An Interview with Bill Joy

Casey Walker: Will you describe nanotechnology and molecular electronics and what can be done with them?

Bill Joy: Nanotechnology is any engineering technology applied to the scale of a nanometer–it can be at the atomic scale and can be with organic or inorganic material. Nanotechnology enables us to manipulate or create any substance that you can imagine. Molecular electronics works much the same way as the electronics we use today, but at a scale that is about three thousand times smaller.

Today, we can certainly recognize that science has created powers that are perhaps beyond our ability to manage well. We had some of that same sense with nuclear power and nuclear bombs, but the particular danger with these new technologies is that many people, perhaps almost everybody, has the ability to use these new technologies to make living or nonliving things with far-reaching consequences. This is something new. We've had a confluence, really, of these incredibly transformative, incredibly powerful technologies with the democratic notion that goes with these technologies being available as information-based technologies. So, for example, genetics is more and more about computers manipulating and helping us to understand a lot of the effect of the transformation that we do. Today we still need to do things using laboratory experiments, but over time, the dream of biologists is to genetically engineer computationally. Nanotech is very much a computational science and robotics certainly is. As these things become more and more computational, the barrier to entry, if you will, for people to make things gets to be very low, and finally disappears completely.

Besides the order of magnitude, and the accessibility, there's also the problem of people thinking that they're in control. Will you describe the problem of these technologies in terms of self-replication?

Genetic and nanotech are part of a spectrum of technologies that can be used to make things that would self-replicate much like a natural disease. Take a flu, for example. Someone sneezes and passes it on to somebody else. The flu basically commandeers part of our bodies and replicates itself. But the flu bug is directing that replication, so, essentially, once it goes from one person to another, it can make more of itself. That's what I mean by self-replication. This is a far different scenario from people in a factory spitting out a bomb that someone else would then set off. These kinds of technologies are like infections spreading on their own. Nanotechnology might be used to make a little machine that could manage itself in the physical world and then make more of itself. And, if you had a robot that could replicate itself, it would be almost like a wild species such as a rabbit. But once we have something that can make more of itself, we have the possibility that it will just spread until it reaches some sort of limit. Just as you can have a flu that becomes epidemic or pandemic, you could also have a creation that is suddenly impossible to get rid of. The destructive technologies that we've had historically were not of this character. Once a person can release something that can self-replicate, the ability to do harm can score far beyond the scale of the initial act.

BILL JOY *is chief scientist of Sun Microsystems, and is a co-founder of the company as well as a member of the executive committee. In 1997, Joy was appointed by President Clinton as Co-Chairman of the Presidential Information Technology Advisory Committee, which provided guidance and advice on all areas of high-performance computing, communications, and information technologies to accelerate development and adoption of infrormation technologies. His cover article, "Why the Future Doesn't Need Us," in"* Wired Magazine, *April 2000, ignited worldwide discussion on the dangers from and ethical issues in the development of new, transformative technologies.*

When you write of Drexler's gray goo or Frank Herbert's white plague, you've also make the point that the original intention may not have been evil, but may well begin from good intentions.

That's correct. We can have disastrous results from bad or good intent and, really, from anyone. It could come from a nation, a company, a group, or an individual, and it could start as an accident, or even from a couple of benign things that come together in an unexpected way. Once we have widespread practices with technologies that are very, very powerful, we're likely to have some accidents. We've been very lucky with nuclear power in terms of not having more accidents, but that's partially because they've been managed very carefully. Management is a key problem with these new technologies.

What can be done now in terms of management?

There are two frames of mind at work. One says, well there's nothing we can do about it, and we should just invest a lot more money in all this technology and try to find solutions to the problems each creates. The other frame of mind says, maybe this isn't the path we should move down and we shouldn't give everybody this kind of power. Now the first path is superficially attractive, because it certainly lets us stay with the whole system that we have of a democratic, unlimited individual pursuit of innovation for these technologies. But I think it ignores the evidence that there are offensive or inappropriate uses of these technologies. For example, in the case of nuclear capability, offense has real advantage over defense for the simple reason that defense has to be perfect. In the case of biotechnologies you cannot defend yourself against all viruses with a single action, you have to stop every single one. And defending against them all, fundamentally, is like having a perfect immune system that can deal not only with what is known but unknown. Impossible.

We're now facing a historical situation. We have a convergence of problems. One kind of problem results unintentionally when our individual behaviors are rational and reasonable but the collective sum of that individual behavior produces an unacceptable outcome. We see this today in the environmental problem. Everybody's consuming a certain amount of materials and energy, but there are a lot of us and it adds up to a whole heck of a lot for the atmosphere and other species' extinctions. But, as much trouble as we're having with that problem, and as important as that problem is, it's also an indirect threat from our collective behavior. Now, with these new technologies, we have as large a scale of threat in terms of global effect, but it presents itself as a direct threat from individual behavior. In this sense our problems with new technologies are more dire. I'm personally not so concerned about some sort of a Mission Impossible or James Bond megalomaniac, as much as I am about some sort of normal business venture that has an unforeseen outcome that is disastrous. And it's those kinds of things that we can't as readily address in terms of policing or managing.

Do you see a political movement capable of constraining the entire sector of technological development?

That would be very rational. Historically, scientists have rejected constraints on the theory that pure science and science in general was good, and people doing pure science shouldn't have to think about social or political concerns. The line between pure and applied science is becoming very blurry, with universities getting involved with companies, and most everything these days is being pursued with much more of an eye toward commercial application. So I think that argument is becoming less and less valid. I like to say science was originally a branch of philosophy, and it's only the modern experimental science with testing hypotheses that we've become divorced from ethical concerns. I think that as uncomfortable as it may be, we have to look at where we are going. If we want to go to a world where everything is possible, then many bad things are going to be easier too. I'm not sure if we'd collectively choose to go there. So my preference is that we look at this larger picture but we don't really have the mechanisms institutionally. We'll have to develop some new mechanisms.

Technological evolution is threatening to take over from what we used to think of as cultural evolution and moving at a rate of about one thousand times faster than cultural evolution. The danger is that the mechanisms that we have in our society for making policy decisions and coming to collective agreements, for culturally expressing some wisdom about these things, is not running at the same speed. How do we respond? If we can all agree that we can get to wherever we want to go and don't need to rush, then maybe the aspects of danger and recklessness can be eliminated. Unfortunately, science and technology is almost a religion. We have to get some control over its ultimate direction.

It's also awfully arrogant to think that we're going to design a new post-biological world. It's fanciful to think that we're going to create some sort of improved silicon-based human and that it's going to be anything at all like us. If we create silicon life forms and let evolution go, which is a very natural process, it won't be human for long. To think that we're going to make humans in a new and improved way seems very unlikely. That doesn't mean we couldn't extend our life-spans substantially in our bodies as we know them. I think we just need to proceed with extreme caution, and we seem to be at the opposite end of the caution spectrum at the moment.

It's a real challenge for us to think on the kind of scale we need to be thinking on. Our humility should equal the danger before us especially when dealing with systems that we understand as little as we do.

Originally published in *Wild Duck Review* Vol. VI No.1 on "The End of Human Nature?"

A complete index of pdf downloads is available at: InstituteforInquiry.org.