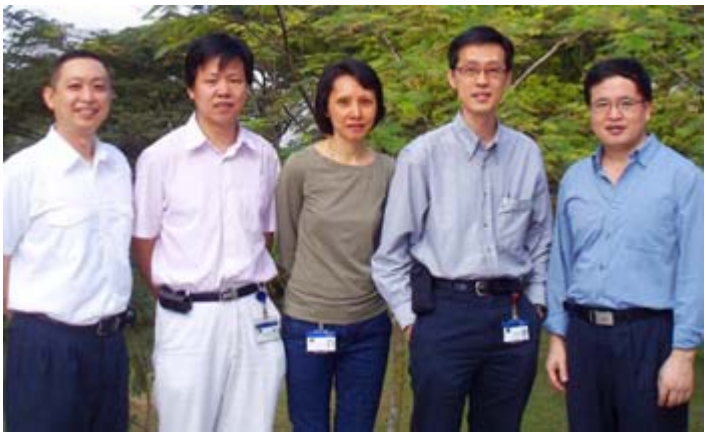
iSGTW INTERNATIONAL SCIENCE GRID THIS WEEK

Feature - Cycle-harvesting around campus in Singapore

At the National University of Singapore (NUS), researchers in the areas of bioinformatics, financial simulation and data-mining had a problem: It was taking days, if not weeks, to run a computation.

The situation encountered by one researcher at the NUS Business School, Dr Ding, was typical. He found that each simulation he needed to do took over ten days to execute on a desktop computer.

Clearly, something had to be done.



The NUS Grid Project Team: (from left) Yeo Eng Hee, Wang Junhong (Project Lead), Grace Foo, Tan Chee Chiang (Manager), Zhang Xinhuai.

Image courtesy of NUS Grid

Then came the NUS Grid (also known as Tera-scale Campus Grid or TCG@NUS), overseen by the Computer Center team at NUS. By using technology similar to that of well-known community grids such as SETI@HOME, FOLDING@HOME and World Community Grid, the team found that they could dramatically cut down computation time.

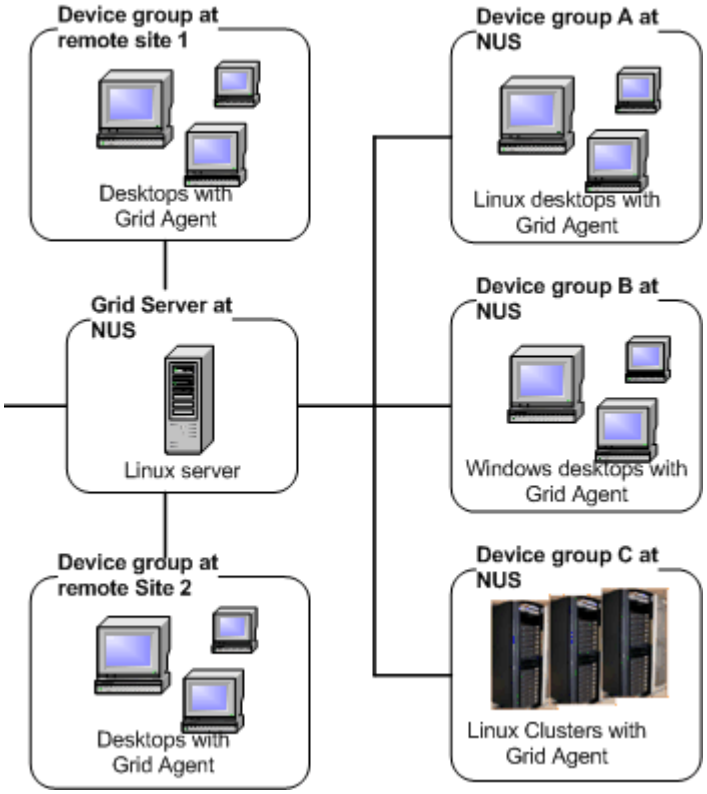
“With the NUS PC grid, we can finish one execution of our programs in several hours,” Ding said. “By running multiple instances on the grid, as required in this project, we were able to complete over a hundred days’ work in just several days.”

Researchers in other areas followed suit, and during the first two years of NUS Grid, investigators running large-scale bioinformatics applications such as BLAST, AUTODOCK, Modeller and HMMER have reduced their computation time from days and weeks to hours. The NUS grid is now being used on everything from financial engineering to data-mining

Divide-and-conquer

NUS Grid works by harnessing idle compute cycles from about ten independent departments across campus, including 1400 PC nodes and 222 multi-core server nodes (with around 1200 cores), which together provide more than 2600 processor cores.

Unused compute cycles are “harvested” and channeled to perform useful tasks. Applications that run on them use the same divide-and-conquer methodology as that of volunteer computing on home PCs; they chop a large computational task into many small tasks that can be performed in parallel on grid nodes.



Schematic view of the NUS Grid.

Image courtesy of NUS.

Although smaller than community grids, NUS Grid's faster campus network, with at least ten times more bandwidth than home PC connections, allows faster dispatch of data-intensive applications. NUS Grid has inter-connected large memory server nodes with up to 16GB of memory each, thus enabling execution of larger memory applications and broadening the range of applications supported.

Team member Tan Chee Chiang says that "We anticipate growth in the multi-core server technology in the coming years, and plan to add more of such computing capability on the NUS Grid. We believe it will offer a greater competitive edge to our research community in enabling new discoveries."

—Updated and reprinted with permission from IT@NUS "Is Grid a solution for researchers?"

Tags: Asia Business Data management Feature Project Profile TeraGrid

Share this page:

Email this page | digg.com | reddit | del.icio.us | StumbleUpon | Ma.gnolia.com

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