

Cluster Computers Speed EM Simulation

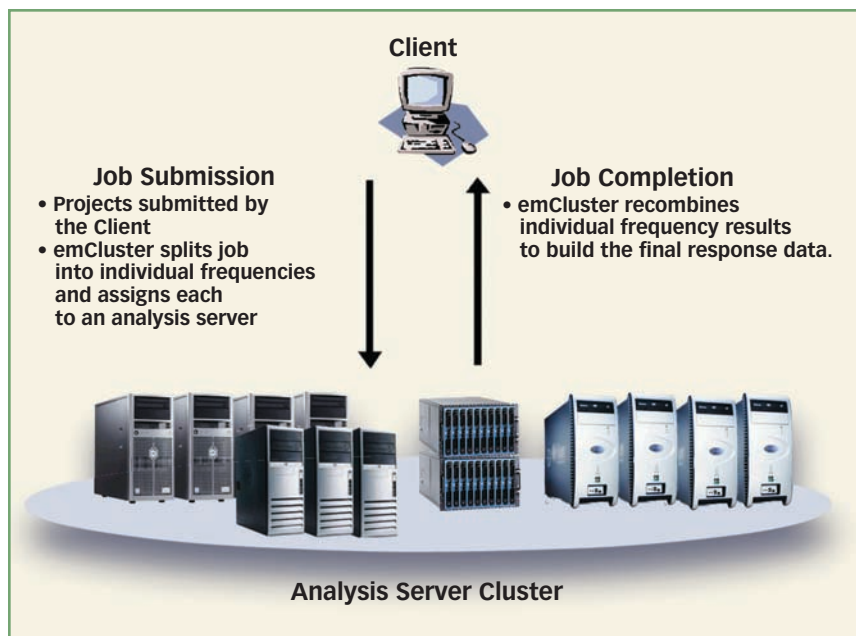
Cluster computers can greatly reduce the processing time of an electromagnetic simulation by combining the processing power of multiple CPUs in a single computing platform.

YUN CHASE
Technical Marketing Manager

BRIAN RAUTIO
Hardware Engineer

Sonnet Software, Inc., 100 Elwood Davis Rd., North Syracuse, NY 13212; (877) 7SONNET (776-6638), (315) 453-3096, FAX: (315) 451-1694, Internet: www.sonnetsoftware.com.

Electromagnetic (EM) simulation has become a powerful tool in modeling high-frequency and high-speed circuits and devices. As the software grows in sophistication, however, computer hardware strains to keep up, often requiring innovative combinations of computer systems to efficiently run modern EM software. One integrated solution is the emCluster[®] Computing approach from Sonnet Software (North Syracuse, NY). It essentially combines the power of multiple computers for faster EM analysis and improved reliability, resulting in shortened design cycles and faster time to market. The emCluster solution can be configured to run on a centralized local cluster. It can also be set to run across a distributed or grid-style computing environment over a Wide Area Network (WAN) or over



1. This diagram shows the typical emCluster process.

the Internet using a Virtual Private Network (VPN) connection. Unlike traditional clusters, it can also run on a heterogeneous environment much like grid computers.

emCluster Computing (Fig. 1) resembles various facets of cluster

computing, distributed computing, and grid computing. Cluster computing makes use of multiple stand-alone homogeneous computers acting in parallel across a high-speed LAN and, in many respects, can be viewed as though

Table 1: Summarizing computer architectures

	CLUSTER	DISTRIBUTED	GRID
NETWORK	LAN	LAN/WAN	WAN/Internet
GEOGRAPHICAL AREA	Centralized	Central/Regional	Global
COMPUTER TYPE	Homogeneous	Heterogeneous	Heterogeneous

it is a single computer. In a distributed computing environment, computers are not exclusively running “group” tasks and are not as tightly coupled as in a cluster computer. When properly configured, distributed computers can utilize computational resources that would otherwise be unused. Distributed computing can ensure delivery of reliable and “always-on” availability of computing resources that would otherwise be impossible.

Grid computing is essentially an evolution of distributed computing where a grid consists of multiple computer resources connected by a network (often the Internet) to solve very large numerical problems. Grid computing can be configured to use idle time on many computers throughout various geographic locations. Such arrangements permit handling of data that would otherwise require the power of expensive supercomputers or would have been impossible to analyze.

Although often confused with cluster computing, grid computers are quite different. For example, clusters are homogeneous while grids are heterogeneous. In addition, grids are spread out geographically while clusters are generally confined to a central location (**Table 1**). Regardless of the architecture being used, various clusters and grids are usually deployed to improve speed and/or reliability over that provided by a single computer, while typically being much more cost-effective (**Table 2**) than single computers of comparable speed or reliability.

