

Analysis of Peer Selection Algorithms in Cross-Layer P2P Architectures

S.M. Saif Shams, Paal E. Engelstad, *Senior Member, IEEE*, and Amund Kvalbein, *Member, IEEE*

Abstract—The large amount of peer-to-peer (P2P) traffic in today's Internet represents a great challenge for most Internet Service providers (ISPs). P2P traffic does not conform to traditional ISP traffic policies, and it makes it harder to perform traffic engineering in the network. It is especially the peer-selection mechanism of the P2P application that influences the traffic patterns and creates an imbalance for the ISP business. Some ISPs benefit from the P2P traffic, while others suffer. There are research efforts proposing cross-layer solutions where the P2P application uses the underlying routing information for the peer selection. These works put a lot of focus on improving the application performance, but take to a little extent into account how the ISPs are affected. In this paper we propose a framework to measure the effects of the peer-selection mechanism, not only on the P2P application performance, but also on the economy of different types of ISPs. We categorize all ISPs into three groups; core, transit and stub, depending on their position in the global Internet hierarchy. Using our proposed framework, we analyze the performance of different peer-selection algorithms.

Index Terms— cross-layer communication, ISP, P2P network.

I. INTRODUCTION

A P2P system is a self-organizing and scalable network of independent entities. It enables easy access and free download of large amount of resources to its users without having a powerful server [1]. Instead, the communication goes directly between the peers, and the peer selection algorithm determines which peers will communicate with each other.

The easy access to resources makes P2P applications extremely popular as a file-sharing technology among a large number of Internet users. The aggregated traffic of popular P2P file sharing applications, such as BitTorrent, Kazaa and eDonkey, dominated the Internet's backbone traffic in 2004 [1-3]. Its ever-increasing popularity opened new opportunities to the global Internet. P2P applications have been mentioned as one of the major reasons for upgrading the end-user Internet access to broadband [4].

However, the popularity of P2P applications also represents a great challenge to the ISPs. Unlike other types of Internet traffic, P2P traffic increase the upstream traffic

for the lower tier ISPs [5], and impose an increasing cost for them [6]. P2P traffic also poses a significant traffic engineering challenge to ISPs, violating traditional Internet traffic rules [3, 7]. This is due to the fact that current peer selection algorithms (neighbor selection algorithms) are network agnostic. These peer selection algorithms rely on application layer routing and do not take the network layer routing and the underlying topology into account [8]. Thus, the P2P application routes data traffic independently at the application layer, i.e. without consulting the network layer. Most of the P2P clients choose their partner peers either randomly or based on some active probing from the top layer. It causes excessive data flow over the network and abrupt changes in traffic pattern resulting in difficulties for ISPs to control traffic flow.

Realizing this negative impact of such an excellent opportunity, some cross-layer peer selection algorithms [9-13] have been proposed that consider network layer routing, and try to ease the austereness of P2P applications over ISPs.

In this paper, we propose a framework to measure the effect of cross-layer peer selection algorithm on the economy of ISPs. We envision that the economical effect of P2P traffic is not similar to every ISP, but rather that it varies for different groups of ISPs. We categorize all ISPs into three tiers reflecting the hierarchical structure of the Internet; core, transit and stub. As shown in the Figure 1, Core ISPs reside at the top of the hierarchy and stub ISPs at the bottom. Transit ISPs are at the middle of these two tiers, that have ISPs above and below of them. This structure has been discussed in the Section III (A). We show that the interest of ISPs in terms of the peer selection algorithm differs based on the location of ISP in the Internet hierarchy. For simplicity, in this paper we consider that there is a one-to-one mapping between an ISP and an Autonomous System (AS), and we use the terms *ISP* and *AS* interchangeably.

Rest of the paper is organized as follows. Section II provides an overview of peer selection algorithms. A framework for modeling of the effects of such algorithms on the application performance and on the economy of the ISPs, follow in the Section III. Our analysis is presented in Section IV, Section V concludes the paper.

