CS-480—Senior Seminar
Synopsis: *Gray Goo*
Spring '04

Jim Rogers

28/Jan

Primary Papers

Peterson, Christine "Taking technology to the molecular level." *Computer*, Vol. 33, No. 1, pp. 46–53, Jan., 2000.

Joy, Bill "Why the future doesn't need us." Wired, (date unknown) `<http://www.wired.com/wired/archive/8.04/joy_pr.html>`, accessed 19/Jan, 2004. Also in *Wired*, Vol. 8, No. 4, (Page numbers missing), Apr, 2000.

Baum, Rudy (with Eric Drexler and Richard Smalley) "Point-counterpoint: Nanotechnology" ACS, (date unknown) `<http://pubs.acs.org/cen/coverstory/8148/8148counterpoint.html>`, accessed 20/Jan, 2004. Also in *Chemical and Engineering News*, Vol. 81, No. 8, pp. 37–42, 1/Dec, 2003.

Dyson, Freeman J. "The future needs us! (Review of *Prey* by Michael Crichton.)" *New York Review of Books*, Vol. 50, No. 2, pp. 11–13, 13/Feb,2003.

Additional Papers

Drexler, K. Eric "Machine-Phase Nanotechnology." *Scientific American*, Vol. 285, No. 3, pp. 74–75, Sept,2001.

Smalley, Richard E. "Of Chemistry, Love and Nanobots." *Scientific American*, Vol. 285, No. 3, pp. 76–77, Sept,2001.

Peterson—Technology at the molecular level, and Joy—The future doesn't need us.

Peterson provides an introduction to the field of nanotechnology, molecular-scale structures built by placing atoms one by one at specific bonding sites. This is in contrast to bulk technologies such as MicroElectroMechanical Systems, in which structures are constructed from bulk crystals using more or less standard IC processing techniques. MEMS structures are larger than nanotechnology structures by a factor of about 10^9 in volume and, being bulk structures, can be tolerant of molecular defects. In nanotechnology, where the unit of construction is individual atoms, such defects are likely to compromise the structure.

The practical distinction between the two is that MEMS is accomplished technology; MEMS devices are actually in production. Nanotechnology beyond the placement of a few dozens of isolated atoms on a crystalline substrate, exists only as simulated structures built using molecular modeling software. This is one of the ways that CS comes to play. It is the utility of such software that raised the prospect of precise atomic constructions to begin with.

The main technical difficulty for nanotechnology, in Peterson's assessment, is thermal noise—at the atomic level molecules are in constant motion, making precise location difficult at best. The general approach to overcoming that challenge has been to make nanomachines extremely stiff, limiting the relative movement within the structure. This is in contrast to the only known functioning precise molecular-level assembly mechanisms, biological systems, which tend to be "floppy" and depend on thermal noise to drive the diffusion on which the mechanisms are based. (Drexler refers to these as "solvent phase" systems.)

A second technical difficulty is the problem of scale. Even to build a MEMS-scale device atom-by-atom would require the precise placement of billions of atoms. If this is going to be feasible, such systems will have to be extremely fast or highly parallel. This is why self-replication is seen to be critical to nanotechnology. The idea is to build molecular assemblers which are capable of building replicas of themselves. These are then allowed to reproduce for some reasonable number of generations. As the population size grows exponentially, the result can be a very large number of assemblers, which can then turn to their intended task.

It is this prospect of potentially self-sustaining replication that drives Joy's concerns. These are based primarily on Drexler's warnings about unconstrained replication of nanobots. Exponential growth is inevitably stunningly prolific. Without safeguards in place, a scenario in which the world is overcome by a "gray goo" of multiplying nanobots is not impossible. The prospect invokes parallels to ecological systems in which the nanobots might have the role of essentially indestructible organisms with unlimited reproductive success.

Peterson, quoting Drexler, points out that such a scenario could be made impossible by adopting safeguards early on. One such safeguard would be to make replication "in vivo" impossible by making it depend on substances unavailable outside a controlled environment. Another possibility would be to make the replication process highly sensitive to errors, making self-sustained replication difficult to achieve.

These controls, as with the scenario itself, seem to be based on parallels with the biotechnology. Peterson and Drexler take the experience with biotechnology to be a model for developing safeguards for nanotechnology. Joy is less sanguine, as are we. We noted that early safeguards analogous to those proposed by Drexler, in particular limiting genetic engineering to organisms that could not survive outside the lab, were abandoned because they are inconsistent with commercial exploitation of the technology. At this point it is impossible to tell what safeguards are or are not in place for genetic engineering because such information is now classified as trade secrets.

