Underlay Awareness in P2P Systems: Techniques and Challenges

Osama Abboud, Aleksandra Kovacevic, Kalman Graffi, Konstantin Pussep, Ralf Steinmetz
Multimedia Communications Lab[1], Technische Universität Darmstadt
{abboud[1], sandra, graffi, pussep, steinmetz}@kom.tu-darmstadt.de

Abstract

Peer-to-peer (P2P) applications have recently attracted a large number of Internet users. Traditional P2P systems however, suffer from inefficiency due to lack of information from the underlay, i.e. the physical network. Although there is a plethora of research on underlay awareness, this aspect of P2P systems is still not clearly structured. In this paper, we provide a taxonomic survey that outlines the different steps for achieving underlay awareness. The main contribution of this paper is presenting a clear picture of what underlay awareness is and how it can be used to build next generation P2P systems. Impacts of underlay awareness and open research issues are also discussed.

1 Introduction

Peer-to-peer (P2P) has lately become a rather surprising Internet paradox. On the one hand, P2P is the most cited reason for upgrading to broadband Internet connection [22]. On the other hand, this traffic has created extra costs for Internet Service Providers (ISPs) due to increased inter-domain traffic and new challenges in traffic management. P2P is a large-scale distributed resource sharing paradigm that leads to enormous amounts of traffic reaching up to 70% of total Internet traffic [27]. Also, due to the distribution of shared content all over the world, P2P traffic is plagued by inter-AS (Autonomous System) traffic nature. Thereby creating large but avoidable costs for ISPs [1].

Underlay awareness is defined as the collection and usage of underlay information in a way to enhance various performance aspects of P2P systems. Without underlay awareness, guaranteed Quality of Service (QoS) is difficult to achieve. This also results in routing inefficiency, larger overheads, bottlenecks, and longer waiting times even for non P2P traffic. Moreover, inter-AS traffic increases, increasing costs for ISPs.

Currently, underlay information is rarely used in most popular P2P systems. In [1], Aggarwal *et al.* show that only less than 5% of peers in the Gnutella network choose neighbors from their own AS.

From the user's point of view, underlay awareness represents an improvement in search and download times. In addition, ISPs are increasingly limiting bandwidth of some P2P applications because of higher costs inflicted by this traffic. By using underlay aware P2P systems, there will no longer be the need for ISPs to resort to such extreme measures. Thus, underlay aware P2P systems are ISP-friendly and allow for better QoS.

The literature on underlay awareness in P2P is not new, however it lacks the big-picture point of view. We propose the first survey that considers studies on underlay awareness and classifies them according to two parameters: techniques used for underlay information collection, and usage of this information for building the P2P overlay network.

Hence, the contribution of this paper is threefold: we present a survey of collection and usage of underlay information with discussion on the different approaches and designs, we propose a taxonomy on how underlay information is collected; and finally, we discuss the impacts underlay awareness has on the Internet.

The paper is structured as follows, Section 2 presents an overview of underlay awareness in P2P systems. Section 3 explains the methods used to collect underlay information. The usage of this information is discussed in Section 4. Section 5 outlines the impacts as well as benefits underlay awareness has on the Internet. Section 6 marks some common challenges for underlay awareness. Finally, we conclude this paper in Section 7 and discuss open issues.

2 Underlay Awareness

The underlay is defined as the substrate on which the overlay resides. Therefore, the underlay abstracts the physical, Medium Access Control (MAC), network and transport layers. When we talk about underlay awareness, we talk about parameters through which the underlay affects the overlay. These parameters, or information, are: *ISP-Location*, *Latency*, *Geolocation*, and *Peer Resources*. Thus, we define underlay awareness as the collection and usage of underlay information to enhance the performance of P2P

[1]This work was funded by the Federal Ministry of Education and Research of the Federal Republic of Germany (support code 01 BK 0806, G-Lab). The authors alone are responsible for the content of this paper.

systems. Collection of underlay information is performed by measuring, calculating and estimating various parameters of the underlay. The collected information is used by the P2P system to build, update and manage its overlay.

It is worth noting that different applications have different QoS requirements, and thus make use of different underlay information to enhance their performance. For example, to reduce ISP costs, more *locality of traffic* is induced through the use of ISP-location information. To enhance the performance of real time communication, latency between direct neighbors in the overlay should be minimized, and so forth.

We now discuss these ideas in more details for different underlay information.