II. OVERVIEW OF PEER SELECTION ALGORITHMS

We will refer to a peer that want so download a block of information as a *leecher*, and a selected peer providing this information as a *seeder*. Thus, a peer selection algorithm (or neighbor selection algorithm) is the mechanism that allows a leecher to select a seeder among the peers in the P2P network.

S. M. Saif Shams, PhD student, Simula Research Laboratory. (phone: +47 67 82 82 00; fax: +47 67 82 82 01; e-mail: saif@simula.no).

Paal E. Engelstad is with SimTel (Simula/Telenor) and holds a professor position at UniK, University of Oslo (e-mail: paal.engelstad@telenor.com).

Amund Kvalbein holds a post-doc position at Simula Research Laboratory, Norway (e-mail: amundk@simula.no).

Four different peer selection algorithms are presented in the following. They are either selection strategies that are in use today, or they are educational corner-cases that illustrate the economic impact of a particular performance objective. These peer selection strategies form the basis of our analysis later in the paper.

A. Random peer selection (BitTorrent)

According to the BitTorrent file distribution system, a peer that wants to download a file (i.e. a leecher) gets a random list of peers (i.e. potential seeders) from an application tracker. It selects its seeder randomly from that list [14]. Sometimes a downloading leecher does not find expected content from its selected seeders and then it seeks another set of seeders with help from the tracker. The basic responsibility of a tracker is to maintain the list of participating seeders and provide support for them. In other words, in this selection algorithm a seeder has been chosen randomly without consulting the network layer.

The main advantage of the random algorithm is that the leecher needs no information about the underlying network. This allows the P2P client to be very simple. However, as we will observe later in this paper, it results in decreased application performance and poor utilization of the network resources, because the number of ISP hops between the leecher and the seeder can be arbitrary.

B. Locality based selection algorithm

A few research works [10, 11, 15, 16, 17] have argued that locality based selection algorithms may improve the performance of P2P applications. However, the definition of locality varies, and a number of different matrices, such as geographical distance, number of hops, AS distance, and bandwidth, have been proposed. Experimenting with several such locality based peer selection algorithms, X. Weng et al. [16] show that AS-based clustering improves the download time of P2P file sharing more than other definitions of locality.

This paper assumes that locality is based on the number of ISP-hops (or AS-hop). In the locality based selection algorithm, a leecher chooses the seeder that is closest in terms of the number of ISP-hops. If several seeders have the same smallest ISP-hop distance to the leecher, then one of these seeders is selected at random.

Unlike the random algorithm, the localized method belongs to a class of algorithms that considers network information for the peer selection [19, 20]. Since a host usually does not get any routing information, the P2P client normally has to probe the network to calculate this information by itself. Although the localized method keeps the number of ISP hops to a minimum and improves the application performance, the active probing leads to a costly network overhead.

C. P4P

There are several research works [9, 10, 13] aiming for cross-layer communication so that P2P applications can benefit from routing information and at the same time allow ISPs to control the P2P traffic. This is a class of selection algorithms that also considers network information but needs no probing. Instead, a cross-layer architecture is used to communicate network information to the P2P application.

P4P [9] is an architecture for cross-layer cooperation between the application layer and the network layer. This cross-layer solution allows for communication between the ISP and different P2P applications. In the P4P architecture, there is a central entity called the *appTracker* and entities in each ISP, referred to as *iTrackers*. The *appTracker* collects network information from the *iTracker*, in the terms of the distance between every pair of peers. Based on this information, the *appTracker* does the peer selection on behalf of the P2P application. The *appTracker* selects a set of seeders for a leecher considering its point of presence (PoP), the distance from other peers, the background traffic on each link and the resilience to peer failure. In order to choose a set of seeders for a particular leecher, it follows some rules (in the given order):

1. The *appTracker* selects up to an *Upper-Bound-IntraPID* fraction of m seeders that reside at the same PoP as where the leecher also resides. The default value of the *Upper-Bound-IntraPID* is 70%. Here m represents the number of seeders required for a leecher, and the *Upper-Bound-IntraPID* is a constant given by the P2P application.
2. The *appTracker* selects up to an *Upper-Bound-InterPID* fraction of the m seeders from the same ISP as where the leecher also resides. The default value of the *Upper-Bound-InterPID* is 80% including the seeders selected in rule 1 (i.e. Step 1).
3. The *appTracker* selects some other seeders from outside the ISP to reach a total number of m seeders. To select those seeders, first the *iTracker* evaluates each seeder, considering the distance between the leecher and the seeder, the link quality, the link priority, the business agreement, and so forth. The *appTracker* gets the evaluation of each seeder from the *iTracker*. The evaluation of the seeders is aggregated into a *link priority* value per seeder.

