WHITE SPIKES: Gallium-arsenide nanowires grew out of golden seeds, not all of which sprouted. The wires reflect the crystalline structure of the silicon base, despite the different spacing of the materials' lattices.

als. Both groups grew the nanowires from gold particles deposited on a silicon substrate.

The Swedish group, led by Lars Samuelson at Lund University, started by depositing an aerosol of gold particles. The Dutch team, comprising researchers from the Philips Research Laboratories in Eindhoven and the Kavli Institute, and led by Erik Bakkers of Philips, obtained their particles by heating a gold film 0.2 nm thick on the substrate until it melted into droplets. Initially, they deposited the droplets on germanium. Now they have put them on silicon as well [see photo, "White Spikes"].

Having obtained a substrate dotted with gold nanoparticles, each team then grew the nanowires by introducing its device into an atmosphere containing III-V semiconductor vapor at high temperature and pressure. The semiconductor material diffuses into a gold droplet, and a semiconductor nanowire starts to grow. A wire 20 nm thick can reach a length of 20 micrometers, long enough to function as an optical waveguide, for example, says Delft's van Dam. The growth is epitaxial; that is, important features of the crystal structure of the silicon or germanium substrate, such as the orientation of the lattice planes, are replicated in the nanowire. This ensures that electrical contact is maintained between the nanowire and the substrate.

Peidong Yang of the University of California, Berkeley—a chemist who has been investigating III-V nanowires for the past five years—says that this advance will markedly accelerate the development of new optoelectronic devices. Lincoln Lauhon, a materials scientist at Northwestern University, in Evanston, Ill., agrees, pointing out that the development "gives you a great deal of flexibility. You can mix materials with different functionalities."

Still, the road to commercialization of viable devices will have its share of potholes and sharp curves. Bakkers expects it will be at least two years before any devices reach the market. High on his wish list: a III-V nanowire that forms the channel of a transistor in which the gate is placed around the wire—a gate-around transistor, suitable for high-frequency applications above 150 gigahertz. "This is the ideal configuration: a channel with high mobility of electrons and a gate with very good capacitive coupling, resulting in high transconductance [current-carrying capability]," he says.

The results from both research groups clearly show that hybrid microcircuits, incorporating the desirable properties of the III-V compounds with those of cheap and ubiquitous silicon substrates, might soon find an important niche in electronics after all.

—ALEXANDER HELLEMANS

IBM

IBM Reclaims Supercomputer Lead

But stay tuned—supercomputers are getting faster, at an even faster rate

Late last year, the news that two heavyweight computers were racing for the title of world's fastest left an excited supercomputing community in suspense. The competition came to an end in November, when the Top500 supercomputer ranking project released its biannual report.

IBM's Blue Gene/L [see photo, "Your Move, Garry"], to be delivered this spring to the U.S. Department of Energy's Lawrence Livermore National Laboratory in Livermore, Calif., took the top spot, with a performance of 70.72 teraflops (trillion floating-point operations per second). Silicon Graphics' Columbia, built for NASA and named after the space shuttle lost in 2003, came in second, at 51.87 teraflops. The two machines displaced Japan's famed Earth Simulator to third place after that 35.86-teraflop computer had reigned supreme for two and a half years [see table, "The Top Three Supercomputers"].

"This was tremendously important," says Tarek El-Ghazawi, a professor of electrical and computer engineering at George Washington University, in Washington, D.C. IBM and Silicon Graphics Inc.—SGI—came up with promising new designs, he says, showing it's possible to achieve dramatic increases in speed without sending costs through the roof. Blue Gene/L cost roughly a quarter of what was spent on the Earth Simulator, and Columbia just an eighth.

The new breed of supercomputers brings technology advances that may ultimately trickle down to a variety of high-

YOUR MOVE, GARRY: IBM's Blue Gene/L, now the fastest computer on earth, is cousin to Deep Blue, the chess machine that beat World Champion Garry Kasparov in 1997.

NEWS

performance computers, thus benefiting not only big-bucks buyers like the Energy Department and NASA but many other organizations in need of serious computing horsepower. That market is worth US $6 billion worldwide, in which IBM and SGI compete with Cray, HP, NEC, and Sun, among others.

