A Measurement Study on AS-aware P2P Streaming Strategies

Bo Liu, Yansheng Lu
School of Computer Science and Technology
Huazhong University of Science and Technology
Wuhan, China
Email: {newpoo, lys}@smail.hust.edu.cn

Yi Cui, Yuan Xue
Department of Electrical Engineering and Computer Science
Vanderbilt University
Nashville, TN, 37235
Email: {yi.cui, yuan.xue}@vanderbilt.edu

***Abstract*—This paper presents a measurement study of AS-aware P2P streaming solution over real-world Internet topology. By using the accesses of nodes of PlanetLab testbed, we create a detailed AS-level map including the end-to-end path of all nodes, as well as the relationship of all involved ASes. Based on this map, we evaluate AS-aware P2P streaming solutions with different goals, i.e., minimizing AS hops and maximizing AS revenue, as well as the random P2P streaming solution which has no AS-awareness built in. Our findings suggest that AS-awareness can help existing P2P solution to significantly decrease load on Internet, and improve upon the uneven financial distribution among ASs. By comparing the performance of AS-aware and random P2P solutions, we also point out the necessity to tradeoff between the goals of achieving fairness among peers and optimizing AS-related performance such as intra-AS traffic and financial fairness.**

I. Introduction

Peer-to-peer (P2P) applications are getting increasing popularity. A few examples include bulk content distribution [13], Voice Over IP [15] and streaming systems [12] [14], all of which have been proved by the commercial deployment of planet-scale systems serving tens of millions of users.

In P2P systems, every peer not only downloads content from other peers, but also makes use of its upload bandwidth to serve other peers. There is great diversity of different peers both in terms of geographic distance and Internet topology, which introduces tremendous amount of inter-ISP traffic. Such traffic often causes great financial loss to ISPs with active P2P users, which motivates them to regulate or even ban P2P traffic. To alleviate the tension between ISPs and P2P users, many research [1], [2], [3], [4] have been conducted to bring ISP-friendliness into the current P2P systems.

However, what is still missing is a comprehensive study of the impact of P2P solutions on the ISPs, which themselves interconnect into a complex network of autonomous systems (AS). Such a study is impossible without a detailed and up-to-date AS-level map revealing the connection among ASes and their financial relationships. While mature techniques have been practiced to infer AS relationship, we still need to know the exact end-to-end path two peers connect with each other, which often traversing multiple ASes. Such knowledge can be only obtained when one is able to access a large number of peer machines.

In this paper, we present our initial investigation into this problem. With the aid of PlanetLab testbed[8], we obtain the access to more than 600 nodes around the globe, and use the information on their all-pair end-to-end path to construct a detailed AS-level map. Based on this map, we evaluate AS-aware P2P streaming solutions with different goals, i.e., minimizing AS hops and maximizing AS revenue, as well as the random P2P streaming solution which has no AS-awareness built in. Our findings indicate that AS-awareness can help existing P2P solution to significantly decrease load on the network, and improve upon the uneven financial distribution among ASs.

The rest of this paper is organized as follows. Related work is presented in Sec. II. Sec. III presents the evaluation methodology, i.e., how we obtain the AS-level map, which our evaluation is based on, and the streaming strategies we choose to study. Sec. IV evaluates streaming strategies on a series performance metrics including network cost, AS financial gain/cost, and peer load. Finally, we conclude the paper in Sec. V.

II. Related Work

We summarize previous works in two key areas related to our research: Internet AS relationships inferring and AS-aware P2P solutions.

There are many works on inferring Internet AS topology. The work by Gao [6] is the first comprehensive study targeted on this topic. By analyzing BGP table entries, it finds valley-free property of Internet AS path, and further identifies relationships between neighbor AS pairs. Spring et al.[5] present a technique for mapping the router-level topology of an ISP or a focused portions of the Internet, which use only end-to-end traceroute measurements. Dimitropoulos et al.[7] introduce some heuristics to address the problems of inferring peer-to-peer and sibling relationships. They also validate the inferred AS relationships. Though there are many works on inferring Internet AS relationships, the data set of inferred AS relationships is still incomplete as we have seen in Sec. III.

