

that converge in a “pocket” around the targeted device. Together, they generate enough RF power for a metamaterial receiving antenna to harvest. The device—say, a cellphone—is tracked in real time, and the beam paths are continuously updated with its location. The microwaves can bounce off surfaces, so as the device moves, the microwaves can reach the receiving antenna even if it’s not in a WattUp transmitter’s direct line of sight.

Energous recently commissioned a performance evaluation from Underwriters Laboratories. UL verified that, under ideal conditions, a WattUp transmitter is capable of delivering microwaves to up to four devices simultaneously. The amount of power the beams deliver is dependent on distance: 4 watts within 1.5 meters, 2 W within 3 meters, and 1 W within 4.6 meters. Energous wouldn’t comment on the end-to-end efficiency of the system during this test, except to say that its eventual goal is 25 percent efficiency. While it has yet to receive the approval of the U.S. Federal Communications Commission, Energous says it plans to make a splash this month at CES 2016 with an announcement of substantial efficiency improvements and integration into near-market consumer devices.

Energous has a competitor in Ossia, based in Redmond, Wash., which is also developing an RF power delivery system. But engineers at the University of Washington have taken a completely different approach, showing that a reliable source of wireless power is hiding in plain sight. Vamsi Talla, a UW electrical engineering doctoral student, is the lead author of a

paper describing Power Over Wi-Fi (PoWiFi), which can distribute tiny but useful amounts of power using an ordinary Wi-Fi router.

Your existing Wi-Fi router, transmitting at full strength, emits enough RF power, at least in theory, to operate or gradually charge very small devices at ranges approaching 10 meters. In practice, this isn’t realistic because your router emits power only when it’s sending data packets. Instead, PoWiFi keeps a stream of “superfluous broadcast traffic” flowing so that the router’s power output is kept at the maximum; this flow of nonessential packets is intelligently managed to avoid interference with the user’s actual Web surfing.


As with WattUp’s microwave beams, the amount of power transferred from PoWiFi’s custom routers to its RF harvesting antennas depends strongly on distance. According to Talla, temperature sensors placed more than 8 meters from a router collected enough charge to return data every few seconds. A modified consumer fitness device left within a few centimeters of a router went from zero to a 41 percent state of charge in 2.5 hours. Six PoWiFi routers, tested in the homes of volunteers, had no discernible impact on the quality of the users’ Internet access.

RF isn’t the only charging modality being promised for 2016. Israeli startup Wi-Charge uses beams of infrared light, and Santa Monica, Calif.-based uBeam relies on ultrasound. uBeam’s recent technical disclosures, and the fact that we’ve yet to see an operational acoustical-to-electrical-energy prototype, have left experts with one big question: How can ultrasound—which needs an unimpeded path and attenuates rapidly in the air—lend itself to power transmission in a world filled with stuff like, say, furniture? [See “Experts Still Think uBeam’s Through-the-Air Charging Tech Is Unlikely,” *IEEE Spectrum*, November 2015.]

Wireless tech will probably never match the efficiency and effectiveness of a plug in a wall socket. But even if real-world end-to-end efficiency ends up in the single digits—well below Energous’s hoped-for 25 percent—our guess is that it won’t stop many people from wanting to give wireless charging a try.

—EVAN ACKERMAN

How can ultrasound—which needs an unimpeded path and attenuates rapidly in the air—deliver power in a world filled with stuff like, say, furniture?



## THREE PATHS TO EXASCALE SUPERCOMPUTING

Specialized architecture, millivolt switches, or 3-D memory can get you there—for a price



For most of the decade, experts in high-performance computing have had their sights set on exascale computers—supercomputers capable of performing 1 million trillion floating-point operations per second, or 1 exaflops. And we’re now at the point where one could be built, experts say, but at ridiculous cost.

“We could build an exascale computer today, but we might need a nuclear reactor to power it,” says Erik DeBenedictis, a computer engineer at the Advanced Device Technologies department at Sandia National Laboratories, in Albuquerque. “It’s not impossible; it’s just a question of cost.”

Although nuclear reactors are its forte, the U.S. Department of Energy aims to use a more modest power source when it brings its first exascale supercomputer on line sometime in the 2020s. The goal is to have the machine consume no more than 20 megawatts.

