

Towards a Mobile Peer-to-Peer Service Platform

Wolfgang Kellerer¹, Zoran Despotovic¹, Maximilian Michel¹, Quirin Hofstätter², Stefan Zöls²

¹*DoCoMo Communications Laboratories Europe, Germany
{kellerer, despotovic, michel}@docomolab-euro.com*

²*Munich University of Technology, Institute of Communication Networks, Germany
{quirin.hofstaetter, stefan.zoels}@tum.de*

Abstract

Peer-to-Peer (P2P) technology is regarded as being disruptive for traditional fixed and mobile operators. However, it can also be considered as an opportunity for new ways of service provisioning. On the one hand, P2P concepts allow a more efficient usage of existing platform resources through relying on existing infrastructure including even the customer equipment and thus provide a cheap service platform. On the other hand, P2P technology brings new business opportunities through extending the service portfolio to user provided services and to ubiquitous environments. Along these lines, we outline in this paper the requirements and building blocks for a P2P based service platform for mobile environments. In addition to P2P lookup and information distribution, these building blocks include reliability, controllability, bootstrapping, reputation management, which address operator specific requirements in particular. Mobile networks pose severe challenges to a P2P service platform in terms of heterogeneous devices and access connections and frequent user joins and leaves. We present a hierarchical P2P system as a core component for a P2P service platform addressing those problems.

1 Introduction

Peer-to-Peer (P2P) networks are self-organizing, distributed systems. They are realized as application layer overlay networks relying on whatever physical connectivity among the participating nodes. The basic problem addressed by a P2P network is self-organized distribution of a set of resources among a set of peers enabling their subsequent fast lookup.

Two classes of solutions have been proposed to solve this problem. In *unstructured* networks, peers connect randomly to one another without maintaining any information about what resources are available at their neighbors. For resource lookups, flooding techniques are used in which lookup requests are forwarded through the

network until the requested resource is found or the entire network is explored [1].

At the other extreme is the concept of *structured* P2P networks, based on Distributed Hash Tables (DHTs). In a DHT peers and individual resources are identified by keys from a key space. Given the set of peers, the key of each peer is associated with a partition of the key space such that the peer becomes responsible to manage all resources identified by keys from the associated partition. To forward resource requests, peers form an overlay routing network by maintaining short-range links to all peers with neighboring keys and, in addition, a small number of long-range links to some selected distant peers. Using the routing network established in this way, peers forward resource requests, trying to greedily reduce the distance to the key being looked up. Most of DHTs resolve lookups with a number of messages logarithmic in the size of network, by using routing tables which are also logarithmic in the size of the overlay network [2, 3].

P2P applications have gained an extreme popularity among Internet users. With the increasing availability of data communications, including Internet access in mobile networks, wired network originated P2P applications become available also in mobile environments. We consider the latter as mobile communication systems for data exchange, in which at least one hop of the communication path is accomplished by a wireless link.

For traditional fixed and mobile operators, P2P technology is still regarded as being disruptive. However, looking at new ways of service provisioning enabled by P2P, it can also be considered as an opportunity in two respects. First, P2P concepts allow a more efficient usage of platform resources through relying on existing infrastructure. Second, P2P technology brings new business opportunities through extending the service portfolio to a new range of applications.

The P2P paradigm relies on using edge resources, such as storage or computing power, which are otherwise wasted. As such it pushes the application logic to the network edges, eliminating the need for high performance servers. In cellular mobile systems this may be a great

source of infrastructure cost savings. Applications without strong data availability requirements are particularly good candidates for a P2P implementation. Examples include various community formation applications, exchange of different information (free parking lots, restaurant recommendations, tourist information) and so on.

Sensor networks may also benefit from the P2P paradigm. Under specific circumstances DHT-like organization of the network may be the most appropriate solution [4]. The same applies for ubiquitous environments characterized by an extreme heterogeneity of the involved devices. A P2P overlay can accommodate variations and bridge even diverse network technologies.

In addition, the P2P platform makes the development of applications for heterogeneous environments easier and faster. It provides a common layer on top of various networks abstracting away the underlying heterogeneity. Once a P2P platform is provided (including routing, data management layer etc.) building new applications on top of it becomes easy; in many cases it only involves refining the provided data management functionality and selecting appropriate data abstractions to meet application and performance requirements.

