

Enabling Next Generation Peer-to-peer Services

Fabio Victora Hecht, Burkhard Stiller

Department of Informatics IFI, University of Zurich
Binzmühlestrasse 14, CH—8050 Zürich, Switzerland
[hecht|stiller]@ifi.uzh.ch

Abstract. Peer-to-peer (P2P) applications have the potential to improve network scalability and create a community in which resources are under the shared responsibility of its members. Current popular solutions target specific applications, such as file sharing and VoIP (Voice-over-IP), but important challenges prevent P2P technology from having wider acceptance in other services. Overwhelmed by interdomain traffic generated by P2P applications that build overlays asymmetric to the network topology, ISPs offer at a home user level only connections with asymmetric bandwidth. This measure efficiently hinders P2P applications. This paper introduces a new incentive model for overlay applications and ISPs to collaborate and to create at the same time a win-win situation. The approach combines two principles: (1) ISPs and P2P applications communicate to correct overlay asymmetry and (2) ISPs diminish bandwidth asymmetry by providing users additional bandwidth reserved for intradomain traffic.

1 Introduction, Motivation, Goal

Peer-to-peer (P2P) applications, such as [2], have emerged in the last decade providing many technical advantages over the client/server model; for instance, scalability in respect to resources, fault tolerance, and load balance. Recent surveys show that P2P traffic is currently responsible for more than 50% of Internet Traffic [7]. Applications that use the Peer-to-peer paradigm provide, for instance, file sharing, Voice-over-IP (VoIP), and real-time video distribution, relying on resources from the end-user to form an overlay network without necessary presence of central servers. The success of those applications causes reactions from Internet Service Providers (ISP). Alleging traffic excess, ISPs take several measures to reduce the use of P2P.

From the user point of view, however, the use of P2P applications brings many advantages. Besides these technical advantages listed above, there are other reasons for its popularity, *e.g.*, anonymity, lack of censorship, vast content variety, performance, forming of communities, and lower (usually zero) price.

Typically, ISPs differentiate between two types of customers: (1) end-users, who receive unreliable downstream bandwidth for a low price, and (2) content or service providers, who pay more for a more reliable connection with large upstream capacity. In a P2P model, end-users also become providers, thus negatively affecting ISP revenue. Additionally, traffic costs increase, since P2P applications offer content and serv-

ices not available through other means. Nevertheless, the popularity of P2P suggest that services and content offered by it is a strong factor for users to contract the ISP in the first place.

Taking those conflicting factors into account, ISPs currently take measures to diminish peer-to-peer traffic [5], while not blocking it. Measures include closing certain client ports, using ISP-level Network Address Translation, shaping traffic, and offering only low upstream bandwidth. P2P applications can, at a certain level, work around those measures, except the latter one. For example, Swisscom [9] — market leader in Switzerland — offers for home users ADSL (asymmetric) connections through which end-users may only upload at rates 5% to 33% of their download bandwidth, efficiently hindering peer-to-peer applications. This is due to the fact that P2P applications have mechanisms to enforce that a peer is only able to download at rates it can upload [4]. Facing those limits imposed, some users are migrating to centralized alternatives [7] in order to explore all their download link capacity. ISPs usually do offer an option for symmetric access, but it is aimed at corporate clients and the relevant market price is much higher. At Swisscom, while prices per month currently start at 390 CHF for a 1.200 kbit/s symmetric connection, an asymmetric connection costs only 49 CHF, with 3.500 kbit/s download and 300 kbit/s upload capacities.

Analyzing these arguments outlined in this introductory chapter, the imbalance of incentives at the ISP level is clear. It is crucial, in order to see P2P applications being used at their full potential, that incentives are in equilibrium. The main question is on how ISPs and end users can cooperate to create a win-win situation.

2 Approach

Connections between ISPs are usually established through commercial agreements [6]. Traffic leaving ISP boundaries (interdomain) represents costs for the ISP. intradomain traffic is expected to cost less for ISPs and to be faster for users, since it does not leave the ISP, travels a smaller distance and has less chance of congestion.

