Combining Service-Orientation and Peer-to-Peer Networks

Michael Conrad Jochen Dinger Hannes Hartenstein Marcus Schöller
Martina Zitterbart
Institute of Telematics University of Karlsruhe (TH)
http://www.internetoekonomie.uni-karlsruhe.de

Abstract:
Almost all current P2P applications are built as monolithic systems: they tightly couple application software with an underlying P2P technology. For simple P2P applications like file sharing of instant messaging this does not represent a deficiency. For more general scenarios, e.g., for providing an electronic market place based on P2P technology, flexibility has to be increased. We propose to decouple the application from its underlying P2P network and introduce a mechanism that allows us to dynamically bind an overlay from a given set to a specific application. This paper describes the design of such a system's architecture that combines the service-oriented architecture paradigm and P2P technology. Using these two different technologies in an integrated way gives us flexibility to bind the best fitting P2P network to the executed service. In addition, it allows us to easily add and deploy new services.

1 Introduction

Providing suited distributed services to networked applications still is an important challenge that lacks scalable and flexible solutions which can be provided across todays Internet. Generally, distributed systems can be viewed as a collection of services that are provided by so-called service providers and that are utilized by applications or service users. In this context, the term service-oriented architectures (SOAs) has been coined. Important aspects for such architectures are publishing, discovering and using dedicated services in a distributed environment. Among others, peer-to-peer (P2P) networks are considered as a promising basis for distributed service provisioning. Thereby, services can be very different, e.g., including group communication or distributed storage.

Recently, a variety of concepts for P2P networks has been proposed with the goal of providing scalable services across large-scale networks such as the Internet. Most of the P2P networks, however, have been designed with a specific application in mind such as file sharing. Consequently, they lack the flexibility to be used for the provision of other services. Furthermore, P2P networks are very often designed tightly-coupled with the applications they should support. Although this is a valid step with respect to this certain application, it does not support the general idea of flexible, application-independent service provisioning. Therefore, we propose an architecture that separates the P2P network from the services it should provide to the application.

Our flexible and scalable distributed service architecture is being designed in the context of an application-driven project that is somewhat unusual in the context of P2P networks. We envision highly distributed economic markets that provide very different goods. Specific examples that are considered include multi-utility markets (MUM) and virtual power plants (VPP). Within the MUM scenario few energy suppliers sell their products, e.g. electricity, natural gas or district heating, to a large group of customers. All customers get a transparent view onto this market and can easily select the best fitting bargain for their needs. The second scenario introduces new suppliers of electricity like owners of small to midrange power plants. Such devices can accommodate a single household up to a block of houses with electricity. But during high peek situations the provided power may not be sufficient so that energy must be bought from an external power supplier or during low peek situations electricity can be sold to others. This means that such a market participant switches from buyer to seller and vice versa spontaneously.

In order to be able to serve these different applications we aim at the development of a highly flexible, scalable and technology-independent architecture for distributed service provisioning. Our work is conducted in the SESAM[1] project which is funded by the German Ministry of Education and Research (BMBF). SESAM is a highly interdisciplinary project comprising researchers from computer science, economics as well as law.

The remainder of this abstract focuses on the distributed P2P based architecture that is being developed within the project SESAM. Major design issues are highlighted and a first basic prototype implementation is briefly discussed.

2 Design

A major design goal of the architecture to be developed is flexibility, i.e., distributed services should be seamlessly usable by different networked applications. In contrast to typical service-oriented architectures (SOA), our architecture offers mechanisms that allow multiple peers to cooperate in order to provide a service. We call this cooperation ServiceNet because the service as a whole is not provided through a single system but through a set of peers. From a functional point of view a ServiceNet does not offer more functionality than a single instance, either if the ServiceNet is built by one or hundred peers. But the non-functional aspects, like robustness, availability and so on might differ drastically.

Furthermore many P2P applications have some basic functions, like authentication or data distribution in common. This naturally leads to duplicated functionality on a P2P node if multiple applications are executed in parallel. Such functionality can be refactored to base services provided by a P2P base system.

It is obvious that different services demand different requirements of the P2P network. Therefore each ServiceNet can use its own P2P network, which is suitable for its needs. Hence we need a clear distinction between P2P network and services. To change the P2P network flexibly an interface to access the overlay layer (see fig.1) is necessary. A major

[1] Self-Organization and Spontaneity in Liberalized and Harmonized Markets

innovation of the proposed architecture can be seen in the fact, that P2P networks can be flexibly appointed to network applications. A first approach to such an interface is given in [DZD+03]. However, only structured P2P networks, like DHT-based P2P networks can be accessed via this interface. Unstructured P2P networks, like Gnutella [Ri01] or GIA [CRB+03], will likely require a different interface. This leads to a set of interfaces, which are integrated in our P2P framework. Furthermore, we have to provide a characterization of P2P networks in a way that a service developer can choose a suitable one.

