

NEWS

Making Machines That Make Others of Their Kind

For decades, self-replicating robots have been a roboticist's dream—and a science-fiction writer's nightmare. Yet engineers haven't found a way to create 'bots that beget 'bots

As any sci-fi fan knows, monkeying with robots ultimately leads to mass carnage. From *R.U.R. (Rossum's Universal Robots)*, the play that in 1921 introduced the word "robot," to the battles with the Daleks in the television show *Doctor Who*, to the *Terminator* movies, the tale has been told time and again. Humans (or humanoid aliens) foolishly make robots that reproduce. The self-replicating robots decide that people are a nuisance and set out to exterminate them. This scenario might seem less far-fetched now that robots can make cars and microchips and stalk terrorists from the skies.

Don't panic just yet. Vicious self-replicating machines resembling Arnold Schwarzenegger won't be breaking down doors anytime soon. Anyone mighty enough to kick a toy or topple blocks can overpower today's self-replicating robots, which actually need a lot of help to make something identical to themselves. Self-replication "is fundamental to nature and at the core of evolution, and yet we have no idea how to do it with synthetic systems," says engineer Hod Lipson of Cornell University. "That's always been a sore point for robotics."

A handful of researchers are striving to change that. Working on shoestring budgets and with materials associated more often with child's play than research, they've developed simple robots that can make others like themselves out of a few relatively complex parts. They're defining more precisely what it means for a machine to self-replicate. And some are striving to emulate nature's knack for reproduction. Progress has been modest—stacks of blocks that stack other blocks won't conquer the world—but researchers are optimistic that, at the very least, they may soon better understand exactly what problem they're trying to solve.

All agree that progress has been slowed by a lack of funding, as self-replicating robots serve no earthly purpose—although in theory, they could be useful in establishing a base on the moon or on Mars. "The field is, like, three people," says mechanical engineer Gregory

Chirikjian of Johns Hopkins University in Baltimore, Maryland. Researchers face conceptual barriers as well. "There is a great need to come up with the basic scientific principles" of self-replication, says aerospace engineer Pierre Kabamba of the University of Michigan, Ann Arbor. Still, researchers have taken intriguing steps toward making machines that build copies of themselves.

Easy, in theory

The notion of self-replicating machines stretches back centuries. But the rigorous theory of self-replication emerged in the 1940s and 1950s, when mathematician John von Neumann, who also laid much of the groundwork for modern computing, analyzed the problem.

Von Neumann considered a collection of automata: self-guided cell-like entities that interact according to specific rules. He wondered what tasks a clump of them would have to do to replicate from raw

materials and basic parts. The thing would have to consist of at least three subunits, he figured: first, a set of instructions for making a device; then, a unit that reads those instructions to make a new device; and finally, one that copies the instructions, which von Neumann envisioned as a coded tape.