When building up a dependable mobile P2P service platform, operator specific requirements such as controllability and reliability have to be met in addition to requirements arising from heterogeneous, mobile environments. Both requirements are usually not addressed in existing P2P systems.

2 Challenges of P2P in heterogeneous, mobile environments

Several problems arise for the realization of P2P concepts in mobile networks. We can summarize that it is especially the *heterogeneity* of the mobile environment (networks, devices, users) that constitutes the challenges. Thus with mobile P2P we especially refer to heterogeneity compared with conventional, fixed network P2P systems in which, to the largest degree, homogeneous nodes are assumed. Requirements that have to be considered for P2P in heterogeneous, mobile environments include:

- Reducing the P2P lookup traffic overhead as much as possible, in order to overcome low transmission data rates of mobile devices.
- Addressing the high churn rates due to frequent joins and leaves of nodes.
- Considering the limited resources of mobile devices addressing the heterogeneity of nodes and their distinct device capabilities.
- Handling the increased failure probability of mobile devices, resulting from link breaks, discharged batteries etc.
- Minimizing the traffic generated in the physical network, by taking into account properties of the physical network.

- Providing trust and incentive models to support users' willingness to comply with the protocol rules.

How are these requirements influenced by the underlying mobile network characteristics? We differentiate between cellular networks, hot spots, ad hoc networks and sensor networks.

Cellular networks are one hop wireless networks connecting peers to the fixed Internet via a single wireless link. Mobility management makes mobility issues such as handovers transparent to the application. In this way, data rate restrictions in terms of throughput and delay and device resources are most significant. Mobile phones are mostly switched on for reachability. However, churn rates might still be an issue due to high data transmission cost in cellular networks making users turning off their data applications.

Wireless data network hot spots offer comparatively high data rates to a locally restricted area. Whereas terminal capabilities and data rates are sufficient for P2P systems, mobility is restricted to the area around a hot spot. In other words, mobility is an issue, i.e., user mobility results in a high churn rate.

Wireless LAN technology is also common in Mobile Ad Hoc Networks (MANETs). Here, the transport link between peers is accomplished by multi-hop wireless links using intermediate nodes as relays. Since nodes in MANETs are assumed to be mobile, connection breaks are common whenever adjacent nodes move out of their radio range, which leads to high churn rates. In addition, MANET routing is costly in terms of bandwidth. This problem can be made even worse by a P2P overlay network that does not take the physical network topology into account.

Ad hoc sensor networks can be regarded as systems with extremely limited device resources connected over MANETs as described above. They are only suitable for special types of P2P applications.

3 P2P Service Platform Building Blocks

We propose a two-layer architecture consisting of different building blocks for an operator-grade P2P service platform. Figure 1 depicts our concept. Distinction between two layers, core P2P services and application specific services, is made. Even though this distinction is not new, there are some important differences in comparison with existing architectures [8]. Most notably, we bring in a number of new components (such as the "Controllability & Manageability" and "Adaptability" component) and also envisage different functionalities of a number of other components which exist in proposed architectures.

The P2P Core Services layer provides, on the one hand, fundamental functionalities such as P2P network bootstrapping and efficient resource lookup. On the other hand, it offers the operator a possibility to monitor and control the network operation, manage reliability or adapt the network organization to varying environmental

conditions. All these components are tailored to suit the requirements outlined in Section 2. The P2P core DHT lookup functionality, which we developed in our previous work, is described in Section 4, while the application specific services are discussed in Section 5. In the rest of this section we describe the core components which are new or substantially differ from the existing ones.

