3. Past and Future

Jörg Eberspächer, Rüdiger Schollmeier
(Munich University of Technology)

3.1 Status Quo: Networks (Over)Filled with Peer-to-Peer Traffic

Within the last few years, starting with the introduction of Napster in May 1999, the disruptive technology of Peer-to-Peer networking has encountered an enormous growth. Today the traffic caused by Peer-to-Peer networks represents a significant portion in the Internet. For example in the German Research Network (Deutsches Forschungsnetz DFN) Peer-to-Peer causes up to 60 percent of the traffic [210]. Similar trends can be observed in other networks e.g. in the Abilene backbone [42]. As we can observe in Figure 3.1, at the beginning of 2002 the traffic caused only by the signaling traffic of Peer-to-Peer applications (no user-data-transfers included) already amounts to 50 percent of the total traffic volume (see Figure 3.1).

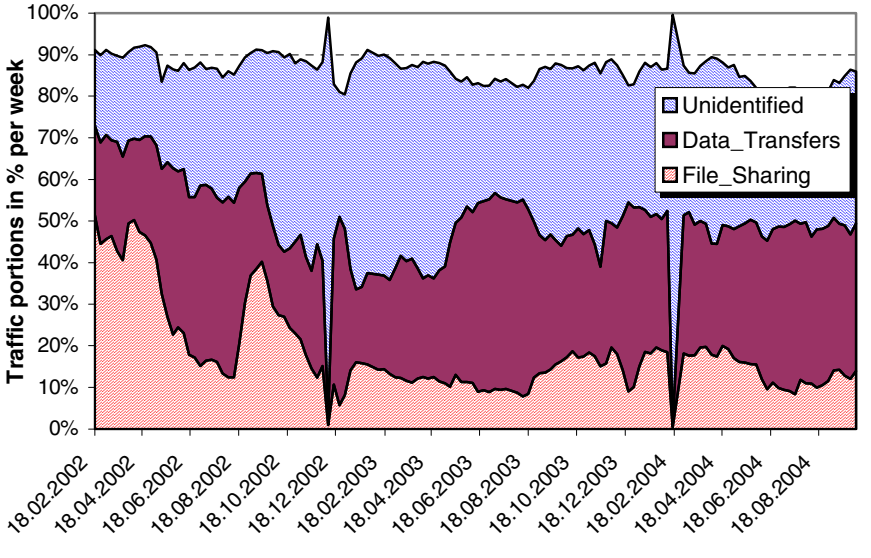

Fig. 3.1: Portions of traffic measured per week in the Abilene Backbone from 18.02.2002 until 18.07.2004 (peaks at 18.12.2002 and 18.02.2004 result from measurement errors. source:[42])

R. Steinmetz and K. Wehrle (Eds.): P2P Systems and Applications, LNCS 3485, pp. 17-23, 2005.

Until the end of 2004 the amount of Peer-to-Peer traffic decreased down to a value of approximately 15 percent. This might point to an increasing efficiency of the Peer-to-Peer protocols, since the signaling traffic is reduced or to a decreasing usage of Peer-to-Peer applications. However if we also have a look at the unidentified traffic and the traffic identified as data-transfers, we can observe that these volumes are increasing and that the total amount of traffic stemming from these three sources stays at a constant level of nearly 90 percent. Analyzing the networking techniques of different Peer-to-Peer applications in more detail this could also indicate that Peer-to-Peer applications are "going underground", i.e. they use TCP port 80, so that they can, on the port level, not be distinguished from common data transfers. Further on more and more Peer-to-Peer applications use so called port hopping, meaning that they change frequently their communication port during run time and can thus not be identified as file sharing applications on the port level. Thus the amount of unidentified traffic and data transfers increases and the amount of identified Peer-to-Peer traffic decreases, while the total amount stays at a constant level of approximately 90 percent.

Hence, Peer-to-Peer communication plays a dominant role in todays networks and is also proliferating into many new application areas. In this chapter, we will have a look at the development of Peer-to-Peer applications in the last few years, analyze the development of the capabilities of the user terminals and finally consider possible directions that development of Peer-to-Peer technology might take in the future.

