The Inner Limits

By Chris Jozefowicz

Are computers approaching the furthest extent of their power?

"If you were living in the year 3,000 B.C.E., how much change would you see in your lifetime?" asks Doug Burger, a computer scientist at the Microsoft Corporation. "You probably wouldn't see any." Change happened very slowly in ancient times.

"In the 1700s or 1800s, change was measured in decades," Burger continues. "In the 20th century, it was measured in somewhere between decades and years. Now we're seeing things change in years to months." Faster, smaller computers, smart phones, and other electronic gadgets drive a continuing revolution in how people work and play.

Will that spectacular rate of change keep going into the future? Not in the same way, says Burger. As computer scientists continue to make the guts of electronic devices smaller and quicker, they've begun to push against some limits of what may be possible.

Tiny Switches

The part of a computer's guts that does all the computing is run by many tiny *transistors.* "Transistors are effectively like little light switches," Burger says. When a transistor is on, electricity flows through it. When it's off, the current stops.

A modern computer has hundreds of millions to billions of transistors packed together to form *integrated circuits,* or *computer chips.* Lightning-quick changes in the complex pattern of switches in the "on" and "off" positions enable a computer to do computations and store information in its memory. "For 40 years we've been making transistors smaller, smaller, and smaller," Burger says. They've become faster, too, while needing less energy. The results have been amazing. The latest smart phone is faster than a refrigerator-sized, 5-ton supercomputer from the 1970s. "It's been magical for us," Burger says.

Lagging Efficiency

Rapid change is still the norm for computer technology, but some rates of change are slowing. Scientists can pack more and more transistors onto chips, but computer chip *efficiency* is lagging. Efficiency is a measure of how something performs compared with how it might perform if its energy use were ideal.

Every tiny transistor wastes a little energy when it's on. That waste takes the form of heat created by the electric current flowing through the transistor. If you've touched a hot cellular phone or felt the warm air blowing out of a computer's vent, you've sensed the wasted energy.

iStock
Smart phones and tablets use computer chips that still have room to get faster. Such devices run on limited programs, so their chips aren't as complex as the ones found in PCs.

The rate of change in chip efficiency is slowing because designers are making less progress reducing how much energy transistors consume—and waste. If the newest chips were to consume all the energy they need to power every one of their transistors, they would get too hot. The chips' parts would melt and fuse. The computers would spit out incorrect results.

To prevent that from happening, some parts of the newest computer chips must remain off while other parts are on. Researchers call the unpowered parts *dark silicon.* (Silicon is the material that most transistors are made of.) When silicon is dark, a chip isn't working as hard as it could, however.

Burger and some colleagues in computer science have studied the problem of dark silicon. Their conclusion: The pace of computer innovation will slow in the coming decades. If computers were to continue evolving at the breakneck speed of the past 40 years, we might expect their performance to improve by at least 30 times by 2024. Burger's calculations suggest that we should expect a fraction of that improvement. "I think we're likely to see something between 3.7 times and 7.9 times," he says.

Twilight of the PC?

Burger's prediction hardly fazes David Patterson, a computer scientist at the University of California, Berkeley. "My position is that we need to innovate," he says.

One innovation is *multi-core processing.* Many of today's computers already have it. Their computer chips are effectively several chips linked together. Future designs might link even more chips to keep speeding things up. Each chip would be specialized to perform a certain function when needed.

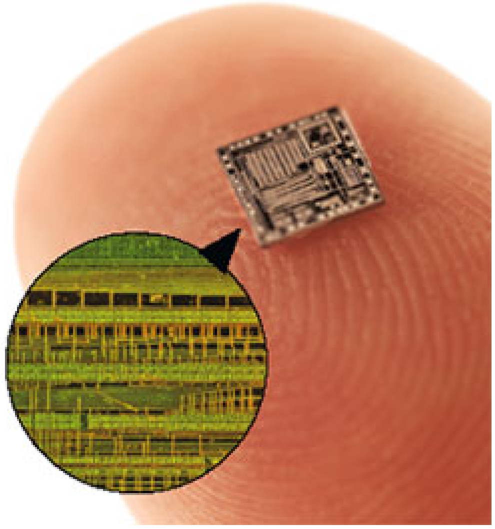

Lucidio Studio Inc./Getty Images; Nick Botto/Getty Images
A computer chip. Left: A close-up of the intricate, microscopic circuitry on a computer chip

Patterson compares multi-core processing to having a car with many engines, each one designed for driving at a different speed or on a different road condition. Some parts remain off while the only parts needed are on. "It doesn't matter if silicon is dark or light," Patterson says. "What matters is if it can continue to deliver performance into the future."

Burger's predictions about dark silicon apply mostly to personal computers (PCs), adds Patterson. But consumers are increasingly turning away from PCs to smart phones and tablets for many of their computing needs. Those devices don't use the same kinds of chips PCs do. "Smart phones and tablets have room to get faster," Patterson says.

Patterson imagines a future in which people carry or even wear many small computers. Such devices might be slower than top-of-the-line PCs, but each one would have its own task. And the devices would communicate and work with one another in a network, almost like a multi-core processor. Researchers call that type of sharing *cloud computing.*

"This is such a fast-moving field," Patterson notes with optimism. "When you start talking about what's going to happen past 10 years, it's hard to worry."