MIT Genius Stuffs 100 Processors Into Single Chip

Call Anant Agarwal's work crazy, and you've made him a happy man.

Agarwal directs the Massachusetts Institute of Technology's vaunted Computer Science and Artificial Intelligence Laboratory, or CSAIL. The lab is housed in the university's Stata Center, a Dr. Seussian hodgepodge of forms and angles that nicely reflects the unhindered-by-reality visionary research that goes on inside.

Agarwal and his colleagues are figuring out how to build the computer chips of the future, looking a decade or two down the road. The aim is to do research that most people think is nuts. "If people say you're not crazy," Agarwal tells Wired, "that means you're not thinking far out enough."

Agarwal has been at this a while, and periodically, when some of his pie-in-the-sky research becomes merely cutting-edge, he dons his serial entrepreneur hat and launches the technology into the world. His latest commercial venture is Tilera. The company's specialty is squeezing cores onto chips — lots of cores. A core is a processor, the part of a computer chip that runs software and crunches data. Today's high-end computer chips have as many as 16 cores. But Tilera's top-of-the-line chip has 100.

The idea is to make servers more efficient. If you pack lots of simple cores onto a single chip, you're not only saving power. You're shortening the distance between cores.

Today, Tilera sells chips with 16, 32, and 64 cores, and it's scheduled to ship that 100-core monster later this year. Tilera provides these chips to Quanta, the huge Taiwanese original design manufacturer (ODM) that supplies servers to Facebook and — according to reports, Google. Quanta servers sold to the big web companies don't yet include Tilera chips, as far as anyone is admitting. But the chips are on some of the companies' radar screens.

Agarwal's outfit is part of an ever growing movement to reinvent the server for the internet age. Facebook and Google are now designing their own servers for their sweeping online operations. Startups such as SeaMicro are cramming hundreds of mobile processors into servers in an effort to save power in the web data center. And Tilera is tackling this same task from different angle, cramming the processors into a single chip.

Tilera grew out of a DARPA- and NSF-funded MIT project called RAW, which produced a prototype 16-core chip in 2002. The key idea was to combine a processor with a communications switch. Agarwal calls this creation a tile, and he's able to build these many tiles into a piece of silicon, creating what's known as a "mesh network."

"Before that you had the concept of a bunch of processors hanging off of a bus, and a bus tends to be a real bottleneck," Agarwal says. "With a mesh, every processor gets a switch and they all talk to each other…. You can think of it as a peer-to-peer network."

What's more, Tilera made a critical improvement to the cache memory that's part of each core. Agarwal and company made the cache dynamic, so that every core has a consistent copy of the chip's data. This Dynamic Distributed Cache makes the cores act like a single chip so they can run standard software. The processors run the Linux operating system and programs written in C++, and a large chunk of Tilera's commercialization effort focused on programming tools, including compilers that let programmers recompile existing programs to run on Tilera processors.

The end result is a 64-core chip that handles more transactions and consumes less power than an equivalent batch of x86 chips. A 400-watt Tilera server can replace eight x86 servers that together draw 2,000 watts. Facebook's engineers have given the chip a thorough tire-kicking, and Tilera says it has a growing business selling its chips to networking and videoconferencing equipment makers. Tilera isn't naming names, but claims one of the top two videoconferencing companies and one of the top two firewall companies.

An Army of Wimps

There's a running debate in the server world over what are called wimpy nodes. Startups SeaMicro and Calxeda are carving out a niche for low-power servers based on processors originally built for cellphones and tablets. Carnegie Mellon professor Dave Andersen calls these chips "wimpy." The idea is that building servers with more but lower-power processors yields better performance for each watt of power. But some have downplayed the idea, pointing out that it only works for certain types of applications.

Tilera takes the position that wimpy cores are okay, but wimpy nodes — aka wimpy chips — are not.

Keeping the individual cores wimpy is a plus because a wimpy core is low power. But if your cores are spread across hundreds of chips, Agarwal says, you run into

problems: inter-chip communications are less efficient than on-chip communications. Tilera gets the best of both worlds by using wimpy cores but putting many cores on a chip. But it still has a ways to go.

