning in these areas seems to work best. One can't master everything relevant to so broad a field.

Fields of Research

Nanotechnology will mean complete control of the structure of matter, building complex objects with molecular precision. It doesn't exist yet, because we don't have molecular assemblers yet. Work related to nanotechnology accordingly falls into two broad areas: the study of nanotechnology itself (which must remain theoretical, for the time being) and research on enabling technologies leading toward assemblers and nanotechnology (which can be theoretical in part, but which also has an experimental, developmental component).

The theoretical study of nanotechnology involves exploratory engineering work in any of several of areas. It includes basic studies in nanomechanical engineering (the study of molecular machines) and nanoelectrical engineering (the study of molecular and atomically-precise nanometer scale electronic

systems). It also includes studies of complex systems, such as assemblers, replicators, and nanocomputers. More broadly, it includes studies of non-nanoscale applications, such as large systems built by teams of assemblers.

Because we lack the tools to do real nanotechnology today, these theoretical studies amount to building castles in the air. Accordingly, there is little funding for such efforts and frequent skepticism about their value. Nonetheless, such studies can be pursued with intellectual discipline, yielding firm results and a better understanding of our choices as a society. They have been my main focus and have spawned the current interest in nanotechnology—including the interest in giving these theoretical castles hardware foundations.

Studying what can be done with assemblers yields more foresight than it does progress; working to develop assemblers yields more progress than it does foresight. Inevitably, more resources will go into development than into theory, because technology development will yield practical, short-term results on the way to long-term objectives. It makes no practical sense to try to build an assembler today, but it does make sense to build tools today that will make it easier to build assemblers tomorrow. These tools are termed "enabling technologies."

Promising enabling technologies fall into several familiar categories. These include:

- protein engineering (involving efforts to develop techniques for designing molecular devices made of protein)
- general macromolecular engineering (involving efforts to develop techniques for designing and synthesizing molecular devices made of more tractable materials)
- micromanipulation techniques (involving efforts to extend the technology of scanning tunneling and atomic force microscopy to chemical synthesis, and then to the construction of molecular devices)

These approaches have differing strengths and weaknesses. Protein engineering can draw on a host of examples and prototypes from nature, and can exploit existing self-replicating machines (bacteria) to make products cheaply—a major consideration, where short-term payoffs are concerned. General macromolecular engineering avoids the major problem with protein engineering (proteins, not having been designed for designability, are hard to design), but at the cost of moving away from natural prototypes and requiring more expensive chemical synthesis techniques for making near-term products (thus reducing the potential

market). Micromanipulation techniques promise to ease design problems by allowing direct construction of molecular objects, but they suffer from higher costs: a chemical reaction typically makes many trillions of molecules at once, while a manipulator would make but one, hence manipulator-made products can be expected to cost trillions of times more, dramatically reducing the potential market. Also, as of this writing [1988], micromanipulation has not achieved even a single chemically-specific step in molecular synthesis, while chemists have built specific molecules containing thousands of atoms.

All the above areas bear watching, and all will be pursued to some extent, regardless of which ultimately proves to have the biggest payoff. Hybrid approaches, combining techniques from several of these areas (e.g., micromanipulation of molecular tools), seem promising. Finally, improved computational modeling of molecular systems is a generic enabling technology, relevant to all these approaches.

Nanotechnology is fundamentally a branch of engineering.

Background Fields: Molecular Science and Technology

There are, as yet no college curricula aimed specifically at preparing students for work in nanotechnology, although there are numerous new programs that have been created in particular subfields. My own course at Stanford many years ago provided at best an overview of the field. Rather than seeking courses (and books, and journals) in nanotechnology, one should seek courses (etc.) in the broad field of molecular science and technology.

Unfortunately, there are, as yet, few (if any) schools that treat molecular science and technology as a unified field. (A note to curriculum reformers: developing a program having this focus makes sense in terms of current science and technology, and would provide a natural home for early studies in nanotechnology.) Students aiming to gain a solid background in areas important to nanotechnology should be prepared to shop around from department to department. The following section lists some of the important topics and some of the departments in which they are frequently taught.

To study science and technology in a serious way, one must have an adequate background in mathematics. Basic calculus is essential, and differential equations and linear algebra are widely used. Problems in

nanotechnology vary widely in the mathematical sophistication required for their solution.