When used with emCluster, Sonnet analysis can be N times faster than executing on a single computer where N equals the number of computing resources within the cluster. As a result, the analysis time for large projects can decrease from days to just hours.

With faster analysis, designers can opt to design their circuits more compactly without concern for extensive EM analysis time.

Even with only two computers (N = 2), the time saved during analysis is twofold. However, for a typical Sonnet analysis project using Adaptive Band Synthesis (ABS), a cluster made of 10 computers (N = 10) offers a good balance between improved speed, reliability, and setup cost. The cost of a dedicated cluster computer is comparable to that of a UNIX workstation about 10 years ago. Cluster computers can be purchased as ready-made systems or assembled from off-the-shelf parts. Each computer within the cluster must have a license for LSF, cluster management software from Platform Computing (www.platform.com) and a license of Sonnet Suites Professional™ Release 10.53. The hardware itself can be fairly commonplace personal-computer (PC) equipment. emCluster interfaces directly with LSF and it also allows clients to be located anywhere around the globe as long as the client is able to establish a network connection to the cluster using VPN or the like.

Traditional EM analysis on a single computer requires each frequency to be analyzed one after another in series. However, with Sonnet emCluster in a cluster computing environment, the entire frequency sweep is intelligently split into as many single-frequency jobs as required by Sonnet’s EM analysis engine. Immediately, these single frequency jobs are individually assigned by emCluster to be processed in parallel on the resources that are available on the cluster. emCluster automatically schedules the jobs for analysis on simulation servers only if the servers meet the job requirements (such as available RAM). This intelligent scheduling by emClus-

ter significantly reduces the overall wait time and ensures delivery of “always-on” availability of resources.

Once all of the frequencies are analyzed by the cluster, they are recombined by emCluster to build the final response data. The process is transparent to the user except for the significant decrease in analysis time and fast but accurate results. Even a cluster of two computers results in a 2× benefit in time saving. When used with a 10-computer cluster (a 10× time savings), a 10-hour project is reduced to 1 hour.

emCluster can also be used to increase the efficiency of an Adaptive Band Synthesis (ABS) sweep. ABS is an interpolation method that provides a fine resolution response for a frequency band requiring only a small number of EM analysis points. Sonnet performs a full analysis at a few discrete points and uses the resulting internal, or cache, data to synthesize a fine resolution band. There are two basic methods of accomplishing an ABS sweep with emCluster. The first method, called the Automatic Algorithm, allows the emCluster to determine at which discrete data points to run a full analysis and run them parallel on the cluster. In the second or User-Defined method, you may use the Frequency Sweep Combinations to

Table 2: Reviewing typical component costs for a fast cluster

COMPONENT	US DOLLARS
3.2-GHz Pentium 4	275
4-GB RAM	375
40-GB hard drive	50
2U rack-mount case	100
300-W power supply	50
Motherboard	150
Total (per computer)	1000

When used with emCluster, Sonnet analysis can be N times faster than executing on a single computer where N equals the number of computing resources within the cluster.

define linear sweeps or single discrete frequency points at which to run an analysis before attempting the ABS sweep and running them parallel on the cluster.

To demonstrate the effectiveness of emCluster, consider the example of a hybrid power splitter (Fig. 2). This 3-dB in-phase power splitter was implemented in GaAs technology and analyzed with conformal meshing on thick metal to model the compact spiral inductor. The total footprint of the splitter occupies an area that is less than 0.5×0.5 mm. Such curved-line structures have traditionally presented a challenge for EM field solvers, which are trying to fit a large number of analysis subsections into the curved structure (with computer time eaten by the increasing number of subsections). Sonnet's conformal meshing technology analyzes the curved structure with a practical number of subsections that still delivers high accuracy.


Using a single 1.8-GHz Pentium 4 PC with 2 GB RAM, analysis of the power splitter at 10 frequencies using a fine mesh for high accuracy required 9 hours and 41 minutes. Analysis with emCluster running on a 10-node cluster system based on the same exact Pentium 4 computers specifications required just 59 minutes (Figs. 3 and 4). emCluster automatically divided the job into 10 frequencies and intelligently assigned (based on cluster specifications and availability of resources) each individual fre-

Table 3: Sizing up typical component costs for a low-power cluster

COMPONENT	US DOLLARS	QUANTITY	TOTAL US DOLLARS
Motherboard	237	10	2370
1.8-GHz Pentium M	309	10	3090
2 x 1-GB RAM	189	10	1890
40-GB hard drive	50	10	500
200-W PSU	50	10	500
Case	100	10	1000
Heatsink	20	10	200
LED lights	10	10	100
Ethernet cables	4	15	60
Network switch	55	1	55
12U rack	400	1	400
Power distributor	85	1	85
		Total	10,250

Table 4: Checking typical component costs for blade servers

COMPONENT	US DOLLARS	QUANTITY	TOTAL US DOLLARS
Blade enclosure	1499	1	1499
Patch panel	1172	1	1172
2.4-GHz 1M 1P blade	2499	16	39,984
Blade sleeve server	199	8	1592
		Total	44,247



2. This compact hybrid power splitter fabricated on GaAs was used to compare the speed on a Sonnet EM analysis on different computer systems. The emCluster system showed about 10 times the speed of a single computing resource.