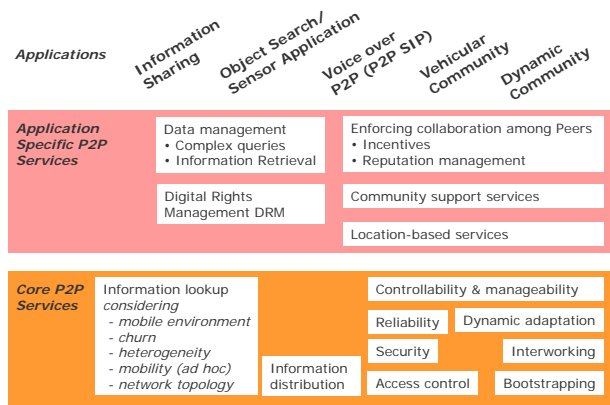


Figure 1: Two-layer design of the P2P Service Platform

Operator-grade services call for components that are usually not found in Internet scale applications, characterized by low guarantees. The operator, endeavoring to provide high service quality, must be able to monitor the status of the provided P2P application and react appropriately to any undesired changes. This is why we introduce the Controllability component in our architecture. As just hinted, there are two main functions that such a component has to realize. One is monitoring of the network. It involves either estimating the current status of the network via a number of operator controlled peers or central monitoring based on, for example, traffic sampling and analyzing. The second function provides a network control capability. For example, if the churn rate becomes too high the operator may need to react by injecting a number of more stable peers into the system. Or, operator controlled peers may initiate a change of the used replication strategy.

These steps do not assume any change in the underlying P2P protocol, i.e., parameters of a specific protocol may change, but not the protocol itself. However, it is sometimes better to change the protocol. For example, when the network size declines and churn rate is high it might be better to replace the currently employed DHT organization with an unstructured one. This is the primary task of the “Dynamic adaptation” component.

P2P network bootstrapping has been a subject of study in the P2P research. However, the emphasis has been put on sequential peer arrivals so far. Fast, parallel construction of the network has only recently drawn attention in the research community [5]. Such a problem becomes particularly relevant in various mobile

applications, such as formation of ad hoc communities at various events. Intimately related to this problem is the problem of efficient merging of two overlay networks. We imagine an application in which various local groups of users form different overlays (using MANET as the underlying network layer), which need to be merged when the groups meet and discover one another [9].

The “Reliability” component also plays an important role in provisioning operator-grade services. The content should be available at any time. This also means that it must be possible to look it up even in case of frequent peer failures. Replication has been recognized as a possible solution to the problem. But, there are many ways to replicate the content and the indices. One of them might take into account relevant characteristics of the replica peers such as their online time, failure probability etc. This information can be maintained in the network itself. Thus the network operation gets a flavor of the chicken and egg problem. It depends on the used replication policy, but the replication itself relies on proper functioning of the network.

4 Core P2P Services Lookup Component

Figure 2 shows a simple hierarchical P2P architecture designed specifically for heterogeneous environments. In this architecture high performance nodes (ultrapeers U) form a DHT, while low performance nodes (leafnodes L) are made free of DHT routing but still participate in the system [6]. This architecture offers a number of important advantages in comparison with flat DHT designs.




Figure 2: A simple hierarchical DHT

First, it fits better heterogeneous environments as depicted in Section 2. If the nodes differ substantially in their capabilities or communication costs, then nodes with cheaper costs should take a larger fraction of the total communication overhead. It is easy to see that the total network operation costs become smaller in this case.

Second, even when there are no substantial differences among the nodes, it pays off to employ such a hierarchical design. [6] argues that the total network traffic increases as the system from Figure 2 becomes flatter, i.e., the more

nodes move to the top level DHT. Thus, it is better to have just a fraction of the nodes operating in the DHT and a non-zero fraction of the nodes operating as leafnodes. The optimal ratio of these two classes of nodes is determined by the characteristics of the DHT nodes, i.e., how much load they can take. The key to understand this is the observation that this hierarchical architecture actually presents a tradeoff between centralized systems, which minimize the network traffic but cannot be realized in practice because no single node can take the entire load generated in the system, and a totally decentralized system (a flat DHT), which incurs the highest operation costs and can be realized.

The architecture from Figure 2 is not the only one possible. [7] mentions a number of other possible hierarchical designs, e.g., one in which leafnodes operate within a new DHT instead of relying completely on their top level DHT proxies. This disburdens the top level nodes to a certain degree and puts more effort to the leafnodes. We are currently evaluating these alternatives in the context of the cost model of [6]. Our preliminary findings say that the architecture from Figure 2 outperforms the other architectures.

