University of Würzburg
Institute of Computer Science
Research Report Series

P2P-based Network and Service Operation for the Emerging Future Internet

K. Tutschku

Report No. 397 January 2007

University of Würzburg
Institute of Computer Science
Department of Distributed Systems
Am Hubland, D-97074 Würzburg, Germany
tutschku@informatik.uni-wuerzburg.de

P2P-based Network and Service Operation for the Emerging Future Internet

K. Tutschku

University of Würzburg,
Department of Distributed Systems,
Am Hubland, D-97074 Würzburg, Germany
E-mail: [tutschku]@informatik.uni-wuerzburg.de

I. Introduction

The Internet has evolved. So far, its main achievement was the interconnection of physical networks by a consistent collection of protocols and technologies. Now, the *Future Internet* is more and more considered a network of applications[1], information and contents. Hence, the new characteristics of the future system can be described briefly as:

- The Future Internet will be viewed as a network of applications.
- It will enable "peer productivity" and becomes an "architecture for participation" [1].
- In particular, the Future Internet will be based on interactive, edge-based applications and overlays, such as P2P content distribution, Skype, MySpace, or YouTube.
- However, it is not yet clear what the next major application in the Future Internet is!

Moreover, recent advances in high speed optical and wireless transmission and virtualization of links and routers have sparked fundamental discussions about how to design the architecture of the Future Internet:

- Is a clean-slate approach required to facilitate new network and application architectures for the Future Internet?
- Would an evolutionary process for designing the Future Internet be more appropriate?
- What kind of network and application features will drive the design of the Future Internet?

Giving a definite answer to these questions is audacious and impossible. However, technological challenges in networking and network applications can be identified and their implications should be considered for the future design.

In addition, peer-to-peer (P2P) based mechanisms have recently proved their capabilities such as scalability, fault tolerance, and self-organization, that they address many of the anticipated challenges of the Future Internet.

Combined with latest advances of directly supporting virtualization in core network elements, one even may think about integrating P2P-based mechanisms deeply in the architecture of the Future Internet, e.g. on transport layer. However, this might require a rethinking of the layering structure of today's networks.

[1]This contribution assumes that the terms "application" and "service" are describing the same item in the Future Internet, which is the application of the network.

The aims of this contribution are *a)* to investigate the technological trends and challenges for the Future Internet and *b)* describe by two examples how these issues might be facilitated in today's or near-future networks by the use of P2P. The paper is organized as follows. In the first part, the Section II sketches a picture of current networking trends. Section III identifies major challenges in the design of the architecture of Future Internet. In the second part, Section IV outlines the concept of P2P. Section V describes a Mobile P2P participation architecture for infrastructure-based public mobile networks and discusses how this architecture addresses the needs of the Future Internet. Section VI suggests a P2P-based automatic service monitoring overlay. Finally, Section VII summarizes the paper.

Part I: Trends and Challenges by the Emerging Future Internet

II. Trends

Today's Internet does not show a general overload situation which would ask for a new network architecture. Actually, it performs very well in some parts, as examples such as P2P file-sharing show. However, some deep limits in its applicability and operation can be observed as well as new ideas in using the system. In addition, recent advances in networking technology have been made which might overcome some of the limitations of today's Internet. In the following, remarkable developments which call for new network and application architectures and for new operation procedures will be summarized.

Services and Applications. The services in classical communication networks, such as ISDN or GSM, are rather platform-dependant. Due to the increased application of abstraction layers, like the Internet Protocol (IP) or overlay techniques, services can be consumed now in a variety of wireless and wireline networks such as ADSL, WLAN, or UMTS. Hence, the transition from single network services to *multi-network services* has occurred.

Classical services are provisioned by network operators. However, the success of P2P file-sharing has blurred the boundary between content providers and consumers. In addition, they showed that edge-based communities can easily design, deploy and offer services. The new services reveal *edge-based intelligence* and form *overlays* with application-specific naming and routing concepts.

Furthermore, users transfer their social behavior increasingly to networks and networked applications. *Social networking web sites* like "YouTube" [2] or "MySpace" [3] with *user-generated content* became tremendously popular. They permit the users to structure the use of the information according to their specific social interests or relationships.

The ubiquity and availability of networked application in today's wired and wireless networks combined with an increasing commercial significance has led to a demand for highly *dependable networks and services*. Hence, automatic resilience, fault management, and overload mechanisms have been introduced on different layers. Examples include dependable overlay services for supporting vertical handovers in mobile networks [4], or multiple source download in P2P content distribution networks.

The success of virtual mobile operators [5] or of the P2P VoIP service Skype [6] has shown that *virtualization of telecommunication services or applications* is no longer an academic concept. For example, Skype replaced and virtualized central indices for user locations by a distributed software running on the end user's client.

High Speed Data Transport. *Advanced optical core networks* using Dense Wavelength Division Multiplexing (DWDM) or hybrid optical network architectures have brought tremendous amounts of flexible point-to-point transmission capacity into core networks [7]. *Fibre-based access technologies*, such as Ethernet Passive Optical Networks (EPON), permit to deliver this capacity to end users at very low cost [8].

