

[Home](#) > [iSGTW 4 November 2009](#) > Feature - Getting GPUs on the grid

Feature - Getting GPUs on the grid

Enhancing the performance of computer clusters and supercomputers using graphical processing units is all the rage. But what happens when you put these chips on a full-fledged grid?

Meet "[Magic](#)," a supercomputing cluster based at the University of Buffalo's [CyberInfrastructure Laboratory](#) (CI Lab). On the surface, Magic is like any other cluster of Dell nodes. "But then attached to each Dell node is an nVidia node, and each of these nVidia nodes have roughly 1000 graphical processing units," said [Russ Miller](#), the principal investigator for CI Lab. "Those GPUs are the same as the graphical processing unit in many laptops and desktops."

That's the charm of these chips: because they are mass-manufactured for use in your average, run-of-the-mill computer, they are an extremely inexpensive way of boosting computational power. That boost comes at a price, however.

"These roughly 1000 processors on each nVidia node are programmed in a synchronous process, basically bringing us back to programming methods of the 1960s," said Miller.

The parallel programs modern supercomputers run are already quite difficult to write. Synchronous programming is a more limited form of parallel programming. "Parallel means doing multiple things at the same time," explained Miller. "Synchronous means doing the exact same thing at the same time."

Synchronous computations could be processing very different sets of data, as long as the algorithm used is identical. For example, an algorithm could instruct two people to kick the object in front of them. If one is playing soccer while the other is learning self-defense, the instruction may be identical, but the context, meaning and effects are quite different. "The job becomes very demanding for a programmer to be able to exploit these roughly 13 000 processors that we have in one rack. But if they can, the returns are huge," said Miller. "We can get roughly 50 teraflops of computing out of one rack of systems."

In a perfect world, scientists could submit their computational jobs to a scheduling application, and the scheduler would take care of finding computing resources. "I want to be able to be lying on a beach in Cancun with my iPhone, and hit a button, and not have to worry about where my data is, or what resources it's using," said Miller. "But there's no way you could submit to the grid and have it assign a GPU for you. That logic is not built into the software stack just yet."

To do so, resource providers would need the ability to specify that they can only handle synchronous



Russ Miller, principal investigator at CI Lab, stands in front of the server rack that holds Magic, a synchronous supercomputer that can achieve up to 50 Teraflops. *Image courtesy of CI Lab.*



[FC](#)
[ste](#)

[Fe](#)
[th](#)

[Fe](#)

[Ne](#)
[Vi](#)
["H](#)

[Im](#)
[th](#)



[iSt](#)

[Re](#)
[Co](#)
[Po](#)

[Ne](#)
[to](#)

[SC](#)
[Co](#)

[Jol](#)



[En](#)
[su](#)
[ne](#)

[en](#)

[Ur](#)



[Ke](#)
[co](#)



computations, and users would need to be able to specify what sorts of resources their computations can exploit.

[Nc](#)

In the meantime, Magic has been hooked up to [Open Science Grid](#) and the [New York State Grid](#) since February. And instead of relying on a high-tech scheduler to assign jobs to the cluster, CI Lab has been relying on much older 'technology' – word of mouth.

4-
Ca

"One of the biggest sources of users we have seen so far is just word of mouth," said Kevin Cleary, the system administrator for Magic. "So getting the word out that on these nodes we do have these massive amounts of power available."

4-
Ca

Once a researcher is aware that Magic is available, he or she can tell the scheduler to submit the job directly to the GPU cluster.

5-
Pr

As other GPU clusters come online, word of mouth may become an impractical solution. In the meantime, however, it is working well for Magic. Said Cleary, "In the past week, nearly 2500 jobs have been run on this cluster with a 98 per cent success rate."

9-
[Mz](#)
Mz

—Miriam Boon, *iSGTW*

9-
Sy

Tags: [Americas](#) [Feature](#) [Infrastructure](#) [Open Science Grid](#)

11

Share this page:

[th](#)
[W](#)
Sv



[Email this page](#)



[digg.com](#)



[reddit](#)



[del.icio.us](#)



[StumbleUpon](#) [Ma.gnolia.com](#)



14
U.

Disclaimer:

These are external Web sites and iSGTW cannot guarantee their security nor endorse their content.

[Mz](#)