A Satellite-based Architecture for Internet P2P Applications

Evangelos Papapetrou
Department of Computer Science
University of Ioannina
45110, Ioannina, Greece
epap@cs.uoi.gr

Niovi Pavlidou
Department of Electrical and Computer Engineering
Faculty of Engineering
Aristotle University of Thessaloniki
54124, Thessaloniki, Greece
niovi@auth.gr

ABSTRACT

This paper proposes an architecture for unstructured p2p Internet applications. The paper focuses on applications that use superpeers. The proposed architecture utilizes satellites for building a backbone network for connecting superpeers. This network provides fully meshed connectivity to superpeers, in order to minimize the complexity of the p2p overlay network and at the same time preserves the flexibility in developing new applications. Technical aspects related to the implementation of the proposed architecture are discussed and useful conclusions are drawn.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design—*Network communications, Network topology, Wireless communication*; C.2.6 [**Computer-Communication Networks**]: Internetworking

General Terms

Computer Systems Organization

Keywords

p2p applications, Internet, satellite

1. INTRODUCTION

Over the last years, p2p applications [3] have gained tremendous popularity over the Internet, compared to the traditional client-server model. More than 60% of the total traffic in the Internet is owned to p2p applications. The p2p communication paradigm has been adopted in the development of a wide range of applications, including distributed computation, Internet service support, database systems, communication and collaboration services (i.e., instant messaging applications such as Aol and MSN), content distribution, internet telephony, etc. Real-life examples of such applications are the project of distributed computation Seti@home, instant messaging applications such as Aol and MSN, internet telephony applications such as Skype, and several content distribution applications such as Napster, Gnutella, Kazaa, Freenet, etc. The are two aspects in the success of the p2p concept. The first one relates to the end users' point of view and concerns the wide range of services, developed based on the p2p concept. On the other hand, from the networking point of view, the p2p communication paradigm provides flexibility and efficiency in developing new services over the complex and in most cases chaotic structure of the Internet. The basic concept of the p2p paradigm is to push functionality to Internet's end users, avoiding in this way the inflexibility related to the network-wide management of core network elements, such as routers. To this end, the p2p architecture introduces the concept of *overlay network* (or *overlays*) to organize its peers in a virtual network that bypasses the core of the Internet. Therefore, overlay networks are implemented by means of transport layer connections between the peers. In this way, functionality of the overlay network is located at the peers. However, the construction and use of overlay networks introduce sub-optimal aspects in the system operation. The first, relates to the difficulty encountered in taking into account the network architecture [13]. Peers connect to each other without taking into account network topology. Thus, overlays, in most of the cases, are not *network-aware*. This results in partially meshed overlay networks, in which one hop breaks down to several network hops, therefore decreasing the network efficiency. To address this problem, the most popular approach has been to derive mechanisms for establishing network-aware overlay networks [13]. However, this leads to frequent overlay reconstructions, thus increasing the complexity. Another disadvantage, existent in state-of-the-art p2p applications, is the problem of *scalability* [3]. The performance of an overlay network decreases with its size, since there is no specific structure in its development. Structured overlay networks have been proposed to address the scalability issue [3]. However, structured overlays involve increased complexity in constructing and maintaining the overlay. As a result, large scale implementations do not exist.

In this paper, a new architecture is proposed to address both the issues of scalability and overlay network awareness involved in the deployment of p2p applications over the Internet. The new architecture, utilizes satellites as fundamental backbone network elements for building the proposed architecture. Satellite systems have been considered in the past for providing backbone connectivity to the Internet and other terrestrial networks [7],[12],[4],[6]. However, in the context of p2p connectivity and overlay networks, the role of satellites can be revaluated on the basis of their inherent capability for broadcasting and its ability to provide efficient backbone connectivity similar to the one provided by the concept of overlays.

The rest of the paper is organized as follows. In Section 2 the investigated problem is addressed in the context of state-of-the-art p2p applications, while in Section the concept of using satellites for building p2p applications, is presented. Then in Section 4 the proposed architecture is delineated and technical aspects are discussed. Finally, in Section 5 useful conclusions are drawn.