Furthermore, infrastructure-based wireless communication has experienced a huge *diversification of radio access technologies* while experiencing a steady increase of capacity. Beyond Third Generation (B3G) wireless networks will comprise highly ubiquitous and very different mobile broadband access technologies such WLAN, HSPA [9], or Mobile WiMax [10].

Network and Service Control and Management. The need for fast responses on failures and the reduction of operational costs (OPEX) led to the development of *autonomous procedures* for network and service operation. Thus, *self-organizing mechanisms* and self-* procedures have been suggested [11]. The algorithms automate, for example, the quick pinpointing of system faults [12] or specific configuration tasks in mobile access networks [13].

The end-to-end control paradigm of TCP/IP networks has decoupled the user and the operator from direct quality feedback. User and network operator are typically not informed about the performance of an application. As a result, *integrated quality feedback mechanisms* have been investigated lately which notify users and operators independently from the application when end-to-end quality degradations occur [14].

Users judge the quality of networks, services, and networked applications more and more by the subjective perception of the performance. Hence, the concept of *Quality-of-Experience (QoE)* has been developed lately [15]. It describes the user's view how usable a service or a networked application is.

III. Technical Challenges

Considering the above mentioned trends, significant challenges arise for the design of the Future Internet which will be discussed next.

Overlays for Participation. A major challenge for the architecture of the Future Internet is the support of *overlays for participation*. Edge nodes should be enabled to form overlays of coordinated communities. In particular, they require mechanisms to define overlays with application-specific name spaces, routing and self-organizing procedures for resource management. Furthermore, if edge nodes want to offer high quality services then they need access to flexibly managed resources, e.g. bandwidth. Thus, resource providers have to offer their supply on small time scales while ensuring the quality.

A further challenge for a participation architecture is the efficient locating and exchanging of *user-generated content*, such as phonebooks, pod- or videocasts. In particular the question where to store the data might deeply influence the future network structure. Centralized storage concepts are easily controllable by operators, e.g. for protecting copyright, and are highly efficient as YouTube has shown. However, they are vulnerable to overload and system faults. Distributed content stores may suffer from synchronization and additional network traffic, but might be more reliable. Hence, scalable, efficient, and controllable edge-based content networking mechanisms are needed for the future Internet. Ideally, these mechanisms should permit to specify the degree of centralization or decentralization at run time.

High Speed Data Transport: Heterogeneity, Mobility, and Core Network Architectures. A major challenge for the Future Internet is the *heterogeneity* of access technologies. For example, future mobile devices will move through a landscape of different wireless access systems, operators and sometimes even use fixed access systems such as DSL for fixed-mobile convergence. Vertical handovers (VHOs) are needed for bridging the heterogeneity between the access technologies and have to executed in very few time. Hence, the architecture of the Future Internet requires *scalable mobility management mechanisms*.

The capability to structure the Future Internet arbitrarily needs *mechanisms for virtual and flexible network configurations, routing architectures and overlays*. Future core network nodes need the capability to support in parallel multiple overlays which form arbitrary and flexible topologies. Each node should be able to configure every of its overlays with a random number of virtual interfaces and virtual edges to other distant core nodes [16]. The virtual edges should expose their performance and any lower layer faults to nodes. In addition, the future core nodes require mechanisms for forwarding and routing traffic along a virtual edge and for distributing the routing information. These mechanisms should also be available and controllable by edge nodes since the Future Internet will not distinguish between the edge and the core of the network.

Future Service Control and Management. Current *self-organization mechanisms for applications and services* are typically designed for end-user constraints where the consumption of network resources is of minor interest and for the optimization of a single objective. In the Future Internet, however, these algorithms need to consider in parallel network resources, multiple stakeholders, and diverse objectives. In addition, the mechanisms should maintain their performance and reliability under adverse network conditions.

Future reliable edge-based applications and services require the provisioning of perusable resources. Hence, *flexible Service-Level Agreements* are needed for negotiating and validating the quality of resource. The new SLAs should address the combination and encapsulation of the provided services, the provisioning on small time scales, and meaningful quality concepts, e.g. QoE.

Offering a universe of diverse services to a large number of users requires the Future Internet to be orchestrated with *scalable monitoring architectures* that survey independently the provisioning of services. Hereby, end-to-end network monitoring has to be achieved while being able to pin-point bottlenecks.

Future Layering and Abstraction Architecture. Today's Internet architecture is largely based on the hourglass concept of the Internet Protocol (IP) where every data is transported over IP and any IP packet is transported over every network. The so-called "IP waist" increasingly constitutes a bottleneck in today's Internet architecture [17].

Internal pressure to the IP waist comes from the increasing complexity and the deficiencies of the IP protocol such as the lack for scalable support for end-to-end quality of service across domains, limited resilience and mobility support for on-going data flows, and the very simple network management protocols. External pressure for the IP waist results, amongst others, from the efficiency and flexibility of application-specific overlays.

Hence, instead of having insufficient layers, such as the IP layers[2], and by-passing them by overlays, a *thinning of the layers* and a *more basic separation of the layers* appears to be needed. This separation should focus on *a)* the application layer (for addressing application needs), *b)* the mediation layer (for network structuring, naming, and routing), and *c)* the transport layer (for reliable and cost-efficient transport), see Figure 1.