3.2 How It All Began: From ARPANET to Peer-to-Peer

Peer-to-peer networking is not new. Basically, it started in the late 1960s with the establishment of the ARPANET. The goal of this physical network was to share computing resources and documents between different US research facilities. In this original system there was nothing like a typical client or a typical server. Every host was being treated equally and one could therefore call this network a first Peer-to-Peer network, although this network was not self organizing and no overlay network was established. Everything matched still to a large extent the physical connection and not virtual connections as we can observe them in today's Peer-to-Peer networks.

Nevertheless, the early "killer-applications" in the ARPANET were typical client server applications, namely FTP and TelNet. Every computer could run a server and client so that every participating node could request and serve content. The missing part however was an instance to inform the participating nodes where which content is provided.

Therefore in 1979 the UseNet protocol was developed. It is a newsgroup application, helping to organize the content and offering a self-organizing approach to add and remove newsgroup servers by the participating users

via a rigorous "democratic" process. However the application itself is still a typical client server application, with "simple" requesting nodes, the clients, and more powerful content providing nodes, the servers.

Approximately 10 years later, around 1990, there was a rush of the general public to join the Internet community. A number of applications were developed, like WWW, email and streaming. Modem connections via the SLIP and the PPP protocol became increasingly popular, with the result that millions of commercial users and customers joined. The basic communication model was the client/server model, with a simple application on the user side (e.g. the Web Browser) which initiates a temporary connection to a well known server, from which the client downloads the requested content and then disconnects again. It is a simple and straightforward model, which provides also the content provider with an easy model to administrate and to control the distribution of content. Further on new security concerns in the Internet had to be taken into account resulting in an Internet partitioned by firewalls.

3.3 The NAPSTER-Story

This was about to change in May 1999. Home users started to use their connected computers for more than just temporarily requesting content from web or email servers. With the introduction of the music- and file-sharing application Napster by Shawn Fenning [437], the users opened their computers not only to consume and download content, but also to offer and provide content to other participating users over the Internet. This phenomenon is best described by the artificial term SERVENT for one node, which is a combination of the first syllable of the term SERVer and the second syllable of the term cliENT.

Comparing the Peer-to-Peer networks, which started with Napster, to the architecture established by the ARPANET we can observe that in contrast to today's Peer-to-Peer realizations, the ARPANET was not self organizing. It was administrated by a centralized steering committee, and did not provide any means for context or content based routing beyond "simple" address based routing. In current Peer-to-Peer networks, the participating users establish a virtual network, entirely independent from the physical network, without having to obey any administrative authorities or restrictions. These networks are based on UDP or TCP connections, are completely self-organizing and are frequently changing their topology, as users join and leave the network in a random fashion, nearly without any loss of network functionality.

Another decentralized and application-layer oriented communications paradigm is Grid computing which became famous with the project SETI-home [557] . It is often compared to Peer-to-Peer as being a more structured approach with the dedicated aim to especially share computing power and

storage space for distributed computations and simulations [217]. Yet, the basic principles in Peer-to-Peer and Grid are similar. However concerning the number of participating users and thus also the traffic volumes Grid computing is taking currently a minor role. Nevertheless it has a high growth potential.

Because of the mostly illegal content shared in the Napster network (content was mostly copyright protected, mp3 compressed music), the Recording Industry Association of America (RIAA) filed in December 1999 a lawsuit against Napster Inc. This was possible, because the Napster network relies heavily on a centralized lookup/index server operated by Napster Inc. This server, which represents a single point of failure in the Napster network could therefore be targeted by the RIAA.

3.4 Gnutella and Its Relatives: Fully Decentralized Architectures

Under the impression of the lawsuit against Napster the company Nullsoft released in March 2000 the Gnutella application as an open source project. In Gnutella the participating peers do not only act as a servent. They additionally take over routing functionalities initially performed in Napster by the Napster server. Thus not only the file exchange and provisioning are completely distributed, but also the content lookup/routing functionality. Thus any central entity and also any single point of failure is avoided. The Gnutella application was taken over by a fast growing development and research community and shortly after its release a variety of similar Peer-to-Peer protocols followed, e.g. Audiogalaxy, FastTrack/KaZaA, iMesh and Freenet [38, 123, 232, 317, 343]. Kazaa is not an open source project and encrypts the signaling traffic between the peers and also to possible centralized elements. Thus the RIAA can not track the peer behavior and as a result can hardly prove any illegal activities on the part of the inventors and operators of the FastTrack network. Other approaches tried to decentralize the Napster server, by distributing the lookup server on several more powerful participating peers [625].