5 Application Specific P2P Services

In order to develop a vast range of applications, several application specific P2P services have to be provided on top of the Core P2P services layer. The data management service provides the necessary data manipulation operations such as inserting, updating, deleting, and querying data. In particular, the complex query management has to be provided. The problem is that DHT's natively support only exact match queries, i.e., the exact value of the key has to be known in order to query for the data item with that key. This is inappropriate in most of the applications. Therefore, support for most common (non-key) attribute queries has to be provided, including the most typical aggregations such as range queries and joins.

Collaboration enforcement through incentives or reputation management becomes important in applications in which potentially large groups of peers provide one another various services. The value of such services is normally low in order to use well established assurance mechanisms (e.g., contractual agreements) but still high enough to reject any of them. Reputation management has been recognized in the research community as an appropriate means to deal with this problem.

A promising class of applications can be built on location awareness. Some location-aware applications are already available (such as a city guide [10] or mobile terminal finder [11]), but we believe that such applications should be built around the same underlying concept and thus deserve support from a common *location-based services* component. As well, both the users who can be positioned in an unambiguous automatic way (e.g. GPS) and those who can have just a descriptive understanding

of the notion of location should be able to use such applications.

6 Conclusion

Peer-to-Peer (P2P) technologies allow for new, resource saving ways of service provisioning. Systems that have been widely discussed mostly support single application classes in fixed network environments. In this paper, we describe the requirements for an operator-grade P2P service platform capable to run in environments characterized by heterogeneity such as the mobile environment. We describe the requirements for P2P in heterogeneous, mobile networks and provide our view of a respective P2P service platform based on a hierarchical P2P lookup service. Furthermore, we introduce new building blocks for an operator-grade P2P service platform such as controllability, manageability, adaptability and reliability.

REFERENCES

- [1] Gnutella. Clip2. The Gnutella protocol specification v0.4 (document revision 1.2). http://www9.limewire.com/developer/gnutella_protocol_0.4.pdf, 2001.
- [2] A. Rowstron, P. Druschel. Pastry: Scalable, distributed object location and routing for large-scale peer-to-peer systems. IFIP/ACM International Conference on Distributed Systems Platforms (Middleware), 2001.
- [3] I. Stoica, R. Morris, D. Karger, F. Kaashoek, H. Balakrishnan. Chord: A Scalable Peer-To-Peer Lookup Service for Internet Applications. ACM SIGCOMM Conference, 2001.
- [4] S. Ratnasamy, B. Karp, S. Shenker, D. Estrin, R. Govindan, L. Yin, F. Yu: Data-Centric Storage in Sensornets with GHT, a Geographic Hash Table. MONET 8(4): 427-442 (2003)
- [5] K. Aberer, A. Datta, M. Hauswirth, R. Schmidt. Indexing data-oriented overlay networks, 31st International Conference on Very Large Databases (VLDB), Trondheim, 30 Aug - 2 Sep, 2005.
- [6] S. Zöls, Z. Despotovic, W. Kellerer. Cost-Based Analysis of Hierarchical DHT Design, 6th IEEE Conference on P2P Computing, Cambridge, UK, 2006.
- [7] L. Garcés-Erice, E. W. Biersack, K. W. Ross, P. A. Felber, G. Urvoy-Keller. Hierarchical P2P Systems. ACM/IFIP International Conference on Parallel and Distributed Computing (Euro-Par), Klagenfurt, Austria, 2003.
- [8] SUN Microsystems. *JXTA v1.0 Protocols Specification*, 2001. <http://spec.jxta.org/v1.0/docbook/JXTAProtocols.html>.
- [9] A. Datta, K. Aberer. The challenges of merging two similar structured overlays: A tale of two networks. IWSOS 2006, International Workshop on Self-Organizing Systems, Passau, Germany, September 2006.
- [10] K. Chevers, N. Davies, K. Mitchell, A. Friday. Experiences of developing and deploying a context-aware tourist guide: the GUIDE project. MOBICOMM 2000, Boston, USA, 2000.
- [11] O2 Germany GmbH & Co. OHG: O2 Handyfinder <http://www.o2online.de/o2/kunden/myo2/handy/handyfinder/index.html>