SERVICE DISCOVERY IN A GRID-BASED PEER-TO-PEER ARCHITECTURE*

M. AMORETTI, G. CONTE, M. REGGIANI, F. ZANICHELLI

Dip. di Ingegneria dell'Informazione – University of Parma (Italy)

Abstract

Peer-to-Peer systems can provide an effective e-business solution and a rich communication medium for work groups in University Campus, Research Lab, Enterprise, Finance environments. The need for coordinated resource sharing and problem solving in large, multi-institutional Virtual Organizations is also the strong motivation behind Grid Computing, which is more concerned on security than Peer Computing but appears quite inflexible for highly dynamic, massive communities.

In our opinion, to fulfil the requirements of community-oriented applications it is necessary to combine elements of both paradigms. In this paper we illustrate the first steps in the design and development of a service-oriented P2P architecture based on the concept of stateful Web Service, i.e. a Grid Service, a network-enabled entity that provides some capabilities. In order to better support highly dynamic scenarios, service providers publish their services as XML advertisements which are then exploited by peers in the discovery phase.

1 Introduction

Recently, there has been considerable interest for a possible convergence of Grid and Peer-to-Peer (P2P) technologies. While both distributed computing models share common objectives, i.e. the organization of resource sharing within virtual communities, they focus on different aspects heavily characterizing the two approaches.

Grid computing mostly emphasizes security and remote resource and service access, as the main goal is to serve the needs of Virtual Organization, i.e. relatively small, closed communities. Conversely, the P2P approach targets highly dynamic, open communities whose aim is content sharing or massive parallel computing. Until now the main application of Peer Computing has been file sharing, but Community Support is emerging as the real killer application. Once endowed with appropriate security mechanisms, P2P systems can provide an e-business solution and a rich communication medium for work groups which belongs to the University Campus, Research Lab, Enterprise, Finance environments.

To fulfill the requirements of educational, research and corporate institutions for scalable, secure and extendable community-oriented applications, the convergence of Grid and P2P paradigms appears promising and worth investigating on account of the complementary strengths of the two approaches [FOSTER03]. The aim of our research activity is to explore the possible extensibility of state-of-art service-oriented Grid technologies in order to cope with P2P-like application scenarios characterized by massive and highly dynamic community of users and service/resources.

Within a service-oriented, massive and dynamic distributed system, service description advertisement and discovery assume paramount importance. Several standards and technologies have been introduced for Web Services and Grid Services, namely Universal Description,

** *This work has been partially supported by a IBM Faculty Award and by the "WEB-MINDS" FIRB project of the National Research Ministry.*

Discovery and Integration (UDDI) and Web Service Definition Language (WSDL)–based Index Services, although their applicability to massive, highly dynamic (but not necessarily business-oriented) virtual organizations appear somewhat questionable due to their mostly centralized nature.
In this paper we illustrate the first steps in the design and development of the Service-oriented P2P Architecture (SP^2A) based on the concept of stateful Web Service, i.e. Grid Service, which is a network-enabled entity that conforms to a set of conventions expressed as WSDL interfaces, extensions, and behaviors. The key idea in this paper concerns the decentralized publication and discovery of Grid Services by means of a P2P-based distribution of XML advertisements.
The paper is organized as follows. Section 2 introduces SP^2A and the technologies on which it is based. The prototype under development is illustrated in Section 3 (service publishing and discovery). Finally, an outline of further work concludes the paper.

