UNIFIED DESCRIPTION AND DISCOVERY OF P2P SERVICES

G. Athanasopoulos, A. Tsalgatidou, M. Pantazoglou
Dept. of Informatics and Telecommunications, National and kapodistrian University of Athens

Keywords: Service Oriented Computing Peer-to-Peer services, Service Description, Service Discovery, Service Invocation.

Abstract: Our era has been marked by the emergence of the service oriented computing (SOC) paradigm. This new trend has reshaped the way distributed applications are built and has influenced current computing paradigms, such as p2p and grid computing. SOC's main objective is to leverage interoperability among applications and systems; however, the emergence of various types of services such as web, grid and p2p services has raised several interoperability concerns among these services as well as within each of these service models. In order to surpass these incompatibilities, appropriate middleware and mechanisms need to be developed so as to provide the necessary layers of abstraction and a unified framework that will obscure a service user from the underlying details of each service platform. Yet, for the development of such middleware and mechanisms to be effective, appropriate conceptual models need to be constructed. Within this paper, we briefly present a generic service model which was constructed to facilitate the unified utilization of heterogeneous services, with emphasis on its properties for the modeling of p2p services. Moreover, we illustrate how this model was instantiated for the representation of JXTA services and present the service description and discovery mechanisms that were built upon it. We regard this generic service model as a first step in achieving interoperability between incompatible types of services.

1 INTRODUCTION

Contemporary application development trends have been highly influenced by the service oriented computing (SOC) paradigm (Papazoglou, 2003). This new software engineering trend has primarily affected the development of distributed systems.

There have been many distinct definitions over what constitutes a service (Vogels, 2003), (Booth, 2004), (Foster, 2002). Nevertheless, all these definitions share some common properties and characteristics (Tsalgatidou, 2005). A service is usually regarded as a self-described software system offering a specific interface (that is described using a specific interface definition language) to its clients who can invoke it via messages over the Internet using predominant protocols and standards such as HTTP, XML, etc.

As it has been articulated by many researchers (Vogels, 2003), (Papazoglou, 2003), services leverage the interoperability among applications. SOC provides an infrastructure, which alleviates the interoperability concerns raised by the underlying platform, programming languages and operating systems used by software applications and systems.

Despite the original hype, the emergence of various types of services such as web, grid and peer-to-peer (p2p) types of services rendered the vision of interoperability elusive. Each of these types of services incorporates its own service model and attributes services with specific properties and characteristics. Heterogeneity is also prevailing among the underlying mechanisms supporting basic functionality such as the description, discovery and invocation of services.

Discrepancies and interoperability problems are not just inter-service type concerns. Such problems are also encountered within services of the same type. WS-I provided the Basic Interoperability Profile (WS-I, 2004) in order to overcome some of the interoperability concerns in the web service computing paradigm. Regarding p2p services, the existence of a plethora of p2p systems and frameworks, which abide by proprietary protocols and models, has aggravated the problem of p2p service interoperability. A comparison on some of the most well known p2p systems and frameworks

Athanasopoulos G., Tsalgatidou A. and Pantazoglou M. (2006).
UNIFIED DESCRIPTION AND DISCOVERY OF P2P SERVICES.
In *Proceedings of the First International Conference on Software and Data Technologies*, pages 203-209
DOI: 10.5220/0001318902030209

such as JXTA (JXTA), Gnutella (Gnutella) and Edutella (Edutella), reveals their diversities on the notion of p2p services. Each of these p2p frameworks and systems incorporates proprietary mechanisms and infrastructure for the description, discovery and invocation of p2p services and attributes p2p services with distinct characteristics and properties according to its own adopted service model.

All these diversities among the existing types of services hinder their widespread utilization and their seamless integration. In order to overcome the obstacles raised by the heterogeneity among the existing types of services one has to be based on a solid foundation. Such a foundation could be a conceptual model, which would provide the necessary basis for the development of appropriate abstractions and mechanisms that could facilitate service interoperability. This was one of the objectives of the SODIUM project (http://www.atc.gr/sodium). A Generic Service Model called GeSMO was developed within the SODIUM project and served as the basis for the development of a set of languages, middleware and tools for the unified discovery and composition of heterogeneous services such as web, grid and p2p services.