One of the major factors that contributes to high cost of ISPs [8] with P2P application regards the latter's method of establishing an overlay network oblivious to the network topology. This situation can be referred to as overlay asymmetry. By building an overlay asymmetrical to the underlay topology, P2P applications waste resources by establishing a high number of interdomain connections. Optimizing overlay would, in this case, be an incentive for both end users and ISPs.

In order to manage this issue, many papers propose optimization on the peer side, not taking the ISP side into consideration [8]. Peers should establish connections preferentially with peers that are closer to them, from a network point of view, according to different metrics. Some papers suggest peers to make measurements — for example, ping, traceroute, a geographical database or IP prefix matching. The use of those techniques causes, however, a significant overhead and can be inaccurate.

[1] proposes an interesting model in which peers query ISPs to check for peers that are located within the same provider. The algorithm improves overall download speed of peers, while lowering ISP cost by reducing interdomain traffic. The work,

however, addresses overlay asymmetry without considering bandwidth asymmetry. The incentive for an ADSL end-user to prefer P2P applications instead of client/server ones is still missing.

ISPs must have an incentive to reduce bandwidth asymmetry without raising price, to allow home users to benefit from P2P. By coping with P2P applications, ISPs will enable users to benefit from new cutting-edge services, for example Internet TV and online gaming, possibly increasing number of customers.

The ideal method for users and ISPs to cooperate for locality exploitation requires actions on both parties. The new newly developed approach combines two principles:

1. Peers explore locality, with the cooperation of the ISP, favouring intradomain connections, and
2. ISPs provide end users with additional reserved upload bandwidth for intradomain traffic only; ideally, the total upload bandwidth is equivalent to the total download bandwidth.

ISPs must provide a helper service for P2P applications to communicate with, obtaining information about peer locality (whether intradomain or interdomain), which can be coupled with useful information, like delay, jitter, packet loss, hop count, and available bandwidth, since applications might have different requirements. For instance, in a file transfer, applications ambition maximization of bandwidth, while, for application layer multicast or voice over IP, low delay and jitter is required, while a constant bandwidth is enough. In this sense, a protocol must be designed for the communication, and P2P applications would have to be modified to use such service.

The P2P helper service can be implemented as an XML-based (eXtended Markup Language) Web Service in a fully decentralized and redundant manner. This technology allows for an easier development over different platforms. The protocol works as follows: the P2P application sends a request message to the helper service, containing the desired application class, and a list of IP addresses from possible peers the application may choose to connect to. Application classes are predefined as part of the designed protocol, and contain a set of characteristics that influence the decisions of the helper service; they may include, for example, different quality levels. The helper service responds with a list containing every peer in the request and a 0-100 integer number, representing a grade for each peer, according to the specified class. Peers from the same ISP will naturally receive a high grade. It is up to the application to use that information in its own benefit.

In general, peers with a higher grade should be preferred, although some randomness may be beneficial for the application to keep its P2P network well connected and more robust. It is important that the grade be an absolute value so peers can cache the value for some time, saving the effort of reevaluating the whole candidate set every time a new peer is discovered. Some metrics might be too volatile to make caching useful, but such inconsistent metrics have very limited usability and should be avoided. Exactly what metrics benefit most each P2P application is an interesting and open research question, together with the method for the ISP to gather such information from its own network.

Technical realization of additional bandwidth for intradomain traffic is trivial. Since ISPs already throttle bandwidth according to several factors, it would be necessary only to add the appropriate rules at the bandwidth broker. By reserving upload bandwidth for intradomain traffic, impact on ISP's costs are foreseen to be minimal, since intradomain traffic has much lower cost if compared to interdomain. Users of mobile phones are already familiar with this scheme, since calls to users within the same company usually cost less. It can also be used as a marketing tool, since users that have several Internet connections are likely to keep loyal to the same ISP and to convince friends to join the same. Further studies in a socioeconomic environment may become necessary at this point.

