Peer-to-Peer Services for Distributed Resource Discovery on Grids

Domenico Talia and Paolo Trunfio

DEIS, Università della Calabria
Via P. Bucci 41c, 87036 Rende (Italy)
{talia,trunfio}@deis.unical.it

Abstract

Grid computing offers an effective way to build high-performance computing systems, allowing users to efficiently access and integrate geographically distributed computers, data, and applications. Several features of today's Grids are based on centralized or hierarchical services. However, as Grid sizes increase, it is necessary to decentralize their services to avoid bottlenecks and ensure scalability. A way to improve Grid scalability and reliability is to adopt *Peer-to-Peer* (P2P) models and techniques to implement decentralized Grid systems. The *Open Grid Services Architecture* (*OGSA*) and the recently proposed *Web Services Resource Framework* (*WSRF*) models provide an opportunity to integrate P2P models in Grid environments, since they offer an open cooperation model that allows Grid entities to be composed in a decentralized way.

A core Grid functionality that could be effectively redesigned using the P2P model is *resource discovery*. Resource discovery is a key issue in Grid environments, since applications are usually constructed by composing hardware and software resources that need to be discovered and selected. In the OGSA framework each resource is represented as a *Grid service*, therefore resource discovery mainly deals with the problem of locating and querying information about useful Grid services.

In Globus Toolkit 3 (GT3) - the current implementation of the OGSA - information about resources is provided by *Index Services*. An Index Service is a Grid service that holds information about a set of Grid services registered to it. A primary function of the Index Service is to provide an interface for querying aggregate views of service data collected from registered services. There is typically one Index Service per *virtual organization* (*VO*). Index Services can be organized in hierarchical tree structures, in which higher-level Index Services hold information about all the underlying resources. However, for scalability reasons, a multi-level hierarchy of Index Services is not appropriate as a general infrastructure for resource discovery in large scale Grids.

To address scalability and reliability, we propose an architecture for resource discovery that extends the model of the GT3 information service to support resource discovery across different VOs in a P2P fashion.

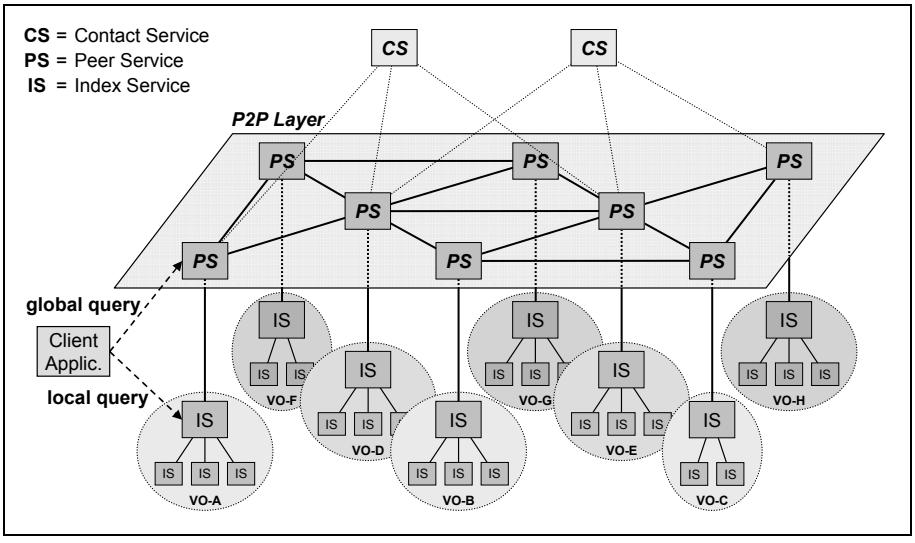

The figure shows the proposed architecture. Some independent VOs are represented; each VO provides one top-level *Index Service* (*IS*) and a number of lower-level Index Services.
A *P2P Layer* is defined on top of the Index Services' hierarchy. It includes two types of specialized Grid Services: *Peer Services* (*PS*), used to perform resource discovery, and *Contact Services* (*CS*), that support Peer Services to organize themselves in a P2P network.
There is one Peer Service per VO. Each Peer Service is *connected* with a set of Peer Services, and exchanges query/response messages with them in a P2P mode. The connected Peer Services are the *neighbors* of a Peer Service. A *connection* between two neighbors is a logical state that enables they to directly exchange messages. Direct communication is allowed only between neighbors. Therefore, a query message is sent by a Peer Service only to its neighbors, which in turn will forward that message to their neighbors. A query message is processed by a Peer Service by invoking the top-level Index Service of the corresponding VO.
A query response is sent back along the same path that carried the incoming query message. To join the P2P network, a Peer Service must know the URL of at least one Peer Services to connect to. A convenient number of Contact Services is distributed in the Grid to support this issue. Contact Services cache the URLs of known Peer Services; a Peer Service may contact one or more well known Contact Services to obtain the URLs of registered Peer Services.
As shown in the figure, a *Client Application* can submit both *local* and *global* queries to the framework. A local query searches for information about resources in a given VO. It is performed by submitting the query to the Index Service of that VO. A global query aims to discover resources located in possibly different VOs, and is performed by submitting the query to a Peer Service at the P2P Layer.
The talk discusses the proposed P2P architecture and *Gridnut* a protocol for P2P Web Services invocation. Finally, research trends and future research work in the area is outlined.