2.1 ISP-location Awareness

ISP-location is the identification of the ISP through which a peer connects to the Internet. ISP-location awareness has attracted a lot of research mainly because of the huge costs inflicted by P2P traffic on ISPs. These costs are avoidable especially because of locality correlated users' searches, whose desired contents are located in the proximity [25, 18, 24]. ISP-location awareness in P2P systems will result in more locality of traffic. Moreover, Inter-AS costs and traffic problems due to congestion and jitter will make ISPs more interested in keeping P2P traffic in their own networks. It is also beneficial for the user, since the user's quality of experience will improve even for non-P2P traffic.

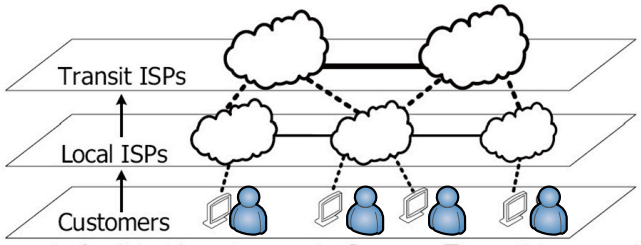


Figure 1. Hierarchy in the Internet. The solid arrows indicate monetary flow, solid lines between ISPs are peer connections and the dashed ones are transit connections [32]

In order to emphasize the impact locality of traffic has on the efficiency of the Internet, we provide here a brief overview about its structure. Roughly speaking, the Internet is built on two types of ISPs: *Local ISPs* that provide connectivity services in limited geographical areas, and *Transit ISPs* that act on a global plane and supply connectivity between local ISPs. Data flow is not free and inter-AS traffic must be paid for. Thus ISPs are ordered in a hierarchy as shown in Figure 1. Transit ISPs charge local ISPs for using the transit links based on the peak rate measured using samples over a months' time [24].

As illustrated in Figure 2, transit traffic costs per Mbps are almost fixed resulting in a proportional increase of costs with more traffic. This leads to high costs due to P2P activity. However, between local or *so-called* peering ISPs, the cost is just that of maintaining the direct link between the two ISPs and is therefore constant. This results in a cost per Mbps that is inversely proportional to the total exchanged traffic.

Therefore, benefit of locality of traffic is twofold: ISPs need to use their transit links less since local peers are preferred by the system, and ISPs are motivated to make more peering agreements with other closely located ISPs. The shift of traffic from transit to peering links due to locality of traffic means that increased P2P traffic does not inflict any additional costs on the ISP.

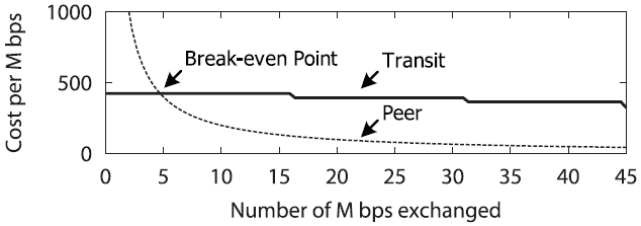


Figure 2. Costs relations [24]

2.2 Latency Awareness

Latency connotes the time it takes for a packet or stream to travel from a sender to the receiver. The importance of latency awareness lies in the fact that delay in information transfer may pauperize the user experience and degrade the QoS especially for interactive applications.

The reduction of network latency and the improvement of QoS are correlated. Nowadays, more and more communication applications need real-time packet delivery, such as VoIP, streaming, live radio or television, music jukeboxes, and on-demand video services. In traditional telephony, long delays occur only with long distance calls or calls with mobile partners. But in the case of packet delivery, the effects of excessive delay have often been overlooked, which results in quality degradation also for geographically close calls.

Awareness of P2P applications on latency is a necessity for next generation IP applications ranging from multimedia content to live communication services.

