

One option for improving energy consumption is to look beyond silicon for new materials that could make transistors more energy efficient.

Demands are high for low-power electronics

By Angela Saini

Never have our personal gadgets been so small and yet expected to do so much. The thirst for energy of our multi-functional cell phones and personal computers has a frustrating impact on their battery lives and also drives up energy usage. It is difficult to compile figures for just how much energy this is in total, but in 2011, researchers at the University of California–Berkeley put a very rough upper estimate of the average power consumed by Internet use globally, including by laptops, smartphones, and data servers, at around 300 gigawatts—nearly two percent of all power. Expert analysis carried out for the *New York Times* newspaper in 2012 estimated that data servers across the world use 30 gigawatts.

Environmental concerns about this, as well as demand from consumers for devices that don't need recharging so often, are pushing the development of electronics that can run on far less power.

But that is not to say that energy savings have not already been made. Koomey's Law, devised by research fellow in energy and finance at Stanford University, Jonathan Koomey, in collaboration with Microsoft and Intel, states that the energy efficiency of computer chips (defined here as electricity used relative to processing power) doubles roughly every year and a half (see graph). It is a trend that has been shown to be relatively stable since the construction of the first general purpose computer in 1946.

"The key driver for efficiency in microprocessors is shrinking the size of transistors," said Koomey. The amount of power used by a transistor is directly proportional to its size. These energy savings have moved almost in parallel with the well-known Moore's Law, which states that the number of transistors per square inch on an integrated circuit doubles every two years.

However, while this rapid pace of miniaturization has transformed the scale and price of gadgets and improved efficiency, as the computer chip shrinks, the leakage current that is lost increases. The smaller chips become, the greater the losses. Even losses as tiny as a few hundred picoamps per transistor become more significant when multiplied by a million or billion transistors on a single chip.

Another problem with scaling down transistors is that it cannot happen forever. Gordon Moore, one of the founders of Intel Corporation and the person behind Moore's Law, suggested this year that it may reach its asymptotic limit within a decade. As a result, researchers and chip designers are looking for other avenues to improve energy efficiency.

Around 80% of the power used by a chip is spent in switching transistors. The minimum voltage at which a transistor will switch is roughly one-eighth of a volt, even though the wires on a chip can operate on as little as 10 millivolts. "Voltage reductions are no longer possible, so semiconductor manufacturers have turned to other techniques, like improving device architecture with multiple cores on a chip and reducing the standby power of chips when they aren't doing any computing," said Koomey. "Those techniques will allow us to continue rapid improvements in processor efficiency for a time, but within the next few decades, we'll need to radically change how we do computing, as we are approaching the limits of what is possible with current technologies."

One option for improving energy consumption is to look beyond silicon for new materials that could make transistors more energy efficient, or which might assist in the development of more sensitive, lower-voltage switches that could succeed conventional transistors. This branch of materials research, however, is very much in its early days.

Arun Thathachary, a researcher with the Department of Electrical Engineering at The Pennsylvania State University, has investigated the possibility of making nanoscale chips out of indium gallium arsenide (InGaAs), a compound semiconductor with higher electron mobility (a key factor in device performance) than silicon. He and his colleagues last year demonstrated that it is possible to harness the improved electron mobility of InGaAs to get a large current in the chip's "on" state, which means that the transistor can operate on a lower supply voltage.

Meanwhile Gianluca Fiori, assistant professor in the Information Engineering Department at the University of Pisa, Italy, among others, has been investigating the potential of ultrathin materials as an alternative to silicon. "The main advantage of thinness is that the gate has perfect control over the potential barrier. If you have better control, you can go for a smaller power supply," he explained. The ultimate thin material, graphene, would be a good chip material were it also a semiconductor, with a bandgap. The fact that it is not means that it can't be used as a transistor or switching element in digital applications. However, graphene, and bilayer graphene, in particular, have instead proven promising for radiofrequency applications.

In 2013, the research arm of IBM demonstrated a silicon chip designed for analog wireless communications, with transistor channels made of graphene. Reported to be 10,000 times faster than the next best graphene chip, it was

Angela Saini, angela.d.saini@gmail.com

nevertheless a functionally basic device, capable simply of receiving and processing signals in the 4.3 gigahertz frequency range. Graphene, therefore, has yet to prove itself before it becomes a standard material for these kinds of applications.