Within this paper we briefly illustrate this generic service model with emphasis on how this model was specialized for the representation of p2p services. Furthermore, we present how this model was instantiated for rendering JXTA services as well as the mechanisms that were developed for the description and discovery of JXTA services. This model has also served as the basis for the development of a p2p service invocation mechanism which will be the subject of another paper.

We have chosen to primarily focus on p2p services since this type of services present the widest range of diversities. Therefore, the application of our approach on this type of services will clearly illustrate its merits and the benefits gained by its usage.

Before proceeding we would like to give a definition on what constitutes a p2p service in order to dispel any misconceptions that may rise. As it has been stated in (Tsalgatidou, 2005), p2p services can be classified either as coarse services providing high level business logic, such as file sharing and instant messaging or as elementary services providing basic p2p network functionality, such as discovery of nodes and resources, message routing and message exchange. As far as coarse p2p services are concerned, these may be defined as "*the provision of resources or the execution of tasks of one or more (temporarily provider) peers on behalf of one or more (temporarily user) peers in a P2P network*" (Gerke, 2003), whereas elementary p2p service may be defined as "*services that support basic functionality in a P2P system, such as discovery of peers or resources, peer membership management, query formulation and routing, etc*".

The goal of this paper is to illustrate the provided generic service model and its extensions for supporting the representation of p2p services as well as how this generic service model was used for the development of appropriate mechanisms supporting the description and discovery of p2p services. In order to achieve its goals the paper is organized as follows: Section 2 illustrates the current state of the art on some of the existing p2p platforms and systems supporting the notion of p2p service. Consequently, sections 3 and 4 present respectively the generic service model and its instantiation for the JXTA service paradigm. Finally section 5 draws our conclusions and future plans on p2p services.

2 CURRENT STATE

A thorough look into some of the most prominent p2p platforms and systems unveils the set of discrepancies among these platforms regarding the properties and characteristics as well as the mechanisms used for supporting basic functionality of p2p services. The results of such an investigation on the JXTA (JXTA), Edutella (Edutella) and Gnutella (Gnutella) platforms are briefly presented below. A more detailed analysis of these problems can be found in (Tsalgatidou, 2005)

- *Support for the Notion of Service*: JXTA is the only framework which inherently supports this notion and enables the development and provision of additional services. Edutella is a platform built on top of JXTA providing a basic set of services, described using web service protocols, which enable the exchange of resources among peers, annotated with meta-data (Nejdl, 2002). Gnutella, on the other hand, is a fully distributed protocol enabling the development of file sharing p2p applications. In contrast to JXTA and Edutella, Gnutella doesn't accommodate the notion of service.
- *Syntactic Service Descriptions*: JXTA utilizes the advertisement construct which is a language neutral meta-data structure represented as an XML document. A service advertisement document conveys the necessary information

that enables a peer to invoke a preconfigured service. However, an advertisement document doesn't convey information related to the interface, operations or messages that are exchanged by a service; all this information is hard-coded into the service and client applications. Edutella utilizes web service description protocols and standards such WSDL or DAML-S. These constructs convey the necessary information that can be used for the on-the-fly invocation of services. Gnutella, in contrast to JXTA and Edutella, doesn't provide any form of service description document.

- *Supporting Mechanisms*: In terms of underlying mechanisms used for publishing, discovering and invoking p2p services, JXTA offers a set of services which facilitate the publishing and querying for available services as well as mechanisms for binding and exchanging messages with selected services. Edutella leverages the underlying mechanisms and services provided by JXTA for the discovery and invocation of the provided set of services. Gnutella on the other hand doesn't accommodate such mechanisms supporting the publishing, discovery and invocation of p2p services.
- *Semantic Service Description*: Among the investigated p2p platforms only Edutella supports the semantic annotation of p2p services descriptions through the use of semantic description frameworks such as DAML-S. JXTA and Gnutella don't accommodate semantic annotations to service descriptions.