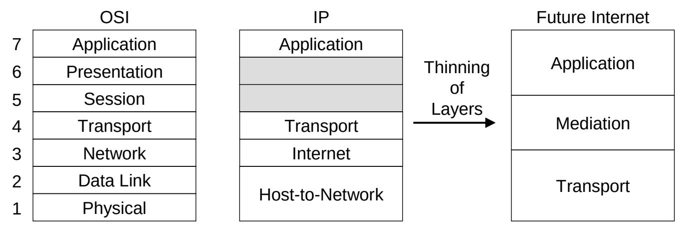

Fig. 1. Thinning of Layers

Within these layers, concepts from application-specific overlays can be applied. For example, for forming *virtual networks* for different applications which are operated in parallel. The overlays can structure their topology more directly to the needs of the application, e.g. for reflecting the relationships of communities, can apply routing strategies which are better suited for an application, and can make efficiently use of cross-layer information.

All in all, future layering concepts should provide advanced programming interfaces for applications in order to let them influence *a)* the (virtual) network structure, *b)* the routing mechanisms, and *c)* the resource management.

As a result, new *overlay and virtualization concepts* are currently under development, like PlanetLab's slices notation [18], [19], or cross-layer visibility as a services [20]. These concepts may implement parts of the Future Internet architecture and if successful, they can be absorbed into the future architecture.

Part II: P2P Overlays for the Future Internet

P2P-based mechanisms have demonstrated their capabilities to form highly scalable and efficient application-specific overlays. Considering the advances of directly supporting virtualization and overlays in future core elements, one may think about transferring P2P mechanisms from application layer to the above suggested mediation or transport layers. Next, the features of P2P overlays are outlined and two examples for P2P-based applications and services for the emerging Future Internet are given.

IV. P2P Overlays

Colloquially, the term P2P is often understood as "not client/server" and is used for a large variety of protocols, mechanisms, architectures and applications. Next, we will outline what is *P2P* and describe its features and *basic functions*.

What is Peer-to-Peer?. In general, a *P2P system* is used to share and exploit resources in a distributed and cooperative way. In detail, equal entities, denoted as *peers*, share resources via direct end-to-end exchanges on application layer. Typical shared resources are disk storage, CPU cycles, or data. The latter one comprises data files such as audio content or measurements as well as *metadata*, like the locations of users or files.

Virtual communication paths are established among peers reflecting their logical relationships. Their collection is called an *overlay*. A *P2P protocol* establishes the overlay and provides basic functions to enable the resource sharing. Functions to join and leave the P2P system are provided as well as functions to insert or retrieve resources. However, the leave function of a P2P system is often rather simple since the P2P concept assumes that any peer may depart from the system without prior notice. In P2P notation, the term *churn* denotes the stochastic process of peer turnover as occurring when peers join or leave the system. The collection of functions of a P2P protocol is called a *P2P service*. An application which is based on the use of a P2P service is referred to as a *P2P application* such as the attractive but disputable[3] P2P file-sharing applications. Due to their forming of overlays, P2P applications and P2P protocols are often referred to as *P2P networks*.

Features of P2P Networks. P2P networks are designed to overcome the drawbacks of the client/server paradigm. They allow peers to leave the system arbitrarily without

[2]Similar considerations can be applied to the OSI layering model since it partly shares the structure and interfaces with the IP layering concept.

[3]As long as controllability and copyright issues are omitted.

compromising the service as a whole. P2P systems support a very large number of peers and stored resources by the use of efficient cooperation strategies, e.g. multiple source download (MSD) and efficient data structures, e.g. Distributed Hash Tables (DHTs). They apply self-organization mechanisms to assign responsibilities to peers, e.g. for load balancing purposes.

P2P systems apply their own application specific addressing and routing. They assign own identifiers[4] to resources and nodes, e.g. hash function values of data files and IP addresses, and perform the routing based on them.

Types of P2P Networks. P2P networks are classified into *pure P2P* architectures, all peers are assumed to be equal, and in *hybrid P2P* systems, where some peers are distinguished from other peers. P2P overlays are denoted to be *unstructured* if the algorithms establish overlay links which do not follow a regular connectivity pattern. In contrast, P2P overlays are said to be *structured* if a predefined but generic organization scheme for the overlay exists, cf. [21].

Basic Functions of P2P Services. For the purpose of a better understanding of the capabilities of P2P systems, the basic functions of P2P services are discussed next. P2P functions can be separated into two classes, functions for *resource mediation* and functions for *resource access control*.

Generally, in P2P system resources are placed in general on arbitrary peers. *Resource mediation* functions are used to locate the resources in the overlay. They vary from centralized concepts such as index servers (e.g. in eDonkey), to highly decentralized approaches such as flooding protocols (e.g. in the Gnutella network), or DHTs (e.g in the Kademlia protocol).

P2P resource access control mechanisms permit, prioritize, and schedule the access to shared resources. A peer providing a resource gets connected to one or more resource consuming peers. At this time the communication takes finally place in a point-to-point manner. Resource access control mechanisms have similarities to conventional content distribution approaches such as multicast protocols or caching entities. The P2P approaches, however, differ from the conventional concepts by their decentralization and their intelligence placed at the network edge.