2.3 Peer Resources Awareness

Peer resources are defined as the set of parameters that estimate different capabilities of a peer. This might include available bandwidth, processing power, harddisk space,

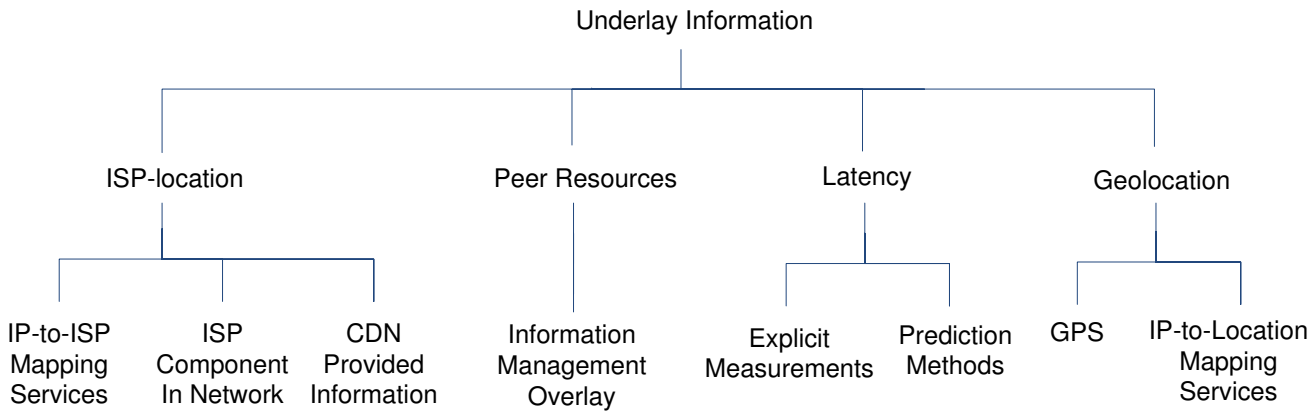


Figure 3. Classification of underlay information and their collection.

available memory, and online times. A P2P system that is aware on peer resources – can thus benefit from an increased performance – since the overlay can be arranged in such a way that different roles in the network are taken by appropriate nodes [11]. The main application area would be in hybrid P2P systems, where some nodes, known as supernodes, take more responsibility in the network. Using peer resources information allows for a more accurate super-peer selection process, and therefore a more stable system.

2.4 Geolocation Awareness

Geolocation is the identification of the geographic location of a peer. Geolocation can also be used to calculate the geographical distances between two different nodes. The geographical distance matches to some extent the transmission delay when both nodes lie within the same network. However, it is important to emphasize that geographical proximity is not necessary correlated with transmission delay, since two peers that are in the same building, could be using different ISPs and thus result in more delay due to traversing several routers and access links.

This underlay information can offer both users and ISPs interesting applications including new business models and solutions. Thus, it enables highly personalized services that could be used to physically locate nearby points of interests (restaurants, addresses, etc). Moreover, life-critical applications can benefit from such system, e.g. emergency services [10].

Geolocation awareness can be applied to construct the P2P network in such a way that peers, which are geographically close, are also closely located in the routing overlay network. One parallel advantage is that this kind of routing may be more scalable since with increasing load, congestion is less probable due to locality of traffic. Challenges faced, when using such an overlay, include routing around dead nodes and operating in low density environments.

Now we start off by presenting how the different underlay information is collected. This would be the first step in building an underlay aware system.

3 Collection of Underlay Information

We now discuss the collection of underlay information as a first step in introducing underlay awareness into P2P systems. A classification of the different methods to collect underlay information is presented in Figure 3.

3.1 ISP-location

ISP-location is defined as the ISP through which a node connects to the Internet. For discovering the ISP of a certain peer, there exists three methods:

IP-to-ISP Mapping Services. The ISP of a certain peer can be discovered simply by using its IP. Since every ISP has a set of well-known IP addresses, mapping every peer to an ISP is straightforward. There exist many IP-to-ISP mapping services, such as [13, 14, 15].

ISP Component In Network. Some solutions [1, 29] aiming at having locality of traffic introduce a new component into the network. This component is situated at the ISP side and helps in making a decision on neighborhood selection of peers. It thus provides information on proximity of peers in the ISP metric space.

CDN Provided Information. Content Distribution Networks (CDNs) are hired to deliver content to end users in order to improve web performance especially during flash crowds. This is achieved by providing content from geographically distinct servers situated at the edge of the Internet near end users. The actual CDN servers which are used for a certain time frame are those which have the least load and shortest paths to the requesting peer. This fact is exploited to infer locality information [5].

3.2 Latency

The latency between any two peers in the system can be either explicitly measured or predicted. We now dis-

cuss these two methods pointing out different advantages and drawbacks.

Explicit Measurements. Latency can be measured explicitly using a simple ping or traceroute technique. This, however, incurs the network with much overhead. These measurement can further congest the network when a group of peers try to measure the latency at once. Typically, these kinds of measurements are used only sparingly, relying mainly on prediction techniques to find latencies between any two peers in the system [35].