Despite the advantages of cross-layer communication in terms of improved application performance and network utilization, there are also hurdles to get it implemented. First, the solution requires a support architectures implemented by the ISPs. Second, the solution needs cooperation from the ISPs, and this involves a business perspective. Thus, a cross-layer solution is not feasible unless it is considered favorable by the ISPs. At the same time, a cross-layer solution must also benefit the P2P application. Otherwise, the application will not have incentives to take advantage of network information, e.g. provided by the ISPs.

D. Customer-oriented selection algorithm

The customer-oriented (CO) selection algorithm is a cross-layer peer selection algorithm proposed in this paper to give more profound understanding of our analyses. It is inspired by the fact that the peer selection algorithm makes an impact on the economy of an ISP. In this peer selection algorithm, ISPs select seeders for their own leechers based on their knowledge of the network topology, with the objective of maximizing their own revenue. Hence, an ISP will first select seeders in one of its customer networks, since this creates traffic that will generate revenue. This kind of intentional traffic redirection is unlikely to occur in a real

III. MODEL

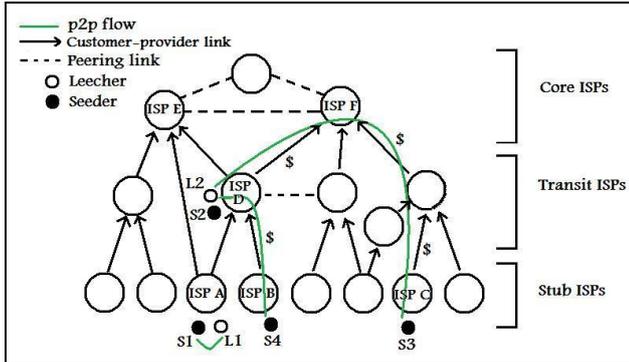


Figure 1: ISP-Level Internet architecture and P2P peer selection.

- L1 chooses S1; ISP A does not pay for this traffic.
- L2 chooses s4 in ISP B; ISP-D gets revenue.
- L2 chooses s3 through ISP F; ISP D pays revenue to ISP F.

networking scenario, since customer ISPs would react negatively to such a practice. While not very realistic, this selection algorithm illustrates an extreme point in the design space for ISP-assisted peer selection algorithms.

There are three possible roles an ISP can have in its relation to another ISP, namely *customer ISP*, *provider ISP* and *peering ISP*. An ISP has exactly one of these roles for each link to another ISP. The term *customer ISP* (or simply *customer*) can be understood by looking at our reference architecture in Figure 1. ISP A, is a stub ISP which is a customer of the transit ISP D and the core ISP E. These two latter ISPs are therefore referred to as *provider ISPs* to ISP A. Likewise, the transit ISP D is also a *customer* of the core ISP E and F. Finally, core ISPs E and F in the figure are peering ISPs to each other. There is a *peering link* between them. A peering link has a relaxed contract agreement for the traffic exchanged on this link, normally implying no payment for the traffic in either direction.

Logically, if a leecher inside an ISP (in the Figure 1, L2 in ISP D) chooses a seeder located in one of the customer ISPs (seeder S4 in ISP B), the P2P traffic may add revenue to the provider ISP (ISP D), since the provider ISP (ISP D) charges the customer ISP (ISP B) based on the actual traffic volumes. As a second best option, if the seeder is located inside the same ISP or within a peer ISP, no revenue is generated (as shown in Figure 1, L1 chooses seeder S1 in ISP A).

