A grid cooperative

FEATURE | AUGUST 17, 2011 | BY ADRIAN GIORDANI

Large-scale scientific facilities such as the LHC at CERN can muster enough support to build large grid computing infrastructures to support particle collision analysis.

Smaller labs face similar challenges, but without the large-scale computing infrastructure.

“To be competitive, they somehow need to collect, manage and analyze . . . tens of terabytes of data — just like the big guys,” said 'father of the grid' Ian Foster, senior scientist at Argonne National laboratory.

OurGrid

Smaller experiments require a smaller grid such as OurGrid, which is both a peer-to-peer (P2P) grid middleware and a desktop grid service. It pools the computing resources of all the computers within a single laboratory or university for researchers to use. OurGrid is the ‘glue’ that brings data, computing power and human expertise together, cheaper than a centralized computing resource according to researchers who use it.

“The jobs best suited for it are small, with processing times ranging from 20-30 minutes, to a few hours,” said Francisco Brasileiro, member of the GISELA grid, based at the Universida de Federal de Campina Grande in Brazil. “Any longer and the probability that jobs will complete is very low,” he said.

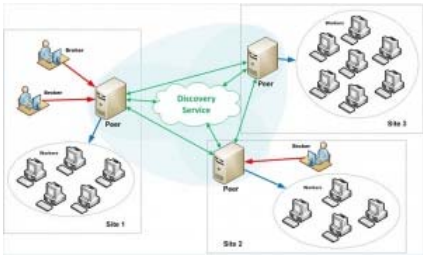

The OurGrid network has three main components: the OurGrid workers (individual computers that run grid jobs), the peer (responsible for managing local workers, interacting with the other peers to donate computer resources and use remote resources), and the broker (the OurGrid software client for users). Peers use a ‘Discovery Service’ to find each other in the P2P network.

Since 2004, OurGrid’s P2P resources have enabled research labs to trade their idle computing time. This trading creates a large computational pool that peers on the network can access to run applications including hydro-meteorological problems, such as understanding the patterns of irregular rainfall in Brazil to help local populations or planning and maintaining road systems.

Picking up the slack

OurGrid optimizes a key phase in a lab’s research cycle. Typically, scientific research is planned, executed and analyzed. Only the execution phase requires intensive computing resources. During planning and analysis, a lab’s computers remain idle. By donating its computing resources, a lab has the ability to exchange idle computer-time with other labs for future executions.

When a job starts processing on OurGrid, trusted software is used to copy any required files to disk. Then a virtual machine is created with sole access to a particular job. After execution, output data is copied to the user, and the virtual machine destroyed.

“In order to protect the resources donated on OurGrid, a strict security model is applied through virtualization technology such as VirtualBox. These virtualized executions have no access to the network. A remote job can be safely executed in any local machine without fear of harm to the local machine or other machines, even if the remote user is granted root access to this virtual machine,” said Brasileiro.

However, he highlighted that confidentiality is still an issue; not just for OurGrid, but all distributed computing: from large grid networks to public clouds. “Unless data can be processed in an encrypted way, it is always possible to have information leaked.”

Image courtesy Wikimedia

Freeriders don't win in 'Network of Favors'

To avoid problems of ‘freeriders’ (peers that consume resources of others without sharing their own), the Network of Favors was set-up. This is an automated P2P reciprocation algorithm that ranks peers based on their resource sharing with other peers. The higher the ranking of a peer, the more priority it receives when competing for computing resources. Peers automatically prioritize donations to other peers to whom they owe the most favors. This is a key motivating factor.

Like any P2P network, the more peers, the better. Even though OurGrid and its community are small - 300 machines and around 30 users - its concept has spawned other P2P grids, including the ShareGrid project.

Massimo Canonico is professor of computer science at the University of Piemonte Orientale, Italy, and founder of ShareGrid. “OurGrid provides a set of decentralized user and resource management mechanisms that eliminate the need and the costs of centralized management,” said Canonico. Cloud implementation is the next challenge in his opinion.

“The next step will be to introduce cloud computing in our platform but it is still work in progress and maybe we’ll switch to some open source cloud platform such as eucalyptus,” he said.

Meanwhile, OurGrid is upgrading. “Getting past the first turning of the wheel has proven to be really hard. We will add around 2,000 extra computer cores, but this will take at least 3 months,” said Brasileiro.

OurGrid is also part of the International Desktop Grid Federation, supported by the EC funded project DEGISCO. DEGISCO’s main aim is to develop mechanisms that allow service grids (e.g. gLite) and desktop grids (e.g. OurGrid) to work together. You can see how both grids complement each other in this educational CGI video.

Average:

Your rating: None Average: 3.3 (3 votes)

About the Author »

Adrian Giordani

RELATED TERMS: GISELA OurGrid ShareGrid VirtualBox grid computing

Comments

ADD NEW COMMENT

Post new comment

Subject:

Comment: *

By submitting this form, you accept the Mollom privacy policy.

SAVE PREVIEW