2 The Service-oriented P2P Architecture and underlying technologies

The SP^2A is designed as a network of peers which are *Service Host Environments* forming secure subspaces of a P2P overlay network. Each of them is endowed with a user interface which allows to explore local services, to publish their advertisements remotely, and to search for remote services. More precisely, the SP^2A architectural model embodies three layers of Grid Services which can be exposed (to the user and/or to the other peers) by a peer, namely System-level Services, Base Services and Application Services. System-level Services are generic enough to be used by all other Grid Services; they are embedded in each peer and transparently perform infrastructural tasks (e.g. service interaction, security). Every peer must be configured with the Base Services (e.g. peer management, peer information retrieval, local services indexing). Moreover, a peer may be specialized in providing particular Application Services (e.g. bandwidth brokering, certification authority, streaming, instant messaging, file sharing, screen sharing), and eventually it can be member of one or more subgroups of the SP^2A network.
Our work is based on two specifications: the Open Grid Service Infrastructure (OGSI) [OGSI], which is the result of recent activities coordinated by the Global Grid Forum (GGF) [GGF], and Project JXTA, Sun Microsystems open P2P initiative [JXTAHOME]. These technologies complement each others: OGSI provides a framework for service-oriented design, while JXTA operates at a lower level than Grid Services providing additional P2P functionalities.
For the development of modular peers we are exploiting the Globus Toolkit 3 (GT3) [GT3], an open source implementation of the *Open Grid Services Architecture* (OGSA) [PHIS], which is a distributed system framework based on OGSI. OGSA also supports a security model that allows a grid-wide sharing of resources, while maintaining local, diverse security policies. For the realization of P2P features, such as connectivity and discovery, our main API is the Java Binding of JXTA.
The general problem we had to address was that users know they need a Grid Service which meets certain requirements, but they have no idea of what its location is. During year 2000 the Universal Description, Discovery, and Integration (UDDI) [UDDI] project was launched, with the intention to define a master directory of publicly available e-commerce services. The focus of the UDDI specification has evolved to support public registries such as the Universal Business Registry (UBR) [UBR], which is currently maintained by four companies (IBM, Microsoft, NTT Com, and SAP), and private registries that may be deployed within a company's own network boundaries. GT3 provides an UDDI-like `IndexService` within which the active Grid Services of a Service Host Environment can be monitored and queried. For example, we could have several `VideoServer` services all around the campus. `IndexService` will allow is to query what `VideoServer` meets the

particular requirements of the users. In the case of our P2P architecture it is unfeasible that each node, which is a Service Host Environment, has to know every other node to perform service discovery through the inspection of their IndexServices. On the other hand, a Centralized Directory Model [CDM] implemented with a unique IndexService providing handles for all the services in the network, would introduce a single point of failure in the system. For those reasons our first effort has been the realization of the PeerManager, which transforms a conventional Service Host Environment in a SP^2A peer. In particular, the PeerManager provides the mechanisms to share and discovery services in the network.

Among related work, the ICENI Grid Middleware [ICENI] also aims to build a service-oriented architecture, although the proposed Jini, JXTA e OGSI implementations are alternative and do not attempt at obtaining a close P2P/Grid integration. Edutella [EDU] aims to provide an RDF-based metadata infrastructure for P2P applications, building on the JXTA framework. Metadata for P2P networks is absolutely crucial, not only for content description but also for service description.

3 Service Sharing and Discovery in SP^2A

In a service-oriented system, service sharing and discovery are the basic features that must be provided to users. In this sense, SP^2A uses the Document Routing Model provided by JXTA. The key idea is to associate a JXTA advertisement to each service that has to be exposed to other peers (*GS-advertisement*). Each GS-advertisement has a “Name” field and a “Description” field, which we use to identify the service and the Service Host Environment (i.e. the peer) which provides it. In the current version of the prototype each peers knows exactly the structure of the service it needs, and the only information it searches for is the Grid Resolver Handle (GSH). GSH are concrete, maybe short-lived addresses such as:

```
http://160.78.28.8:8080/ogsa/services/SP2Aservices/VideoServerFactoryService
```

The previous GSH identifies a VideoServer *factory* service, to which a peer can connect and request that a new (transient) VideoServer instance be created, to use it, and then destroy it.

Service sharing is implemented through the *Shared-Resource Distributed Index (SRDI)* mechanism to index the corresponding GS-advertisement to its rendezvous super-peer. It is important to note that rendezvous super-peers do not cache GS-advertisements, they just maintain an index of GS-advertisements published by their edge peers. Not caching advertisements makes the rendezvous architecture more scalable, and reduces the problem of caching out-of-date GS-advertisements. Each rendezvous has a *Rendezvous Peer View (RPV)*, which is an ordered list of other known rendezvous, and uses a Distributed Hash Table (DHT) function to propagate its GS-advertisement index entries to a subset of its RPV.