Example of a Popular P2P Content Distribution Architecture: eDonkey. eDonkey[5] is one of the most popular P2P applications [22] and provides file-sharing services. Its architecture is shown in Figure 2. In eDonkey, the resource mediation is implemented by specialized *index servers* which store the location of files. eDonkey peers publish and look up shared files using these index servers (cf. "publish edges" and "query edges" in Figure 2). Before an eDonkey peer can download a file, it first gathers a list of all file providers. To accomplish this, the peer connects to an index server and sends a query to it, which returns a list of matching files and their locations. Thus, the index servers handle all of the queries in order to relieve peers from this kind of traffic.

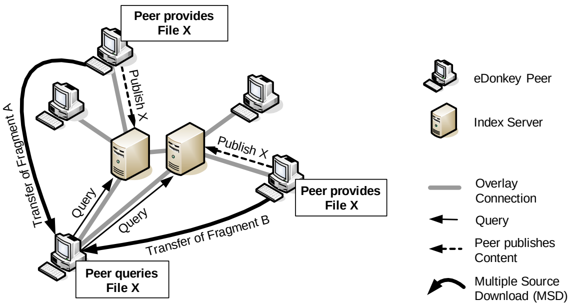

Fig. 2. The eDonkey architecture

eDonkey's resource access control is distributed and located on the peers. It uses the MSD concept, which permits a requesting peer to download different parts of a file in parallel from multiple providing peers (cf. "MSD edges" in Figure 2). In addition, this concept lets peers share fragments of a file before completing it. thus, reducing the dependability on providing peers and speeding up the file transfer.

eDonkey is a hybrid P2P architecture because of the specialized index servers. It is an unstructured P2P system since the overlay imposes no predefined organization scheme, resulting in a random structure.

Example of Efficiently Locating Resources: Kademlia. Kademlia is a DHT-based mediation mechanism focusing solely on looking up data or peers [23]. The basic idea of DTHs is to map keys of resources (e.g. data) and peers into an *id space* by a hash function. Each peer is responsible for parts of the id space. If resources or peers are queried then the ids are used for routing the queries. DHTs apply typically a generic *logarithmic routing scheme* which is implemented by *shortcut links*. Hence, they are classified as structured P2P networks. All-in-all, DHTs use their own addresses, i.e. the ids (= hash values), and their own routing, i.e. query forwarding along shortcut links, to efficiently look up resources.

The operation of Kademlia is outlined in Figure 3. The peer ids are leaves of a binary tree of the id space. Each peer's position is determined by the shortest unique prefix

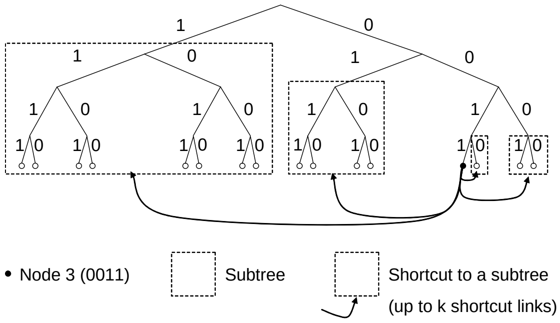

Fig. 3. The Kademlia binary tree

[4]The term "identifier" is abbreviated as "id".

[5]The term "eDonkey" subsumes the original eDonkey 2000 application and its derivatives, e.g. eMule.

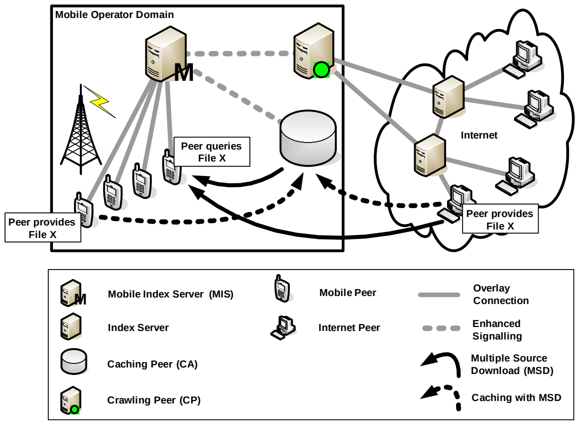

Fig. 4. A MP2P participation architecture for infrastructure-based networks

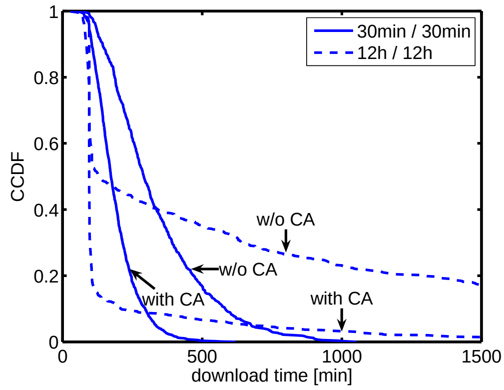

Fig. 5. Download time improvements with the Caching Peer

of its id. A peer divides the binary tree into subtrees of different prefixes that do not contain the peer id. Each peer maintains up to k shortcut links to peers in each subtree. The collection of shortcuts links to a specific subtree are denoted as a *bucket*. In Figure 3, for example, the peer with id "0011" maintains shortcut links to the subtrees with prefix "1", "01", "001" and "000".