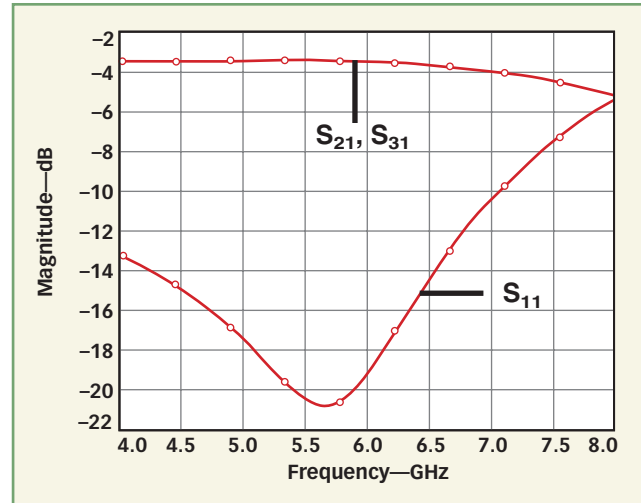
Since the fastest processors draw the most power, a cluster system's power requirements can be considerable.

quency calculation to an available cluster computer.

Build An emCluster

Constructing an emCluster computer system can be as simple as narrowing down choices of hardware and peripheral devices. Using 1U-space (1.75 in. thick and 19-in. wide) rack-mount cases, a cluster can be built in little more space than a conventional computer. Because they can be accessed remotely, things like a mouse, keyboard, monitor, optical drives, and sometimes even hard drives are unnecessary—further reducing cost, space, and power. Network switches and power distribution are also commonly integrated into racks, simplifying external connections. Cluster hardware is designed for reliability. For example, memory errors are minimized with error correction and checking (ECC) RAM. Hard-drive problems can be minimized with RAID arrays. Data sent to a RAID controller is copied and written to two hard drives. If one hard drive crashes, the RAID array continues to use the functional drive until the other is replaced. It then copies the data over to the new drive.

Redundant power supplies can minimize power-supply-related failures. Redundant power supplies consist of two separate power supplies in one enclosure, working together to share the load. If one goes down, then the load is immediately handed off to the




3. This simulated response data for the GaAs hybrid power splitter was generated in less than an hour with the emCluster system.

functioning power supply, without even requiring a computer reboot.

Profiling a user base is essential to establishing the requirements for a cluster system. A profile is based on knowing the number of Sonnet users, typical problem sizes, times of greatest activity, user platforms, and geographic locations of users. The next step is to define a set of environmental specifications, such as power and electrical connections, cooling requirements, and noise considerations.

Once these have been defined, a decision can be made on vendors (using such trade journals as *ClusterWorld* to find companies). It makes sense to classify vendors according to the level of support they can provide once the cluster has been assembled and connected to the network.

The next step is to write a Technical Request for Proposal (RFP) addressed to selected vendors listing specific requirements for the cluster system. The RFP should



4. This 3D plot of surface current density for the GaAs hybrid power splitter was generated with the emCluster system.



5. The first cluster system constructed at Sonnet included 10 nodes and blue LEDs for a touch of style.

include such information as the number of nodes, the type of processors, amount of memory, type of motherboard, type of Ethernet card, hard drive, video graphics card, riser card (if necessary), type of rack mount, power suppliers, operating system, cluster management software, additional software and compilers, network requirements and IP addresses, cluster physical requirements, and system cooling requirements. Of course, costs, delivery time, and warranties will vary, but a quick survey of the New England area revealed a 5-to-6-week delivery from one vendor with two-year warranty on hardware and lifetime technical support on hardware and operating system related issues.

Since the fastest processors draw the most power, a cluster system's power requirements can be considerable. For example, the power draw from ten 3.2-GHz nodes is in the 3-kW range—so multiple 15-and/or 20-A outlets become advisable. Large clusters frequently are in separate rooms, with dedicated air conditioning to control the heat.

It is often more practical to construct a small cluster than to have a vendor assemble it. When assembling a cluster, a number of companies are available to provide

components or assistance, including www.newegg.com (system components); www.zipzoomfly.com (system components); www.outpost.com (components and peripherals); www.frozencpu.com (cooling + customizations); www.microway.com (cluster hardware and integration); www.pricewatch.com (price search engine); and www.shopping.com (price search engine).