Our proposed CO selection algorithm is described by the following rules (in the given order) carried out by the ISPs on behalf of one of its P2P clients (leechers) that is looking for a seeder:

1. The ISP tries to select a seeder in a way so that the data flows through its customer ISP that pays revenue for it. For example, in the Figure 1, a leecher in ISP F choose a seeder in ISP C.
2. If it fails to find such a seeder, it tries to choose a seeder that is one of its own P2P clients (i.e. a seeder located by the same ISP as the leecher)
3. Otherwise, it tries to find a seeder located in one of its peer ISP.
4. Finally, it tries to find a seeder from one of its provider ISP, otherwise it waits until it gets another peer list from the application Tracker.

To analyze how the peer selection algorithm affects the economy of an ISP, we first need to develop realistic models for the underlying Internet topology, the BGP routing decisions, the distribution of the peers in the network and the traffic pricing. These models are presented in the following.

A. The topology of the ISPs in the Internet

We are interested only in the ISP-level Internet topology since the economy of the ISP (in terms of traffic-generated income and variable costs) is affected by the P2P traffic when it crosses the border of its network. We consider the Internet as a collection of ISPs forming a hierarchical structure among them. We adopt the topology model presented in [18], which captures several key properties of the Internet topology, including the hierarchical structure. This model also incorporates the business relationship between ISPs, where ISPs can have either a customer-provider or a settlement-free peering relationship. Customers pay their providers when traffic flows between the two ISPs, while traffic between ISPs that use settlement-free peering is free of charge for both parties.

We categorize all ISPs into three tiers; core, transit and stub. At the top of the hierarchy are the core ISPs, which are connected in a clique. Core ISPs are connected using settlement-free peering. Transit ISPs reside below core ISPs, and have one or more provider ISPs which can be either a core ISP or another transit ISP. At the bottom of the hierarchy, we have stub ISPs. These ISPs do not have any customers, and they buy transit services from one or more transit or core ISPs.

B. BGP routing decisions

The flow of traffic between ISPs is determined by the routing decisions made by the Border Gateway Protocol (BGP) - the interdomain routing protocol used in the Internet. BGP is a policy-based routing protocol, where routes comply with policies based on the business relationship between ISPs. We adopt standard routing policies in our model, where routes learned from customers are advertised to all neighboring ISPs, while routes learned from providers or peering ISPs are only advertised to customers. Routes learned from customers are always preferred over routes learned from peering ISPs, which are in turn preferred over routes learned from providers. This gives what is referred to as valley-free routing, where interdomain traffic flows up, then sideways, then down in the ISP hierarchy.

C. Peer distribution

To evaluate a P2P application, it is important to model how peers are distributed throughout the Internet hierarchy. This translates into determining how many peers should be assigned to each ISP in the Internet topology. In reality, the number of peers will depend on the type of ISP, e.g. whether it has primarily home or office customers, and the size of the ISP. For simplicity, we adopt a model where the peers are uniformly distributed across all ISPs.

D. Pricing model

Customer ISPs typically pay revenue to their providers based on the amount of traffic transmitted each month. The

total amount of revenue collected or paid by an ISP on a provider-client link is normally determined by the 95th percentile of the aggregated peak traffic on that link [1]. We make the fair assumption that the P2P traffic is uniformly distributed in time, so that on average the P2P traffic increases the 95th percentile of the traffic with an amount equal to the size of the P2P traffic. The result is that the P2P traffic considered in our analysis is charged on a per-traffic-volume basis.

Traffic pricing is predominantly determined by the traffic volume on the link. The price per Mbps is considerably higher on a 100Mbps link than on a 10Gbps link. In this paper, we consider a simple pricing model where link pricing depends on the two connected ISPs and the relation between these two ISPs. Since traffic volumes from a stub ISP to its provider ISP is typically smaller than in the core of the network, we consider that stub ISPs pay the highest price for utilizing a link. By the same token, the price on a link from a transit ISPs to a core ISP is considered to be the lowest, since these links are of high capacity. Traffic on a customer-provider link between two transit ISPs is priced between these two. Finally, an ISP does not pay for traffic on a peering link to another ISP. The link prices in our model are not affected by the amount of P2P traffic that is routed over the link according to different peer selection algorithms.

