Visualisation and Navigation in large scale P2P Service Networks

H. Unger, H. Coltzau

Abstract—In Peer-to-Peer service networks, where peers offer any kind of publicly available services or applications, intuitive navigation through all services in the network becomes more difficult as the number of services increases. In this article, a concept is discussed that enables users to intuitively browse and use large scale P2P service networks. The concept extends the idea of creating virtual 3D-environments solely based on Peer-to-Peer technologies. Aside from browsing, users shall have the possibility to emphasize services of interest using their own semantic criteria. The appearance of the virtual world shall intuitively reflect network properties that may be of interest for the user. Additionally, the concept comprises options for load- and traffic-balancing. In this article, the requirements concerning the underlying infrastructure and the graphical user interface are defined. First impressions of the appearance of future systems are presented and the next steps towards a prototypical implementation are discussed.

Keywords—Internet Operating System, Peer-To-Peer, Service Exploration

I. MOTIVATION

WITH the ongoing evolution of the world wide web, the user's role becomes more and more active. In Web 1.0, most users only where information consumers, only few users were able to add, change or remove contents. With Web 2.0 technologies, users can not only add contents to the web, they also have a higher level of interaction than before. Additionally, AJAX and similar technologies allow the creation and use of rich web-applications that no longer require any installation of software on the user's machine (except of a browser). As a consequence, the center of the user's activities moves from his machine to "the net". But the graphical user interfaces of all modern operating systems are still based on the desktop-metaphor developed in the late 1970's, where all documents and applications were organized on a virtual desktop. Intensive use of 3D-technology (e.g. BumpTop[1]) only conceals that users effectively cannot stand up from their virtual desktop and actually "go into the net". Instead, the internet browser is used to provisionally access remote services or applications. But without additional (active) components like JavaScript or Shockwave-Flash, Web-applications can not or only laboriously be used. Since such browser-plugins often miss the possibility to be integrated into a system- or network-wide security concept, the evolving risks need to be managed with additional efforts [5].

Even within the borders of local-area-networks only limited possibilities exist to explore the services offered within the LAN. For searching services outside the LAN, only commercially operated search-engines are available. Such search-engines fail to provide an intuitive way to explore contents and/or services easily. Similarity-searches or complex queries are only available for some specific types of data and only with limited scope. In the end, a user can never be sure that the result of a search-query really fits the user's needs best or at least almost best.

Standards for a consistent definition of services already exist (e.g. WSDL[2]), but most users cannot access such services intuitively or even create and publish their own services. The success of blogs, community-Portals and twitter shows that users have a great interest in actively shaping the contents and services of the WWW.

Consequently, a coherent concept for future operating systems is needed that fulfils at least following requirements:

1. The system dynamically creates a virtual overlay-network containing all services and contents available through the net. Users can access and explore this virtual network easily and intuitively. The system displays only those services and contents the user actually has access to.

2. The interface to services and contents is independent from the user's and the service's current location. There is no need to distinguish between the representation of local-area-network and wide-area-network services.

3. Users can explore the network without the need to use keyword-based searches. They can see other users (avatars) and interact with them.

4. Semantic searches and filters allow to explore the services in a context-sensitive manner. The representation of the network's content is concentrated on the services that are of interest for the user.

5. The graphical representation of the network reflects specific network-parameters for both services and communication channels (current load, estimated waiting-time, bandwidth, etc.).

---

[1] BumpTop is a 3D-Desktop with realistic animation of desktop elemements. s.a.: http://www.bumptop.com/

[2] WSDL=**W**eb**S**ervice **D**escription **L**anguage, s.a.: http://www.w3.org/TR/wsdl20/

6. An adequate service interface allows to access remote services from within applications.

Established features of modern operating systems, e.g. access control, data-security and safety, must be integrated into the system, too.

After giving an overview about the related work in chapter II, chapter III takes a closer look on three central aspects of future operating systems that consequently follow from the requirements above. In chapter IV we show several possibilities, how those aspects can be turned into a graphical user-interface. Finally, in chapter V, open research questions concerning the implementation are discussed.

II. RELATED WORK

Approaches for distributed operating systems already exist for a longer time. Plaice et.al. already presented prototypically a web-operating-system (WOS) [1][2], which made it possible to publish services on the WWW and to access them world-wide. But - alike to the WebOS [3] of Berkeley University - the project has not been carried on through the last couple of years.

