

A conversation with Dr. Eric Drexler, January 23, 2015

Participants

- Dr. Eric Drexler – Academic Visitor, Oxford Martin Programme on the Impacts of Future Technology, Oxford Martin School, University of Oxford
- Dario Amodei – Scientific Advisor, Open Philanthropy Project
- Nick Beckstead – Research Analyst, Open Philanthropy Project

Note: These notes were compiled by the Open Philanthropy Project and give an overview of the major points made by Dr. Eric Drexler.

Summary

The Open Philanthropy Project spoke with Dr. Drexler of as part of its investigation into atomically precise manufacturing (APM). Conversation topics included the feasibility of APM, potential development timelines for APM, potential applications of APM, and options for minimizing potential risks from APM.

Atomically precise manufacturing: distinguishing concepts

Atomically precise manufacturing (APM), is a proposed technology that would be capable of processing simple and inexpensive chemical feedstocks (which would be molecules composed of elements from the upper-right hand corner of the periodic table, including carbon, oxygen, nitrogen, and silicon) into an extraordinarily wide range of high-performance products built with atomic precision by means of arrays of nanoscale devices that guide the motion of reactive molecules. "APM" is roughly synonymous with the older term "molecular manufacturing," and is often associated with "molecular nanotechnology" (a broad and less well defined concept), or "nanotechnology," a term that now often refers to substantially unrelated areas of materials science and nanoscale device fabrication.

It is important to distinguish between manufacturing capabilities discussed:

- In Dr. Drexler's first book, *Engines of Creation*.
- In Dr. Drexler's *Nanosystems*.
- By other proponents of atomically precise manufacturing.

The concepts are overlapping and related, and sometimes a critique of one can be taken as a critique of another, even in cases where that is unjustified.

Engines of Creation aimed to estimate a boundary between what seems likely to be possible in the future and what does not, considering everything that could be made by manufacturing systems that could be made by manufacturing systems...that could be made by manufacturing systems that could be made today. In comparison with *Nanosystems*, it operated with a less conservative standard of proof and discussed a wider range of possibilities.

In contrast, *Nanosystems* aimed to argue more conservatively, concretely, and persuasively for the in-principle feasibility of a range of manufacturing systems with a smaller range of input materials and potential applications.

Dr. Drexler notes that although he used the term "universal assembler" in a section heading in *Engines of Creation*, he did not argue that assemblers could be universal in the strong sense that a Turing machine is universal, and noted that assemblers "will not be able to build everything that could exist". However, some people (such as Dr. Richard Smalley) represented and criticized proposals for mechanically guided assembly as if these called for devices with impossibly strong universality. To be clear, atomically precise manufacturing does not include the concept of a "universal assembler" capable of making any possible object.

Development pathway

Dr. Drexler envisions developing a mature form of atomically precise manufacturing through a series of steps that begins with self-assembling, biomolecular, atomically precise "3D printers" operating in solution and constructing objects by guiding the bonding of biomolecular or other well-defined nanoscale building blocks. Dr. Drexler suggests that such devices could be used to construct more capable devices, leading eventually to nanofactories operating in a dry environment and manufacturing objects that consist of hard materials. This more mature form of atomically precise manufacturing would use small-molecule feedstocks (i.e., molecules composed of elements from the upper-right hand corner of the periodic table such as hydrogen, carbon, oxygen, nitrogen, and silicon) and could, in terms of performance, offer a superset of the range of artifacts that can be made by modern industry. Both proposed kinds of technology are described in greater detail below.

Feasibility of a self-assembling, biomolecular, nano-resolution 3D printer

Dr. Drexler has a concept for a self-assembling, biomolecular, nano-resolution 3D printer operating in solution. The active head of this device (analogous to a printhead) could be moved by linear stepper motors with displacement increments of about a nanometer, operating along three axes and controlled by external optical inputs. Such a device could also be accurate to a resolution of a nanometer (though the system's components would not be stiff enough to enable accurate positioning at the small-molecule length scale, due to thermal fluctuations), and would construct objects out of biomolecular materials. One way the active head of a 3D printer could work is by removing protective groups from active sites on a surface, allowing the feedstock materials in the solution to bind to the selected sites, transported by Brownian motion. The active head would not "pick up and place" molecules (another common misinterpretation of Dr. Drexler's ideas that became a target for many critics).

