

This article is more than **1 year old**

Applied Micro ARM servers all over Open Compute Project

Facebook friends, including Dell, spec out X-Gene microservers

 [Timothy Prickett Morgan](#)

Thu 17 Jan 2013 // 00:41 UTC

OPEN COMPUTE 2013 ARM server chip upstart Applied Micro Circuits was all over the Open Compute Summit hosted by Facebook and its server and data center designing friends in Santa Clara, with several partners flashing devices based on Applied Micro's forthcoming 64-bit X-Gene ARM-based processor. Company president and CEO Paramesh Gopi also gave a brief keynote address that talked a bit about the X-Gene roadmap.

You might have been expecting only one X-Gene chip to be delivered this year, especially considering that this is the first server chip to adhere to the new ARMv8 specification from ARM Holdings, the company that does the basic design of the ARM chip architecture and resells it to anyone who wants to give Intel some grief. But as it turns out, Applied Micro has two X-Gene chips in the works for this year.

The first generation chip has taped out and Taiwan Semiconductor Manufacturing Corp, the fab that X-Gene has chosen to etch its chips, has cranked out samples using its 40 nanometer processes. Gopi told Open Compute Summit attendees in Santa Clara that a second-generation X-Gene system-on-chip (SoC) would come out later this year, and that it would sport 100 gigabits of I/O and have an enhanced fabric compared to the first generation. This would enable X-Gene servers to start being used by enterprises, not just cutting-edge hyperscale data centers like those run by Google and Facebook.

As *El Reg* [confirmed with Gopi last summer](#) at the Hot Chips conference, the X-Gene architecture uses a two-core module as its basic building block (much like Advanced Micro Devices' last two generations of Opteron processors for servers). Each core a four-wide, out-of-order execution unit for integer work, and includes full virtualization support including nested page tables. Each core has L1 and L2 cache memory, and there is a shared L3 cache across the multiple core modules. The elements of the X-Gene SoC are glued together by a coherent network that runs at the full CPU clock frequency; it has a 15 nanosecond latency, delivers

160GB/sec of bandwidth, and has interfaces to network accelerators, 10GE ports, and PCI-Express 3.0 peripherals, and SATA slots.

The first generation X-Gene chips will have four core modules for a total of eight ARM cores running at a top speed of 2.5GHz, with 256GB of memory max across those cores. These chips are implemented in TSMC's 40 nanometer processes. The second generation X-Gene, which will presumably tape out before 2013 ends (which is what Gopi means by delivering it this year) will shrink to a 28 nanometer process, boost the number of cores on the die to sixteen, and push the clock speed up to 3GHz on the cores. That coherent network will also be extended off the die to 64 X-Gene sockets in a glueless manner, allowing for vendors to create a system with 1,024 cores all without a Layer 2/Layer 3 switch. In effect, X-Gene has the switching functionality built into its own network, which is implemented to look like Ethernet when it has to. It is not clear if that 64-socket system will have coherency across cache memory and therefore be able to host a single image of a Linux (and maybe someday Windows) operating system across the nodes. But that sure would be a neat thing.



Applied Micro CEO Paramesh Gopi

The third generation X-Gene, we learned today, is slated for delivery in 2014. The exact feeds and speeds of that chip have not been divulged, but what Gopi did confirm is that its fab will be moving to FinFET technology that is similar to Intel's TriGate 22 nanometer processes used to make its current Xeon and Core processors. These X-Gene chips will have lots of goodies thanks to the shrink, presumably to something 20 nanometers but Gopi did not say and he did not confirm that TSMC would be its fab during this generation when *EI Reg* asked during an interview after the keynote. (ARM Holdings has already partnered with TSMC to get its FinFET tech implemented on ARM cores at and below the 20 nanometer node, and it is highly unlikely that Applied Micro would use another fab or another technology.)

But what Gopi did say is that the future X-Gene chips would offer 10GE, 40GE, and 100GE networking with an aggregate of "fractional terabit" capacity across the chip. Gopi added that these X-Gene-3 chips would offer a "geometric reduction in TCO" compared to the original prior X-Gene chips, which Applied Micro reckons will offer a 67 per cent TCO advantage compared to Xeons in the data center.

The X-Gene chip is in Applied Micro's labs now being tested and will have its production ramped up in the coming quarters, giving partners access to the devices for testing in their own designs. Volume shipments are expected later this year.

Facebook and Dell show off X-Gene microservers

Applied Micro was not the only one talking about the X-Gene, and it is pretty significant that Frank Frankovsky, who is vice president of hardware design and supply chain at Facebook as well as chairman of the Open Compute Foundation, held up an X-Gene microserver that adheres to its new "Group Hug" Open Compute Common Slot specification, during his presentation.




An X-Gene ARM microserver that plugs into the Group Hug backplane

Frankovsky never once said that it was making a commitment to deploying ARM servers in Facebook's data centers, and has not confirmed that for the past two years when it looked like a reasonable possibility. But as a company that writes its own software stack atop Linux and other open source code as well as its own homegrown applications, Facebook certainly could compile all of its code on ARM servers if it wanted to.

Gopi told *EI Reg* that the X-Gene microserver design that it worked on with Facebook was considerably more compact than the one Applied Micro had cooked up and was showing off

back at Hot Chips and as the basis of prototype server and storage designs back in November last year. The feeds and speeds of the Facebook X-Gene board were not divulged, but the word on the street is that the system's design is expected to be given to the Open Compute effort as the Group Hug backplane has been.



Dell has forged an X-Gene ARM microserver code-named Iron

Dell also gave a quick glimpse of an X-Gene microserver design called "Iron" that it has built as a prototype for hyperscale customers to play with. Jimmy Pike, chief architect and technologist at Dell (which used to design and build Facebook's custom servers before the social network decided to do the job itself), showed off the server sled as part of a presentation concerning management interoperability between x86 and ARM servers.

The Iron sled servers slide into Dell's PowerEdge C5000 chassis, code-named "Viking" and sporting "Dragon" Intel Xeon E3 v2 microservers and "Copper" Marvell Armada XP ARM microservers. The Dell Iron server sled looked like it had three two-socket X-Gene motherboards on the sled. ®