BusinessWeek online

FEBRUARY 4, 2005

NEWS ANALYSIS TECH
By Otis Port

Chips That Thrive on Uncertainty

As transistors shrink, consistent performance diminishes. Big problem? Not if Krishna Palem is right about the benefits of unpredictability

In an ideal world, all the transistors on each computer chip would be identical. Moore's Law -- the edict that says silicon power will double every 18 months or so -- could then continue on its merry, exponential way for three or four more decades. Semiconductor engineers would just need to keep trimming the size of transistors by roughly 10% a year.

In the real world, though, things are a lot more complicated. In fact, transistors are like snowflakes: No two are exactly the same. This variation hasn't mattered much so far, but that will soon change.

"In a nutshell," says Shekhar Borkar, director of research for Intel's (**INTC**) Microprocessor Technology Laboratory, "what's happening is that as transistors get smaller and smaller, their variabilty is steadily increasing. In the future," he adds, "two transistors sitting side by side, which you intended to make the same size, could look different electrically."

Looking different electrically is much more serious than a cosmetic blemish. It means haphazard variations in performance. The upshot would be unreliable chips -- and untrustworthy computers, cell phones, navigation systems, and other products built around silicon circuits. But, Borkar says, research at Georgia Institute of Technology holds "great promise" for solving this problem of unpredictable variability.

"LIVE WITH IT." Ironically, the problem is part of the solution, says Krishna Palem, director of Georgia Tech's Center for Research in Embedded Systems & Technology. If making silicon transistors with predictable properties is going to be impossible in 10 years or so, he says, "let's sink into a morass of uncertainty. Let's learn to live with it and see what we can do with unpredictability."

Palem recently unwrapped his first "live with it" prototype chip, and so far tests have confirmed his hunch that uncertainty can be turned into an asset. One benefit is reduced energy consumption. Gulping less energy not only trims heat generation but it also could prolong the recharge cycle of cell phones, handheld computers, and other battery-powered gadgets.

With Palem's chips, cell phones might keep going for weeks, not days. Typical energy savings, he believes, will be around 20%, but it could be much higher if an application could sacrifice speed to gain longer battery life. "There's a trade-off," he explains.

FEWER DOUBLE CHECKS. Similarly, excessive heat is now a major obstacle for speed-demon chips. To avoid melting copper circuit lines, some new chips are fitted with speed limiters. They prevent the chip from crunching numbers as fast as it otherwise could. A lot of this heat stems from today's deterministic approach to chip design, Palem notes. The chip gobbles large amounts of energy to be absolutely certain that each data bit is either a 0 or 1 at every step of a calculation.

However, if a chip can get by without all the double checks to assure absolute certainty, then energy consumption could be slashed --and speed would get a simultaneous boost. That's the notion behind Palem's concept of probabilistic bits, or Pbits. As he puts it: "Uncertainty, contrary to being an impediment, becomes a resource."

Palem figures that Pbit chips could find short-term uses even before they're needed to address manufacturing variability. "The initial applications won't be in general-purpose computing," he says. That's where Intel's chips reign. "Instead, they'll be special-purpose, embedded applications."

GOOGLE'S EXAMPLE. Embedded chips are essentially microprocessors *sans* keyboards. These chips are buried inside TVs and stereos, kitchen appliances, motors, and, these days, most products that run on electricity. They're far more common than typical microprocessors. Indeed, every personal computer has several so-called microcontrollers. They're in the keyboard, the hard-disk drive, and the display. Most cars have a couple dozen embedded processors, and luxury models can have 50 or more, controlling the fuel-injection and braking systems, power windows, and dashboard displays.

So what kind of jobs could tolerate a smidgen of uncertainty? Palem ticks off a rather surprising list: digital-signal processing, medical prosthetics, database searching, robotics, face and voice recognition, computer-aided product design, and even financial and risk analyses. In short, virtually any task for which a pat algorithm doesn't yet exist (face recognition, for example), or those that already involve statistical methods (digital-signal processing in cell phones), or where the so-called problem domain is itself riddled with uncertainty (financial analysis).

"Google (**GOOG**) is a great example of what you can achieve with probabilistic techniques," says Palem. "They have to deal with tons of information but still provide you with a quick answer. So they need a fast search engine to do a huge amount of analysis, using statistical models, pattern matching, and probability searching. And Google does give you an answer quickly. But it's not a 'hard' answer" -- the one specific thing you wanted. Instead, he notes, "you get back this list of possibilities, ranked probabilistically."

WALL STREET APPLICATIONS? Palem is especially optimistic about Pbit-chip prospects for artificial eyes and ears in both robots and people. Computers are notoriously inept at recognizing even simple objects such as chairs, let alone faces. "Probabilistic chips might be able to learn to do a better job," he asserts. "They could also help improve impaired human vision and hearing."

Might such chips bring new insights to Wall Street? "Financial analysis would be extremely well suited for this," Palem predicts. Then he clams up. Whether any Wall Street firms are getting regular briefs on Palem's research, as Intel and IBM (**IBM**) are, he won't say. Wall Street doesn't like people blabbering about technology that promises a competitive advantage.

What is known is that it'll be a while before any of Palem's visions emerge into the light. He figures he needs at least a year to design and validate a more sophisticated chip. But if that chip checks out, it could kick-start a new silicon revolution.