Although the legal pressure on Napster increased further, as some copyright holders, like the hard-rock band Metallica additionally sued Napster, the number of exchanged files in Napster was still growing and reached a total of 2.79 billion files exchanged only within February 2001. However in July 2001 Napster Inc. was convicted and finally had to stop the operation of the Napster server and therefore the service was no longer available.

As we can observe from Figure 3.1 and Figure 3.2, the users of Napster adapted very fast to this situation and switched to other protocols, like Gnutella or FastTrack/Kazaa. In August 2001 for example already 3.05 billion files were exchanged per month via the Gnutella network. The attractiveness of Gnutella did not only result from its distributed structure, but also from

the enhanced protocol, which consists of two hierarchical routing layers [359]. The foundation for this development of Gnutella has already been laid in October 2000 by the presentation of the Reflector/SuperPeer concept. These Peer-to-Peer networks with a second dynamic routing hierarchy are called the second generation Peer-to-Peer networks. As shown by Figure 3.2, even today second generation Peer-to-Peer protocols are widely used. Edonkey2000 and FastTrack are based on such an overlay routing concept [184, 358, 410, 423].

However in May 2003 things began to change again. Applications based on the FastTrack protocol caused significantly less traffic, whereas on the other hand the traffic amounts of e.g. Gnutella or Edonkey increased. In addition, we can observe from Figure 3.2, that the traffic caused by the BitTorrent network increased significantly and caused at the end of 2004 the majority of the traffic [127, 320].

Two main reasons explain this phenomenon. First of all in KaZaA the amount of hardly identifiable corrupted content increased significantly due to the weakness of the used hashing algorithm (UUHASH). Thus users switched

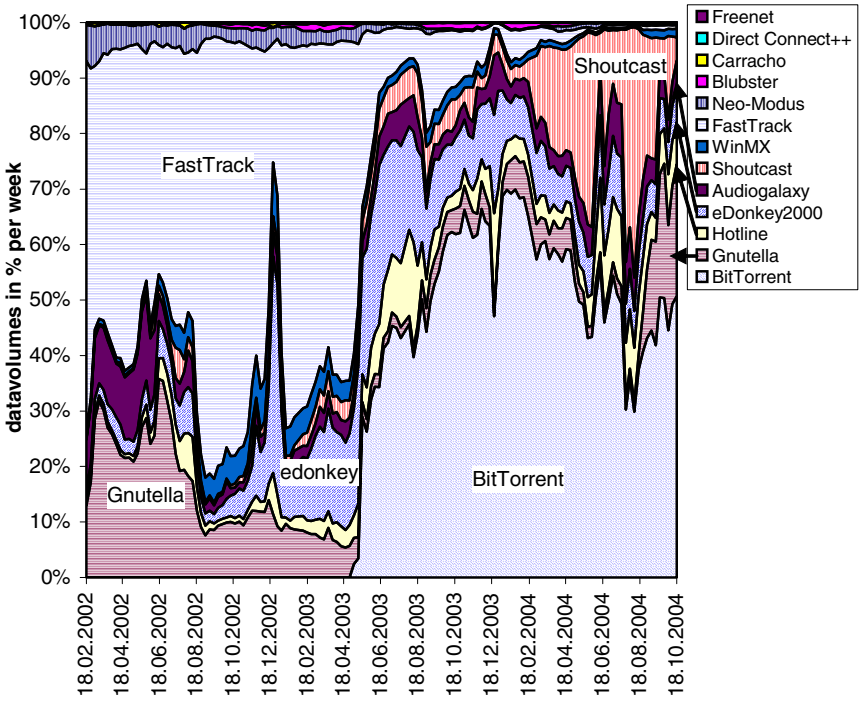

Fig. 3.2: Traffic proportions of the different Peer-to-Peer applications and protocols from the traffic measured per week in the Abilene Backbone from 18.02.2002 until 18.07.2004 (source:[42])

to applications like Gnutella or Edonkey, where the number of corrupted files was significantly smaller. Secondly, upon having a closer look at the traffic caused by the BitTorrent protocol we have to take into account that in contrast to other Peer-to-Peer protocols, in BitTorrent also the traffic caused by the file transfers is part of the measured amount of data given by Figure 3.2.