All these incompatibilities among the investigated types of services render the seamless utilization of p2p services offered by those platforms an arduous task. In order to overcome these shortcomings, developers have to provide customized solutions that facilitate the integration of such services. Nonetheless, such solutions are usually strongly bound to the underlying infrastructure as well as the specific application logic.

An appropriate solution that could be applied to this problem is to develop a generic service model which will specify all necessary constructs. This service model would provide the basis upon which middleware and tools could be built so as to accommodate the necessary abstractions. Such a model is presented in the following section.

3 GENERIC SERVICE MODEL

In this section, we briefly illustrate the structure and the concepts of a Generic Service Model (GeSMO), which can be used as a basis for the development of appropriate languages and middleware that address service interoperability. GeSMO was based on a thorough investigation of the current state of the art on web, grid and p2p services. Specifically with respect to the p2p services, GeSMO has been primarily influenced by the work in JXTA (JXTA) as the latter is one of the very few p2p networks that ardently support the notion of service.

The generic service model was constructed in such a way that it efficiently models all common characteristics of web, p2p and grid services, while at the same time provides for the modeling of the distinct characteristics per service type. The architecture that was selected for the development of the Generic Service Model (GeSMO) is a layered one consisting of *a core layer*, *an extension layer built* on top of the core layer and *a number of layers* orthogonal to the core layer and its extensions. More details on the GeSMO model can be found at (Tsalgatidou, 2005).

As expected, the fundamental element in GeSMO is the notion of service. Services are regarded as self-described software systems, which interact with their clients over the Internet through messages. As we can see in Figure 1, a service description facilitates its clients in identifying the messages that can be exchanged as well as where and how these messages should be sent.

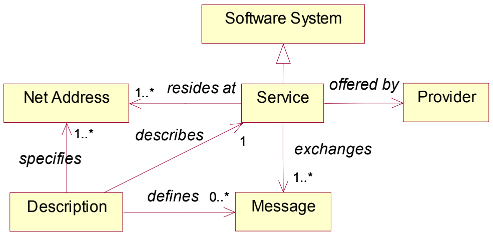

Figure 1: Basic Service Model.

The syntactic elements of a service along with their structure are illustrated in Figure 2. According to Figure 2, a service provides one or more interfaces which consist of the operations that this service offers to its clients. An operation groups a set of messages that are exchanged among a service and its respective clients. Each message consists of a set of elements which adhere to specific data types.

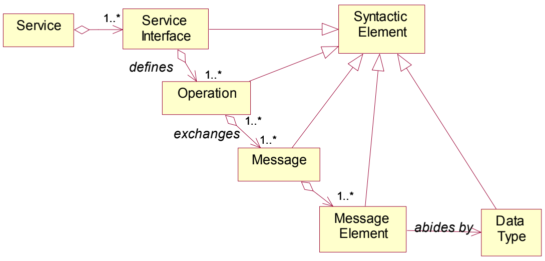

Figure 2: Syntactic Service Model.

3.1 P2P Service Model

The P2P Service Model extends the core concepts of the generic service model with additional elements that are needed for the description of p2p services. Although the extensions for the p2p service type were primarily influenced by the JXTA platform, they are still generic enough and thus applicable to other types of platforms.

According to Figure 3 the concepts that were added are that of P2P Service, Peer, Peer Group and PSDL (P2P Service Definition Language) description. A p2p service represents a service that is provided by a peer or a set of peers in a p2p network. A p2p service may be a coarse service or an elementary service according to the classifications provided in (Tsalgatidou, 2005).

A peer represents a node of a p2p network, which is able to communicate with other peers and provide them with services. A peer group represents the logical group of peers that may be formulated in a p2p network.