Since 2012, DeBenedictis has been working with the IEEE Rebooting Computing initiative and the International Technology Roadmap for Semiconductors on a practical way to get to exascale. He sees three technologies that will lead the way forward: millivolt switches, 3-D stacked memory chips, and specialized processor architectures. Each involves some trade-offs, particularly in the realm of software specialization and because it will be difficult to predict when the technologies could be deployed. —JEREMY HSU

**STACKS UP NICELY:** A 3-D memory stack from Tezzaron Semiconductor boosts speed and reduces power draw.

### THE MILLIVOLT SWITCH

In accordance with Moore's Law, engineers had hoped that operating voltages for silicon transistors would continue to fall along with transistor sizes. Instead, the once steady improvement in operating voltages has remained stuck at close to 1 volt for the past decade.

"We could theoretically lower the operating voltage from where it is now—about 800 millivolts—to a few millivolts," says Eli Yablonovitch, a computer engineer and director of the National Science Foundation's Center for Energy Efficient Electronics Science at the University of California, Berkeley. Such a "millivolt switch" would be a radical solution to supercomputing's energy efficiency problem. But it remains an elusive goal.

Nobody knows how a millivolt switch might become a reality, and worse—from the standpoint of building exascale computers in the near future—nobody knows when. But Yablonovitch's center has chosen to focus on four possibilities. First, there's a device called a tunnel field-effect transistor, or tunnel FET: Switching it on makes the electronic barrier between the two sides of the transistor so thin that electrons are more likely to "tunnel" through it [see "The Tunneling Transistor," *IEEE Spectrum*, October 2013]. A second possibility is tiny nanoelectromechanical switches that would conserve power because little current can leak across an open switch [see "MEMS Switches for Low-Power Logic," *Spectrum*, April 2012]. Third, nanophotonics could enable faster, lower-power communication between switches by replacing physical wires with light signals. And fourth, logic circuits based on nanomagnetism could pave the way for lower energy, nonvolatile circuits [see "Better Computing With Magnets," *Spectrum*, September 2015].

Yablonovitch coined *millivolt switch* to illustrate the far-reaching potential

of "an energy version of Moore's Law," he says. Sandia's DeBenedictis thinks it won't have quite as huge an impact as Moore's Law did, but even so, he estimates that a millivolt switch could improve efficiency by 10 to 100 times.

### MEMORY IN 3-D

The U.S. Department of Energy's vision of a 20-MW exascale supercomputer would reserve 30 percent of the power budget for memory. That power budget goal will require memory-stacking architecture that improves on 3-D stacked DRAM, which puts towers of memory chips closer to a computer chip's processor units.

Such an arrangement means shorter interconnects between the units—and less power lost to capacitance and resistance when transferring bits of data. Even better, they also mean quicker communication [see "Memory in the Third Dimension," *Spectrum*, January 2014].

Stacked memory has already caught on in the broader computing industry. Before 2011, tech giants merely tinkered with it as "science fair projects," says Robert Patti, chief technology officer and vice president of design engineering at Tezzaron Semiconductor, which has worked on U.S. government supercomputing projects. But since that time, big chipmakers have all developed and

commercialized different versions of 3-D stacked memory.

For its part, Tezzaron has pioneered a method of designing and manufacturing 3-D memory stacks that boosts memory performance, enables a higher density of interconnects between each memory chip layer, and reduces manufacturing costs.

"In a lot of ways, 3-D integration allows you to have your cake and eat it too," Patti says. "You get speed improvement and power reduction at the same time."

Exascale supercomputing may need to rely less on volatile DRAM—computer memory that holds data only while powered—and move toward more nonvolatile yet fast storage such as resistive RAM, Patti says. A new architecture based on 3-D memory stacking would also require some software retuning as memory access latencies change. But once that's done, using more memory to upgrade a 1-exaflops supercomputer to 2, 3, or 4 exaflops wouldn't result in a system that consumes much more power, says DeBenedictis.

### SPECIALIZED CHIP ARCHITECTURE

Adding specialized cores to computer processors designed to run only a few specific applications might help save on supercomputing's future power budget. Each part of the chip would run at full power only when its particular function is needed and would be powered down when it's not.

But such specialized architectures would also require programmers to tailor their software to make use of the added efficiency, and that would constrain what a particular supercomputer could perform efficiently.

"If you want something else, you will have to replace the chip with one of a new design," says DeBenedictis. He believes that such specialized architectures will likely become a part of exascale supercomputing. As the cost of making chips continues to go down, building a general-purpose supercomputer from specialized chips with dedicated functions might become reasonable.

### EXASCALE TRADE-OFFS

The road to an exaflops supercomputer won't be smooth. The millivolt switch, for example, would dramatically reduce power draw. But how to make one, and when it would be ready, is anybody's guess.