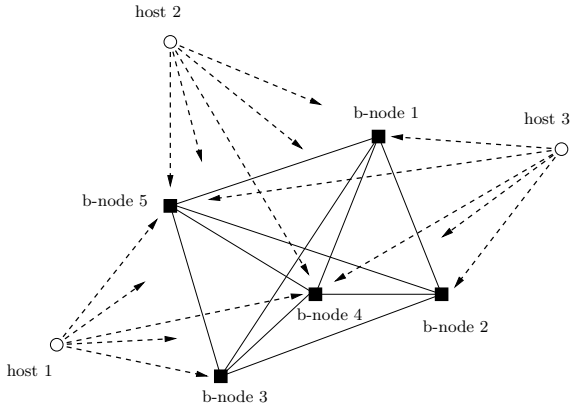


Figure 4. Proposed architecture for the Internet coordinate system (five beacon nodes and three ordinary hosts) [20].

Prediction Methods. Using prediction methods, it is only required that each node in the system measures latencies to just a small set of other nodes (typically its neighbors or landmarks) [16]. Many methods exist that give a solution to this problem, the most prominent one is Vivaldi [7]. Landmark based prediction methods estimate the latency between two arbitrary peers without direct measurement using a set of well-known landmarks. In [20], Lim *et al.* propose such an architecture that uses GPS-like triangulation techniques, where the so-called beacon nodes perform the functionality of a satellite transmitter as shown in Figure 4.

3.3 Geolocation

There are numerous techniques to harvest geolocation information. These techniques are divided generally into two classes:

The first class is based on inferring the geolocation from a satellite positioning system such as GPS, Galileo, or GLONASS [12, 18]. Typically, the UTM (Universal Transverse Mercator) coordinate system is used to represent this geolocation [12]. This method is usually used for effectively locating points of interest.

Another class of methods to infer the geolocation of peers is based on IP-to-Location mapping. For this, there exist a plethora of commercial [13, 14] and non-commercial [15] mapping services. This method is less accurate and thus gives only a rough geographical area in which a peer is (most probably) located.

Here again, ISPs can also provide geolocation information about users connecting through it to the Internet, since each ISP knows the addresses and exact locations of all of its customers.

3.4 Peer Resources

The most interesting solution for collecting peer resources is based on an information management overlay [11]. This overlay is used to generate statistics on the P2P system, which enables resource-based peer search and thus different parameters can be retrieved to build a resources-aware P2P system.

Next, we present how the different underlay information can be used to enhance the performance of the P2P system.

4 Usage of Underlay Information

ISP-location information is usually used to improve the neighborhood selection of the peers in such a way so as to reduce inter-AS traffic. To explain how ISP-location is used to build the system, we consider as an example the solution proposed in [1]. This solution introduces an ISP component into the network called the oracle. The oracle is queried for locality information about the peers. Mainly, it just considers ISP-location-based ordering of peers to avoid inter-AS traffic. As presented in Figure 5, the use of the oracle correlates both the overlay and the underlay, constraining peers into their own networks. The process of choosing neighbors from the same ISP is known as *biased neighbor selection*. Figure 6 shows the impact of biased neighbor selection on the topology of the P2P system. Clustering of the network based on ISP boundaries is noticed. Here there is a mini-