Applications of the biomolecular 3D printer

Early stages of this kind of biomolecular 3D printer will likely be less important than work in synthetic biology, but in the long-term the impact of the hard, more advanced manufacturing systems it enables could potentially be greater.

One distinct aspect of this research direction over applications of synthetic biology is that stepper motors can be directed to move a specified number of steps, offering precise control of a kind that is not possible for motor proteins (e.g., myosin) transporting materials along protein (e.g., actin) filaments.

These products might include more complex structures than are currently produced in synthetic biology and DNA nanotechnology, including:

- Structures incorporating a greater diversity of active molecules.
- Packages with more fine-grained differential reactions to input from the cellular environment.
- Packages able to remain active inside a cell and interact with the cytoplasmic RNA content, modifying the cell's state without stimulating the cell's defensive autophagy and achieving results similar to genetic engineering without altering the genome.

Progress in this direction could also produce better instrumentation for monitoring tests of these products.

"Layered" packages

By guiding the assembly of molecular building blocks, these biomolecular "3D printers" could create layered packages, capable of interfacing with the cellular environment via sensors and actuators, and a protected inner mechanism to perform computations. Such packages could be safely delivered to cells, perhaps using next-generation exosome technologies, carrying a payload able to sense the state of the cell (e.g. recognizing the active messenger RNA in the cell) through binding and unbinding molecules, and applying present DNA-based computational methods to responding differentially (e.g. synthesizing a nucleic acid, exposing an active site that causes degradation of a particular molecule).

The consequences of such packages would be similar to an advance in small interfering RNA (siRNA) technology, and would avoid the delivery issues that siRNA faces. Layered packages may not have much impact on the issue of off-target effects (although Dr. Drexler suggests that this research direction could potentially create more discriminating packages than siRNA).

These results could in principle be achieved through other techniques (e.g., RNA interference or CRISPR), but this kind of biomolecular 3D printer would significantly shorten development cycle times, speeding up innovation in the area.

Comparison to rational drug design

In drug design, trial-and-error methods tend to have more success than rational drug design using computational chemistry. Dr. Drexler suggests that that this kind of biomolecular 3D printer could significantly improve the effectiveness of computational design approaches by allowing researchers to design, easily customize, and almost immediately test new chemical packages.

Nanosystems with hard materials and their applications

Dr. Drexler describes a continuous gradient of technological advances between in-solution, entirely self-assembling “soft” nanotechnology (in which the design aspect is difficult, while actual fabrication is relatively simple, exploiting current techniques for macromolecular synthesis and self-assembly) and “hard” nanotechnology (for which design is simple, though fabrication is not yet possible). Some discussion of this technological gradient is included toward the end of *Nanosystems* and in an appendix of *Radical Abundance*.

At a more mature stage in the development of APM, Dr. Drexler envisions desktop-sized (or larger), programmable “nanofactories” which would use earth-abundant materials (e.g., molecules composed of elements from the upper-right hand corner of the periodic table, including carbon, oxygen, nitrogen, and silicon) as inputs/feedstocks, and use arrays of molecule-binding nanoscale devices (using gearboxes, motors, and so on to implement positioning/transport mechanisms) to guide the motion of reactive molecules from the feedstocks to assemble objects and machines defined by data files (as in 3D printers today). Such a nanofactory would be capable of placing its inputs in controlled configurations in controlled sequences, providing a degree of control of chemical synthesis that cannot be achieved by means that rely on the diffusion of molecules in solution. Drexler suggests that systems of this general kind could produce a superset of the range of products that can be made by modern industry. The physical principles, mechanisms, and potential system architectures of such nanofactories are examined in greater detail in *Nanosystems*.

Many useful engineering materials are high-strength, high-stiffness, and low-density. These properties could be significantly improved, beyond current capabilities, by nanofactory-based production of strong covalent solids (e.g., diamond, graphene, silicon, aluminum oxide, silicon dioxide). Metals, semiconductors, photovoltaics, and strong, insulating solids could be manufactured by nanofactories using the earth-abundant materials described above. Atomic precision would be especially beneficial for semiconductors, and for covalent solids, which are weakened by flaws in their molecular structure.