A PSDL description provides information on what a p2p service does and how it can be invoked. This concept was added so as to facilitate the description of p2p services with additional information elements. Hence, a PSDL document doesn’t aim at replacing the existing description constructs provided by the underlying platform, but rather to annotate them with additional elements that these constructs lack. Furthermore, PSDL descriptions provide the basis upon which mechanisms and middleware could be built so as to support the more efficient discovery and utilization of p2p services. More info about the PSDL description may be found in (Tsalgatidou, Athanasopoulos 2006)

The aforementioned concepts represent a basic set of elements that may be used for the description of p2p services. However, when it comes to modeling p2p services offered by specific platforms, such as JXTA or Edutella, additional elements or specializations of the aforementioned concepts might be needed. Such specializations for the JXTA platform as well as a description and discovery mechanism which were based upon this model are presented in the next section.

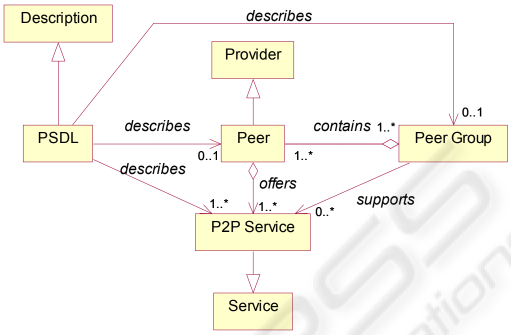

Figure 3: P2P Service Model.

4 JXTA SERVICE PARADIGM

JXTA (JXTA) is a set of open, generalized peer-to-peer protocols that allow network connected devices to communicate and collaborate as peers. JXTA protocols are independent of any programming language and up to now there have been several implementations for different environments. In general, we may regard JXTA as a framework that can be used for the development of p2p networks.

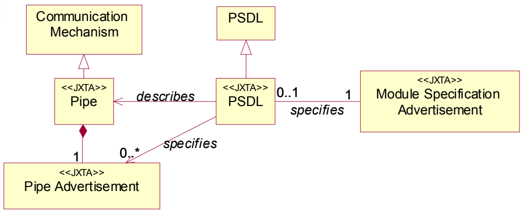

Figure 4: JXTA Service Model.

The provided JXTA protocols standardize the services used for: a) Discovering peers, b) Organizing peers into peer groups, c) Advertising and discovering network services, d) Handling peer communication and e) Handling peer monitoring. With respect to the development of services, JXTA offers constructs such as Pipes, Module Specification Advertisements and Pipe Advertisements. Pipes are communication mechanisms used for handling the exchange of messages among peers facilitating thus the interaction of service requestors and service providers. Pipe Advertisements and Module Specification Advertisements are specializations of the advertisement construct which are used for publishing information about pipes and services respectively.

The specializations of the generic service model provided for representing the concepts of the JXTA platform are illustrated in Figure 4. An extended PSDL description for the JXTA platform describes the pipes that may be used for exchanging messages with a service and provides references to their respective Pipe Advertisements. A p2p service's Module Specification Advertisement provides reference to the PSDL description document of that service. This association is accomplished through the use of an extension element (SURI element) that is inherently supported by a Module Specification Advertisement.

4.1 JXTA Service Description

A JXTA service description document such as the one presented in Figure 5 is an instantiation of the PSDL description concept that is illustrated in Figure 4. As it can be seen in Figure 5, PSDL is an extension of WSDL 1.1 (Christensen, 2001) which has the same abstract part as a WSDL document and an extended concrete part. The concrete part contains additional elements to accommodate the concepts described in Figure 3 and Figure 4. The WSDL 1.1 elements that were extended so as to accommodate the extra concepts presented in Figure 3 and Figure 4 are the *Binding* and *Port* elements. The *Binding* element has been extended via the *jxta:binding* sub-element (see Figure 5) which declares that this is a binding to a JXTA p2p service. Moreover, the *jxta:binding* element provides information on the type of interactions that might take place via this binding i.e. *send, receive* or *send-receive* interactions as well as the patterns of these interactions i.e. *synchronous* or *asynchronous* interactions. The *Port* element on the other hand has been extended with the *jxta:provider*, *jxta:pipes* and *jxta:moduleAdv* sub-elements. The *jxta:provider* element conveys information describing the provider of a p2p service which might be either a specific peer or a peer group (JXTA). The *jxta:pipes* element lists the pipes that support the exchange of messages with the service. For each pipe element that might be used for interacting with a service the pipe advertisement of that pipe as well as the interface accessible through that pipe are specified. Finally, *jxta:moduleAdv* element provides reference to the service's Module Specification advertisement. A PSDL document provides all the necessary information for identifying the functionality a JXTA service offers as well as how it can be invoked.