Note that this simple pricing model does not take into account the changing traffic patterns in the Internet, and hence does not capture how the traffic volume on each link is affected by the reduction or increase in P2P traffic caused by different peer selection algorithms. We argue that our pricing model still captures the most important aspects of the economic impact of interdomain P2P traffic.

For a link between:	Price (per flow)
two peering ISPs	0
a transit ISP and a core ISP	1
a transit ISP and a transit ISP	3
a stub ISP and a transit ISP	10
a stub ISP and a core ISP	10

IV. ANALYSIS

A. Simulation Setup

We have implemented a snap-shot model for peer selection algorithms that calculates P2P traffic over each ISP-ISP link. The model is based on the framework outlined in Section III and considers a set of peers (leecher and seeder), an underlying ISP-level network topology, and the BGP routing information. Using this model, we have implemented the four peer selection methods discussed in Section II. We perform our experiments on an ISP-level topology consisting of 100 ISPs, where 400 P2P clients (200 seeders and 200 leechers) are uniformly distributed across all ISPs (i.e. according to the *Per-ISP* peer distribution method mentioned in Section III). The majority (80%) of ISPs are stubs, while 14% are transit and 6% are core. To reduce the random effects of the topological properties, we execute our simulations on ten different ISP-

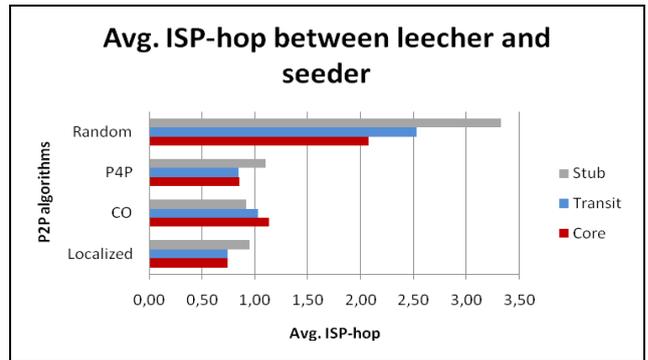


Figure 2: Average distance in ISP-hops between a leecher in different groups of ISPs and its seeders.

level topologies. The path of a P2P flow is determined by the outcome of the BGP route selection process.

Our snap-shot model allows us to measure inter-ISP traffic as a function of a particular peer selection algorithm, and to illustrate the different interests of different types of ISPs. However, since our model only captures the P2P traffic flow at a particular moment, it is unable to capture certain dynamic aspects of a peer selection algorithm, such as the robustness against peer failure or the nature of content distribution over a time span. Such evaluations will be part of our future work.

Since the business agreements between ISPs are generally not publicly available, it is hard to develop an accurate pricing model for today's Internet. We consider a simple pricing model that conforms with our discussion in Section III. We classify all links between ISPs into five different groups and allocate a fixed price for the unit traffic volume to each group of links. The type of link and price allocated for those groups are given in the Table I. Generally, the price of a unit data on a link is inversely proportional to the size of the link, according to our discussion in Section III.

As we are interested only in the flow of revenue between ISPs, we ignore the revenue from end users (like a DSL user) who pay a fixed price to their ISP for the Internet connectivity. Hence, in our model, the revenue of stub ISPs (which do not have ISP customers) will always be negative, because the fixed revenue that the stub ISPs collects is not taken into account. For the same reasons, we do not consider the network maintenance cost of an ISP, since this is also assumed independent of the traffic volumes.