3 Challenges

Setting up such new service raises important questions and challenges on both sides. One concern regards customer and strategic data protection by the ISP. The helper service does not need to provide information on a detailed level, but instead general application classes are present in the request, and only an absolute grade is part of the reply. The right trade off must be found not to reveal information that are sensitive about customers, but that are still useful for applications.

By supporting P2P traffic, it is also feared that total traffic could increase. Although preference will be given to intradomain traffic, it is possible to reach a point that interdomain traffic will also increase. Further investigation is required to find this exact point.

What is considered a major issue regards legality of content. Anonymity, provided by some P2P applications to support censorship free information, is also used by some to offer a range of illegal content — mostly due to copyright infringement. Yet, anonymized networks are not necessarily the focus of P2P networks, as it is possible to use authentication methods as in a client/server approach if wanted.

4 Simulation Scenario

In order to validate the proposed approach, it is imperative to define a realistic scenario and perform simulations and measurements. Ideally, the solution should be deployed in an ISP and implemented on one or more popular P2P applications. Results must include several measurements. On the user side, quality of experience (QoE) is a key metric — depending on the application chosen, it may include start-up or buffering time, delay, and total bandwidth. On the helper service, the overhead created by introduction of the proposed protocol, together with its scalability and fault-tolerance can be measured. Benefits to the ISP must also be accounted, comparing volume of intradomain and interdomain traffic, preferentially from a cost perspective.

In addition to making measurements on a P2P application with and without the proposed mechanisms, it would be useful to also compare it to classical client/server (C/S) distribution paradigm. The results could be an incentive for ISPs to actually

support P2P, assuming that a customer has the choice of obtaining content from either a P2P or C/S system. Possibly, it can be shown that, with locality aware mechanisms in place, hindering P2P traffic may not be a good decision.

5 Preliminary Conclusions

Although it is believed that the approach presented in this paper represents a major improvement over the state of the art due to the combined approach, extensive and realistic simulations must be done to prove its assertions. A more detailed protocol specification must be developed, with predefined service classes. Although there are different P2P applications, they share similar concepts. Since there are marketing issues involved with the approach as well, research on their evaluation must be performed.

Acknowledgment. The authors appreciate many discussions with Dr. David Hausher on this paper's topic.

References

1. V. Aggarwal, A. Feldmann, and C. Scheideler: Can ISPS and P2P users cooperate for improved performance? ACM SIGCOMM Computer Communication Review, Vol. 37, No. 3, pp. 29-40 July 2007
2. Bit Torrent. <http://www.bittorrent.com>. Last visited: February 22, 2008
3. S. Brecht (P. Racz, B. Stiller, eds.): Investigation of Application-level Routing, University of Federal Armed Forces Munich (UniBwM), Information Systems Laboratory IIS, Master Thesis, 2003
4. B. Cohen: Incentives build robustness in BitTorrent. First Workshop on the Economics of Peer-to-Peer Systems, Berkeley, California, U.S.A., June 2003.
5. F. Constantinou, P. Mavrommatis: Identifying Known and Unknown Peer-to-Peer Traffic. Fifth IEEE International Symposium on Network Computing and Applications (NCA 2006), Cambridge, Massachusetts, U.S.A., 2006, pp. 93-102.
6. G. Huston: Interconnection, Peering, and Settlements; The Internet Protocol Journal, Vol.2, No. 1, pp. 2-16, March 1999
7. T. Karagiannis, A. Broido, N. Brownlee, K.C. Claffy, and M. Faloutsos: Is P2P dying or just hiding? IEEE Global Telecommunications Conference, 2004 (GLOBECOM'04), vol. 3, pp. 1532-1538, Dallas, Texas, U.S.A., December 2004.
8. H. Le, D. Hong, and A. Simmonds: A self-organising model for topology-aware overlay formation. 2005 IEEE International Conference on Communications (ICC 2005), Vol. 3, Seoul, Korea, May 2005, pp. 1566-1571.
9. Swisscom. <http://www.swisscom.com>. Last visited: February 22, 2008
10. M. Zink, A. Mauthe: P2P streaming using multiple description coded video. 16th Euromicro Conference, Sicily, Italy, August/September 2004, pp. 240-247.