Another ultrathin option is single-layer molybdenum disulfide, which has the benefit of being a semiconductor with a wide bandgap, so it can be used to construct transistors with high on-off ratios and high voltage gain. Xiangfeng Duan, a professor of chemistry and biochemistry at the University of California–Los Angeles, has been part of a team to develop transistors made from molybdenum disulfide. Its major drawback, he says, is that it has an intrinsically low electron mobility relative to graphene. Duan suggests it would be best suited to low-power, high-frequency applications.

A different answer to the transistor problem, developed by Sapan Agarwal (now at Sandia National Laboratories, New Mexico) and Eli Yablonovitch at the University of California–Berkeley, is a piezoelectric transformer that fits on top of a transistor, ramping up the voltage exactly where it is needed, while keeping the power supply lower in the wires. The potential reduction from current energy levels, per operation, is considerable: theoretically, power consumption could drop by four orders of magnitude.


So far, though, all these options are in the research stage, and none are yet an immediate challenge to conventional silicon-based chips. Esther Rodriguez-Villegas, an associate professor in low power electronics at the Department of Electrical and Electronic Engineering at Imperial College London, said, “The problem is that silicon is a really good material. Other materials that are coming bring other problems to solve.”

Even so, the largest technology firms are betting on a low-power revolution. As well as being crucial to the future of the mobile internet, newer applications for which low-power chips are necessary include mobile swarms of wireless sensors, such as weather-monitoring or mapping drones, and body-centered networks, such as wearable electronics or health monitors and drug delivery systems that are implanted inside the body.

Rodriguez-Villegas is developing wearable sensors and implantable technology for use in health care. “The sensors have to be reliable and function for a long period of time on little power. Many of these systems have to be continuously monitoring,” she said. (The title image shows a microchip developed by Rodriguez-Villegas and colleagues that carries out complex real-time signal processing of brain signals from an electroencephalograph. It has an overall power consumption of under 10 microwatts and is less than six square millimeters.)

In January 2015, Intel announced plans for a hardware module it named Curie based on the Intel Quark system-on-a-chip architecture. This module can operate on low power for “always on” wearable products in which it says silicon has never before been used, “from rings, bags, bracelets, pendants, fitness trackers to even buttons,” according to the company’s press release.

When it comes to wearable electronics, materials researchers are also tackling the energy problem from a different direction: by changing the power source. Zhong Lin Wang, professor and chair in materials science and engineering at the Georgia Institute of



Graph of Koomey's Law showing the implications of historical trends in the electrical efficiency of computing. Reprinted with permission, <http://doi.ieeecomputersociety.org/10.1109/MAHC.2010.28>.

Technology, is exploring how to generate small amounts of current by harnessing the tiny vibrations of everyday life, such as the flow of blood over a sensor implanted inside the body. He has developed piezoelectric nanowires using zinc oxide, which he believes could be used to help design a nanoscale power source in the future.

The need for low-power chips is not limited to small devices, but may in the future also extend to the world’s sprawling data farms, which store an increasing proportion of our information. According to a Data Center Efficiency Assessment produced by the Natural Resources Defense Council (NRDC) in 2013, US data centers consumed power equivalent to the output of 34 large coal-fired plants—enough to power all the households in New York City for a year, twice over.

Unlike our personal gadgets, however, which rely on constantly improving technology, data farms face different energy challenges. Data center operators have been slow to upgrade old and wasteful power distribution equipment. Much energy is also wasted, according to the NRDC assessment, because equipment is powered while doing no work: the average server operates at between 12 to 18 percent of full capacity. Server capacity is the equivalent of a computer’s hard drive, and operators provide capacity depending on the workload they expect, such as file requests and data transfers. Capacity management is therefore crucial.

Even though energy efficiency is quickly rising up the agenda in both personal electronics and data farming, in an information age that has for so long depended on silicon, a wholesale change to new materials in computer chips is unlikely to happen very soon—not the least because few have shown much promise. “There have been advances in new materials, but none have been proved at scale yet, and many details have yet to be worked out,” said Koomey. □