2. NETWORKING PERSPECTIVES OF P2P APPLICATIONS

Recently there has been a peer-to-peer (p2p) trend in developing applications and services over the Internet. The popularity of this approach stems from the need of Internet's end users to efficiently share data without the need of a server, as in the client/server model. The p2p concept represents the flexibility in sharing and distributing information across the Internet. However, there are some technical issues related to this approach.

2.1 Overlay Networks and related issues

In order to accomplish direct communication between peers, p2p applications introduce the concept of *overlay* network (or overlay). Overlay networks represent an application layer networking of peers. In today's Internet, connections between peers are implemented at the transport layer and in most cases by means of TCP connections. However, other solutions, such as UDP, are possible. The reason for this methodology lies in the difficulty to manage the network layer of the TCP/IP stack across the whole Internet. Making the networking layer (i.e., the IP) aware of the characteristics of the application layer would require a world-wide upgrade and management of routers, which is not an easy task due to the Internet's size and chaotic structure. The concept of overlay represents the trend to push network functionality to its edges (i.e., the end hosts) in order to facilitate quick and flexible deployment of new applications.

Nevertheless, the tradeoff for quick and flexible application development is the overheads introduced by the use of a transport layer protocol, such as the extra headers imposed to data and other signaling information exchanged by the transport protocol. Moreover, a path within the overlay network consists of several subsequent transport layer connections, thus aggravating the imposed overheads. For example, in the case of an overlay that is constructed by means of TCP, the overheads for acknowledging data in a multihop path within the overlay, are imposed in each hop. Furthermore, the transport layer of each intermediate node should parse packets , which induces additional overhead to intermediate nodes. Furthermore, there are additional overheads related to the construction and the maintenance of the overlay network. For example, consider the case that TCP is

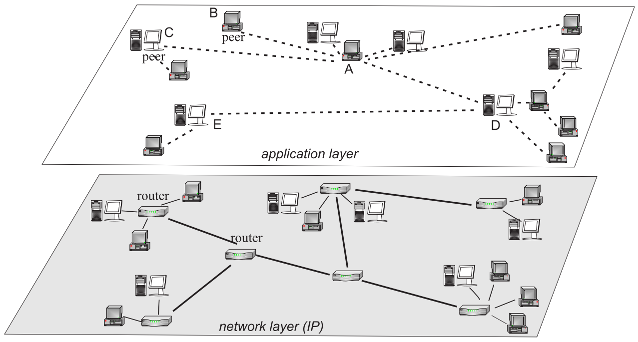

Figure 1: An example of an overlay network

used for the construction of the overlay. There is overhead in setting up connections of new peers that enter the overlay, as well as in maintaining them. On the other hand, TCP connections are highly susceptible to temporary link failures and congestion conditions. As a result, in real-life networks, the reconstruction of the overlay is frequently required. Finally, overlays should provide transport services that fit the profile of the exact p2p application. For example, the best choice for real-time applications would not be an overlay constructed by means of TCP since in such a case acknowledgements are not required.

2.2 Application layer p2p connectivity and implications

Besides the role of overlay networks, another important aspect of the performance of p2p applications relates to the application layer networking of peers, which serves the flexibility in exchanging data. To this end, the construction of overlay networks depends only on application layer semantics and rarely takes into account the actual structure of the network. However, this approach introduces overheads in the operation of the underlying network. In several cases the semantics of the application layer mandate that some peers have to connect with other peers, which are distant in terms of network connectivity. Let us consider for example the case of the overlay network depicted in fig. 1), where node B is connected to node A. As a result, an application layer connection may span several network hops, resulting in waste of resources in intermediate links and in degradation of performance due to frequent congestion conditions within the Internet. Alternatively, in the previous example, node B could connect to node C, which is in close proximity. The aforementioned problem is known as the *overlay network awareness* problem and affects the efficiency of communications within the overlay. Apart from traditional point-to-point communications, p2p applications are usually make use of multicasting. For example, p2p telephony applications require multicasting over the overlay for supporting conferencing. The efficiency of such communication schemes raise considerable skepticism. Consider for example the case that node A wishes to multicast a data packet to nodes B,C,D and E in figure 1. In this case, node A should originally transmit three packets to nodes B, C and D and then node D should retransmit the data packet to node E. This results in eleven transmissions between routers of the underlaying network. On the contrary, if multicasting is performed at the network layer, delivery to all interested nodes

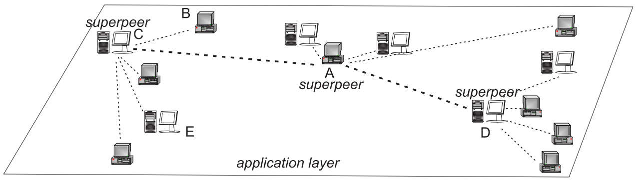

Figure 2: An architecture with super-nodes.

may be accomplished with only five transmissions.