SecondLife [4] is a centralized platform providing a virtual 3D-world, where users can acquire own properties and cover them with buildings. By the help of a proprietary scripting-language, interactive services can be offered. Even data from external servers can be integrated. Interaction with other users is possible. For trading purposes, an own currency exists, the Linden-Dollars, which can be exchanged towards US-Dollars. Nevertheless, there is a variety of disadvantages: the system brakes down frequently due to server overload and limited bandwidth. Further more, the burglary of identities and their misuse is a problem as well as the theft of virtual money. A basic disadvantage is that the user has no influence on the reliability and accessibility of the services offered because of the chosen client-server architecture. Above that, insecurity exists about property rights, which makes the commercial adoption difficult.

Schuster et.al. [6] developed a concept of a 3D-Web, which bases on a completely decentralised architecture. However, this virtual world only consist of loosely coupled regions, which are connected through so-called gates which are comparable to links in the current WWW. This way, the net may be explored, an intuitive orientation, however, seems to be possible only in networks with only a few smaller regions. Possibilities to search for services are not being discussed in the respective publications.

A strongly occupied and coherent presentation of the virtual world can be realized through distributed structuring algorithms presented by Sukjit and Berg [9][10], which base upon the natural example of cell growth. Using only local knowledge, the algorithms create a regular grid-like structure and assign a position on the grid to each member of the network. The algorithm is robust against failures of peers and independently corrects arising errors.

In the area of MMOGs (Massive-Multiplayer Online Games), [7] and [8] show that displaying the the area of interest (AOI) for a specific user works significantly better in distributed virtual environments using P2P- architecture, than it does in client-server architectures. By the help of Voronoi-charts they identify the close neighbourhood of a (mobile) peer. Because of this they are able to inform only the effected users when changes occur. This concept, however, refers only to the activity of the user in the net and to the arrangement of the results emerging. It does not contain a concept for how also the virtual world might be created in distributed manner.

III. CENTRAL ASPECTS

A. Services

In the context of this article, the definition of a service is deliberately week. A service is understood to be every kind of function, application or content, the user can access through the WWW. This includes, of course, w3c webservices but is by far not limited to them.

B. Changing the GUI-metaphor

Although many users have high bandwidth internet access, which allows to push aside the differences between local and remote service execution, the desktop has not been developed considerably. Therefore, it is almost not possible to present all resources the user has access to in a clearly arranged manner. In real life, we are used to move and orient ourselves in cities, supported by maps, if necessary. Consequently it seems appropriate to make use of city-metaphors for the graphical user interface, also. In the approach discussed here, a single computer (peer) in the service network is represented through a virtual building, which can be entered by the user. The buildings (peers) are connected by streets (the network). The user can move on these streets with his avatar, see and interact with other users. The outward appearance of a peer (the façade) shows a general representation of the services offered by that specific peer. In order to make the navigation more easy, the streets are arranged orthogonal. In contrast to the work of Schuster et.al. [6], the virtual world is presented on a coherent and dense grid.

The user has the possibility to explore and to orient in the net without necessarily having a precise goal. All aids of orientation known from the real-world, such as finger boards, distinctive buildings etc., can have an equivalence in the virtual world and through this may help to make the orientation considerably easier.

C. Adaptive Views

The more participators there are, the more confusing the virtual environment will be. Therefore it is of central importance for the user to find an appropriate arrangement of the peers and services that are of interest for him. By the help of keywords and their combinations the user shall be able to hide or to bring out peers and services. Each service will be described by the provider within an appropriate ontology so that the system can create semantic maps of services, where

similar services are arranged in the same neighbourhood. This is how the user gains a complex overview over all accessible services, which are of interest for him. If necessary, the search can be narrowed if, for example, only results of a certain geographical radius, of a certain provider or of certain features are presented. At any time, filters whose use instantly forms a new landscape, can be turned on or cut of. Especially through the use of meta-informations a similarity-search is possible. The user requests a list of all services which are, in certain aspects, alike to an already known service. In case the services offer geographic coordinates, the user can e.g. limit the search to his current whereabouts.