4.2 JXTA Service Discovery

In this section, we will briefly describe the mechanism we implemented for the discovery of JXTA p2p services. Due to the lack of space, and the fact that this is not the main purpose of this paper, implementation details have been left out. As it is shown in Figure 6, the course of service discovery in general involves the following steps: a) Specify service requirements in the form of a query; b) Submit the query to the appropriate broker, registry or network, in order to retrieve service description

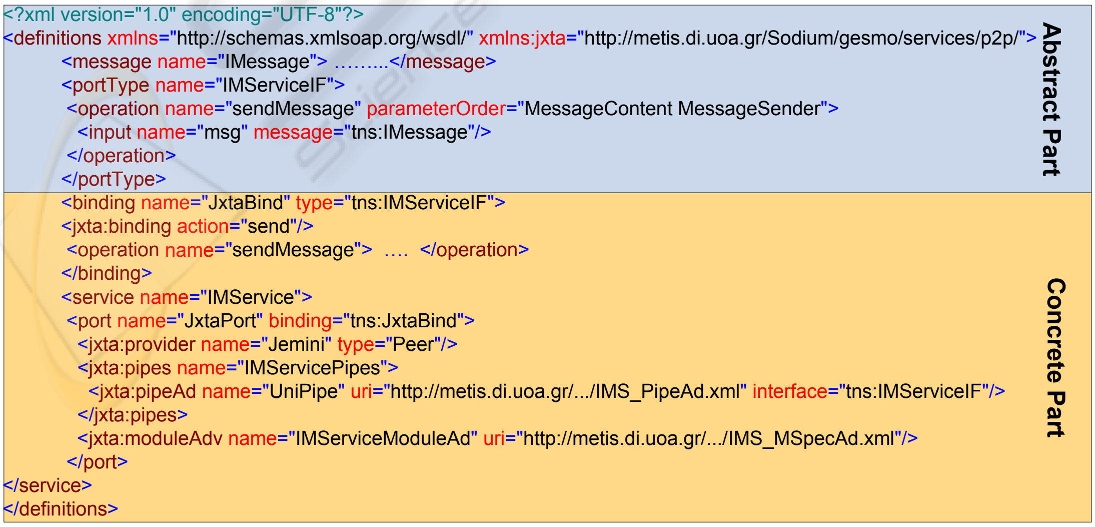

```
<?xml version="1.0" encoding="UTF-8"?>
<definitions xmlns="http://schemas.xmlsoap.org/wsdl/" xmlns:jxta="http://metis.di.uoa.gr/Sodium/gesmo/services/p2p/">
        <message name="IMessage"> ………</message>
        <portType name="IMServiceIF">
         <operation name="sendMessage" parameterOrder="MessageContent MessageSender">
           <input name="msg" message="tns:IMessage"/>
         </operation>
        </portType>
        <binding name="JxtaBind" type="tns:IMServiceIF">
        <jxta:binding action="send"/>
         <operation name="sendMessage">  ….  </operation>
        </binding>
        <service name="IMService">
         <port name="JxtaPort" binding="tns:JxtaBind">
           <jxta:provider name="Jemini" type="Peer"/>
           <jxta:pipes name="IMServicePipes">
             <jxta:pipeAd name="UniPipe" uri="http://metis.di.uoa.gr/.../IMS_PipeAd.xml" interface="tns:IMServiceIF"/>
           </jxta:pipes>
           <jxta:moduleAdv name="IMServiceModuleAd" uri="http://metis.di.uoa.gr/.../IMS_MSpecAd.xml"/>
         </port>
</service>
</definitions>
```

Figure 5: Instant Messaging Service PSDL description.

advertisements; c) Match the contents of the query against those advertisements; d) Wrap matching services in a message and send it to the requester, as response to his/her request.