There have been several proposals on AS-aware P2P solutions. Bindal et al.[1] propose biased neighbor selection mechanism to reduce inter-ISP traffic, which requires no dedicated central servers. Karagiannis et al.[2] show locality-aware peer-to-peer can significantly alleviate the induced cost at the ISP.

The work by Ren et al.[3] confirm the benefits gained by peer-relay in VoIP, and then propose an AS-aware peer-relay protocol for P2P VoIP system. Huang et al.[4] confirm the benefits of peer-assisted VoD, and show that locality-aware P2P solution can reduce inter-ISP traffic. However, there is little work on the evaluating impact of AS-aware P2P solutions based on real world Internet topology.

III. Methodology

A. Obtaining AS-level Map

Our experiment is conducted over real-world Internet topology. We construct an AS-level map on the PlanetLab testbed[8]. Fig. 1 illustrates this process step by step.

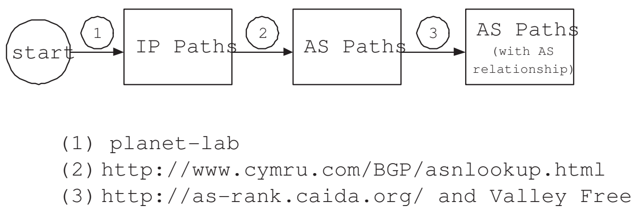

Fig. 1. Step-by-step Process of Obtaining AS-Level Map

First, we run traceroute between each pair of PlanetLab nodes to obtain the IP-level end-to-end path between them. We then assemble these paths into an IP-level map. Some nodes are eliminated from the map since the traceroute program fails to return the IP-level path over them. The final IP-level map consists of two kinds of nodes, 340 end hosts which are PlanetLab nodes, and more than 13000 routers which represent the IP substrate interconnecting the end hosts.

Second, we convert each IP path obtained in step one into an AS path. For each IP address shown up in the IP-level map, we find its AS number through public AS-lookup service such as the one run by the CYMRU team[9]. Since an end-to-end path linearly traverses multiple ASes, we aggregate consecutive IP addresses with the same AS number into a single AS node. In this way, we condense an IP-level path into an AS-level path, and further transfer the IP-level map into an AS-level map. A small number of nodes are further eliminated due to the failure of AS lookup.

Third, we mark the AS-level map with AS relationship data provided by CAIDA[10]. Each adjacent AS pair must have one of the following four kinds of relationships: Provider-Customer, Customer-Provider, Peer-to-Peer, and Sibling. Each customer AS should pay its provider AS for for both inbound and outbound traffic. Traffic across peer-to-peer ASes or sibling ASes are usually free. 70% of the AS pairs in our AS-level map are identified via the CAIDA dataset.

To identify the relationship of the remaining AS pairs, we apply the "valley-free" property proposed in [6]. In brief, if we represent our AS-level map in a hierarchical structure where every AS is positioned lower than its provider, higher than its customer, and at the same level with its siblings and peers, then any AS-level path should not form a valley, i.e., the path should start with zero or more customer-provider pairs, then zero or more peer-to-peer pairs, finally zero or more provider-customer pairs, and sibling pairs can exist in any place of an AS path. Assuming all AS-level paths follow this property, we improve the percentage of identified AS pairs to 90%. Finally, we eliminate the unidentified AS pairs from the map.