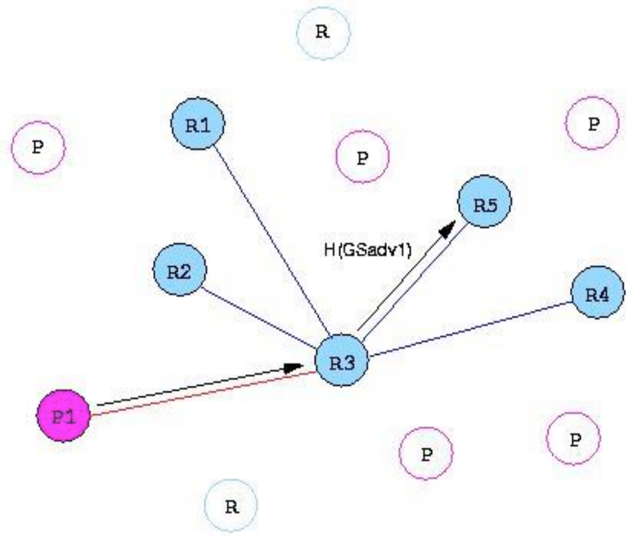

Fig.1 - Edge peer P1 publishes an advertisement to its rendezvous R3; R3 propagates the index entry of GSadv1 to R5 (according to the DHT function).

For example figure 1 describes the situation in which edge peer P1 publishes GSadv1 to its rendezvous super-peer R3. For clearity, we enumerate the five involved rendezvous (R1 to R5). The RPV of R3 contains R1, R2, R4 and R5. Suppose that the DHT function applied to the index entry of GSadv1 (H(GSadv1)) returns R5. So, R3 pushes the index entry to R5.

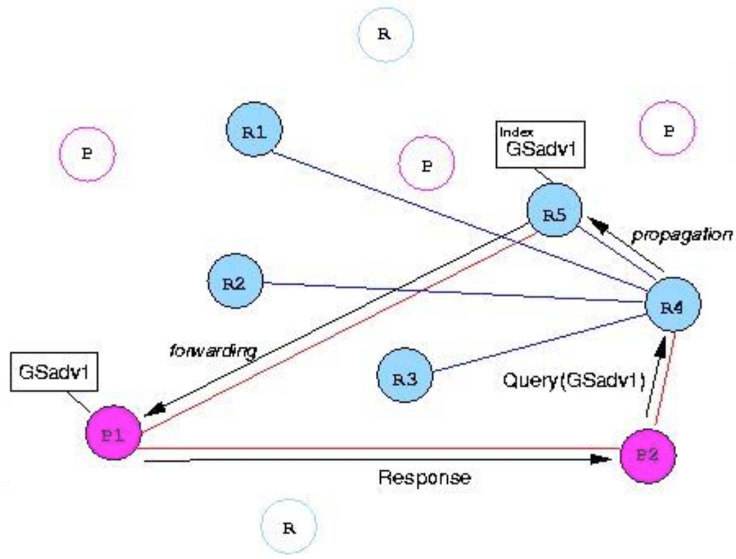

Fig.2 - Edge peer P2 issues a GSadv1 query to its rendezvous R4; R4 does not index Gsadv1, but it has the same RPV of R3 so it propagates the query to H(GSadv1) = R5 which has the index entry.

To understand the discovery mechanism implemented in the PeerManager, please refer to figure 2. Edge peer P2 issues a query for a GSadv1 (e.g. search by name) to its rendezvous R4, which looks if it has one or more suitable index entries. If local search fails, R4 uses the SRDI mechanism to establish the rendezvous to which propagate the query. If R4 has the same RPV of R3, the DHT function applied to GSadv1 returns R5. R4 propagates the request to R5, which has an index entry for GSadv1 pointing to P1. Finally, R5 forwards the query to P1 which then responds to P2. If the query issued by P2 leads to more than one GSH, the desired service can be found through the inspection of the IndexServices of the matching edge peers.

