

[News Blog](#)

August 21, 2007 9:30 AM PDT

Live from Hot Chips 19: Technology and software directions

Posted by [Peter Glaskowsky](#)

Font size

Print

E-mail

Share

This is the sixth in a series of posts from the [Hot Chips](#) conference at Stanford University. The previous installments looked at [process technology](#), [multicore](#) designs, [IBM's Power 6](#) efforts, [Vernor Vinge's keynote address](#), and [Nvidia](#). Other CNET coverage may be found [here](#). This is sort of an experiment for me; I usually prefer to have time to review my work before I publish it. If you see anything wrong, please leave a comment!

We began Tuesday morning with a session on assorted technology developments.

The first talk was from Sun Microsystems, about the company's Proximity chip-to-chip interconnect technology. Today, to put multiple chips in a package--a common technique in high-end servers, for example--each chip will be individually connected to the package substrate through conductive metallic bumps; then printed wiring on the substrate provides the electrical connection between the bumps. Proximity uses capacitive coupling instead, substituting flat electrodes for the bumps. Chips are overlapped within the package, bringing the electrodes of one into close proximity to the electrodes of another.

Sun claims this method greatly increases the communications bandwidth between chips. It measures bandwidth in terms of the data rate for a given area of chip surface, quoting the potential of Proximity at 10 terabits/s per square mm, about 10 times the bandwidth density that can be achieved using conventional substrate bonding.

There are drawings and more detailed explanations on Sun's Proximity Web pages ([here](#)).

So the interesting question is, how useful is this technology? IBM's Power6 (which I mentioned in yesterday's Hot Chips blog post [here](#)) uses the conventional solution and achieves very high packaging density. As I described, each Power6 chip has 300GB/s of off-chip bandwidth. Would Proximity technology have produced a better result? It's possible, I suppose, but it isn't exactly obvious to me.

Sun's focus is more on networking hardware than servers. The company is looking to enable network switches with potentially thousands of ports and many terabits/s of bandwidth. But the market for such things is very small, and there are other ways--albeit possibly less efficient ways--to build big switches.

We probably just haven't reached the point where Proximity's high bandwidth density is really needed. By the time we get there--when we can put 16 or more high-performance CPU cores on a single chip--it's likely that traditional chip interconnects will also run at much higher speeds than they do today.

But it's still nice to see companies like Sun investing time and effort in basic R&D. That's been one of the nice things to watch over at Sun for the last few years. Although the company was in trouble for a while there, between the collapse of the telecom business and the bursting of the dot-com balloon, it never shut down its R&D efforts. Sun's investments during that time are now paying off in the success of its [Throughput Computing](#) initiative and other areas. One hopes that if there is another economic upheaval in the computer industry, other companies will follow Sun's lead.

Up next was a talk from [T-RAM Semiconductor](#). T-RAM is another alternative memory technology for integration into high-speed logic chips such as microprocessors. When such chips need on-die memory--for large caches, for example--they usually use SRAMs with six transistors per bit.

SRAMs are good, but provide low storage density. Many companies have been working on alternatives for a long time now. [MoSys](#), for example, has been selling its 1T-SRAM for many years. That product name refers to the use of one transistor per bit--while still maintaining the functionality of an SRAM array, internally the MoSys technology is actually based on DRAM. Innovative Silicon ([Z-RAM](#)) and [Renesas](#) also offer high-density single-transistor memory technology for use on logic chips.

T-RAM's approach uses a thyristor in place of a transistor. A thyristor is basically a three-junction semiconductor component, as compared with two-junction transistors. The extra junction allows thyristors to latch into a conducting state, thus storing a bit. You can read the rest of the theory on the company's Web site. But the bottom line is that T-RAM technology is denser than SRAMs. That translates to more bits in the same area, or alternately the same amount of storage on a smaller chip.

T-RAM claims to be faster than its direct competition with comparable density. Whether that's good enough for the company to succeed is not clear to me. MoSys has been able to sell its technology to quite a few customers, but like other SRAM alternatives, it still isn't widely used on mainstream microprocessors.

The final presentation in the session was from a group of Stanford students proposing a new method of improving software security. The proposal (named [Raksha](#) after the Hindi word for protection) is to add another bit to each data item stored in the computer that tracks the trust level of the data item. Untrusted items are marked "tainted," according to the presentation. When tainted data items are processed, the taint is propagated to the results.

The presentation described how Raksha can be implemented in hardware and supported by operating systems. Raksha imposes a 12.5 percent storage overhead--one extra bit per byte stored--and about 7 percent overhead in computational logic as well.

The Stanford team has implemented a prototype Raksha system in hardware, running Linux on a modified Sparc processor (which is interesting in itself--it's an open-source processor design called LEON3 from [Gaisler Research](#)).

The presentation shows how the Raksha team tested their system and demonstrated that it achieves its goals. But this talk raised more questions in my mind than it answered. It's easy enough to say that some piece of data is or isn't "trusted," but what does that really mean? And what's the practical consequence of saying some data item isn't trusted?

Untrusted data must still be processed--enabling such processing is the essence of the Raksha project, after all--and the results of this processing will eventually be presented to users. How will they know whether to trust the results? The user's trust in the result

is a function of the user's trust in the input data. Raksha tracks this at a coarse level--trusted vs. untrusted. But people use a more subtle mental model of trust that Raksha doesn't emulate. I think I'll have to give this more thought after Hot Chips, and maybe I'll come back to it here later.

In the meantime, the next session is on wireless communication...stay tuned to Speeds and Feeds.

Originally posted at [Speeds and feeds](#)



Peter N. Glaskowsky is a technology analyst for The Envisioneering Group. He is a member of the [CNET Blog Network](#), and is not an employee of CNET. [Disclosure](#).

Topics: [Security](#), [Processors](#)

Tags: [Hot Chips](#), [Sun](#), [Proximity](#), [T-RAM](#), [Raksha](#)

Bookmark: [Digg](#) [Del.icio.us](#) [Reddit](#) [Yahoo! Buzz](#)

Related

From CNET

[Intel completes 32-nanometer chip development](#)

[Allegations, denials of 'bad' Nvidia chips in Apple MacBook](#)

[IBM to start-up: Industry vet responds to recession](#)

From around the web

[LG Shows First Modem Chips For 4G Phones...](#) Wired

[Intel Pushes Ahead With Tiny Circuitry](#) Wall Street Journal

[More related posts](#) powered by