2.3 Data-centric communication

The p2p concept introduces a *data-centric* communication paradigm. This means that the set of peers that exchange information is determined based on application layer characteristics rather than being predetermined between an origin/destination pair. For example, in p2p file sharing applications the peers that are going to exchange data are determined by the p2p network based on the file that is requested. In many applications, especially those build for sharing content, the actual exchange of information is preceded by the routing of a query message through the overlay, in order to locate the peer or peers that the requested data reside [3]. Usually a flood-based technique is utilized for routing the query to all or part of the peers. Routing of queries induces significant overhead to partially meshed overlay networks, usually used by today's p2p applications in the Internet. There have been several proposals for minimizing the cost of a query [3], however, scalability is still an issue, especially in the context of the increased popularity of such applications. An approach that eliminates the problem, introduces the concept of *structured overlay* [3], [10]. In this case the overlay is built such that a limited number of hops is required for locating data without the use of flooding. However, the tradeoff for structured overlays is the increased cost for maintaining them, especially in a network of frequently changing conditions. Instead, the most common approach in the Internet for minimizing the cost of query routing is the hierarchical construction of overlays [3]. According to this approach besides simple peers, there are also superpeers, which have increased communication and computational capacity. All simple peers connect to some superpeer and the overlay is completed with connections between superpeers (see fig. 2). This kind of architecture, although minimizing overhead, does not solve the problem.

3. THE SATELLITE-BASED ARCHITECTURE APPROACH

Following the discussion of the previous section, it is clear that the performance of state-of-the-art p2p applications is encumbered by their current architecture. Two are the major drawbacks, with the first being the construction of the overlay without considering the underlying network. The second one relates to the large size and the partially meshed topology of the overlay, which raise scalability issues. Such inefficiencies are closely related to the current architecture of the Internet, therefore the most effective solutions should involve the development of a new architecture. In order for such an architecture to be efficient, it should:

- minimize the use of overlay networks in order to evade the problem of network awareness; efficient resource utilization should be attained for data transport in order to enhance cost-effectiveness,
- provide fully meshed connectivity between most, and if possible, all peers; this approach on one hand targets at providing efficient resource utilization, and on the other at eliminating inefficiencies caused by the large number of peers.

Satellites emerge as the best candidates for fulfilling the aforementioned requirements, since they can provide an efficient and cost-effective solution due to their inherent characteristics, such as the large scale coverage and the broadcasting capability. In particular, satellites may provide direct broadband connections between peers in order to establish a fully meshed connectivity. However, the incorporation of satellites in a satellite-based architecture for p2p applications, in order to meet the specified requirements, is not straightforward due to the plethora of networking options provided by state-of-the-art satellite technology. Furthermore, some additional requirements are of essential importance for the success of such an architecture:

- World-wide coverage: although satellites provide large scale coverage, in the context of the Internet, it is essential to provide services to peers world-wide. In this way, the cost-effectiveness of the proposed architecture is also enhanced.
- Flexibility: while minimizing the use of overlays, the advantage of flexibility in developing new applications and services, should be maintained in order to realize a competitive alternative to the terrestrial Internet. In this context, flexibility is still interpreted as transparency to IP.
- On-demand connectivity: the nature of p2p applications is characterized by frequent variations of peers' participation. Consequently, providing services on an on-demand basis is fundamental for supporting p2p protocols.
- Low complexity: since the target is to derive an architecture for world-wide p2p applications, low complexity of the system is a key feature for supporting flexibility. Furthermore, simplicity should also characterize the user interface in order to enhance the competitiveness of such a solution.