Important requirements for the first cluster computer built at Sonnet included low power, low heat, and low fan noise (in case the system was on display at a trade show). The logical starting point for a low power cluster is a notebook central processing unit (CPU). Sonnet's cluster is based on 1.8-GHz Intel Pentium-M processors. Despite the low clock speed, Pentium-M processors offer more instructions per clock (IPC) and more cache memory than traditional Pentium 4 processors. The net result is higher performance. A single 1.8-GHz Pentium M is comparable to a 3.4-GHz Pentium 4 in a Sonnet benchmark. Multiplied by 10, this yields serious computer power: 40 Gigaflops of numerical performance. To complement the performance of the CPU, each system is loaded with 2 GB of 400 MHz PC-3200 RAM. The i855gme motherboard was used in the first Sonnet cluster system—one of only three desktop motherboards on the market today with support for Pentium M processors.

A 40-GB hard drive and a 200-W power supply complete each cluster node. Blue light-emitting diodes (LEDs) were added to the system for aesthetic reasons (**Fig. 5**), but these turned out to be practical when they helped to reveal a power-supply problem (when one set of the LEDs failed to light). **Table 3** shows the total cost of materials (about \$10,000 US).

The Pentium M processors con-

tribute greatly to the overall low-power consumption of the Sonnet cluster. A 1.8-GHz Pentium M consumes only 21 W compared to a 3.4-GHz Pentium 4 at 100 W. The power draw on a full Pentium M computing node measures 55 W under load. The 10-unit cluster draws less than 600 W, well within the limits of traditional 110-VAC, 15-A outlets in the US. Less power draw, of course, also means less heat, which means less cooling is required. With some slight modification, the traditional six high-rpm fans were replaced with only two low rpm fans, which produced virtually no audible noise.

Of course, this may be the first cluster system based on Pentium M processors. It is so unique that Intel has purchased one of the systems for evaluation. And since there is no dedicated Pentium M cluster hardware, the desktop hardware available required considerable customization to work with the Pentium M units.

For example, being a mobile CPU, a heat spreader, which normally protects the CPU core from damage, is not used so as to save weight and space. This places the CPU approximately 2 mm lower in the socket. This put the CPU core 1 mm below the heat sink, and thus no physical heat sink-CPU contact was possible. The motherboard included a special low seating heat sink because of the heat sink alignment issue, but that heat sink was approximately 1 cm too tall for a space-saving 1U case. Therefore, Sonnet's hardware specialist designed a custom heat sink retention mechanism that holds the heat sink lower and in reliable contact with the CPU.

Blade servers are the pinnacle of today's low to moderate cost cluster hardware. Highly optimized, blades reduce space and power draw by sharing resources. Extremely thin, they are hot-swappable

stored side by side in an enclosure that shares power, networking, and storage. Each actual blade only needs to have a CPU, motherboard, and RAM. The high non-recurring cost of the enclosure, however, only starts to make blades cost effective in quantities above about 40. **Table 4** illustrates a 16-unit blade that occupies about 6U (six 1U servers) of rack space.

Special cluster hardware isn't always necessary to run a cluster. Many facilities already have an existing network of PCs, which

can be used as a cluster. Even just two spare computers on a network can run Sonnet twice as fast. Cluster computing in this way requires no additional hardware cost, and otherwise idle computers are put to productive use. It should be noted that using standard PCs in a cluster takes it away from other tasks. In addition, the slowest PC in a cluster will essentially limit the speed of the other computers in the cluster.

In short, the new emCluster Computing from Sonnet Software

offers many benefits but most importantly it offers faster EM analysis along with improved reliability, which in the simplest of terms, translates to shortened design cycle and faster time to market for the end user. **MRF**

FOR FURTHER READING

"Cluster Computing," Wikipedia,
http://en.wikipedia.org/wiki/Computer_cluster
John F. DeFord, "Electromagnetic Modeling on Computer Clusters." Simulation Technology & Applied Research, Inc., available from RFGlobalnet,
[http://www.rfglobalnet.com/content/news/article.asp?docid=\(cc665bff-3f7d-426e-9951-4fbfa937a327\)](http://www.rfglobalnet.com/content/news/article.asp?docid=(cc665bff-3f7d-426e-9951-4fbfa937a327))

Copyright © 2005 by Penton Media, Inc.



Sonnet Software, Inc
100 Elwood Davis Rd.
North Syracuse, NY 13212 USA
Phone: 315.453.3096
Fax: 315.451.1694
Web: www.sonnetsoftware.com