Of course, the access to services of a peer can still be made via URL, consisting of a protocol(service-)name, a unique identifier for the peer and service as well as other parameters of the service. Additionally, a unique service identification based on through tags or through applied filters and parameters is possible. The more the search is narrowed, the more a certain peer is identified by the parameters of the search.

All tables and figures you insert in your document are only to help you gauge the size of your paper, for the convenience of the referees, and to make it easy for you to distribute preprints.

D. Visualization of Network- and System-Parameters

Besides content criteria also the availability of services is of importance for the user. Peers, whose services have a large load already, may not be able to respond further requests in the requested quality. It is also possible that the available bandwidth between the user's computer and the hosting peer does not fulfil the needs of either the service or the user. Therefore it is of high interest for the user to be able to estimate such network parameters beforehand. Here, the user shall not be flooded with loads of numeric data but instead, the parameters shall be reflected in each peer's appearance and therefore be understood immediately and intuitively.

The load of a service can be easily recognised by the amount of avatars, who apply for the service at that time. Since the avatars can be recognized by other users, they can get an immediate impression of current load. In case avatars are modelled with physical features (i.e. they consume space), a "natural" limitation of the number of users, who can use the service at the same time, can be configured.

More examples for the reflection of network features in the appearance of virtual objects can be found: The broadness of a street can vary dynamically representing the quality of the connection (bandwidth, latency). The load of a peer can be modelled by using colour saturation, e.g. peers with a high load would appear grey. The same parameter could alternatively be shown through transparency which shows highly loaded peer only dimly.

Following this thought, also parameters who are not directly connected to features of the network could be represented intuitively.

Table 1 shows a range of possible classifications for

TABLE I
REPRESENTATION OF NETWORK PARAMETERS

APPEARANCE	Indicated Network Parameter
Broadness of the street	Available Bandwidth
	Efficiency
Height of a building	Avaialibiliy/Efficiency
	Number/Quality of Services
	Similarity in Semantic search
Transparency	Avaialibiliy/Efficiency
Color saturation	Similarity in Semantic search
Size of building area	Number/quality of services
Elevation (mountains, valleys)	Frequently Used Services/Popularity
	Number/quality of services
Traffic lights	Load balancing facilities
Barriers	

different parameters and their possible appearance the virtual world.

Some parameters can also melt into the arrangement of the peers. Peers, whose load is rather high at the moment can be faded out, appear transparent or distant. An exhausting discussion about all possibilities one could think of is hardly possible and will be topic of a following research paper.

IV. GRAPHICAL REPRESENTATION

The next graphics give a first impression of how the network surrounding could be presented to the user. By the help of an overview map (figure 1) an extensive orientation in the service landscape is possible. Only the rooftop of each building is shown, which shows a representative image of the service available on that peer. To make it easier to orientate, the streets go orthogonally. One could think of showing the building areas of each peer in either the same or variable size. The streets could either be managed by the peers or by the system as a whole (figure 2). This consequently means that one peer is either having the responsibility for one crossing (model a) or for one area which is framed by streets (model b), whereupon the streets are being added through the user's computer.

Fig. 1 Overview map

Both models have advantages and disadvantages. In the first case, the user has more freedom to arrange buildings on his building area. Additionally, the presentation of network parameters with reference to street features is easier. Nevertheless it is a disadvantage that using this method peers

might inhibit the alleyway to other peers and thereby isolate whole parts of the network. The second model avoids this problem but demands a higher effort in localizing and assigning objects (e.g. avatars), who are on the streets themselves.

Since the system shall refer to the users needs, it will not be possible for the service provider to buy building ground, as it is possible in SL. Nevertheless, a system wide ranking-algorithm may be able to emphasise peers of particular importance and provide them with more space to build upon. Still, the algorithm must be robust against manipulation, as is can be done with methods of search-engine-optimisation (e.g. Sybill-attack, Linkfarms etc.). In the end it has to stay in control of the users to take advantage of these accentuations.

model a   model b
Fig. 2 Responsibilities of the peers

Figure 3 shows exemplary the change of the map affected by the filter the user employs. It is obvious that peers with similar services are presented in the same areas of the map. This makes it easier to compare the services directly. In order to make the results even more comparable, it is possible without further ado to list the features of peers in a table, using another GUI-component.