There’s also a limit to how wimpy your cores can be. Google’s infrastructure guru, Urs Hölzle, published an influential paper on the subject in 2010. He argued that in most cases brawny cores beat wimpy cores. To be effective, he argued, wimpy cores need to be no less than half the power of higher-end x86 cores.

Tilera is boosting the performance of its cores. The company’s most recent generation of data center server chips, released in June, are 64-bit processors that run at 1.2 to 1.5 GHz. The company also doubled DRAM speed and quadrupled the amount of cache per core. “It’s clear that cores have to get beefier,” Agarwal says.

The whole debate, however, is somewhat academic. “At the end of the day, the customer doesn’t care whether you’re a wimpy core or a big core,” Agarwal says. “They care about performance, and they care about performance per watt, and they care about total cost of ownership, TCO.”

Tilera’s performance per watt claims were validated by a paper published by Facebook engineers in July. The paper compared Tilera’s second generation 64-core processor to Intel’s Xeon and AMD’s Opteron high end server processors. Facebook put the processors through their paces on Memcached, a high-performance database memory system for web applications.

According to the Facebook engineers, a tuned version of Memcached on the 64-core Tilera TILEPro64 yielded at least 67 percent higher throughput than low-power x86 servers. Taking power and node integration into account as well, a TILEPro64-based S2Q server with 8 processors handled at least three times as many transactions per second per Watt as the x86-based servers.

Despite the glowing words, Facebook hasn’t thrown its arms around Tilera. The stumbling block, cited in the paper, is the limited amount of memory the Tilera processors support. Thirty-two-bit cores can only address about 4GB of memory. “A 32-bit architecture is a nonstarter for the cloud space,” Agarwal says.

Tilera’s 64-bit processors change the picture. These chips support as much as a terabyte of memory. Whether the improvement is enough to seal the deal with Facebook, Agarwal wouldn’t say. “We have a good relationship,” he says with a smile.

While Intel Lurks

Intel is also working on many-core chips, and it expects to ship a specialized 50-core processor, dubbed Knights Corner, in the next year or so as an accelerator for supercomputers. Unlike the Tilera processors, Knights Corner is optimized for floating point operations, which means it’s designed to crunch the large numbers typical of high-performance computing applications.

In 2009, Intel announced an experimental 48-core processor code-named Rock Creek and officially labeled the Single-chip Cloud Computer (SCC). The chip giant has since backed off of some of the loftier claims it was making for many-core processors, and it focused its many-core efforts on high-performance computing. For now, Intel is sticking with the Xeon processor for high-end data center server products.

Dave Hill, who handles server product marketing for Intel, takes exception to the Facebook paper. “Really what they compared was a very optimized set of software running on Tilera versus the standard image that you get from the open source running on the x86 platforms,” he says.

The Facebook engineers ran over a hundred different permutations in terms of the number of cores allocated to the Linux stack, the networking stack and the Memcached stack, Hill says. “They really kinda fine tuned it. If you optimize the x86 version, then the paper probably would have been more apples to apples.”

Tilera’s roadmap calls for its next generation of processors, code-named Stratton, to be released in 2013. The product line will expand the number of processors in both directions, down to as few as four and up to as many as 200 cores. The company is going from a 40-nm to a 28-nm process, meaning they’re able to cram more circuits in a given area. The chip will have improvements to interfaces, memory, I/O and instruction set, and will have more cache memory.

But Agarwal isn’t stopping there. As Tilera churns out the 100-core chip, he’s leading a new MIT effort dubbed the Angstrom project. It’s one of four DARPA-funded efforts aimed at building exascale supercomputers. In short, it’s aiming for a chip with 1,000 cores.

www.wired.com [1]

Posted by Janine E. M0oney, Editor

January 23, 2012

Source URL (retrieved on *12/12/2013 - 10:20pm*):
http://www.wirelessdesignmag.com/news/2012/01/mit-genius-stuffs-100-processors-single-chip

Links:
[1] http://www.wired.com