The study of physical systems is founded on physics. A knowledge of basic classical mechanics and electromagnetism is essential, as is a knowledge of at least the rudiments of quantum mechanics. Anyone aiming to do any sort of sophisticated work in chemistry and molecular machines can benefit from deeper knowledge of quantum mechanics; anyone interested in molecular electronics should make quantum mechanics a chief focus of study. "Quantum mechanics" is a broad area, however. The quantum mechanics of interest here is not quantum electrodynamics, quantum chromodynamics, or superstring theory, but the garden-variety quantum mechanics of electrons in matter, the sort studied by chemists and solid-state physicists. Both quantum chemistry and solid state physics are topics of great relevance to nanotechnology.

As with mathematics, so with physics: problems in nanotechnology vary widely in the sophistication needed for their solution.

Nanomachines and nanoelectronic devices are often greatly influenced by thermal noise. To understand its effects, one needs knowledge of thermodynamics and of statistical mechanics. Thermodynamics deals with the flow of energy and heat in matter in bulk; its principles constrain all physical systems and its subject matter is regarded as a prerequisite for the study of statistical mechanics, which describes much the same territory in a more detailed, molecular fashion. These topics are often taught in chemistry and physics departments.

Nanotechnology can be viewed as an outgrowth of chemistry, the leading science in the field of molecular devices and molecular manipulation. Anyone planning serious work in nanotechnology should seek at least a basic background in chemistry, focusing on its structural, molecular aspects. Those interested in assemblers and molecular mechanical devices should study organic chemistry, and those interested in the chemical-synthesis path to nanotechnology should study synthetic organic chemistry, and learn the arts of the chemistry lab.

Many specific fields have special relevance. Chemical kinetics and reaction transition-state theory is of special relevance to assembler theory. Molecular mechanics is fundamental to any sort of molecular machine design. Studies in materials science (often considered closely allied to chemistry) are also of value; materials scientists consider the mechanical

behavior of larger systems of bonded atoms than chemists typically contemplate.

Biology is the leading science in the study of existing molecular machines. Here, biochemistry is central: enzyme reaction mechanisms provide examples of what many nanomachines will need to do; the folding of proteins and the self-assembly of protein systems provide examples of how complex first generation molecular machines may be made. Familiarity with these fields is of considerable importance to anyone interested in enabling technologies.

Although nanotechnologists will need a thorough grounding in relevant scientific principles, nanotechnology is fundamentally a branch of engineering. To work as an engineer, one must learn to think as an engineer, and that means studying (and doing) design. Nanosystems will be systems, and so the principles of systems engineering apply. Many nanosystems will be mechanical, and so the principles of mechanical engineering apply. Studies in solid mechanics, system dynamics, mechanisms, and control theory all are relevant to both nanotechnology and enabling technologies. Engineering departments often teach more specialized topics of relevance to nanotechnology, such as VLSI circuit design (relevant to nanocomputer design) and microfabrication (relevant to possible enabling technologies). The principles of conventional electronic circuit design are applicable to moderately large nanoelectronic systems, and the principles of quantum electronics are applicable to the smallest systems.

Software systems will be vital to nanotechnology and to enabling technologies along the way. A basic introduction to computers and software will be of value to anyone in any area of science or technology. Those interested in software related to nanotechnology should pay special attention to numerical simulation methods for molecular mechanical and quantum electronic systems, and to the design of programs for highly parallel computer systems, since this is the direction hardware will be moving in the coming years. Parallel systems will help designers develop nanotechnology, and nanocomputers will later be used to build massively parallel (trillion processor and up) computer systems. Finally, if powerful systems are to be useful in molecular design, they will need to be accessible through fast, clean, intuitive interfaces that let designers see and manipulate model molecules.

Learn the fundamentals of molecular science and technology.

Levels of Knowledge

"In short, to do good work in nanotechnology, one must master everything relevant to the physics, chemistry, and engineering of molecules, from quantum mechanics to advanced software architectures." Fortunately, this isn't true. Of course, the more you know, the better you'll do (within limits—studying mustn't completely displace doing), but one can't master everything relevant to so broad a field.

What one can and should do is try to master some areas and know a lot about the others. Real molecular devices can do many different things: they can vibrate, pull apart, shake apart, deform, transform, photolyse, or pop from state to state—any of these behaviors can occur in a simple mechanical part, and any can make it fail. Real physical systems will do something when used, and if what they will do is strikingly different from what you think they will do, then the work you're doing may be a waste of time for you and for anyone who listens to you. It's much better to be right about what will work, and this means knowing enough to steer clear of potential problems.

It makes sense to think in terms of three levels of knowledge about a field:

1. Knowing what a field is about—knowing what sorts of physical systems and phenomena it deals with, and what sorts of questions it asks and answers.
2. Knowing the content of a field in a qualitative sense—having a good feel for what sorts of phenomena can be important in what circumstances, and knowing when you need answers from work in that field.
3. Knowing how to get those answers yourself, based on personal mastery of enough of the field's subject matter.