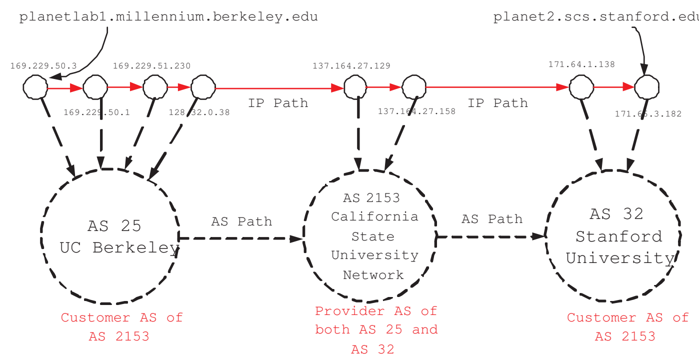

Fig. 2. A Processing Example

In Fig. 2, we illustrate the above steps by a sample PlanetLab pair "planetlab1.millennium.berkeley.edu" and "planet2.scs.stanford.edu". In this example, the IP path consists of eight IPs, which locate in three different consecutive ASes. These ASes are UC Berkeley(AS 25), California State University Network(AS 2153) and Stanford University(AS 32). The AS 2153 is provider of both AS 25 and AS 32. Therefore, for traffic from AS 25 to AS 32 or vice versa, both of them will be charged by AS 2153.

B. Streaming Strategies

All strategies are evaluated by the same setup as follows. We choose a YouTube[11] server as the source of the streaming. We then create a request sequence where each peer, i.e., PlanetLab node in our AS-level map, is scheduled to joined the streaming at different times. The size of the media file is 483 blocks sized 32K bytes, a random video file taken from the YouTube website. The strategies are as follows:

1) *client-server* (CS) strategy, where the server handles the requests from all peers. The AS-level path from each client to the server is obtained the same way as described in the previous subsection.
2) *random P2P streaming* (Random P2P) strategy, where each peer randomly selects its parents from earlier-joined peers, and download the same amount of data from each parent. The maximum numbers of children and parents per peer are upper bounded by a Degree Constraint parameter (DC).
3) *minimum-AS-hop P2P* (Min-AS-Hop) strategy, where each peer selects, among the earlier-joined peers, the one with the shortest AS-level path. If a tie exists among multiple candidates, the peer will make all of them to be its parents, and download the same amount of data from each parent. The intuition here is that the fewer number of ASes a P2P connection has to go through, the less inter-ISP cost it will incur. In light of the fact that a peer

is not able to serve too many children, we also test the degree-constrained version of the same strategy, where the maximum number of children per peer are upper bounded.

4) *maximum-AS-revenue P2P* (Max-AS-Revenue) strategy. Here, each peer selects, among the earlier-joined peers, the one which brings the maximum revenue to the AS it belongs to. In other words, the peer has the highest preference on the beginning AS hop to be provider-customer, which brings revenue to the ISP it belongs to. The second preference is the type of peer-to-peer or sibling, which brings neither revenue nor charge. The least favored option is customer-provider, which makes the ISP pay. The tie-breaking method is the same as the Min-AS-Hop strategy. Similar to the case of Min-AS-Hop, we also test the degree-constrained version of the Max-AS-Revenue strategy.

IV. Evaluation

We evaluate the performances of aforementioned streaming strategies over the following metrics, network cost in terms of IP hops and AS hops, financial gain/cost on each AS, and the load on each peer.