WHAT INNOVATIONS catapulted Blue Gene/L and Columbia to the top?

crunching a tiny piece of a massive computational task. After five years of work, IBM engineers managed to cram 16 384 chips into 16 refrigerator-sized cabinets, then installed 5 types of networks that link the processors together in different ways and guarantee that none of them starve for data.

"It's just a maniacal focus on power, and it's a maniacal focus on integration," says Tilak Agerwala, vice president of systems at IBM Research and an IEEE fellow. At one-fourth of its planned size, Blue Gene/L has twice the computing power of the Earth Simulator, yet it's 1/50 the size and consumes 1/14 the power required by the Japanese computer, Agerwala says.

THE TOP THREE SUPERCOMPUTERS

SYSTEM	APPLICATIONS	BUILDER	ARCHITECTURE	PERFORMANCE (teraflops)	COST (US $, millions)
Blue Gene/L* Lawrence Livermore National Laboratory, Livermore, Calif.	Materials science, nuclear stockpile simulations	IBM	32 768 processors; 8 terabytes of memory; 28 terabytes of disk storage; Linux and custom operating system	70.72	100
Columbia NASA Ames Research Center, Moffett Field, Calif.	Aerospace engineering, simulation of space missions, climate research	SGI	10 240 processors; 20 terabytes of memory; 440 terabytes of disk storage; Linux operating system	51.87	50
Earth Simulator Earth Simulator Center, Yokohama, Japan	Atmospheric, oceanic, and earth sciences	NEC	5 120 processors; 10 terabytes of memory; 700 terabytes of disk storage; Unix-based operating system	35.86	350–500**

SOURCE: Top500 list (November 2004) and companies * Current configuration ** Industry estimate

IBM takes particular pride in its building block: a microchip containing not one but two microprocessors, based on the company's PowerPC processor family. It added memory, communications functions, and extra circuitry to speed up floating-point operations. Even so, the chip barely covers a fingernail and consumes a mere 10 to 15 watts.

It's important that the building blocks be small and energy efficient, because they work together by the thousands, with each one crunching a tiny piece of a massive computational task.

SGI, in Mountain View, Calif., isn't too shy to boast of its achievements, either. For Columbia's design, it took 20 of its Altix systems, high-performance computers powered by 512 Intel Itanium 2 processors, and connected them with an industry-standard InfiniBand networking system, getting them all up and running in only 120 days, says Dave Parry, a senior vice president at SGI.

The 10 240-processor machine began operation by adding one batch of processors at a time, without disrupting other systems, says Walt Brooks, chief of NASA's advanced supercomputing division at the agency's Ames Research Center, in Moffett Field, Calif. Simulations of space shuttle missions that used to take two to three months to generate are now running in hours, he says. The machine will raise NASA's overall computing power 10-fold.

A PRODIGIOUS PERFORMANCE on the part of the new supermachines hasn't stopped many in the supercomputing field from asking whether they will really stay ahead of new contenders for long.

For one thing, Blue Gene/L and Columbia—and other top-ranking supercomputers, for that matter—have distinct architectures, perhaps an indication that there's currently no single winning strategy for building these ultrafast computers. In fact, just a few years ago, the field "was very monochromatic, with very few architectures," says IEEE Fellow Jack Dongarra, a professor of computer science at the University of Tennessee at Knoxville and one of the organizers of the Top500 project. "Now there is a much richer collection of machines," he adds.

Current designs range from off-the-shelf clusters, which are basically large collections of cheap computers hooked together with some standard high-speed network, to highly customized machines like the Earth Simulator. It was built by Japanese computer maker NEC Corp. out of specially designed—and especially costly—components such as vector microprocessors, which perform some operations in parallel and thus faster than standard microprocessors.

IBM and SGI chose a mix of commodity and custom architectures—a strategy that seems to have proved successful, especially for Blue Gene/L, whose monstrous final 130 000-plus-processor configuration will achieve 360 teraflops. Dongarra predicts that this machine will lead the Top500 list for at least several more editions.

"But in the end," he hastens to add, "it's not really a question of speed, it's a question of what kind of science, what kinds of new ideas, what kind of deeper understanding do we get by using this equipment." The number of teraflops is just "a trophy," Dongarra says. "And we all know that trophy won't last forever."

—ERICO GUIZZO