B. Effect on the application performance

We calculate the distance between a leecher and its seeder in terms of number of ISP hops with different peer selection algorithms. We assume that in general, a lower ISP hop distance gives a better transmission quality (shorter delays and better throughput) [12]. Figure 2 shows the average distance between the leecher and the seeder for each group of ISPs with the different peer selection algorithms discussed in Section II. It shows that all three cross-layer peer selection algorithms perform better than the random peer selection in terms of number of ISP hops. This is because all cross-layer peer selection algorithms try to prevent a leecher from selecting a seeder far away from it. Since the locality based peer selection algorithm always prefers the seeder closest to the leecher, it has the minimum hop distance in every tier among all peer selection

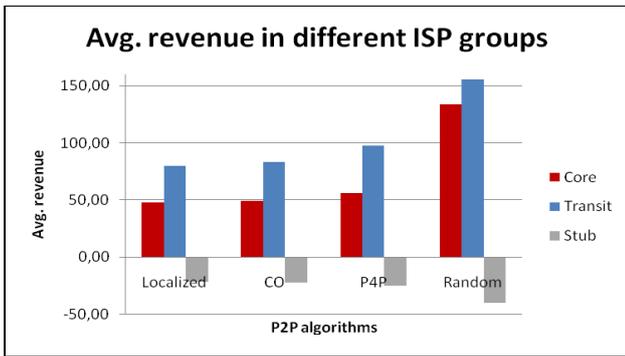


Figure 3: Average revenue collected by ISPs in different tiers

algorithms. Leechers in stub ISPs experience longer paths compared to peers in ISPs at upper tiers in all peer selection algorithm except CO.

In CO, transit and core ISPs exhaustively search its customer tree for a seeder and this leads to an increased average ISP-hop distance for them. For stub ISPs, using CO is nothing but using localized peer selection since they do not have any customer ISP below them. Even though CO prioritizes economic benefits over application performance, the application performance in CO is still much better than the application performance in random peer selection.

P4P tries to increase reliability by choosing at least 20% of seeder from outside of the network, and hence it has a marginally higher average ISP hop distance than the localized one. As a whole, for any type of ISP, a cross-layer peer selection algorithm that tries to localize the traffic provides better performance than the random peer selection. This is good news, since it indicates that cross-layer mechanisms normally will be a benefit to the P2P application.

C. The impact of peer selection algorithms on ISP revenue

The effect of peer selection algorithms on the economy is not uniform on all ISPs; it depends on the position of that ISP in the global Internet hierarchy. The amount of P2P traffic in an ISP depends on three parameters; number of seeder inside an ISP, number of leecher inside an ISP, and transiting P2P traffic passing through it. ‘Transiting P2P traffic’ refers to the traffic that passes through an ISP but is originated and terminated to other ISPs. With the uniform random peer distribution over the network, every ISP contains on average equal number of peers (seeders and leechers) inside it (e.g. on average, two seeders and two leechers in an ISP in this experiment). However, the amount of transiting traffic depends on the hierarchical structure of the Internet and peer selection algorithm used by the P2P application. When a leecher randomly chooses a seeder, the probability of choosing a seeder from other ISP becomes close to 1. e.g. $(1-2/200)$ for our simulation setup. This results in a large amount of inter-ISP traffic, giving a high volume of transiting traffic for transit and core ISPs. Hence, transiting traffic is always the main source of revenue for them. Since stub ISPs has no (or negligible) transiting P2P traffic through them, it pays for every p2p traffic ends in it. Figure 3 shows the average revenue collected by ISPs in different tiers when all ISPs are using a particular peer selection algorithm.

We observe that random peer selection mostly benefits

core and transit ISPs at the expense of stub ISPs. Core ISPs usually carry higher amount of transiting traffic compared to the transit ISPs. However, as shown in the figure, Core ISP earns less revenue than a transit ISP in our model, because transit ISP sell bandwidth to stub ISP at much higher rate than what it buys from a core ISP. The localized and P4P peer selection algorithms try to select a seeder for a leecher inside the same ISP. This strategy reduces inter-ISP traffic, decreasing the transiting cost of stub ISP from by 40%. However, these peer selection algorithms do not bring overall benefit to transit and core ISPs. Their revenue decreases compared to what they earn in case of random peer selection, since transiting traffic through transit or core ISP decreases dramatically. When every ISP shifts from using random peer selection to P4P, the transiting traffic reduces 82% for core ISPs and 65% for transit ISPs.