The USQL Engine (Tsalgatidou, Pantazoglou, 2005), our proposed service search engine prototype, implements the aforementioned generic service discovery mechanism. The USQL Engine is an extensible framework, capable of applying service discovery against heterogeneous service registries and networks, in a unified manner. As its name implies, it makes use of the Unified *Service Query Language (USQL)* for the formulation of incoming service requests and outgoing service responses. USQL abides by the principles of GeSMO, and allows requesters to express their requirements by making use of syntactic, semantic, as well as quality criteria. A detailed description of the language is given in (Tsalgatidou, 2006). Like GeSMO, both USQL and the USQL Engine have been developed within the context of the SODIUM project.

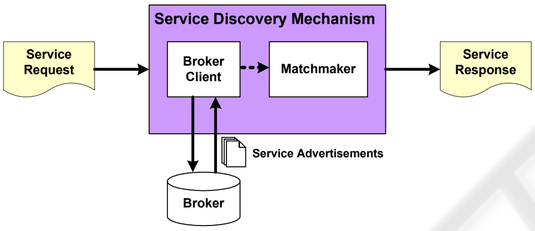

Figure 6: Perception of Service Discovery.

The USQL Engine retains its abstraction from heterogeneous service types and their related technologies by establishing a flexible plug-in based architecture. More specifically, the framework provides extension points for discovery agents and document handlers. The former are responsible for accessing and querying the various registries and networks by utilizing their existing discovery mechanisms, while the latter are responsible for parsing the various heterogeneous service advertisements hosted in those registries and networks.

In the case of JXTA service discovery, the USQL Engine has been extended with a *JXTAAgent* plug-in. This component utilizes the JXTA Discovery service, in order to access a peer group and retrieve the advertisements it contains. Provided that these advertisements contain a link to an external PSDL description document, the agent employs the *PSDLHandler* plug-in in order to parse these service descriptions and supply the *USQL Matchmaker* with input. The JXTA services that were found to meet the requester's requirements are included in the resulting USQL response, which contains adequate information for their immediate invocation. The whole procedure is displayed in Figure 7, which may be perceived as an instantiation of the generic mechanism shown in Figure 6.

In the current implementation of the JXTA agent plug-in, the USQL Engine needs to become a member of a peer group, before applying service discovery. In a highly volatile environment, such as a p2p network, and given the asynchronous nature of advertisement retrieval in JXTA, this approach may not always be the best one. A more sophisticated approach would break the plug-in into two loosely coupled parts: one part which would be a peer in the peer group of interest, caching all service advertisements, and another part acting as a remote control of the peer, asking for service advertisements. Such an approach is within our plans for future work. Nevertheless, the USQL Engine framework is flexible enough to be able to accommodate both approaches, as well as other potential architectures. For instance, the USQL Engine being exposed as a JXTA service, it could be used as the default service discovery mechanism within the boundaries of peer groups and networks, thus significantly contributing to the JXTA infrastructure.

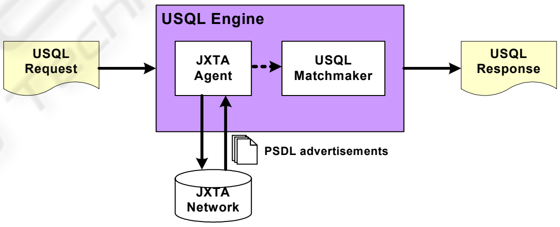

Figure 7: JXTA Service Discovery Process implemented by USQL Engine.

5 CONCLUSIONS

Service Oriented Computing has reformed the way distributed applications are built. P2P along with grid computing have been heavily affected by this new engineering trend. Although, grid computing ensued a straightforward approach in accommodating the service model through the (Foster, 2002) and (Foster, 2005) specifications, p2p computing has not ensued a similar approach.

The incompatibilities among the existing p2p platforms in supporting the notion of service are hindering the widespread utilization of p2p services as well as their interoperation with other types of services. In order to surpass the difficulties raised by the discrepancies among the existing types of services as well as of the existing p2p service platforms a model describing the necessary concepts needs to be provided. Upon this model appropriate mechanisms and middleware can be built so as to provide the necessary abstractions.