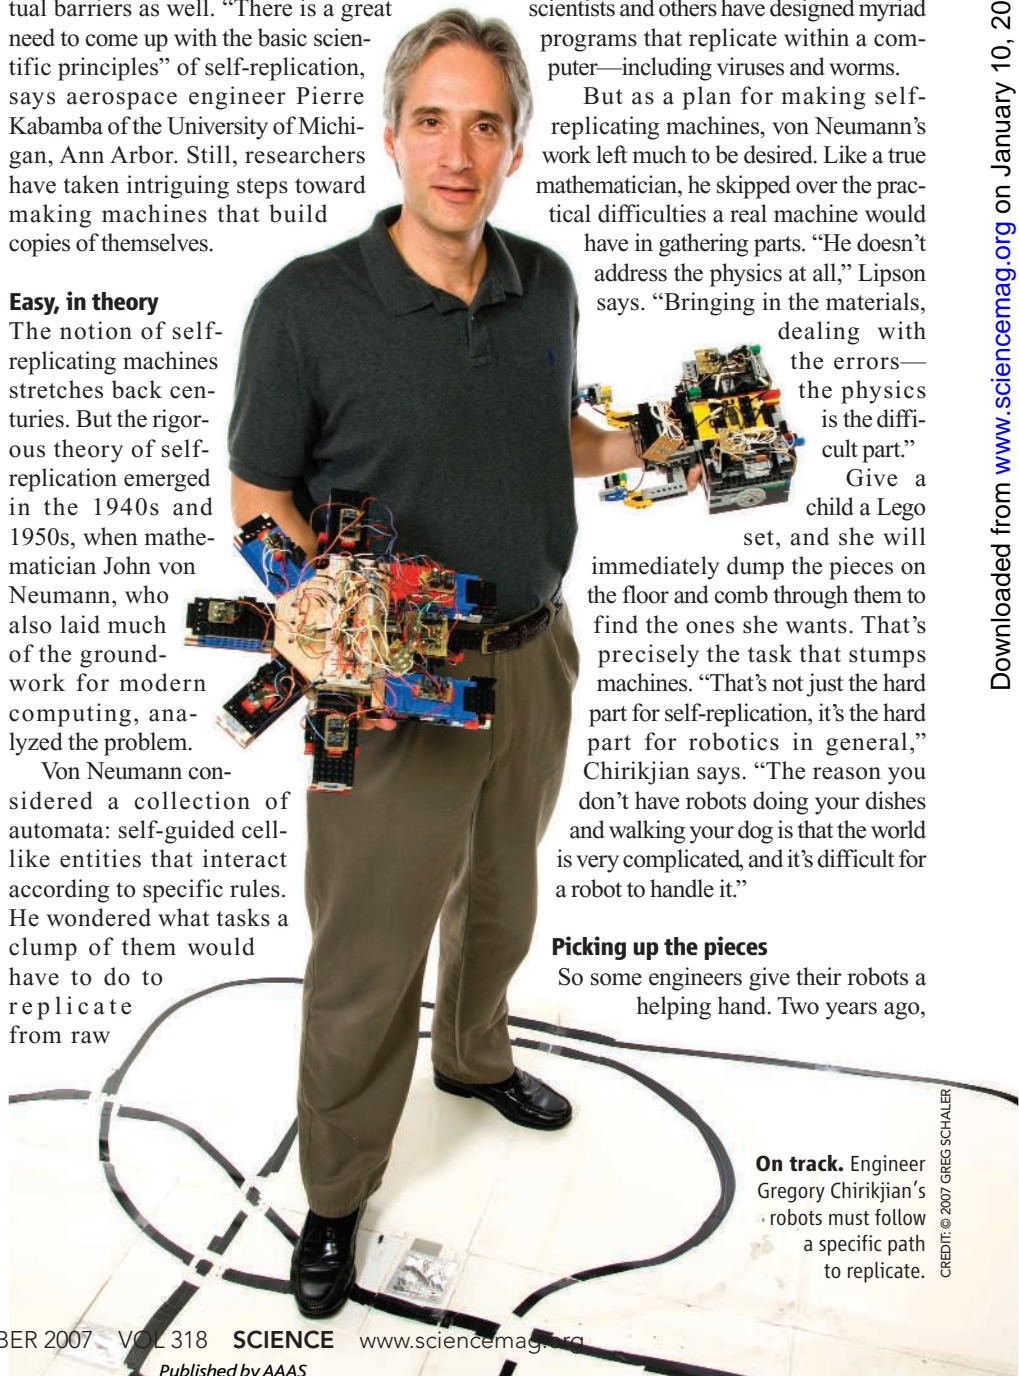
This agglomeration would read the tape, make its progeny, and pass a copy of the tape to its offspring. The scheme bears a striking resemblance to biology, in which cells replicate by reading and copying tape-like molecules of DNA, the structure of which was discovered after von Neumann cooked up his ideas. Spurred by von Neumann's work, computer scientists and others have designed myriad programs that replicate within a computer—including viruses and worms.

But as a plan for making self-replicating machines, von Neumann's work left much to be desired. Like a true mathematician, he skipped over the practical difficulties a real machine would have in gathering parts. "He doesn't address the physics at all," Lipson says. "Bringing in the materials, dealing with the errors—the physics is the difficult part."

Give a child a Lego set, and she will immediately dump the pieces on the floor and comb through them to find the ones she wants. That's precisely the task that stumps machines. "That's not just the hard part for self-replication, it's the hard part for robotics in general," Chirikjian says. "The reason you don't have robots doing your dishes and walking your dog is that the world is very complicated, and it's difficult for a robot to handle it."

Picking up the pieces

So some engineers give their robots a helping hand. Two years ago,



On track. Engineer Gregory Chirikjian's robots must follow a specific path to replicate.

Downloaded from www.sciencemag.org on January 10, 2008

CREDIT: © 2007 GREG SCHALLER

Lipson and colleagues unveiled programmable blocks measuring 10 centimeters across. Each consisted of two pyramid-shaped halves that could swivel against each other, and each block could grip others using magnets on its faces. Wriggling like a drunken hula dancer, a stack of four blocks could assemble a second stack, if new blocks were fed in at the right place and times, the researchers reported in the 12 May 2005 issue of *Nature*.

Although one stack of blocks does form another, it still seems a far cry from a fully self-replicating robot. Instead of some basic part, each cube is itself a fairly sophisticated robot. And the contorting tower requires plenty of human assistance to help it locate the additional blocks. To produce something truer to the spirit of self-replication, Lipson is now experimenting with simpler cubes measuring only 500 micrometers wide that jumble together randomly in a fluid. “What is the smallest building block from which we can make everything?” Lipson says. “That’s the crucial question.”