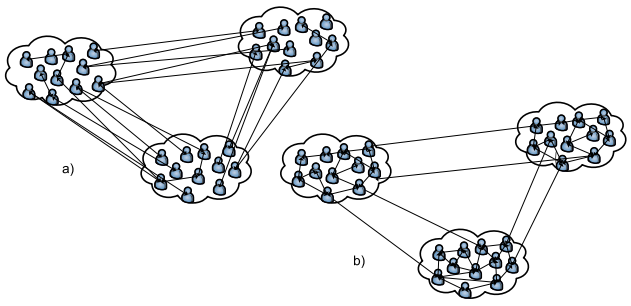
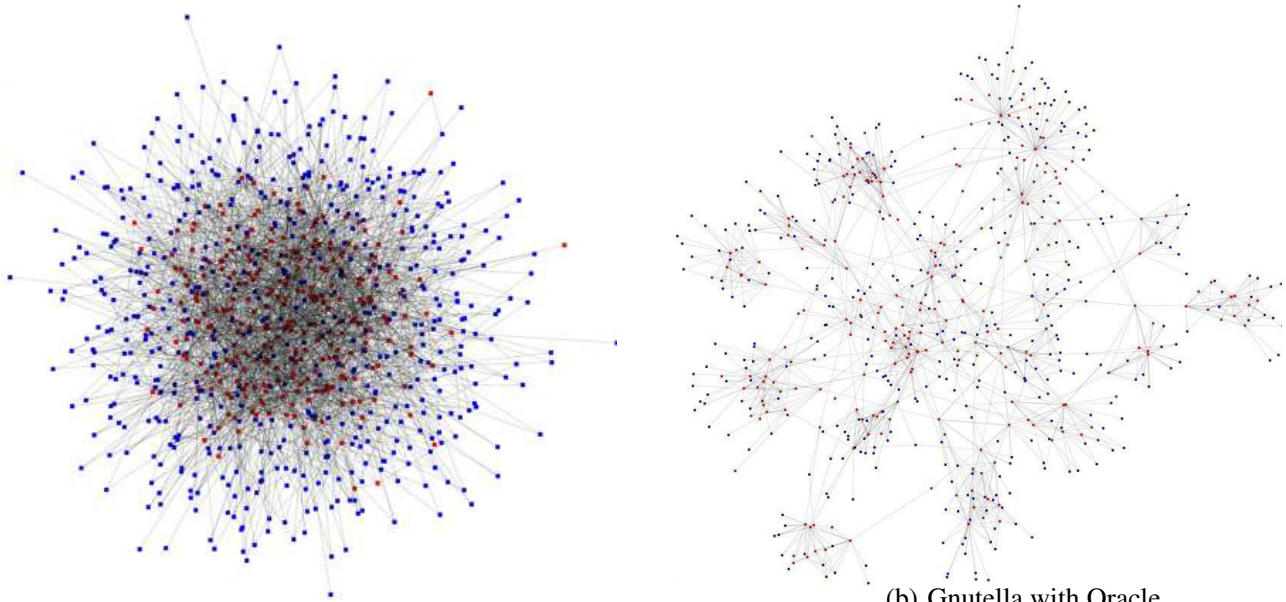


Figure 6. (a) Uniform random neighbor selection and (b) biased neighbor selection.



(a) Unbiased Gnutella

(b) Gnutella with Oracle

Figure 5. Comparison of original Gnutella overlay topology with the one after applying the Oracle [1]

mal number of inter-AS connections necessary to keep the network connected.

Introducing locality to P2P traffic is not restricted to content distribution. In [17], Kaune *et al.* extend the routing algorithm of Kademlia to reduce inter-AS traffic due to the distributed hash table-lookup algorithm.

Latency information is used to build a latency aware overlay. In [33], Yinzhe *et al.* apply this concept on a structured overlay. Both content identifiers and latency information are processed together using a special hashing function called *Geographically Scoped Hashing* to produce the final peer and content identifiers.

Geolocation information is used to build an overlay where neighboring peers are geographically close. In [2], Araujo *et al.* focus on location-constrained queries and information dissemination based on geographical information. In [19], Kovacevic *et al.* present a hierarchical tree-based P2P system that enables geolocation-based overlay operations.

Peer Resources information is used to construct P2P systems with shorter search and download times, higher responsivity and stability. This is achieved by arranging the overlay in such a way that different roles in the network are taken by appropriate nodes [11].

The most prominent underlay aware solutions, classified by underlay information, are listed in Table 1.

ISP-location	BNS [3], TSO [31] , Ono [5], LTM [21], CAT [32], Brocade [36], Plethora [9], Mithos [28], MBC [35]
Latency	gMeasure [34], Genius [23], eCAN [30], Leopard [33], Hop-based Proximity [8], Proximity in DHTs [4], Landmark-based proximity [26]
Geolocation	Globase.KOM [18], GeoPeer [2]
Peer Resources	SkyEye.KOM [11], Bandwidth Aware [6]

Table 1. Various underlay aware systems.

5 Benefits and Impacts on the Internet

We now present an analytical study of the influence that underlay aware P2P systems would have once fully implemented. Classified by scopes, impacts are presented trying to add a business point of view to the merely technical one. In addition, legal issues and challenges are presented.

5.1 Impact on Users

The most important impact on P2P users is better QoS. This includes low delays, successful searches, short download times, etc (depending on which underlay information is used).