We also discussed the possibility of a scenario developing in which the parallel with biological systems extended to evolution. What is to prevent random variations in the structure of nanobots (including, presumably, any built-in programs) from producing nanobots that were marginally more efficient than their siblings. Were the mutation inherited by the offspring of that nanobot, even a single mutated individual might rapidly come to completely dominate the population. If the mutation did not interfere with the intended function of the devices, this might be harmless. On the other hand, if all it did was prevent the intended function from being carried out, it could be catastrophic. Since it seems unlikely that the intended function would be an evolutionary benefit (indeed, one would presume it would include self-destruction on completion) advantageous mutations might well interfere with it.

In this way evolution might not be necessarily limited to biological organisms. While the coincidence of both increased efficiency and inheritability may be unlikely in the extreme, the fact that evolutionary-model process are currently exploited in some software systems makes it not implausible that it might be designed in. Moreover, as with biological evolution, the staggering size of the population makes even exceedingly rare events reasonably likely.

Drexler and Smalley—Point/counter-point

This is a series of open letters written by two of the major players in the nanotechnology field, presumably debating the feasibility of the "molecular assemblers" on which the self-replication model is based.

Smalley points out that the fact that one is working in the arena of physical chemistry makes macroscopic scale mechanisms useless as models. The assemblers themselves must be molecular mechanisms subject to the same laws as the devices they are assembling. This leads to two problems: the sticky fingers problem—since picking an atom up is a matter of atomic bonding, what then breaks that bond—and the fat fingers problem—the devices which will be used to precisely position atoms, being themselves atomic level devices, will need devices to precisely position the atoms that comprise them, and so on. It is not at all clear that there is actually room for all this structure without blocking the paths to the active sites.

Drexler claims that these finger-based arguments are irrelevant, that no proposed system uses what he terms "Smalley fingers". The editor for *C&E News* has included an image, courtesy of Drexler, which could be taken to show how the sticky fingers problem might be solved. The mechanism, however, certainly seems to have the form of a "finger" and, beyond the active point of the device, is rendered without detail, suggesting the kind of bulk material that one might assume from macro-scale models (and certainly begging the fat fingers question).

The initial observation of the class was that neither Drexler nor Smalley seem to be listening to each other. As the image might suggest, neither of them seems to be addressing the technical issues that the other is raising. In fact, the debate seems to be a much more personal one. Drexler accuses Smalley of dismissing his work by misrepresenting it. Smalley accuses Drexler of scaring the children with disaster scenarios based on science fiction.

A not entirely charitable reading might see Drexler aiming to attract funding for his research on the grounds that he is addressing the potential risks of the technology responsibly (presumably in contrast to others in the field) and Smalley accusing Drexler of scaring away funding for the entire field on the basis of hopelessly overstated claims (presumably in contrast to Smalley's own claims). A more charitable reading might see the debate as an example of what Peterson refers to a "disciplinary disconnect" with Drexler taking the more abstract, system level view and exuberant rhetoric characteristic of CS and Smalley taking the more concrete, detail focused view and more modest rhetoric characteristic of Chemistry.

Dyson—Future needs us

Dyson's article is a review of an actual work of science fiction, Michael Crichton's *Prey*, in which he imagines something like Drexler's doomsday scenario. He points out some reassuring physical inconsistencies in Crichton's story, in particular the fact that speed through a fluid scales with size, leaving nanobots unable to exceed about one-tenth of an inch per second. It was pointed out, however, in our discussion that larger organisms can be regarded as conglomerations of nanobot-scale organisms, and that these have no problem exceeding that speed. But, as it is not likely that the gray goo scenario will depend on nanobots having the characteristics imagined by Crichton, these questions are not really relevant.

The bulk of the review, however, is dedicated to a response to Bill Joy's call for specific measures to limit the risk of nanotechnology. While Dyson seems to be happy with the idea of internationally coordinated safeguards—indeed, like Peterson, he views the biotechnology experience as a reassuring historical precedent—he takes issue with Joy's call for restrictions on the dissemination of knowledge on which nanotechnology is based. Drawing parallels to attempts to control the dissemination of dangerous ideas in the 17th century by limiting the publication of books, he questions whether knowledge, in itself, can be dangerous and whether it is possible to relinquish the pursuit of knowledge in any case. As was pointed out in class, restriction of information about technology can prevent the development of effective safeguards. The experience of the cryptology community serves as an example. The only reliably effective way of insuring the security of an encryption algorithm is to expose the algorithm to public scrutiny. Perhaps we would be as reassured by the safeguards on genetic engineering technology if those safeguards had not be moved into the realm of secret information.