Figure 3 shows an interesting fact in case of CO peer selection. Although transit and core ISPs are redirecting traffic towards its customer network to get more revenue out of it, they fail to increase their revenue compared to random peer selection. The reason behind their failure is that stub ISPs dominate the Internet and when all stub ISPs start using CO by restricting p2p flows inside the ISP, transit and core ISP loose a major portion of transiting traffic which is the primary source of revenue for them in case of random peer selection.

D. Cross-layer peer selection algorithms for early adaptors

In our second experiment, we focus on the advantage of using peer selection algorithms based on cross-layer communication for the early adaptors. We start from a situation where all ISPs use random peer selection, and gradually increase the fraction of ISPs using an alternative method. For each point, we have repeated the experiment 10000 times in order to avoid random effects of the selection of ISPs and the peer distribution.

Our first experiment showed that cross layer peer selection algorithm benefits stub ISPs. Since this group constitutes the great majority of ISPs, it is interesting to see the effect on the economics as an increasing number of ISPs start using this method. Figure 4 shows how the early adoption of P4P has positive effect on the revenue collection of stub ISPs. The first stub ISP to adopt P4P reduces its traffic cost by more than 12%, by guiding its leechers to find seeders from inside the ISP, if available. Note that an ISP can only influence the choices of its leechers; i.e. Peers in other ISPs may still connect to seeders in this ISP as before. However, since many internal seeders will be busy serving local leechers, this traffic will also be reduced, explaining the non-linear shape of the curves in Figure 4. As all ISPs adopt P4P, the cost reduction for stub ISPs increases to 15%. Stub ISPs that do not convert to P4P are hardly affected. P4P reduces the revenue of transit and core ISPs, for the reasons explained above. Figure 4 shows that for early adaptors of P4P, this decrease is somewhat smaller than for transit and core ISPs that stay with random peer selection, since they are able to influence the selection process of their own leechers.

We repeat the same experiment for an increasing number of ISPs using other cross layer algorithms that shows the same trend of revenue collection for early adopters except

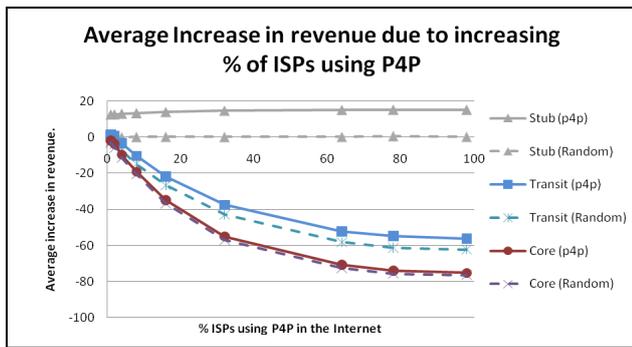


Figure 4: Increase in average revenue collected by ISPs in different tiers.

the fact that localized and CO are more positive to stub ISPs and less attractive to core and transit ISPs.

From the Figure 4, it is evident that early adaptors of cross layer strategies will reap much of the benefit from these algorithms from day one, even while a majority of ISPs use random peer selection. The full benefit will only be reached when all ISPs have made the change. ISPs using cross layer peer selection algorithm will always be more profitable than those that use random peer selection.

V. CONCLUSION AND FUTURE WORK

This paper finds that ISPs at different levels in the Internet hierarchy have different interests for using cross-layer peer selection algorithms. Cross-layer architectures like P4P or localized peer selection may improve the economy of stub ISPs as well as the performance of P2P application but decreases transiting traffic for transit or core ISPs causing less revenue for them. However, once all stub ISP start using a cross-layer peer selection algorithm, it is better for other ISPs at the higher layer to use it instead of using traditional random peer selection. Cross-layer architectures allow an ISP to influence the peer selection algorithm of the P2P applications inside its network. The results presented in this paper indicate that stub ISPs, in particular, have much to gain from adapting such methods. Cross-layer P2P algorithms primarily benefit ISPs that get their revenue from hosting users of P2P applications, for example ISPs with large number of residential broadband users, rather than from transiting traffic for other ISPs.