A resource is stored at the peer whose *XOR difference* between the resource id and its id is minimal. When a peer looks up an id, it checks which subtree has the longest matching prefix with the id and forwards the query to α randomly selected peers from the bucket of that subtree. Each peer returns a bucket of a smaller subtree closer to the queried id. Then, the procedure starts recursively from the peer originating the query. Because of the prefix matching scheme, a lookup query is resolved in $O(logN)$ hops.

When a peers joins the Kademlia network, it fills its bucket by querying peers in different subtrees. A Kademlia peer maintains its buckets by recording received queries. In this way, the Kademlia algorithm facilitates the self-organization of the overlay.

V. A Mobile P2P Participation Application

The Future Internet will be a system for ubiquitous participation and will not distinguish between wireless and wireline users. Hence, a participation architecture has to *a)* level the heterogeneity of the peers and users, *b)* support mobility, and *c)* permit adaptivity and controllability.

A Mobile P2P File-Sharing Architecture. A P2P-based participation architecture to exchange user-generated content in mobile and wireline environments, denoted as a *MobileP2P (MP2P) file-sharing architecture*, was suggested in [24], [25]. The architecture facilitates the above stated requirements and is based on the hybrid P2P concept of eDonkey. The original eDonkey architecture is enhanced by three additional entities which are placed in the operator's domain, see Figure 4: *a)* the mobile index server (MIS), *b)* the caching peer (CA), and *c)* the crawling peer (CP).

The MIS is an enhanced eDonkey index server. It implemented parts of the resource mediation function in the MP2P architecture. The MIS tracks frequently requested content, triggers the caching peer to fetch it, and forces the mobile peers to download files from the CA by returning it as the major source. It steers how often a file is transmitted over the air interface. Furthermore, the MIS's decision which files to store on the CA imposes adaptivity and self-organization in the MP2P architecture. The content is automatically placed on the CA if performance is needed. Content for which efficiency is not needed as well as downloaded files remain on the peer. Hence, the MP2P application can still be denoted as a P2P application. Last but not least, the MIS can enforce copyright control by deregistering inappropriate contents.

The CA is a modified eDonkey peer. It implements parts of the resource access control. The CA appears to be an ordinary peer to other peers, but distributes popular content only. Since the CA is located in the wireline part of the network it replaces mobile-to-mobile file exchange by wireline-to-mobile communication, thus reducing the uplink data volume of mobile peers. Furthermore, the CA levels the heterogeneity of peers by serving wireline peers with high capacity and mobile peers with reduced throughput. The CA permits downloads even when original providing peers are off-line, thus supporting the mobility of users. In addition, it avoids the multiple exchange of the files across operator domains, thus reducing interdomain traffic.

The CP supports also the resource mediation, i.e. the locating of files. It searches for sources of files on behalf of mobile peers, thus reducing the mediation traffic on the air interface. In addition, it can perform the task even if the requesting peer goes off-line, reinforcing the support of peer mobility in terms of churn.

Performance of the Architecture. The performance enhancement by the CA is outlined next as an example for the efficiency of the suggested MP2P participation architecture. A comprehensive investigation is available in [25].

Figure 5 depicts the impact of the CA on the download time for files of 5 MBytes size when GPRS is used as the transmission technology. It depicts the complementary cumulative distribution function (CCDF) of the file download time with and without the CA and for different churn times, i.e. different average on-line/off-line times t_{on}/t_{off}. The churn time of peers reflects the effect of their mobility.

Figure 5 shows that the application of the CA results in significantly decreased download time CCDFs and low probability of high download times. The example demonstrates that the additional infrastructure entity, the CA, decreases the download time and is even able to reduce the effects of mobility.

Possibilities for the Future Internet. Although the suggested MP2P participation architecture is based on a today's P2P concept it addresses already future network challenges such as heterogeneity, mobility support, self-organization and user participation. The transport of large size contents will be facilitated by the advances in high capacity wireless and wireline transmission.

Furthermore, the MP2P participation architecture may give guidelines for the mechanisms available in the Future Internet. The future network should directly support the easy set-up of the additional infrastructure elements. Moreover, the future core elements, i.e. the future access points or routers, should assign the infrastructure elements and the peers to a specific instance of the participation overlay. Thus, achieving a closed but well-controllable virtual community. In the participation overlay, the peers (i.e. user peers and infrastructure elements) can set-up virtual links and set the routing tables on the core elements for efficient content distribution, with the infrastructure elements may having superior rights. Thus, the participation architecture would address the needs of the user, the applications, and the future operators.

VI. P2P-based Edge Service Monitoring

The provisioning of high performance applications by edge-based communities in the Future Internet requires monitoring mechanisms that are directly located on these nodes. Hence, distributed, scalable, and automatic monitoring mechanisms are needed.

Advantages of Edge-based Service Monitoring. When developing a distributed and edge-based monitoring infrastructure the question why not simply using a centralized concept is inevitable.

