

Ciara Technologies rethinks supercomputer design

12/8/2004 5:00:00 PM - Quebec company uses commodity components to make system affordable

by Poonam Khanna

Ciara Technologies says it wants to bring supercomputing to the masses and announced Tuesday the launch of a supercomputer it says rethinks the way high-performance computing (HPC) clusters are designed.

Ciara also announced Tuesday plans to make an HPC benchmarking centre available starting in Q1 of 2005 for businesses, governments and academic institutions in need of computing power.

Ciara's VXR-3DT can scale from 16 up to 19,440 Intel Extended Memory 64 Technology Xeon processors, providing more than 140 teraflops of computing power, said to Patrick Scateni, the VXRACK co-creator and director of business development for high-performance computing at Saint Laurent, Que.-based Ciara.

Each motherboard has two Xeon processors and an integrated switch. Two motherboards are paired together and the switch on each is connected to the four processors on the board, providing two routes to each processor. If one switch fails, the network fabric can still connect to the four processors through the second switch. Traditional supercomputers are "poorly designed," Scateni said. The hierarchical design of conventional HPC means if a switch breaks down, all of the processors that connected to it become unavailable, he said.

"If I lose a switch at a point on the layer, I lose everything below it. That's not a great thing," he said.

HPC clusters also have unneeded components that can increase the chance of failure, he said. The 3DT, which runs on the Linux operating system and has native storage, has no video, no USB, and no extra SCSI, Scateni said. It does, however, have two ports for feature expansion.

The computer uses commodity components, such as the Xeon processors and InfiniBand, keeping its cost low, Scateni said.

The company's goal is to offer "50 per cent of the speed at 20 per cent of the price," Scateni said.

The benchmarking centre will give users affordable access to up to 640 Intel Xeon processors on a VXR-3DT, Scateni said, adding Ciara also plans on offering free CPU cycles to its local community in order to promote R&D.

Ciara worked with **Intel** and **Raytheon Co.** to develop the cluster. Raytheon developed the scheduler software.

The problem with tier-two HPC manufacturers is that they're hardware manufacturers, and there's no common software to make the clusters run, said James Ballew, high performance computing architect at Dallas-based Raytheon. This means setting up an HPC cluster is a project that can take months, he said.

"The thing that made PCs happen really was the availability of a common software system — Microsoft's — that everybody could use. It was open to anybody that built a PC," he

ROI
ASAP.



IBM eServer
xSeries systems



said. "Well the thing that's missing at this level is that there's a lot of pieces available, but there's no standard integrated software that these manufacturer companies could just plug into. We saw this as something we needed to help make happen. The idea here is to have a system you can install on the user side, and the software is there to make it work. And the users can go right to work on it instead of it being a computer science project for the next six months."

Several processors co-operate to solve a problem and in a machine with a 3-D structure, you have to take into account the topology of the machine and the topology of the problem when you assign tasks to a processor, Ballew said.

For example, computational fluid dynamics problems are normally expressed either 2-D or 3-D. If you're splitting a problem up between two processors on a 2-D computer, you'd want to put pieces of the problem between two nearest neighbours to shorten the time it takes for them to communicate with each other, because some problems are tightly coupled, Ballew said.

"The overhead of communicating is very high."

Problems placed on a 3-D system like Ciara's have to be reconfigured depending on their nature, Ballew said.

"What you do is rearrange the geometry of the problem to fit inside the geometry of the processing array. For instance, if you had a one-dimensional problem, if all of the processors want to communicate equally to each other, then you start in the centre of the array and then assign the processor to the next open location that is the shortest distance to the centre — like winding up a ball of twine. Every wrap is the closest it can be to the centre, and you just go out from there."

With 2-D problems, you can fold up in a number of ways to fit in a 3-D array, he said.

A 3-D problem you'd try to fit in as it is, Ballew said.

But if the natural state of a 3-D problem is 3x6x3, and the processors available in the array aren't configured in the same way, then you rotate the problem to try to fit it in, he said.

Keeping the board at a reasonable size was a real challenge, said Hugues Morin, a field sales engineer at Intel Canada Ltd. in Montreal.

While the board removed all unnecessary components, integrating the two chipsets and switches made real estate a challenge, he said.

"What's unique in this design is the amount of redundancy," said Doug Cooper, country manager for Intel of Canada in Toronto.

Comment: info@itbusiness.ca