A. Network Cost

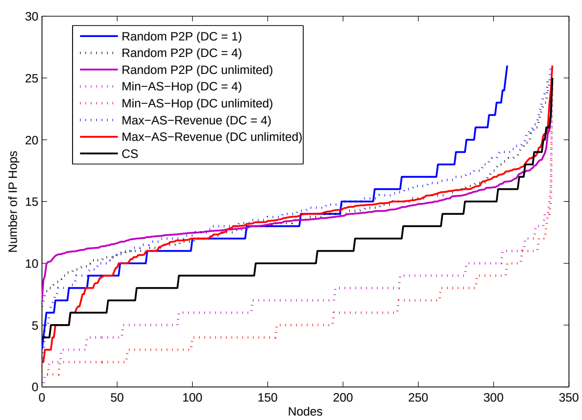

Fig. 3. Average Network Cost in Terms of IP Hops

Fig. 3 shows the sorted curve for each peer the average network cost in terms of IP hops. In any P2P streaming strategy, a peer might download blocks from different peers. Therefore, the average cost is calculated as the summary of number of IP hops traveled by all blocks divided by 483, the total number of blocks. Note that in CS strategy, since the each peer downloads all blocks from the server, its curve actually reveals the distance from the server to each peer. Also note that for all P2P strategies, namely Random, Min-AS-Hop, and Max-AS-Revenue, they render the same result when the degree constraint is 1 (DC$=1$). This is because each newly-joined peer has no choice but to choose the previous peer as its parent. Thus, all strategies end up with a P2P relaying chain of the same sequence.

Not surprisingly, the Min-AS-Hop strategy delivers the best performance among others. In particular, the average IP hop reaches the minimum when no degree constrain is enforced, about 2 hops less than the case when DC$=4$. On the other hand, Random and Max-AS-Revenue both perform worse than the CS strategy, indicating that the average distance among peers is higher than the distance between the server and all peers.

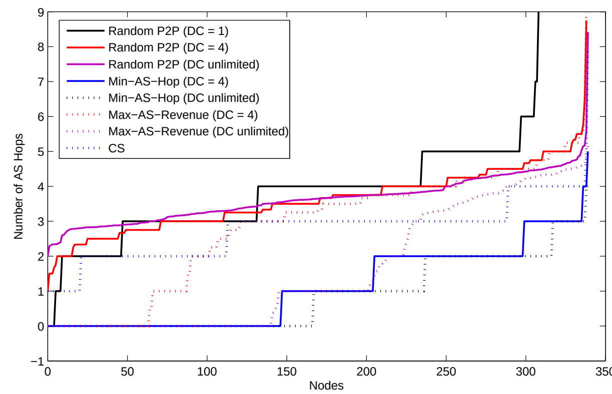

Fig. 4. Network Cost in Terms of AS Hops

Fig. 4 shows the network cost in terms of average AS hops, which is calculated the same way as average IP hops. It largely demonstrates the same trend as observed in Fig. 3. The Min-AS-Hop strategy performs the best. In particular, when the degree constraint is unlimited, around half the traffic is contained within the same AS, mostly due to the fact that each participating institution of the PlanetLab testbed must maintain at least two nodes, making each AS to contain at least two PlanetLab nodes. When DC$=4$, the same value is reduced to 40%. In terms of maximizing intra-AS traffic, Max-AS-Revenue returns comparable performance with Min-AS-Hop when the degree constrain is unlimited. Interestingly, Max-AS-Revenue outperforms the CS strategy, contrary to what is shown Fig. 3, possibly because it prefers downward AS paths.

From both figures, we have the following common observations. First, Min-AS-Hop consistently returns the best performance, due to its design goal. Second, the Random strategy performs worse than all other streaming strategies, confirming the fact that P2P solutions not aware of AS topology will unnecessarily increase the load on Internet, even worse than the CS strategy.

B. AS Financial Gain/Cost

We measure the financial gain/cost on each AS by number of blocks, assuming the same charging rate is practiced by all ASes. Positive value indicates financial gain, while negative value indicates otherwise.

Fig. 5 shows the sorted curve for each AS. From it we discover that among all strategies, min-AS-Hop involves min-