Dr. Drexler does not expect such nanofactories to be developed in the next 10 years, though it could possibly be achieved in that time given a well-organized, properly focused research effort. Some aspects of such a project might run counter to the intuitions of chemists and biologists, and would require a multidisciplinary systems-engineering style of organization that is not found in small laboratories pursuing research in specialties within the molecular sciences.

Returning to biomedical applications, Drexler suggests that at this level of fabrication technology, it would become possible to build devices on a submicron scale with a sufficient computational capacity to perform complex operations in a mammalian cellular environment (e.g., a 1 GHz, 32-bit computer with a few megabytes of storage can be made as small as a cubic micron, which is approximately one one-thousandth the volume of a human cell).

Computers and chips

Modern computer chips use a wide range of chemical elements (for example, about one-third of all non-radioactive elements are used in Intel's latest processor), but these are often chosen to enable present day fabrication technologies (e.g., diffusion barriers to avoid mixing of adjacent materials during high-temperature processing, low-melting point solders made of complex mixtures of element to depress freezing points, and so on). Dr. Drexler believes that a set of far fewer elements and compounds would be sufficient to produce the needed conductors (e.g., carbon nanotubes, copper), semiconductors (e.g., silicon, silicon carbide, diamond) and insulators for comparable but smaller scale chip technologies. There is, however, no need to restrict the elements used to a small set. Transition metal oxides, for example, could be used in chips if the elements were available as inputs (typically in the form of simple molecules containing the desired elements).

Transition from solution-based, soft nanosystems to dry, hard nanofactories

In a nanofactory that does not transport its feedstock materials in solution, transport of molecules must instead be done mechanically, and this would require considerable complexity and progress along the technological gradient mentioned above. Working in a dry environment would require applying other methods to perform the functions currently performed by solvents. For example, via Brownian motion, solvents provide transport without requiring transport structures, as well as motion with six degrees of freedom (though it cannot be controlled). An intermediate step towards dry, hard nanofactories would involve mechanically controlling the position of materials along spatial axes, and using Brownian motion to achieve correct orientation.

On the path toward APM, it will likely be most effective to take advantage of the benefits of solvents for as long as possible, because of the complexity and hence difficulty of performing these functions by mechanical transport.

Dr. Drexler has not written extensively about the transition from solvent to non-solvent APM in particular, but remarks that mechanical transport can be introduced in a solution environment, and then use of solvents can be reduced as convenient.

Specialization and generalization of nanofactories

For the sake of efficiency, the mechanisms used in nanofactories would need to be specialized to handle the particular chemical makeup of different input materials. Specialization of nanofactory processing systems would be necessary to enable efficient and high-throughput APM production systems.

However, macro-scale nanofactories could be made general-purpose by incorporating many specialized nanoscale subsystems. Such a nanofactory could use a wide range of input materials and produce a wide range of products, including the components of more nanofactories. The nanofactory would be limited to products specified in its programming language.

Although it would be most convenient to use separated stocks of specific input materials for a nanofactory, even people living in areas without industrial infrastructure (such as rural regions of Africa) could use the earth-abundant materials from their environment as feedstocks for nanofactories. A chemical processor could extract the desired elements from feedstocks containing them and put them in water solution as inputs to nanofactory units. A nanofactory could use atmospheric nitrogen, hydrogen, oxygen and carbon as inputs.

These operations would of course require substantial energy inputs, which Drexler suggests could be small multiples (a factor of five or ten?) of the thermodynamic energy requirements. These requirements are on the same chemical energy scale as conventional operations, such as production of metals from oxide ores.

Building a nanofactory using another nanofactory

A nanofactory could manufacture another nanofactory, most likely by producing several separate pieces that would fit together into a new nanofactory. Dr. Ralph Merkle has designed one potential method for this.

On the smallest scale inside a nanofactory, nanoscale mechanisms of tens to hundreds of nanometers in size would transport molecular building blocks at a speed in the range of centimeters per second. These basic blocks would be brought together to react and form larger components.

At a transport speed of 10 cm/s, with parts spaced 1 nm/apart, a molecular processing device in a nanofactory could perform about 100 million reactions per second. A system able to construct macroscale products would require large arrays of processing subsystems that contain long chains of such devices, and would require sequences of larger devices to combine small atomically precise building blocks to form larger components. Making allowance for all this, and de-rating the small devices by a factor of 100, to 1 million operations per second, would still allow a nanofactory to build enough parts to build another meter-scale nanofactory in about 1,000 seconds, given adequate energy inputs and cooling.