If one has enough knowledge at levels (1) and (2) in enough fields, then one can steer clear of problems in those fields while doing work in a related field where you have knowledge at level (3). And this is a good thing, because knowledge at levels (1) and (2) takes far less time to acquire. But to make proper use of knowledge at levels (1) and (2) requires a harsh discipline: attempt to assume the worst about what you don't know. Don't assume that a poorly-understood physical effect will somehow save your design; do assume (until finding otherwise) that it may utterly ruin it. Without this discipline, you'll become an intellectual hazard. With it, you'll be able to make a real contribution.

Ideas about real systems must somehow be disciplined by reality.

Modes of Learning

How can one get this sort of general knowledge of a field? Courses can help, but they tend to focus on mastery of a narrow range of knowledge, rather than familiarity with a wide range of knowledge. One can gain this familiarity by reading magazines and journals that offer broad coverage of science and technology: good choices include *Science*, *Nature*, *Science News*, *Scientific American*, and *IEEE Spectrum*. Another good tactic is to skim a wide range of books on the new books shelf of a science library, on a regular basis, and to do likewise with a wide range of technical journals.

To do all this properly requires the discipline to read what you don't understand—despite the school-induced reflex which says "Oh, no! I don't understand, so I'll fail the test—maybe I should drop this subject!" By reading what you don't understand, you gain a sense of the patterns of the field—the terms and abstract relationships, the kinds of problems being addressed, and the kinds of knowledge required to understand more. And this adds up to an important sort of understanding. Later, this familiarity makes it much easier to consult the literature: one knows which disciplines deal with what problems, and what one needs to study to gain a deeper understanding. Also, it fills your mind with questions, so that you can later recognize the answers and have your mind seize them more firmly.

For a thorough grounding in a basic field, classes can be excellent. If classes aren't available, textbooks can often serve well, especially if you work many of the problems.

In any evolving, interdisciplinary field, you must learn to learn from books and journals. Learn to use libraries (as horrible as they are, compared to more accessible internet publishing systems). Learn to read skeptically—it is a rare book or journal that doesn't have a few serious errors, and occasionally publications are utter bilge, especially in interdisciplinary fields (which too often lack any discipline at all).

Finally, tackle problems. If you can find a professor doing good, interesting work, consider becoming an apprentice researcher. If not (or in addition), pursue technical problems that interest you. The best way to learn is to seek answers to questions that interest you,

and there is no other way to make an original contribution.

Learn to criticize ideas, especially your own. Most new ideas are wrong or inadequate. If you don't reject most of your ideas promptly, then you're almost surely fooling yourself, and if you also spread them, you're almost surely polluting the intellectual world. But if an idea really seems to stand up under testing, try filling in more details, and criticizing it again.

Get criticism from others. Learn to present ideas in discussions, papers, and talks, and listen to the responses, especially from people who know relevant fields. If they disbelieve your idea and tell you why, either understand and refute their criticism, or consider working on a different idea. If they look at you oddly and change the subject, consider whether you are perhaps overlooking a really big, basic problem—are you really familiar with the relevant fields? If they disbelieve you at first, but can be persuaded, congratulations! You've probably got hold of something interesting, perhaps even new and important.

Always remember that ideas about real systems must somehow be disciplined by reality. Experimental work brings its own discipline from nature, if the experimenter uses good technique. This discipline is direct and hard to escape. Theoretical work, in contrast, must be disciplined by knowledge of experimental results and natural law; this discipline doesn't impose itself, it must be sought out and largely self-applied. To be a careful thinker, try to understand things in more than one way: if you get the same answer from physical calculations and from analogies to known machines and from analogies to biology, then you're probably right. If all you have is a rough analogy or a crude calculation, you may well be wrong.

Seek out weaknesses in ideas, and build only on ideas that pass rigorous tests, or you may see the foundations of your thinking later crumble and dump a year's (or a decade's) work into the trash. Beware of those who have neither experimental results nor a theoretician's voluntary discipline; expect them to spout great streams of plausible nonsense, unconstrained by reality. Don't become one of these, even if you find that many (ignorant) people are intrigued and entertained by your wilder imaginings.

In short, learn the fundamentals of molecular science and technology. Survey other relevant knowledge. Learn to learn from books and journals. Pursue problems, think critically, and learn more. Design and calculate or experiment. Publish your contribution and add to the world's knowledge.