First, the benefits of distributed concepts can be exploited. Central monitoring units might fail due to faults or attacks. A distributed framework achieves increased *reliability*. It can avoid single points of failure and can still operate even when parts of it fail, assuming it is properly designed. In addition, a distributed framework provides for *scalability*. The workload can be shared among the distributed entities, i.e. each entity contributes a fraction of the required processing power, storage, and transmission capacity.

Second, an edge-based framework permits true *end-to-end based views* on service performance. A central network monitor can test the performance only from its location to a client, e.g. by sending ping messages or by probing the bandwidth. It cannot actively check the performance of a direct connection between two clients without additional software, cf. Figure 6a). Hence, the status of direct paths remain unknown. In contrast to this, a distributed concept can initiate tests locally between any two end points, cf. Figure 6b). The results can be reported to or queried by any other location if necessary, e.g. by a central manager.

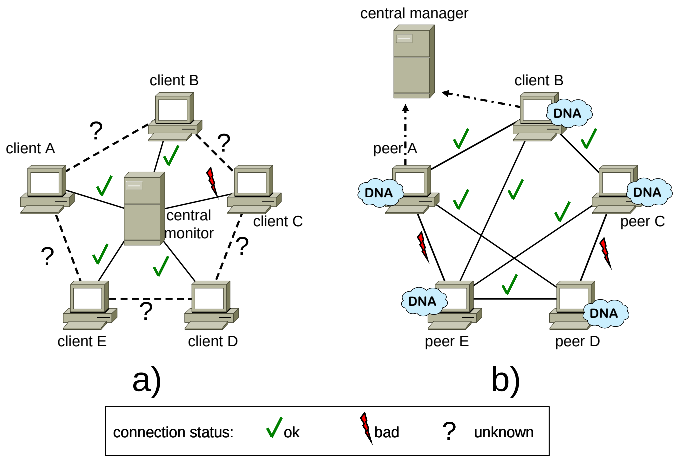

Fig. 6. Central monitoring vs. DNA Overlay

The P2P-based DNA Overlay. The application of an edge-based monitoring infrastructure, however, requires an automatic coordination of distributed entities. In addition, it is expected that the the network addresses of Future Internet nodes may change frequently, e.g. due to node mobility or roaming users. Hence, a scalable and self-organizing monitoring overlay using network independent addresses is required. These requirements are met by P2P mediation mechanisms using DHTs. Therefore, the DHT-based DNA (Distributed Network Application) overlay was suggested [26]. The DNA is a peer, i.e. an edge node, which runs the DNA software. A DNA can execute arbitrary tests implemented by plug-ins. The main purpose of the DNA overlay is to maintain the connectivity between the peers and to find any other peer in a reasonable amount of time, e.g. in order to execute tests on respectively between between nodes or to retrieve data stored on nodes. The P2P overlay can still operate even when a fraction of the peers fail.

The *Kademlia* algorithm has been chosen as the basis of the DNA overlay since it offers an appealing set of features:

1. *Short query paths:* Due to the symmetry (i.e. $d(x, y) = d(y, x)$) and the unidirectionality (i.e. for any id x and arbitrary distance $s > 0$, there is exactly one point y such that $d(x, y) = s$.) of the XOR metric.
2. *Parallel queries:* Timeouts on the forwarding path do not necessarily delay the search process. Thus, guaranteeing faster and more reliable searches under high churn rates.
3. *Flexible shortcuts*: In [23] shortcuts were chosen by the time of last contact to obtain more reliable connections. However, they can be chosen by any criterion like

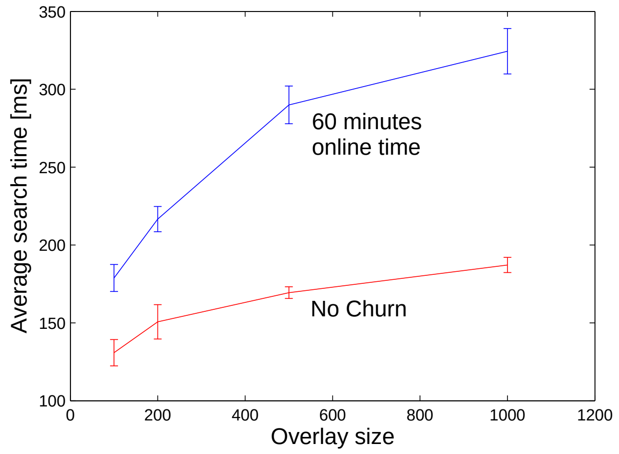

Fig. 7. Duration of a search as a function of the overlay size

trustability or latency. Thus, achieving trade-off between reliability or speed.
4. *Low periodic traffic:* Kademlia uses almost no periodic overhead traffic. Configuration information spreads automatically as a side-effect of key lookups.

Performance of the Architecture. The performance of the DNA overlay is outlined by the scalability of the search times with regard to the network size and churn times, cf. Figure 7. A comprehensive investigation of the architecture is available in [26]. Figure 7 depicts the simulation of an overlay of 1000 peers assuming a exponential distributed network transmission time with a mean of 50 ms per hop. The concave curves shows that the system does indeed scale. In an overlay without churn, the average search time stays below 200 ms. If churn is present then the search times increases but the overlay still scales.

