

Speeding EDA Design Cycles

Synopsys is a world leader in delivering semiconductor design software and professional services for chip manufacturers. The company's products enable semiconductor, computer, communications, consumer electronics and other companies to increase productivity and streamline the process of moving from design through manufacturing to finished silicon. Founded in 1986, Synopsys is continually advancing chip design capabilities for semiconductor companies around the world.

Synopsys maintains its leadership not only by continually upgrading its software products, but by delivering outstanding customer service that includes troubleshooting customer production jobs as well as consulting on appropriate computing platforms for customers' design-for-manufacturing work. Among its portfolio of solutions, Synopsys offers a product called CATS™, which prepares semiconductor design data for photomask manufacturing. With installations in virtually every photomask manufacturing facility worldwide, CATS has become the de facto standard for mask manufacturing, inspection and direct-write-on-wafer data preparation.

But as designs become more complex and design file sizes increase, it takes longer and longer to process the files. To stay ahead of the curve, Synopsys has invested in world-class cluster computing technology. Thanks to solutions from Voltaire, Mellanox Technologies, DataDirect Networks, and Cluster File Systems, Synopsys has been able to build and scale its computing performance to remain at the forefront of semiconductor design and support for the foreseeable future.

Massive Files Need High-Performance Infrastructure

In a world where the number of components on a chip still doubles about every two years, semiconductor makers relentlessly move to smaller and denser designs to improve the functionality of their products—placing billions of transistors in complex architectures ever closer together—often requiring complex data modifications and enhancements to ensure manufacturability.

CATS analyzes semiconductor designs and then 'fractures' the data, breaking down each of the design elements into fundamental structures and layouts to be reproduced on a photomask and eventually onto the layers of semiconductor material at the chip. Depending on the device complexity, this process can increase design file sizes from hundreds of megabytes to tens of hundreds of gigabytes. In fact, according to Naji Bekhazi, Director of CATS R&D at Synopsys, as designs continue to cram more into a given space, 2008 will bring terabyte design file sizes, and file sizes will increase to as much as six terabytes by 2011.

For Synopsys, ever-expanding file sizes require an ongoing investment in greater computing horsepower. Initially, the company built compute clusters using NFS-based storage systems on a Gigabit Ethernet interconnect, but it found that these technologies severely limited scalability and processing performance. According to Glenn Newell, senior staff IT Architect at Synopsys, the NFS-based system worked well up to about 32 CPUs, but above that it began to bog down due to latency in the interconnect and file system. "At some point, says Newell, "performance tails off as you scale. NFS and

Gigabit Ethernet works with 32 CPUs, but when you go up to 64 CPUs, the performance actually drops because the bandwidth becomes a problem."

Another part of the problem was that the NFS-based storage couldn't scale properly. "With a traditional file system like NFS, once you get to the maximum a file server can deliver, you can't get more out of it," says Newell. "You can add more servers, but because CATS is trying to read one single file, it doesn't help the program to have more file space on different volumes—the file space needs to be in one location. Even clustered NFS servers don't help—you need a parallel file system where you're striping the file across multiple file servers to get the aggregate performance."

In an effort to overcome these challenges, Glen Newell's team partnered internally with the CATS R&D team to spec and build a next-generation, high-performance compute cluster optimized for CATS and for customer mask data preparation applications. The goal was to try to overcome the performance bottlenecks of past systems by looking at other cluster computing solutions.

Initially, Glen's IT team built additional test clusters using another parallel file system, and it tried proprietary optical-based interconnects. However, the team wasn't pleased with these solutions' performance either. Ultimately, Newell realized that there were three requirements for a true solution: a high-speed, low-latency interconnect, a parallel file system, and high-performance storage.

A Massively Scalable Interconnect

Tackling the interconnect first, the Synopsys team selected a Voltaire® Grid Switch™ ISR 9024 10 Gigabits/second InfiniBand switch to connect all of the servers in the cluster. The Grid Switch features 24 InfiniBand ports that can support speeds of 10 to 20Gbps per port. By using InfiniBand technology's remote direct memory access (RDMA) that bypasses system processors during I/O and writing directly to disk, the wire-speed, non-blocking Voltaire switch delivers latency of less than 420 nanoseconds—some 20 times faster than Gigabit Ethernet and four times faster than 10-Gigabit Ethernet.

During design and deployment of the new clusters, technical staff from both Voltaire and Mellanox (the InfiniBand silicon and HCA supplier) helped Synopsys tune the switch and interconnect configurations for optimum performance. "Voltaire and Mellanox really impressed us with their proactive suggestions and their willingness to help us build a system that exceeded our expectations," says Newell.

