


“By 2020, most computing and storage will be in the cloud, and we'll stop calling it the cloud.”

GOOGLE ENGINEERING DIRECTOR PETER MAGNUSSON AT THE IEEE TECHNOLOGY TIME MACHINE 2012 SYMPOSIUM



**TICK TOCK:** Cyclos Semiconductor tested its clock-energy recycling technology on an ARM processor.

PHOTO: CYCLOS

clock tree and thus minimizes timing variations. But this approach comes at the cost of power. Clock meshes are essentially big capacitors; injecting and pulling charge through all the added metal consumes a lot of energy.

About 10 years ago, researchers began eyeing a seemingly simple fix: Wire up a number of inductors to the mesh. The result is an LC oscillator that naturally bounces between two voltages as it alternately stores energy in the mesh and in the inductors. A chip that uses such an oscillator to make a “resonant clock” can rely on smaller, less power-hungry circuits to drive its clock distribution network. They’ll be needed only to prime the oscillator and inject a little energy during each cycle to keep it running, like pushing a swing.

Cyclos, a start-up based in Berkeley, Calif., and Ann Arbor, Mich., has been in stealth mode since 2009, when the company reported a proof-of-principle test showing that it could use a single inductor, located off-chip, to recover as much as 85 percent of the energy in the clock distribution network of a 200-MHz ARM microcontroller chip.

Fully integrating the technology into a high-performance chip presented a few challenges. For one thing, says Cyclos cofounder Marios C. Papaefthymiou, an electrical engineering and computer science professor at the University of Michigan,

## Good Timing

A method for recycling clock energy is making its way into high-performance chips

A LONG-STUDIED strategy for recycling the energy used to clock computer processors is now starting to make its way into commercial chips. The approach—dubbed resonant clocking—has been integrated into Advanced Micro Devices’ new AMD A10-4600M processor. The chip, known by the code name Trinity, debuted in May and is now shipping in desktop and laptop computers.

Taking a cue from the pendulum, the technology uses inductors to periodically store and release the energy that would ordinarily be dissipated through the clock distribution network on a

chip. According to AMD and Cyclos Semiconductor, which licensed the technology to the processor firm, the approach can cut down on the power needed to clock the chips by up to 24 percent. Alternately, the technology can be used to boost the top speed of the 4-gigahertz-capable chips by 100 megahertz without requiring more power.

Finding ways to save clock power, which can consume 30 percent of a high-performance chip’s energy, is tricky. Chips need the steady heartbeat of a clock to drive computation. And the clock’s signals must be distributed widely, to tens of thousands—

or even hundreds of thousands—of locations.

Chipmakers have traditionally used “trees” to distribute clock signals. These employ many branching wires and networks of small circuits to pick up and propagate the clock signal. But trees don’t work well for chips that are meant to operate at speeds of roughly 2 GHz and up, where slight variations in wire lengths, transistor properties, and power supply can significantly alter the time it takes for a clock signal to reach two different spots on a chip. To prevent timing-violation errors, chips are run at less than their maximum speed capabilities.

Designers of high-performance chips have largely solved this problem by adding a wire grid—often called a clock mesh—which shorts out the endpoints of the

# update

the team had to find a way to incorporate spiral loops of inductors into the already quite dense layers of interconnects that carry power and data around the chip.

“The way the power grids are designed, they’re full of loops,” says Papaefthymiou. “The inductors are close to power supply lines and close to ground lines, and the closer they get to these lines, the more they interfere with them.” The team found a way to minimize interference by making the inductors a bit less ideal, winding them around in a spiral but also up and down, through the two top metal layers of the chip. The team’s scheme for AMD’s Piledriver CPU design—which forms the heart of the company’s new Trinity chips—uses 92 of these 100-micrometer-wide inductors, spread out over each dual-core processor module. The basics of the design were reported earlier this year at the IEEE International Solid State Circuits Conference.

Although AMD saw clock power drop by a quarter, the potential power savings could be much greater, says Matthew Guthaus, an associate professor of computer engineering at the University of California, Santa Cruz. His team has shown that by customizing the size of inductors and their location on the chip, it may be possible to reduce clock power by as much as 90 percent.

Most of the savings would come from reducing the length of wiring between the inductors and the flip-flops at

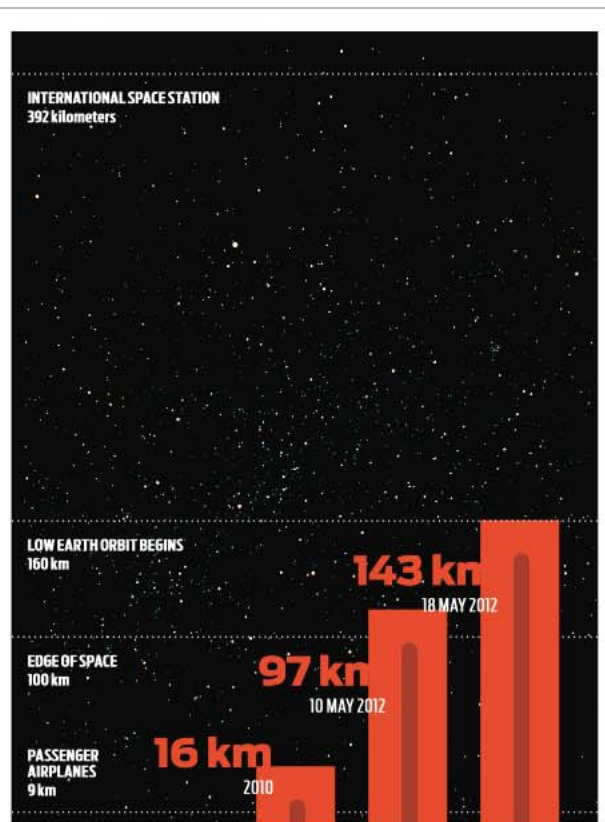
the end of the clock network, which are not distributed uniformly on the chip. Guthaus’s team presented such a design approach in June at the Design Automation Conference in San Francisco.

Despite this potential, it’s still unclear how far resonant clocking can be extended. Not all chips use clock meshes, and inductors can consume a lot of real estate. What’s more, although a resonant clock doesn’t have to be run directly at the resonant frequency of the LC circuit, its efficiency as an energy recycler goes down when the clock is run significantly faster or slower. That reduces the utility of one tried-and-true power-saving strategy—temporarily running chips at slower speed.

“Driving far from resonant frequency won’t save power, and at some point [the circuit] won’t work at all,” says Phillip Restle, a member of the research staff at the IBM Thomas J. Watson Research Center in Yorktown Heights, N.Y. Cyclos has inserted a switch that allows the AMD chip to turn the resonant part of the clock on and off. But Restle, who performed some of the earliest work on resonant clocks in microprocessors, says the switch isn’t a perfect fix because it adds to the power a chip consumes.

Still, Restle says this isn’t the last we’ll see of this clocking approach. “I consider this early days for resonant clocking,” he says. “There are more rabbits to pull out of the hat.”

—RACHEL COURTLAND



## Beam Me Up

European scientists led by the University of Vienna’s Anton Zeilinger say they’ve set a new record for the mysterious but potentially important phenomenon of quantum teleportation. The basis for a communications scheme that can’t be hacked, quantum teleportation reproduces the quantum properties of one particle—a photon in this case—on another, distant particle. Scientists have been in a race to find out just how distant they can make these particles. Progress has been rapid, and the prize is nearly within reach: If you can establish a quantum connection between the ground and an orbiting satellite, you can, in principle, establish an impenetrable, globe-spanning network.

GETTY IMAGES