Measurement Concept. While the DNA overlay facilitates mainly the maintaining of the virtual infrastructure, measurement mechanisms for assessing the end-to-end quality are needed. These mechanisms can be used as plug-ins to the DNA. Such a plug-in may be the proposed *Network Utility Function (NUF)* [27]. The NUF combines the observed network utility at the inlet and the outlet. It captures the damping effect of the network onto user-perceived quality from an end-to-end perspective. The NUF is highly intuitive due to its mapping to a simple value between 0 and 1. Hence, it is first step toward efficiently evaluating the QoE of future applications.

Possibilities for Managing the Future Internet. Major possibilities for managing the Future Internet are based in the self-organization capabilities of a P2P-based service monitoring overlay. The automatic operation reduces operational expenses and may adapt to the heterogeneity of peers, i.e. manage the roaming of users. By using P2P mechanisms, monitoring overlays can be set up on-demand by communities of edge nodes for specific applications. Thus, permitting managed and surveyed communities services. The coordinated edge-based approach can inform third parties, such as operators. Thus, it may bridge the gap between pure end-to-end control and network control when edge systems experience service degradations. Ideally, the support of service monitoring overlays should be included into the edge and core of the Future Internet, e.g. the node should directly be placed in overlays.

VII. Conclusion

The Future Internet is no longer a collection of links, routers, and protocols. It will be viewed as a network of applications, information, and contents and becomes an architecture for participation. Hence, *intelligent edge-based applications and services* will dominate the Future Internet. These applications and services will be typically implemented in an abstract way as *overlays*. Additional challenges come from recent advances in networking technology such as high speed optical networking, wireless transmission, or virtualization of links and routers.

In order to address these challenges new *methodologies for implementing and operating overlays* are needed. In particular, new mechanisms are required which permit edge-based overlays to structure their topology, to define their routing scheme, and to manage their resources independently.

The capabilities of P2P-based overlays for solving some of the control tasks of the Future Internet and its applications have been demonstrated by two examples: a Mobile P2P participation architecture and a P2P-based service monitoring infrastructure. The examples underline that application-specific structuring and routing of the P2P overlays should be supported directly by the future network nodes.

Moreover, the pressures from the efficiencies of overlays on the conventional layering model of IP and OSI currently initiate a rethinking of these models. A *thinning of the layers* and a *more basic separation of the layers* appears to be needed. This separation should focus on a split into three layers: *a)* the application layer (for addressing the application needs), *b)* the mediation layer (for network structuring, naming, and routing), and *c)* the transport layer (for reliable and cost-efficient transport).

Different international and national research initiatives and projects for evolving today's system into the Future Internet, such as GENI [28], IKT 2020 [29], or VINI [16], have been started lately. Hence, it can be expected that specific parts of the architecture of the Future Internet, such the virtualization of router and links, will soon be available in real world networks. Thus, the question whether an evolutionary or a clean-slate design is necessary for the Future Internet might be answered quickly by operating networks which absorb the future traffic.

Acknowledgement: The author would like to thank Prof. P. Tran-Gia, F.-U. Andersen, A. Binzenhöfer, T. Hoßfeld, M. Fiedler, M. Menth, and D. Staehle for stimulating discussions.

References

[1] C. Anderson, "People Power," *Wired*, p. 132, Jul. 2006.

[2] YouTube Inc., "YouTube Fact Sheet," 2006, Available at http://www.youtube.com/t/fact_sheet.

[3] MySpace Inc. , "MySpace.com," 2006, Information available at http://www.myspace.com/.

[4] T. Hoßfeld, S. Oechsner, K. Tutschku, F.-U. Andersen, and L. Caviglione, "Supporting Vertical Handover by Using a Pastry Peer-to-Peer Overlay Network," in *Proceedings of the 3rd IEEE Int. Workshop on Mobile Peer-to-Peer Computing (MP2P'06)*, Pisa, Italy, Mar. 2006.

[5] D. Varoutas, D. Katsianis, T. Sphicopoulos, A. Cerboni, S. Canu, K. Kalhagen, K. Stordahl, J. Harno, and I. Welling, "Economic viability of 3g mobile virtual network operators," in *Proceedings of 3G Wireless 2002*, San Francisco, USA, May 2002.

[6] Skype Technologies S.A., "Skype Homepage," Information available at `http://www.skype.com/`, 2006.

[7] C. Gauger, P. Kuhn, E. Breusegem, M. Pickavet, and P. Demeester, "Hybrid Optical Network Architectures: Bringing Packets and Circuits Together," *IEEE Communications Magazine*, vol. 44, no. 8, 2006.

[8] P. Chanclou, S. Gosselin, J. Palacios, V. Alvarez, and E. Zouganeli, "Overview of the Optical Broadband Access Evolution: A Joint Article by Operators in the IST Network of Excellence e-Photon/One," *IEEE Communications Magazine*, vol. 44, no. 8, 2006.

[9] UMTS Forum, "HSPA: High Speed Wireless Broadband - From HSDPA to HSUPA and Beyond," 2006, White Paper; Information available at `www.umts-forum.org`.