The first new compute cluster used 32 AMD dual-core Opteron™-based servers managed with Lustre™ cluster file system software. The storage solution was direct-attached SATA Disks.

Since then, Synopsys has built two additional clusters, one with 32 nodes and another with 50 nodes. These two clusters share a common network using a Grid Director™ ISR 9288 96-port 10 Gigabits/second InfiniBand switch. The two clusters share 32 terabytes of DirectData Networks storage. The largest cluster includes 50 nodes (48 of which are dual-

socket nodes populated with AMD dual-Opteron processors) that combine the computing power of 200 compute cores.

"Basically, what we're doing is using supercomputer technology to deliver high-performance storage, and this was what led us to InfiniBand, since it's used by almost 27 percent of the top 500 high performance computing sites," says Newell. "Those same sites are what led me to Lustre as the parallel file system."

Giving CATS More Agility

With the InfiniBand-based interconnect, a parallel file system, and a high-performance storage appliance, the performance improvement has been dramatic. "CATS uses distributed processing," says Newell. "The file sizes are too big to be contained in memory, so they have to be in storage. CATS has to access this file and process it and write it back to disk, and this puts a tremendous strain on the disk farm and on the network. This is where we've seen that InfiniBand has helped a lot. The wider bandwidth has allowed more CPUs to communicate at the same time."

Before using the current clusters, Synopsys customer support teams needed days or even a week to process a customer job for analysis. With the InfiniBand-based cluster, the process now takes about 8 hours.

"It's basically ten times faster," says Newell. In addition, storage performance has increased about 17 times over the previous NFS-based storage solution.

Another InfiniBand feature that has improved operations is superior error handling. Any error in data transmission across the interconnect or in the storage system can end up costing hundreds of thousands of dollars for a new mask while delaying a new product's appearance on the market. InfiniBand's RDMA technology and ultra-low latency slash the error rate when compared with Ethernet technologies. The high-performance DirectData Networks storage appliance also reduces errors by striping data across multiple drives to eliminate errors due to failures in a single disk.

What the speed improvement means for Synopsys customers is that they can fracture files for mask preparation far more quickly, which enables them to process more jobs in a given period of time to increase revenue and productivity.

Supporting Future Growth

Today's speed marvel is tomorrow's also-ran technology, so the Synopsys team continually keeps pushing for higher performance. According to Naji Bekhazi, customer needs are tied to business cycles as well as increasing file sizes.

"Customers want to be able to fracture a design in less than eight hours, and that's our goal," Bekhazi says. "If customers can fracture a 500-gigabyte file in less than eight hours today, then they'll need to fracture a 1-2-terabyte file in less than eight hours next year. We have to keep improving the cluster's performance."

Synopsys uses 50 dual-core servers to process 600-gigabyte files in 16 hours today, but with a goal of processing 1-terabyte files in eight hours within a few months, it will be beefing up processing performance on the server side. "Today we're throwing more than 200 cores at the problem," says Bekhazi, "but in the future we'll be using hundreds of servers and over a thousand cores to reduce the job cycle time." The effort will be helped with the emergence of quad-core processors.

Fortunately, the InfiniBand interconnect has plenty of capacity to handle several years' worth of upgrades to server processors and cluster enhancements. "We can always go to the double data rate (20 Gigabits/second) InfiniBand," says Newell, "and Mellanox is telling us that they'll have 40-gigabit interconnect silicon available in 2008, so we're confident in having a technology path that will get us down the road for a few years."

Today, the high-performance computing clusters, high-performance storage, and massively scalable InfiniBand interconnect at Synopsys give it the power and scalability to continue rapid product development and outstanding customer support. With plans for additional processors and higher InfiniBand throughput in the works, the company is well positioned to increase the performance of CATS while reducing processing times for its customers.

Voltaire
6 Fortune Drive
Billerica, MA 01821
www.voltaire.com
800-865-8247

Mellanox Technologies
2900 Stender Way
Santa Clara, CA 95054
www.mellanox.com
408-970-3400

Synopsys
700 E. Middlefield Rd.
Mountain View, CA 94043
www.synopsys.com
508-263-8006

DataDirect Networks
9351 Deering Avenue
Chatsworth, CA 91311
www.datadirectnet.com
800-837-2298

Cluster File Systems, Inc.
4800 Baseline Road
Boulder, CO 80303
www.clusterfs.com
415-928-3633