In future we intend to explore dynamic behaviors of different cross-layer p2p algorithms in realistic network topology and find the financial effect of them on ISPs. It is also interesting to see how those algorithms spread the content over the network and also how they perform in case of node or path failure.

REFERENCES

- [1] R. Steinmetz and K. Wehrle, "P2P Systems and Applications". Springer Lecture Notes in CS, 2005.
- [2] Thomas Karagiannis, Pablo Rodriguez and Dina Papagiannaki, "Should Internet Service providers Fear peer-Assisted Content Distribution?", IMC'05, Berkeley, 2005.
- [3] Light Reading, "Controlling P2P Traffic," http://www.lightreading.com/document.asp?site=lightreading&doc_id=44435&page_number=3.
- [4] T. Mennecke, "DSL Broadband providers Perform Balancing Act," <http://www.slyck.com/news.php?story=973>.
- [5] P. Rodriguez, S.-M. Tan, and C. Gkantsidis, "On the feasibility of commercial, legal P2P content distribution," SIGCOMM Comput. Commun. Rev., vol. 36, no. 1, pp. 75–78, 2006.

- [6] W. B. Norton, "A business case for isp peering," 2001. [Online]. Available: <http://www.equinux.com/>
- [7] R. Keralapura, N. Taft, C. Chuah, and G. Iannaccone, "Can ISPs Take the Heat from Overlay Networks?," in *HotNets*, 2004.
- [8] V. Aggarwal, S. Bender, A. Feldmann, and A. Wichmann, "Methodology for Estimating Network Distances of Gnutella Neighbors," *GI Jahrestagung - Informatik 2004*.
- [9] H. Xie, Y. Yang, A. Krishnamurthy, Y. Liu, A. Silberschatz, "P4P: provider Portal for Application", SIGCOMM, Washington, USA, 2008.
- [10] V. Aggarwal, A. Feldmann, and C. Scheideler. "Can ISPs and P2P systems cooperate for improved performance?" ACM CCR, July 2007.
- [11] Bindal et.al., "Improving Traffic Locality in BitTorrent via Biased Neighbor Selection," IEEE ICDCS, 2006.
- [12] Bo Liu, Yi Cui, Yansheng Lu, and Yuan Xue, "Locality-Awareness in BitTorrent-Like P2P Applications", IEEE transactions on multimedia vol. 11 no.3, April 2009.
- [13] D. P. Pezaros, L. Mathy, "Explicit Application-Network Cross-layer Optimisation", The 4th International Telecommunication NETworking WorkShop (IT-NEWS), Italy, 2008.
- [14] BitTorrent. <http://www.BitTorrent.org/>.
- [15] B. Ruan, W. Xiong, H. Chen, D. Ye, "Improving Locality of BitTorrent with ISP Cooperation", International Conference on Electronic Computer Technology, 2009.
- [16] X. Weng, H. Yu, G. Shi, J. Chen, X. Wang, J Sun, W. Zheng, "Understanding Locality-awareness in peer-to-peer Systems", International Conference on Parallel Processing – Workshops, 2008.
- [17] C. Tian, X. Liu, H. Jiang, W. Liu, Y. Wang, "Improving BitTorrent Traffic Performance by Exploiting Geographic Locality", Global Telecommunications Conference, 2008. IEEE GLOBECOM 2008.
- [18] A. Elmokashfi, A. Kvalbein, and C. Dovrolis. "On the Scalability of BGP: the roles of topology growth and update rate-limiting", CoNext 2008, ed. by ACM, 2008.
- [19] Y. Chu, S. Rao, and H. Zhang, "A case for end system multicast," in *Proceedings of ACM Sigmetrics*, pp. 1–12, June 2000.
- [20] D. Andersen, H. Balakrishnan, F. Kaashoek, and R. Morris, "Resilient Overlay Networks," in *Proceedings of 18th ACM Symposium on Operating Systems Principles (SOSP)*, pp. 131–145, October 2001.