Impact on	Parameter	ISP-location	Latency	Geolocation	Peer Resources
Users	Download time	++	o	o	++
	Delay	o	++	+	o
ISPs	ISP OAM	++	o	o	o
	ISP Costs	++	o	o	+
Both	New Application Areas	o	+	++	o
	Resilience	++	++	o	+

Table 2. Impact of underlay awareness on Internet users and ISPs. Table legend; ++: big effect, +: small effect, o: neutral.

However, in solutions based on ISP-provided information, this will only happen as long as this does not contradict with the ISP's own interests. Moreover, an important issue here is that users must be able to trust ISPs with private information, since in some solutions information about location (ISP-location and geolocation), resources, and activity are used by the ISP to enhance the performance of the overlay. It is worth noting that most users will not use a service that involves the ISP knowing about certain P2P activity.

5.2 Impact on ISPs

ISPs are mainly interested in having underlay aware overlays in order to reduce inter-AS costs. Thus they support mostly ISP-location based underlay awareness. Since locality of traffic would bind P2P flow to the local network, ISPs can save costs and edge resources, avoid bottlenecks, and enhance internal network management.

Moreover, a wider business market as a result of underlay awareness is expected. The different types of underlay awareness information allow for individualized and focused solutions depending on applications or even users with specific requirements. For example, using geolocation information from the ISP, new location-based services can be realized.

Different underlay information can be used for different applications to optimize application specific QoS requirements. One interesting possibility is that ISP can sell specific resources for applications in order to enhance their performance, and thus, ISPs can gain money, while providing a better QoS for applications.

5.3 Impact on P2P System Providers

The main impact which is foreseen for the people behind P2P system development is some joint ventures with ISPs which might result in better cooperation and less fighting. However, this is still a very challenging task mainly because of legal issues.

5.4 Impact on the Quality of P2P Systems

As discussed throughout this paper, benefits from underlay awareness are numerous and promising. However more research should go into the direction of impacts on different performance aspects of P2P systems. This includes search completeness, overlay connectivity, and robustness especially against churn. This and a general study about the introduced overhead due to underlay awareness remain open issues.

A summary of the impacts of underlay awareness on the Internet in general and on different QoS parameters in particular is presented in Table 2.

6 Open Issues and Challenges

Any solution aimed at being underlay aware has to be able to handle some challenges in order to properly introduce underlay awareness into the Internet.

Asymmetric Node Selection and Long Hop

Asymmetric node selection is one of the most common problems found in latency-aware systems since such solutions rarely consider it. This problem occurs when an incoming peer searches for the closest peers from the overlay. It can happen that the path from node A to node B is the shortest for node A, but at the same time the path from node B to node A is not the shortest for B. Therefore, the asymmetry of peer selection results in less precise underlay measurements.

The *long hop problem* on the other hand occurs in hop-based latency aware systems which do not take message delays into consideration. Thus, one single hop may represent a big distance in terms of delay and hence, the medium through which the physical signal travels introduces a time penalty, even with active locality awareness.

ISP Internal Information

Some proposed solutions [1, 29] are based on the cooperation of ISPs furnishing internal information such as routing

policies, costs or traffic flow management. In many cases the recipient would be an external node, not controlled by the ISP, which could be malicious. This is an important issue to consider, since an ISP will not provide such information unless there exist clear incentives or benefits to do so.

Mobile Support

Mobile users are envisioned to increase rapidly in the next years. For this, underlay mechanisms that support low footprint mobile devices should be incorporated. In turn, some underlay provided information such as ISP-location and latency no longer apply because of continuous variation, or at least this might introduce additional overhead to any *mobility-aware* P2P system in case of low-scale mobility.

Legal Issues

This is the most controversial aspect of P2P applications since the most popular use of this technology is file sharing. This situation may be hindering the incorporation of the P2P technology to major IT companies.

In [3] Bindal *et al.* suggest the use of an ISP component which would manage P2P traffic by scanning the packets and taking the role of the tracker. This could situate the provider in a delicate situation due to privacy issues.

7 Conclusion

In this paper we analyzed the state-of-the-art in underlay awareness for P2P systems. We provided a taxonomic classification of the literature on underlay awareness in P2P systems according to two main categories, collection and usage of underlay information, and the impacts and benefits it triggers. The studied underlay provided information were: ISP-location, latency, geolocation, and peer resources.