Preparation and Safety

There are useful preparatory steps that could be taken in parallel with development work on paths toward APM-level fabrication technologies, including:

- Developing a design language and compiler for APM devices.
- Applying quantum-chemistry modeling to explore relevant mechanically constrained transformations of molecules.
- Applying computer modeling to design larger molecular assemblies, creating an early “parts catalogue” for APM products.
- Creating concrete designs for potential APM products (e.g., computers, solar cells, etc.).

Safety

There is no technical research agenda for the safe development of APM, such as exists for, e.g., artificial intelligence (AI). Dr. Drexler suggests that the nature of the technologies (essentially small-scale chemistry and mechanical devices) creates no risk from large scale unintended physical consequences of APM. In particular the popular “grey goo” scenario involving self-replicating, organism-like nano-structures has nothing to do with factory-style machinery used to implement APM systems. Dangerous products could be made with APM, but would have to be manufactured intentionally.

The most effective way to ensure that nanofactories are not used to make dangerous products is to build nanofactories that are only capable of producing a narrow range of products (“restricted” nanofactories), which could potentially be expanded carefully over time. Ensuring that only safe nanofactories are created would be mainly an institutional problem. It would be very difficult to manually alter an appropriately designed restricted nanofactory to create a nanofactory capable of making dangerous weapons without having tools with the same potentially concerning capabilities as a general-purpose nanofactory.

Potential focus areas for an APM policy community

It would be beneficial for the conversation on APM policy and strategy to start relatively early, though it is hard to imagine the ideas gaining political traction without first having a more widespread understanding of APM and a positive assessment of its feasibility of APM. In a context with this basic understanding, topics to consider might include:

Arms control

A community focused on APM-related policy could work to explore potential frameworks for arms control and how they could be implemented. Dr. Drexler believes that, to avoid needless risks, nations would be well advised to pursue cooperative and transparent APM development programs. Access to nanofactories could likely be controlled by policies like the ones we use for nuclear weapons, but current policy does not offer an effective way to restrict independent development of nanofactories. “Nano-solutions for the 21st Century,” which Dr. Drexler co-authored with Dennis Pamlin, is the best resource on this issue, and includes a comparative analysis of nuclear and APM arms control considerations. *Radical Abundance* also addresses this issue to some extent.

While arms control is important, Dr. Drexler is concerned that an arms-control centered discussion of APM-related policy would frame the technology inappropriately, when it would be better to focus on potential multilaterally beneficial, problem-solving applications of APM in areas such as energy, climate change, agriculture, clean water, medicine, and computation. Dr. Drexler thinks it would be beneficial for people with experience in these areas to be involved in the early stages of APM policy.

He notes that APM applications in these areas (particularly in resources, climate, and economic development) could reduce at least some of the major pressures that

are now expected to drive 21st century conflicts, and that this situation reframes basic questions of national interest. He argues that the temptation to consider APM in the context of conflicts that it will make obsolete could drive incoherent and needlessly risky national policies, compared to a more integrated and coherent approach to understanding the prospects and their implications for national interest.

Economic impact

The policy community could also explore the potential economic impact of APM, especially its effect on employment and which jobs humans would likely still perform after APM becomes available. Similar economic questions arise in the context of AI development and general developments in robotics.

Impact on AI development

APM may impact AI development by dramatically increasing computational resources. For example, APM could potentially produce arrays billions of 1 GHz processors, each a few cubic microns in size and requiring less than one microwatt of power, at an extremely low cost. Conservatively, Dr. Drexler estimates that APM could reduce the power required for equivalent computation by a factor of roughly a million. The energy consumption required for computation in such devices may conceivably approach the Landauer limit. The main limit in constructing large, fast arrays would be adequate cooling.

Potential impact of increased surveillance

Dr. Drexler thinks APM will likely extend the trends toward dense surveillance that are already being driven by conventional nanoscale electronics. A general exploration of policies intended to ensure good quality of life in a world with dense surveillance could have implications for APM policy.

*All Open Philanthropy Project conversations are available at
<http://www.givewell.org/conversations>*