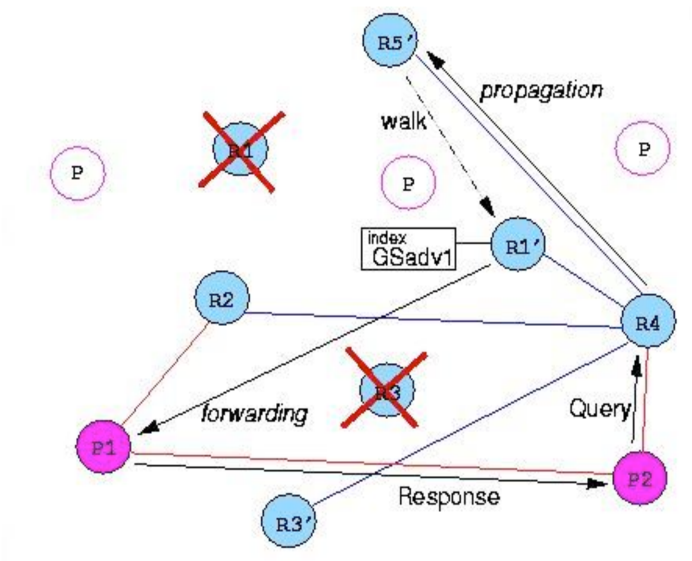

Fig.3 - R1 and R3 leave the network, while two other rendezvous connect to R4. In R4's RPV, R5 becomes R1' and the newly connected rendezvous become R3' and R5'.

In the previous example, we assumed that the RPV is the same on R4 and R3. This is not true, in general, because in a fluctuating network it is difficult to maintain a consistent RPV. Referring to figure 3, suppose that R1 and R3 leave the network, P1 connects to R2, and R4's RPV changes (R5 takes the place of R1 and two new rendezvous join). When R4 applies the DHT to GSadv1 it finds R5' which has no index entries for GSadv1. In this case, an alternative mechanism is used to continue searching: a limited-range walker is used to explore the RPV from the initial DHT rendezvous target (refer to [JXTA2] for the details).

4 Conclusions and Future Work

We have illustrated the basic features of our service-oriented P2P architecture, in particular connectivity, service sharing and service discovery. An experimental evaluation of the prototype and of the DRM algorithm provided by JXTA is necessary and will also be performed by simulating a network of several thousands of nodes, possibly investigating alternative discovery algorithms to embed in the peer prototype.

Our future work will concern other aspects of the service-oriented P2P architecture, such as the very critical issue of providing security and QoS to P2P systems. The often massive scale and the high dynamics of peer arrival/departure seem to make traditional enforcing techniques quite inapplicable. However, the campus-wide P2P computing grid we envision appear to be less critical than an Internet-wide one and probably amenable to the application of techniques for providing service level agreement among peers.

Bibliography

[FOSTER03] I. Foster, A. Iamnitchi. *On Death, Taxes, and the Convergence of Peer-to-Peer and Grid Computing*. 2nd International Workshop on Peer-to-Peer Systems (IPTPS '03), 20-21 February 2003, Berkeley, CA, USA.

[OGSI] S. Tuecke, K. Czajkowski, I. Foster, J. Frey, S. Graham, C. Kesselman, T. Maguire, T. Sandholm, P. Vanderbilt, D. Snelling. *Open Grid Services Infrastructure (OGSI) Version 1.0*. Global Grid Forum Draft Recommendation, 6/27/2003.
[GGF] Global Grid Forum, http://www.ggf.org
[JXTAHOME] Project JXTA, http://www.jxta.org
[GT3] The Globus Alliance, http://www.globus.org
[PHIS] I. Foster, C. Kesselman, J. Nick, S. Tuecke. *The Physiology of the Grid: An Open Grid Services Architecture for Distributed Systems Integration*. Open Grid Service Infrastructure WG, Global Grid Forum, June 22, 2002.
[JXTA2] B. Traversat, A. Arora, M. Abdelaziz, M. Duigou, C. Haywood, J.-C. Hugly, E. Pouyoul, B. Yeager. *Project JXTA 2.0 Super-Peer Virtual Network*.
[UDDI] *The Evolution of UDDI – UDDI.org White Paper*. The Stencil Group, July 2002.
[UBR] Universal Business Registry, http://www.uddi.org/register.html
[CDM] D. Milojicic, V. Kalogeraki, R. Lukose, K. Nagaraja, J. Pruyne, S. Richard, S. Rollins, and X. Zhichen. *Peer-to-Peer Computing.* Technical report, Hewlett-Packard Company, March 2002.
[ICENI] N. Furmento, J. Hau, W. Lee, S. Newhouse, J. Darlington. *Implementation of a Service-Oriented Architecture on top of Jini, JXTA and OGSI.* UK e-Science All Hands Meeting conference, September 2003.
[EDU] W. Nejdi, B. Wolf, C. Qu, S. Decker, M. Sintek, A. Naeve, M. Nilsson, M. Palmer, T. Risch. *Edutella: a P2P networking Infrastructure Based on RDF*. WWW2002, May 2002, Honolulu, Hawaii, USA.