The use of satellites for providing communication services has been proposed in several cases in the past [4],[6]. However, in many cases the acceptance of such proposals has been limited due to the existence of more cost-effective terrestrial solutions. In the case of p2p applications however, the major advantage of satellites compared to terrestrial solutions is the broadcasting capability and the large scale coverage. However, the wide acceptance of a satellite-based topology depends strongly on cost-effectiveness which can be achieved by fulfilling the aforementioned characteristics.

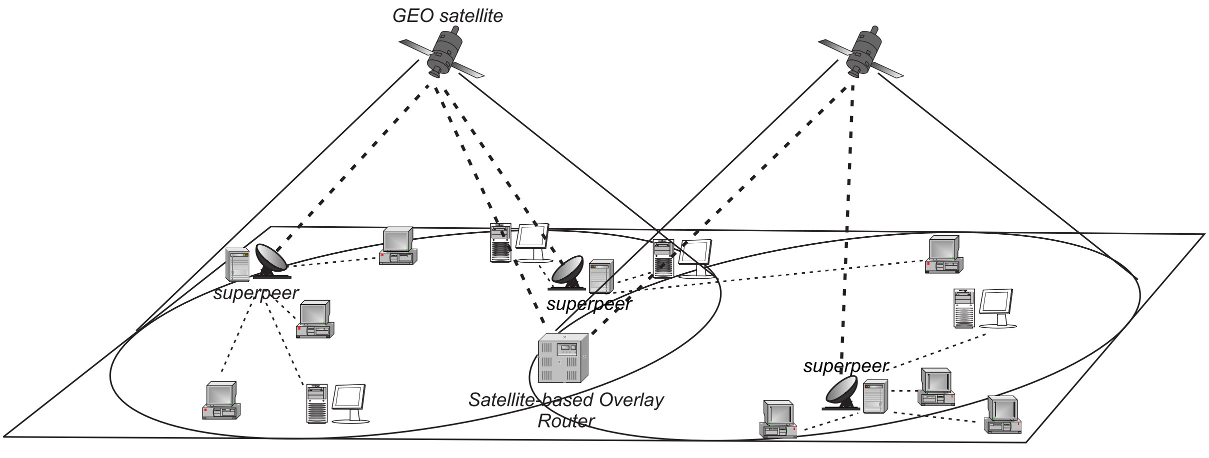

Figure 3: The proposed architecture

4. PROPOSED ARCHITECTURE

In this section we propose a new architecture for building p2p applications over the Internet. To this end, we focus on unstructured p2p systems, which are widely used in the Internet. Furthermore, we are especially interested in architectures that incorporate superpeers. In real-life systems, superpeers are frequently used to handle the vast number of peers and reduce the complexity of the overall system. *The key concept of the proposed architecture is the introduction of satellites for providing connectivity between superpeers.* To this end, we aim at forming a satellite backbone for providing quick and high speed communication services to superpeers. The proposed architecture is presented in fig. 3. In this architecture, superpeers are connected by direct connections through satellites. Instead, simple peers still use transport layer networking (e.g., TCP connections) for connecting to superpeers, through the terrestrial infrastructure. The motivation for using satellites only for connecting superpeers is twofold. On one hand, the equipment cost for establishing communication through a satellite is limited only to some peers. Although this cost is minimized due to recent technology advances, it could still be considered significant for simple users. On the other hand, superpeers can be viewed as traffic aggregation points. Therefore, the proposed architecture allows of the efficient utilization of high speed satellite connections, augmenting in this way the cost effectiveness of the proposed solution. It must be noted that satellites could be used in an optional basis. There is no actual limitation for a hybrid system, in which superpeers may be connected also through terrestrial infrastructure. The capacity for modular deployment of the proposed solution represents a great advantage for its wide acceptance.

An integral part of the proposed architecture is the space system, which consists of the satellites and the *Satellite-based Overlay Router (SOR)*. As far as communication satellites are concerned, there are more than one possible solutions such as GEO (geostationary), MEO (medium) or LEO (low) earth orbit satellites. However, the requirements for low complexity and world-wide coverage, render GEO satellites the most appropriate solution. Furthermore, GEO satellites may provide services to greater areas, making possible traffic aggregation at a greater extent. Finally, the DVB-RCS technology [1],[2], which is used in GEO satellite networking, is well developed and constitutes another motivation for selecting GEO satellites. As far as the SOR is concerned, its functionality is to extend the coverage of the proposed architecture. Although GEO satellites may provide large scale coverage, in order to achieve global support, SOR may be used for simple forwarding of data amongst satellites. As a result, the routing capabilities at SOR may be minimized along with the complexity of the proposed architecture. Moreover, the use of only one satellite for supporting parts of an existing p2p application, is also possible.