Chirikjian also began with robots that assembled others from a few complex chunks. Starting in 2002, he and his students began experimenting with robots made of Lego bricks. At first, they built remote-controlled vehicles that could be broken into a few components. When placed in a pen, one robot could push the components of another together—a crude form of self-replication, given that the guts of a robot lay mostly in the one component containing the computerized controller.

Since then, Chirikjian and his students have striven to make their robots more autonomous and to assemble them from simpler parts. They developed a system of optical sensors that allows a robot to follow a colored stripe to find various parts. They have simplified the robots by replacing the central controller with cruder electronics distributed throughout the pieces. Recently, the researchers demonstrated a self-replicating robot made of six fairly simple modules, and Chirikjian and a grad student are working on one consisting of 100 pieces.

Chirikjian’s robots look more or less self-sufficient, but they do not truly forage for parts. Rather, they depend on a track to guide them. Chirikjian says that he is working to eliminate the track. But he notes that even biological systems depend on their environment to reproduce. “If you take the DNA out of the environment of the cell, it’s no longer self-replicating,” he says.

Doing what comes naturally

Given the challenges of step-by-step, or deterministic, assembly, some researchers



Basic parts? A stack of Hod Lipson’s cubes stacks more cubes, but each is itself a complex robot.

are opting for chaos instead. Rather than making their robots fetch pieces, they’re relying on random collisions to bring parts to the robots in efforts that mimic the mingling of biomolecules in cells.

For example, as a graduate student at the Massachusetts Institute of Technology (MIT) in Cambridge, materials scientist Saul Griffith developed smart tiles that can latch onto one another as they glide and jumble on an air table. Whether two tiles latch depends on how they are already connected to other tiles. When the tiles were properly programmed, a chain of them could form another chain, Griffith and colleagues reported in the 29 September 2005 issue of *Nature*. “In many respects, self-replication is just a party trick,” says Griffith, now president of Makani Power in Alameda, California. “You don’t even need much logic.”

The random, or “stochastic,” approach may have a key advantage. Ironically, jumbling huge numbers of pieces together should be easier than putting them together one by one, says engineer Eric Klavins of the University of Washington, Seattle, who has developed a similar set of triangular tiles. “If you want to do self-replication with billions of parts, you’re not going to get away with determinism,” he says. The stochastic approach presents its own challenges, however. For example, researchers must figure out how to form larger useful structures from the pieces while preventing them from glomming together in undesirable ways.

A few molecular biologists are even pushing to develop artificial cells. For such research, the emphasis is a little different, says Jack Szostak of Harvard Medical School and Massachusetts General Hospital in Boston. In chemistry, self-replication is fairly common, as any chemical that catalyzes its own production does it. Szostak and colleagues are striving for

something more. “What we’re trying to do is to develop a self-replicating chemical system that can evolve,” Szostak says.

For the membranes for his artificial cells, Szostak employs molecules called lipids, which can form fluid-filled shells. Within the shells, he hopes to store a length of DNA, RNA, or a related molecule that can store coded information, replicate, and mutate. The researchers have already shown that they can make the shells grow and divide—by forcing them through a small pore—and they are working on the material to store within the shells.

Researchers have a long way to go, however. For example, molecular biologists have been searching for a strand of RNA called a ribozyme that can catalyze the replication of itself and other strands. Such a ribozyme would have to churn out strands a couple of hundred chemical letters long, but so far the best candidate can string together only about 20 letters. “Twenty years ago, I thought this would be a 20-year project,” Szostak says. “Maybe it still is.”

Waiting for the Terminator

Where research on self-replication will lead remains unclear. Some say that practical considerations will inevitably force researchers toward biomolecular systems. “Self-replicating robots are going to be made out of biomolecules long before bulldozers start copying themselves,” Griffith says. Others say it’s not so clear that self-replication in synthetic biology is easier than in mechanical robotics. “You’re comparing two very difficult things,” says molecular biologist David Bartel of MIT. “So which one is more difficult may not matter.”

Meanwhile, some say that the concept of self-replication needs a rethink. Researchers have thought that a system is either self-replicating or it isn’t, Lipson says. But given that even biological systems rely heavily on their environment, it seems there are different shades of self-replication. Both Lipson and Chirikjian have developed mathematical tools to quantify them. Using them, researchers might analyze a system to figure out how to make it more self-replicating, Lipson says.

Of course, employing such scales, one might argue that self-replicating robots already exist. Machines are typically made by other machines these days, albeit with plenty of help and guidance from humans. So perhaps the entire industrial enterprise constitutes a swarm of self-replicating robots. That seems plausible. But it also seems to be a disappointingly long way from the grand vision of machines that don’t need people. Maybe that’s a good thing.

—ADRIAN CHO