ISP-location information helps in realizing solutions that keep P2P traffic inside the ISPs' networks in order to decrease inter-AS costs. An important challenge here is not to degrade the performance of the overlay because of ISP-based clustering of the network. Latency-aware systems aim at reducing the communication delays between peers by choosing shortest paths. This is of particular importance for services based on multimedia distribution or live communication. Geolocation information can be used for next generation point of interest localization services, in which the P2P-system always has up-to-date information about the location of its peers. Peer resources information is used to construct P2P systems with better efficiency since different roles in the network are taken by appropriate nodes.

Underlay awareness has many impacts on the Internet. We pointed out the possible improvement of QoS for users, as well as better data routing management and reduction of costs from inter-AS traffic for ISPs. Overlay providers may find new business models with the possibility of joint ventures with ISPs and offer for user added values. Underlay awareness is thus an important property that must be considered in next generation P2P systems.

Some solutions were based on collecting and using underlay information at the ISP. Here, we have an extra ISP-owned component that affects and manages the overlay in many ways. ISP-provided information can thus alter the network based not only on ISP-location, but also on other policies and mechanisms of the ISP. An open issue here remains in trust and in sharing internal information.

Another open research issue is the development of a general architecture for underlay awareness in which different underlay information can be collected and used. Thus an underlay awareness framework is the definitive next step in implementing underlay awareness in the Internet.

New technologies in multimedia distribution [18] such as on-demand content distribution, Internet TV and VoIP services require the switch to P2P to have lower costs. But without underlay awareness, this might never be an easy switch. Nonetheless, through the synergy of efforts of ISPs, content providers, and P2P system developers, a more reliable Internet and cheaper services might not be so far off.

References

[1] V. Aggarwal, A. Feldmann, and C. Scheideler. Can ISPs and P2P users cooperate for improved performance? In *ACM SIGCOMM Computer Communication Review, Volume 37, Number 3*, pages 23–40, July 2007.

[2] F. Araujo and L. Rodrigues. Geopeer: A location-aware peer-to-peer system. *Network Computing and Applications*, pages 39–46, 2004.

[3] R. Bindal, P. Cao, W. Chan, J. Medved, G. Suwala, T. Bates, and A. Zhang. Improving traffic locality in BitTorrent via biased neighbor selection. In *Proc. of the ICDCS 2006 Conference*, 2006.

[4] M. Castro, Druschel, P., Hu, Y., and A. Rowstron. Exploiting network proximity in peer-to-peer overlay networks. Technical report msr-tr-2002-82, 2002.

[5] D. Choffnes and F. Bustamante. Taming the torrent: a practical approach to reducing cross-ISP traffic in Peer-to-Peer systems. *SIGCOMM Comput. Commun. Rev.*, 38(4):363–374, 2008.

[6] A. P. da Silva, E. Leonardi, M. Mellia, and M. Meo. A bandwidth-aware scheduling strategy for p2p-tv systems. *Peer-to-Peer Computing, IEEE International Conference on*, 0:279–288, 2008.

[7] F. Dabek, R. Cox, F. Kaashoek, and R. Morris. Vivaldi: a decentralized network coordinate system. *SIGCOMM Comput. Commun. Rev.*, 34(4):15–26, 2004.

[8] P. Danielis, S. Kubisch, H. Widiger, J. Schulz, and D. Timmermann. A conceptual framework for increasing physical proximity in unstructured peer-to-peer networks. In *IEEE Sarnoff Symposium*, pages 226–227, 2008.

[9] R. A. Ferreira, S. Jagannathan, and A. Grama. Locality in structured Peer-to-Peer networks. *J. Parallel Distrib. Comput.*, 66(2):257–273, 2006.

[10] K. Graffi, A. Kovacevic, K. Wulffert, and R. Steinmetz. Echop2p: Emergency call handling over peer-to-peer overlays. In *P2P-NVE*, pages 58–67, Dec 2007.

[11] K. Graffi, A. Kovacevic, S. Xiao, and R. Steinmetz. Skyeye.kom: An information management over-overlay for getting the oracle view on structured p2p systems. 2008.

[12] M. Grewal, L. Weill, and A. Andrews. *Global Positioning Systems, Inertial Navigation, and Integration*. Wiley-Interscience, 2007.

[13] IP2Country. http://www.ip2country.net/, 2008.

[14] IP2Location. http://www.ip2location.com, 2008.