4.1 Technical Aspects

Before describing the mechanisms of the proposed architecture, it is important to discuss some technical issues regarding its feasibility. As mentioned previously, a key design decision is to use GEO satellites with DVB-RCS, which is a well established technology. Furthermore, in order to establish mesh connectivity, we adopt the use of regenerative satellites with on-board processing (OBP) capabilities [2]. Moreover, the proposed architecture aims at being an integral part of the Internet. To this end, superpeers may act as gateway stations, as far as DVB-RCS is concerned. Furthermore, from the technical point of view, our proposal utilizes the IP technology over DVB-RCS. The issue of deploying IP over DVB networks has been extensively investigated [9], [8], [5], however some architecture specific issues must be addressed.

While IP connectivity is a necessity, a major objective of the proposed architecture is to preserve the flexibility of developing and deploying new application layer protocols. This suggests that increased transparency to IP is desired. To this end, all peers within a satellite footprint should be organized in one IP subnet. Since we assume GEO satellites, this results in a predefined number of IP subnets (one per satellite). Furthermore, the lack of satellite mobility with respect to the Earth's surface, simplifies the management of such IP subnets. In order to establish network-wide communication, the SOR should be capable of routing data through subnets. However, the required routing functions are of minimum complexity, since there are only two options for incoming packets. The one is that the packet is destined to a superpeer within the IP subnet, while the other is that the packet passes through. Moreover, for passing through packets there is only one option; to be forwarded toward the SOR that provides connectivity to the next satellite. In other words, each SOR should maintain two entries for routing data; one for destinations within the IP subnet and one for the next IP subnet. This is represented in fig. **??**. From the DVB network point of view, the SOR may be implemented as a double gateway station (one per satellite). The described routing function is necessary for interconnecting the gateway stations.

In the process of forming and maintaining IP subnets, the IP addressing of superpeers is an important mechanism. Addressing may follow a fixed assignment, however in order to serve the on-demand nature that characterizes the participation of superpeers to the p2p application, an efficient solution would be to use the DHCP protocol. Such a server could reside at the SOR or the NCC of the DVB-RCS network. Another important issue relates to the address of the IP subnet itself. Besides using assigned IP addresses, satellite network operators may use private IP addresses. In such a case, superpeers must also implement Network Address Translation (NAT) in order to establish connectivity to the terrestrial IP network. Finally, an important mechanism that relates to the addressing of superpeers is the translation required between IP and DVB layer-2 addresses. Such mechanisms have been proposed in the literature [5].

4.2 P2P Application Mechanisms

In this subsection we describe the mechanisms required for establishing the functionality of a traditional p2p application on the proposed architecture. On-demand connectivity is a fundamental feature of p2p traditional applications. Thus, a peer is required to join the overlay network by sending a join message to a peer which is already in the overlay. In the proposed architecture, peers are also required to join the overlay. However, the procedure of joining is not the same for superpeers and simple peers. A superpeer, after establishing IP connectivity, can be considered to have joined the backbone of superpeers. On the contrary, simple peers may use the traditional approach for connecting to one of the superpeers. In traditional p2p applications, the selection of a superpeer is frequently determined by application layer specifics. However, the advantage of the proposed mechanism is that peers may use network layer metrics (for example the hop count or the delay) for selecting the superpeer to which they wish to connect. According to this approach, given a peer, its distance (in hops) to every superpeer is the same and equals the count to the closest superpeer plus the one hop through the satellite backbone. On the contrary, the distance to superpeers through the terrestrial infrastructure is, in principal, greater. Clearly, the network awareness of the overlay is enhanced, in the context of the proposed architecture.