3 Basic communication architecture of SESAM

Derived from the design issues, the architecture depicted in figure 1 has been developed. It basically shows the differentiation of three major 'layers': communication layer, overlay layer and service layer.

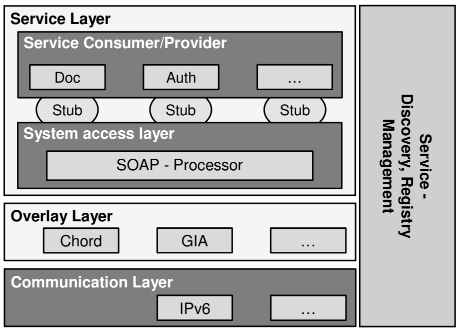

Figure 1: Communication Architecture

- The *Communication Layer* provides an abstraction from actual network protocols, like IPv6 based networks. This is similar to the objective of JXTA [Go02].
- The *Overlay Layer* makes various P2P networks available to services. E.g. Chord [SMK+01] might be included as structured DHT-based P2Ps and GIA as unstructured one. This layer is responsible for handling overlay tasks like initializing procedures. Access to this layer is offered through one or more interfaces that are similar to [DZD+03]. These interfaces encapsulate specific P2P networks.
- The *Service layer* consists of a *SOAP Processor* and the services themselves. The SOAP processor is responsible for translation of SOAP messages to programming language objects. *Service Consumers* and *Provider* are integrated through a stub/skeleton mechanism. This mechanism facilitates service developers to write services and use other ones. They can benefit from the P2P networks without knowing them in detail. Through the use of Web Service techniques this layer also offers type safety and programming languages independence. Two base services depicted are the document service (*Doc*) and the authentication service (*Auth*). The document

service provides a 'document pool' that can be used for publishing and finding documents. The authentication service is used to authenticate participants and provides multiple mechanisms like password- or smartcard-based authentication.

Complementing these three layers, there is a vertical block, which includes service management functions. It includes methods for finding, publishing and binding of services. In traditional SOAs the service registry is a central entity. In our approach the registry is distributed and benefits from the integrated P2P networks (see [FM04]). Because of that a distributed service discovery is included that uses itself a ServiceNet.

Our architecture is easily extensible as indicated through the dots in the figure. This means that we can build up on new networks through the communication abstraction, integrate new P2P networks in the overlay layer and finally extend our system through new services.

Based on this architecture a first prototype has been developed and tested. Within the prototype implementation we make use of Apache Axis (see http://ws.apache.org) as Web services implementation and SOAP processor respectively. This prototype also includes a first version of the document service that allows us to publish and find electricity offers.

4 Conclusion

We argued that application logic should be separated from underlying P2P network techniques in order to derive a general and flexible P2P framework for provisioning of distributed services. In our proposed architecture we combine ideas from service-oriented architectures with P2P concepts. Our prototype has shown that architecture realization is feasible. Nevertheless, we still have various open tasks to tackle in the future, e.g., provisioning of security mechanisms or a characterization model for P2P networks. Evaluation of appropriate unstructured overlay networks is also 'work-in-progress'.

References

[CRB+03] Chawathe, Y., Ratnasamy, S., Breslau, L., Lanham, N., und Shenker, S.: Making gnutella-like P2P systems scalable. In: *SIGCOMM*. Karlsruhe. 2003. ACM.

[DZD+03] Dabek, F., Zhao, B., Druschel, P., Kubiatowicz, J., und Stoica, I.: Towards a common API for structured peer-to-peer overlays. In: *2nd Int. Workshop on P2P Systems*. 2003.

[FM04] Forster, F. und Meer, H. D.: Discovery of web services with a p2p network. In: *International Conference on Computational Science*. pp. 90–97. 2004.

[Go02] Gong, L. Project JXTA: A technology overview. 2002.

[Ri01] Ripeanu, M. Peer-to-peer architecture case study: Gnutella network. 2001.

[SMK+01] Stoica, I., Morris, R., Karger, D., Kaashoek, M. F., und Balakrishnan, H.: Chord: A scalable peer-to-peer lookup service for internet applications. In: *SIGCOMM01*. San Diego, California, USA. 2001.