[15] IPGEO. Online: http://www.ipgeo.com/, 2008.

[16] M. Kaafar, L. Mathy, C. Barakat, K. Salamatian, T. Turletti, and W. Dabbous. Securing internet coordinate embedding systems. *SIGCOMM Comput. Commun. Rev.*, 2007.

[17] S. Kaune, T. Lauinger, A. Kovacevic, and K. Pussep. Embracing the peer next door: Proximity in kademlia. *In Proc. of IEEEP2P 2008*, pages 343–350, 2008.

[18] A. Kovacevic, O. Heckmann, N. Liebau, and R. Steinmetz. Location awareness - improving distributed multimedia communication. *Special Issue of the Proceedings of IEEE on Advances in Distributed Multimedia Communications*, 2008.

[19] A. Kovacevic, N. Liebau, and R. Steinmetz. Globase.KOM - a p2p overlay for fully retrievable location-based search. In *IEEE P2P 07*, 2007.

[20] H. Lim, J. C. Hou, and C.-H. Choi. Constructing internet coordinate system based on delay measurement. *IEEE/ACM Trans. Netw.*, 13(3):513–525, 2005.

[21] Y. Liu, X. Liu, L. Xiao, L. Ni, and X. Zhang. Location-aware topology matching in P2P systems. In *IEEE INFOCOM*, 2004.

[22] T. Mennecke. IPGEO, DSL Broadband Providers Perform Balancing Act, 2005. Online: http://www.slyck.com/news.php?story=973.

[23] N. Ning, D. Wang, Y. Ma, J. Hu, J. Sun, C. Gao, and W. Zheng. Genius: Peer-to-peer location-aware gossip using network coordinates. In *International Conference on Computational Science (3)*, volume 3516 of *LNCS*, pages 163–170. Springer, 2005.

[24] W. B. Norton. A business case for ISP peering. Online. http://pages.cs.wisc.edu/˜akella/cs640/f06/readings/norton.pdf. 2002.

[25] A. Rasti, R. Rejaie, and W. Willinger. Understanding the effect of P2P overlay on the AS-level underlay. In *CoNEXT '07: Proceedings of the 2007 ACM CoNEXT conference*, pages 1–2. ACM, 2007.

[26] S. Ratnasamy, M. Handley, R. Karp, and S. Shenker. Topologically aware overlay construction and server selection. *IEEE INFOCOM'02*, June 2002.

[27] H. Schulze and K. Mochalski. Data about P2P, VoIP, Skype, file hosters like RapidShare and streaming services like YouTube. *Ipoque, Internet Study 2007*, November 2007.

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

[29] H. Xie, A. Krishnamurthy, A. Silberschatz, and Y. R. Yang. P4P: Explicit communications for cooperative control between P2P and network providers. *P4PWG*, 2007.

[30] Z. Xu, C. Tang, and Z. Zhang. Building topology-aware overlays using global soft-state. In *ICDCS '03: Proceedings of the 23rd International Conference on Distributed Computing Systems*, page 500. IEEE Computer Society, 2003.

[31] G. Xue, Y. Jiang, Y. You, and M. Li. A topology-aware hierarchical structured overlay network based on locality sensitive hashing scheme. In *In Proc. of UPGRADE '07*, pages 3–8, New York, NY, USA, 2007. ACM.

[32] S. Yamazaki, H. Tode, and K. Murakami. CAT: A cost-aware BitTorrent. In *LCN '07: Proceedings of the 32nd IEEE Conference on Local Computer Networks*, pages 226–227. IEEE Computer Society, 2007.

[33] Y. Yu, S. Lee, and Z.-L. Zhang. Leopard: A locality aware peer-to-peer system with no hot spot. In *NETWORKING*, pages 27–39, 2005.

[34] X. Zhang, J. Liu, Q. Zhang, and W. Zhu. gMeasure: A group-based network performance measurement service for peer-to-peer applications. In *In Proc. of GLOBECOM*, pages 2528–2532, 2002.

[35] X. Zhang, Q. Zhang, and W. Zhu. Measurement-based construction of locality-aware overlay networks. In *IEEE ICC*, 2004.

[36] B. Zhao, Y. Duan, L. Huang, A. Joseph, and J. Kubiatowicz. Brocade: Landmark routing on overlay networks. In *Proceedings of 1st International Workshop on Peer-to-Peer Systems (IPTPS)*, pages 34–44, 2002.