One of the most interesting mechanisms, employed especially in resource sharing p2p applications, is the discovery of available resources that reside in some peers in the overlay. In unstructured systems, resource discovery is initiated by the requesting peer by transmitting a message that describes the requested resource. This message is propagated through the overlay in a flood-based fashion, which raises scalability concerns. The proposed architecture minimizes the cost of performing a resource discovery. A node that wishes to perform a resource discovery, forwards its relevant message to the connected superpeer. The superpeer is responsible for forwarding the message to all other connected peers and to the rest of the superpeers. The existence of the peer-superpeer hierarchy minimizes the cost of the discovery on the terrestrial network, while the respective cost on the superpeer backbone is eliminated by the use of the satellite. This is because only one transmission is required for delivering the query message to all superpeers in the IP subnet. As a result, the bandwidth consumption is minimized and scalability is achieved in the satellite backbone network.

Finally, the actual data transfer between peers is also of interest. Peers may exchange data, either in an one-to-one manner or by multicasting. As far as, point-to-point communication is concerned, peers in the proposed architecture may use either the satellite network or the terrestrial infrastructure. However, the proposed architecture provides an advantage for multicast communication of peers. Establishing IP multicast over satellites has been investigated in the literature [11].

5. CONCLUSIONS

In this paper we have presented a new architecture for building unstructured p2p applications that use superpeers. The proposed architecture utilizes satellites for building a backbone network that provides mesh connectivity to superpeers. The backbone network enhances the network awareness of the p2p overlay, minimizing in this way the consumption of network resources. Furthermore, the high speed connectivity between superpeers benefits several functions of a p2p application, such as resource discovery and data transfer. Finally, the proposed architecture preserves the flexibility in developing and using new application protocols over the Internet.

6. ACKNOWLEDGMENTS

This work was carried out in the framework of the EU funded Network of Excellence "SatNEx II" (contract No. IST-027393).

7. REFERENCES

[1] Etsi en 301 790 (v1.3.1), digital video broadcasting (dvb); interaction channel for satellite distribution systems, March 2003.

[2] Etsi tr 101 790 (v1.2.1), digital video broadcasting (dvb); interaction channel for satellite distribution systems; guidelines for the use of en 301 790, January 2003.

[3] S. Androutsellis-Theotokis and D. Spinellis. A survey of peer-to-peer content distribution technologies. *ACM Comput. Surv.*, 36(4):335–371, 2004.

[4] B. Evans, M. Werner, E. Lutz, M. Bousquet, G. E. Corazza, G. Maral, R. Rumeau, and E. Ferro. Integration of satellite and terrestrial systems in future multimedia communications. *IEEE Wireless Communications*, 12(5):72–80, 2005.

[5] G. Fairhurst and M.-J. Montpetit. Address resolution mechanisms for ip datagrams over mpeg-2 networks. Internet Draft, March 2007.

[6] Y. Hu and V. O. Li. Satellite-based internet: a tutorial. *IEEE Communications Magazine*, 39(3):154–162, 2001.

[7] A. Iera and A. Molinaro. Designing the interworking of terrestrial and satellite ip-based networks. *IEEE Communications Magazine*, 40(2):136–144, 2002.

[8] M.-J. Montpetit, G. Fairhurst, H. Clausen, B. Collini-Nocker, and H. Linder. A framework for transmission of ip datagrams over mpeg-2 networks. RFC 4259, November 2005.

[9] H. Skinnemoen, R. Leirvik, J. Hetland, H. Fanebust, and V. Paxal. Interactive ip-network via satellite dvb-rcs. *IEEE Journal on Selected Areas in Communications*, 22(3):508–517, 2004.

[10] I. Stoica, R. Morris, D. Karger, M. Kaashoek, and H. Balakrishnan. Chord:a scalable peer-to-peer lookup service for internet applications. In *Proc. of SIGCOMM 2001*, 2001.

[11] Z. Sun, M. Howarth, H. Cruickshank, S. Iyengar, and L. Claverotte. Networking issues in ip multicast over satellite. *International Journal of Satellite Communications and Networking*, 21(4-5):489–507, 2003.

[12] T. Taleb, N. Kato, and Y. Nemoto. Recent trends in ip/ngeo satellite communication systems: Transport, routing, and mobility management concerns. *IEEE Wireless Communications*, 12(5):63–69, 2005.

[13] M. Waldvogel and R. Rinaldi. Efficient topology-aware overlay network. *SIGCOMM Comput. Commun. Rev.*, 33(1):101–106, 2003.