[10] WiMAX Forum, "Mobile WiMAX Part I: A Technical Overview and Performance Evaluation," 2006, White Paper; Information available at `www.wimaxforum.org/news/downloads/`.

[11] IBM Research, "Autonomic Computing," 2006, Information available at http://www.research.ibm.com/autonomic/.

[12] D. Patterson, A. Brown, P. Broadwell, G. Candea, M. Chen, J. Cutler, P. Enriquez, A. Fox, E. Kiciman, M. Merzbacher, D. Oppenheimer, N. Sastry, W. Tetzlaff, J. Traupman, and N. Treuhaft, "Recovery-Oriented Computing (ROC): Motivation, Definition, Techniques, and Case Studies," UC Berkeley Computer Science Technical Report UCB//CSD-02-1175, Mar. 2002.

[13] S. Oechsner, T. Hoßfeld, K. Tutschku, F.-U. Andersen, and L. Caviglione, "Using Kademlia for the Configuration of B3G Radio Access Nodes," in *Proceedings of the 3rd IEEE Int. Workshop on Mobile Peer-to-Peer Computing (MP2P'06)*, Pisa, Italy, Mar. 2006.

[14] C. Elster, D. Raz, and R. Wolff, "Autonomous End-to-End QoS monitoring," in *Proceeding of the IEEE/IFIP Workshop on End-to-End Monitoring Techniques and Services (E2EMON)*, Nice, France, May 2005.

[15] D. Soldani, M. Li, and R. Cuny, Eds., *QoS and QoE Management in UMTS Cellular Systems*, John Wiley & Sons, LTD, Chichester, 2006.

[16] A. Bavier, N. Feamster, M. Huang, L. Peterson, and J. Rexford, "In VINI Veritas: Realistic and Controlled Network Experimentation," in *Proceedings of the ACM Sigcomm 2006 Conference*, Pisa, Italy, Sep. 2006.

[17] T. Roscoe, "The End of Internet Architecture," in *Proceedings of the Fifth Workshop on Hot Topics in Networks (HotNets-V)*, Irvine, USA, Nov. 2006.

[18] L. Peterson, T. Anderson, D. Culler, and T. Roscoe, "A Blueprint for Introducing Disruptive Technology into the Internet," *Computer Communication Review*, vol. 33, no. 1, Jan. 2003.

[19] T. Roscoe, *Peer-to-Peer Systems and Applications*, chapter 33. The PlanetLab Platform, Springer, Berlin, 2005.

[20] R. Kompella, A. Greenberg, J. Rexford, A. Snoeren, and J. Yates, "Cross-layer Visibility as a Service," in *Proceedings of the Fourth Workshop on Hot Topics in Networks (HotNets-IV)*, College Park, Maryland, Nov. 2005.

[21] J. Wu, Ed., *Theoretical and Algorithmic Aspects of Sensors, Ad Hoc Wireless, and Peer-to-Peer Networks*, Auerbach Publications, Boca Raton, 2006.

[22] CacheLogic Research, "P2P in 2005," available at `http://www.cachelogic.com/home/pages/research/p2p2005.php`, 2005.

[23] P. Maymounkov and D. Mazieres, "Kademlia: A peer-to-peer information system based on the xor metric.," in *Proceedings of the 1st International Workshop on Peer-to-Peer Systems (IPTPS 2002)*, Cambridge, MA, Mar. 2002.

[24] J. Oberender, F.-U. Andersen, H. de Meer, I. Dedinski, T. Hoßfeld, C. Kappler, A. Mäder, and K. Tutschku, "Enabling Mobile Peer-to-Peer Networking," in *Proceedings of the First Int. Workshop of the EURO-NGI Network of Excellence on Wireless Systems and Mobility in Next Generation Internet, LNCS 3427*, Dagstuhl, Germany, Jan. 2005.

[25] T. Hoßfeld, K. Tutschku, F.-U. Andersen, H. de Meer, and J. Oberender, "Simulative Performance Evaluation of a Mobile Peer-to-Peer File-Sharing System," in *Proceedings of the Next Generation Internet Networks (NGI2005)*, Rome, Italy, Apr. 2005.

[26] A. Binzenhöfer, K. Tutschku, B. auf dem Graben, M. Fiedler, and P. Arlos, "A P2P-based Framework for Distributed Network Management," in *Proceedings of the Second Int. Workshop of the EURO-NGI Network of Excellence on Wireless Systems and Network Architectures in Next Generation Internet, LNCS 3883*, Loveno di Menaggio, Como, Italy, Jan. 2006.

[27] M. Fiedler, S. Chevul, O. Radtke, K. Tutschku, and A. Binzenhfer, "The Network Utility Function: A Practicable Concept for Assessing Network Impact on Distributed Services.," in *Proceedings of the 19th International Teletraffic Congress (ITC19)*, Beijing, China, Aug./Sep. 2005.

[28] GENI Consortium, "GENI - Global Environment for Network Innovations," Information available at `http://www.geni.net/`, 2006.

[29] German Federal Ministry of Education and Research (BMBF), "Förderprogramm IKT 2020," Information available at `http://www.bmbf.de/`, 2006.