In this paper, we presented a generic service model called GeSMO, which was developed within the SODIUM (http://www.atc.gr/sodium) project for facilitating the modeling of web, grid and p2p services. The extensions that were provided for the description of p2p services as well as an instantiation of the p2p service model for the JXTA platform were also presented. The extensions provided for JXTA were the basis for the development of appropriate mechanisms supporting the description and discovery (Tsalgatidou, Pantazoglou, 2006) of JXTA services.

Despite that currently the presented p2p service model was only instantiated for the description of JXTA services, its concepts are generic enough so as to enable the modeling of other types of p2p services such as Edutella services or Gnutella services. Furthermore, although we have presented mechanisms which facilitate the description and discovery of JXTA services, additional mechanisms can be provided so as to enable the invocation of these services.

Our future plans include the instantiation of this model for other types of p2p service platforms e.g. Edutella, and the development of mechanisms which will enable the invocation of p2p services. Further future plans include the extension of the generic service model so as to accommodate other emerging types of services such as sensor services.

ACKNOWLEDGEMENTS

This work has been partially supported by the Special Account of Research Funds of the National and Kapodistrian University of Athens under contract 70/4/5829 and by the European Commission under contract IST-FP6-004559 for the SODIUM project.

REFERENCES

A. Tsalgatidou, G. Athanasopoulos, M. Pantazoglou (May 2006). A P2P Service Description Language Specification: Technical Report. Retrieved from http://www.di.uoa.gr/~gathanas/TR/psdl-1.0-spec.pdf

A. Tsalgatidou, M. Pantazoglou, G. Athanasopoulos (2006, February 3). Specification of the Unified Service Query Language (USQL), Technical Report. Retrieved from http://cgi.di.uoa.gr/~michaelp/TR/usql-1.0-spec.pdf

A. Tsalgatidou, G. Athanasopoulos, M. Pantazoglou, et al. (2005, June 23). Generic Service Model Specification: Technical Report. Retrieved from http://cgi.di.uoa.gr/~gathanas/TR/gesmo-1.0-report.pdf

I. Foster, et al. (2005, January 29).The Open Grid Services Architecture, Version 1.0. *Informational Document, Global Grid Forum (GGF)*. Retrieved from http://www.gridforum.org/documents/GWD-I-E/GFD-I.030.pdf

WS-I, (2004, August 24). Basic Profile Version 1.1. Retrieved on 8, May 2006, from http://www.ws-i.org/Profiles/BasicProfile-1.1.html

D. Booth, et al. (2004, February 11). Web Services Architecture. *W3C Working Group Note*. Retrieved on 26, April 2006, from http://www.w3.org/TR/2004/NOTE-ws-arch-20040211/

Werner Vogels (2003). Web Services Are Not Distributed Objects. *Journal of IEEE Internet Computing*, Vol 7, No 6, 59-66

M.P. Papazoglou and D. Georgakopoulos (2003). Service Oriented Computing (special issue). *Communication of the ACM*. Vol 46, No 10, 24–28

I. Foster, et al. (2002, June 22). The Physiology of the Grid: An Open Grid Services Architecture for Distributed Systems Integration. *Open Grid Service Infrastructure WG, Global Grid Forum*.

J. Gerke, et al. (2003 April). An Architecture for a Service Oriented Peer-to-Peer System (SOPPS). In Praxis der Informationsverarbeitung und Kommunikation (PIK) 2/03, p.90-95.

W. Nejdl, et al (2002, May 7-11). EDUTELLA: A P2P Networking Infrastructure Based on RDF. WWW2002. Honolulu, Hawaii, USA.

E. Christensen, et al. (2001, March 15). Web Services Description Language (WSDL) 1.1. Note, W3C. Retrieved from http://www.w3.org/TR/wsdl

JXTA, Retrieved from http://www.jxta.org/

Gnutella. Retrieved from http://www.gnutella.com

Edutella. Retrieved from http://edutella.jxta.org