Today one of the major drivers for Peer-to-Peer is certainly the exchange of all kinds of content (mp3 encoded music, videos, DVDs, pornographic content, software, ...), free of charge and administration. However there do already exist attempts to use this new and successful networking paradigm for other applications and to develop a business case out of it. One first promising approach is to use the Peer-to-Peer overlay network as a kind of distributed phone book to provide means to establish IP based voice communication channels between participants without any centralized instances. One approach, Skype, allows calls free of charge within the Internet any and charges users only at gateways to other fixed telephony systems [567].

To increase the reliability of such a system it is necessary to adapt the overlay routing schemes to reflect the characteristics of a significantly decreased replication rate. Every user is available only once instead of several times as is the case for a common mp3 encoded music file. Further on, it is also necessary to establish a call fast and to receive a dedicated answer from the overlay network, whether a certain user is currently available or not. Therefore research on the third generation of Peer-to-Peer networks started around 2001 [526, 575]. The third generation Peer-to-Peer networks and protocols are mainly characterized by using a proactive routing algorithm based on Distributed Hash Tables (DHTs).

Yet the question why Peer-to-Peer came up in 1999 and developed so rapidly to the major application in the Internet at 2002 has still not been answered completely, though one reason is certainly the possibility to receive copyright protected content for free.

3.5 Driving Forces Behind Peer-to-Peer

In the following we want to have a look at the development of the physical and technical capabilities of the used physical networks and the participating terminals. Regarding the data rates at the access level around 1997/98 the first broadband connections for residential users were available via cable modems which allow data rates of up to 10 Mbps. Beginning in 1999 DSL and ADSL connections became available for the public with data rates of up to 8.5 Mbps. Compared to 56 kbps modem connections the available data rate for the private user was plentiful at the beginning of 2000. Further on the deregulation of the telecommunication markets showed already first effects, as the ISPs and telecommunication network providers reduced their tariffs

significantly due to increased competition. Thus not only a significantly increased access data rate became available but additionally at comparably low prices, e.g. flat rates.

In a similar manner the storage space and the processing power of common end user computers evolved. In 1992 the average size of hard disks was around 0.3 Gbyte and developed within 10 years to sizes of several 100 Gbyte. Regarding the processing power personal computers in 1992 were available with clock frequencies of 100 MHz, whereas in 2004 computers with more than 3 GHz are commonly available. Thus the computers available since the beginning of 2000 have capabilities comparable to those of servers a few years earlier. Resulting the technical prerequisites to operate a personal computer as a high performance server and a client at the same time for reasonable prices were available, when the first Peer-to-Peer networks appeared. Additionally we could observe since the end of the 1990s a general trend towards self organizing networks also in the mobile area (mobile ad hoc networks, MANET) [483, 548].

This also resulted in more and more intelligence distributed over a whole network and pushed to the place where it is demanded, i.e. at the edge of the networks. Currently it appears that this trend is still unbroken, as e.g. the application of Peer-to-Peer networks to telephony application or the traffic amounts caused by Peer-to-Peer applications show. A number of other future application areas of Peer-to-Peer networking certainly also include self organizing collaborative environments, context and location based services in conjunction with mobile networks, especially MANETs, Peer-to-Peer media streaming networks or the self organization of active network environments. Therefore certainly a number of open research issues, e.g. how to provide security, trust and authentication or accounting and access control in such a distributed environment to provide the basis for business in the area of Peer-to-Peer have to be solved. Further open problems which have to be addressed also include reliability, availability, load-balancing, QoS and network organization, as well as cross layer communication especially in mobile environments.

There is no doubt about it: Peer-to-Peer technology is yet in its infancy, and will play a key role in next generation networks!