Fig. 3 Views and perspectives
a) Perspective=Printing

b) Perspective=Communication

A first impression on how the exploration of the network could look like is presented in figure 4. The user can explore his environment and while he is doing so he notices the buildings representing the peers. If he finds an interesting service, he can enter the according peer. The inner appearance of the peer offers possibilities to access the services

V. RESEARCH QUESTIONS AND OUTLOOK

The authors believe that the system discussed here will only be accepted if the users and administrators stay in control over their hard- and software. Therefore, the service maps must be generated in a totally distributed (P2P) manner. Additionally to keeping the administrative control to the owner, P2P systems are also known for their scalability and stability.


Fig. 4 Network neighborhood

Hence, the questions to be answered concentrate on how adaptive maps of virtual service areas can be produced in a totally distributed environment. For the centrally organised case, interesting solutions already exist. A well known example is the use of neural networks to create semantic maps of pictures, music and other contents (*Self-Organized-Maps, SOMs*). An equal solution for the distributed case does not exist, yet.

The system has to be robust enough to compensate temporary failures and to limit their impact on the user to locally bordered effects. Mechanisms have to be found which, on demand, hand over a part of the functionality of the damaged peer to its neighbour until it is on-line again. As an interesting side effect, possibilities for load-balancing may result from this.

Furthermore the question arises, how three dimensional representation of a peer from a certain perspective can be transferred to several users in an acceptable amount of time, regarding limited bandwidths. Distributed caches could be installed, which would have to be semantic-aware, accurate and shall deliver latest information only, of course.

In contrast to SecondLife, burglary of ones identity shall not be possible. This means that a suitable trust-management has to be found. Some solution, basing upon delegation, already exist (e.g. [11]), those should be adopted if applicable.

REFERENCES

[1] John Plaice, Herwig Unger: *The Web Operating System – WOS.* Distributed Communities on the Web, Springer-Verlag, LNCS 1830, Berlin Heidelberg New York, 2000

[2] John Plaice, Peter Kropf: *WOS Communities - Interactions And Relations Between Entities In Distributed Systems.* Distributed Computing on the Web (DCW'99), University of Rostock Press, Rostock, Germany, 1999

[3] Amin Vahdat, Paul Eastham, Chad Yoshikawa, Eshwar Belani, Thomas Anderson, David Culler, Michael Dahlin: *WebOS: Operating System Services for Wide Area Applications.* Proc. 7th IEEE Symposium on High Performance Distributed Systems, 1997

[4] http://www.secondlife.com, Stand vom 30. Mai 2009

[5] A. A. Noureddine, M. Damodaran: *Security in web 2.0 application development.* Proceedings of the 10th international Conference on information integration and Web-Based Applications & Services (Linz, Austria, November 24 - 26, 2008)

[6] Daniel Schuster, Moritz Biehl, Thomas Springer, Jörg Müller: *Eine offene, skalierbare Architektur für 3D-Web als Informations- und Kommunikationsplattform der Zukunft.* Gemeinschaften in neuen Medien (GeNeMe); 2008

[7] S. Krause: *A case for mutual notification: a survey of P2P protocols for massively multiplayer online games.* In Proceedings of the 7th ACM SIGCOMM Workshop on Network and System Support For Games (Worcester, Massachusetts, October 21 - 22, 2008).

[8] S. Hu, G. Liao, G.: *Scalable peer-to-peer networked virtual environment.* In Proceedings of 3rd ACM SIGCOMM Workshop on Network and System Support For Games (Portland, Oregon, USA, August 30 - 30, 2004)

[9] Panchalee Sukjit, Daniel Berg, Herwig Unger: *A New, Fully Decentralized Grid Generation Method.* Proc. NCCIT Conference Thailand, 2009.

[10] Daniel Berg, Panchalee Sukjit, Herwig Unger: *Grid Generation in Decentralized Systems.* Angenommen für IDNS Conference, Klagenfurt, Österreich im Juli 2009

[11] Abhishek Agrawal, Douglas J. Brown, Aditya Ojha, Stefan Savage: *Bucking Free-Riders: Distributed Accounting and Settlement in Peer-to-Peer Networks.* Technical Report, CS2003-0751, UCSD, June 24, 2003