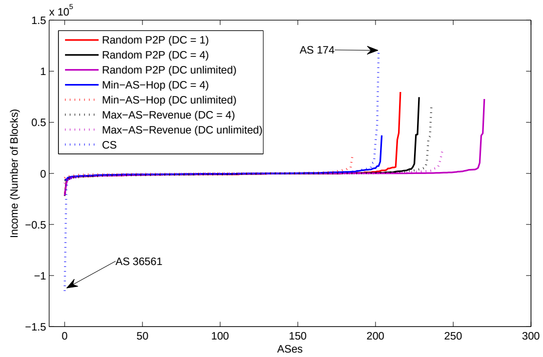

Fig. 5. Financial Gain/Cost in Terms of Number of Blocks

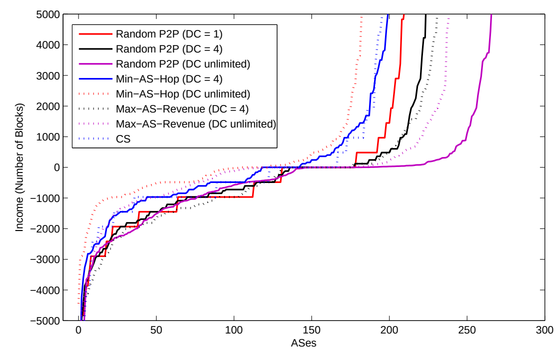

Fig. 7. Financial Cost of middle range ASes

imum number of ASes. The CS strategy comes in second, followed by Random when DC$=4$, and Max-AS-Revenue. Random strategy with no degree constraint involves the maximum number of ASes. We also observe that, for all strategies, the financial gains/costs among all ASes are highly uneven. A small number of ASes gain or lose significantly, while the majority of ASes stay slightly up or down 0. In what follows, we focus on the left (Fig. 6), the middle (Fig. 7) and the right (Fig. 8) parts of this figure.

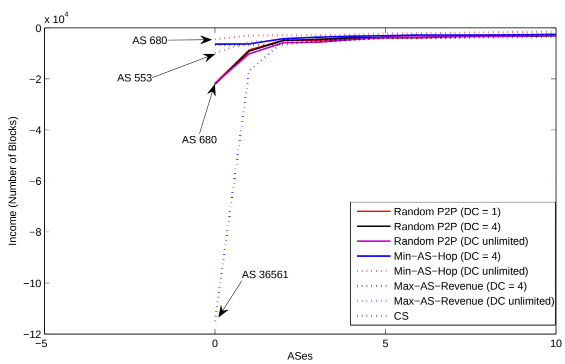

Fig. 6. Financial Cost of First 10 ASes

Fig. 6 shows the CS strategy's income distribution to be the most uneven of all. AS 36561 pays more than 80 percents of the total cost of all ASes in debt, Not surprisingly, this AS is "Youtube Inc.", which connects to many other ASes as a customer or a peer. On the positive side, any P2P streaming strategy can reduce this number by from 82% to more than 97%. Again, Min-AS-Hop produces the maximum reduction, followed by Max-AS-Revenue, then Random.

In Fig. 7, we show ASes in the middle range, where the crossing points from ASes in debt to ASes in profit clearly show. It is an obvious fact that this is zero-sum game, where total gain equals total cost. However, we find that for all strategies, the number of ASes in debt is slightly more than those which profit. This is consistent with the current Internet hierarchy, where the provider charge its customers on both inbound and outbound traffic.

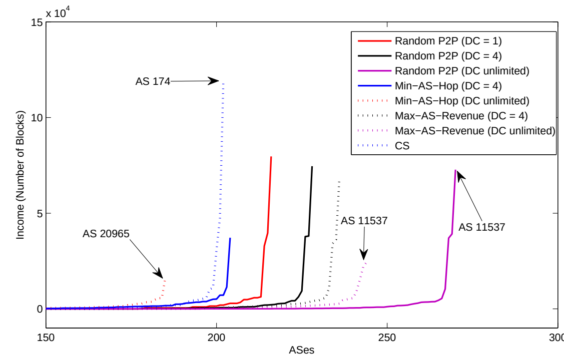

Fig. 8. Financial Cost of Last 150 ASes

Fig. 8 show the last 150 ASes, all of which profit. Among all strategies, CS creates the highest profit for a single AS. It is AS 174, "Cogent Communications", which is a provider of a lot of other ASes. Max-AS-Revenue fails to achieve so, although its goal is to maximize the revenue for the AS each peer belongs to. We conjecture that this strategy introduces the competition, causing the profit to distribute more evenly among all ASes. When the degree constraint is dropped, more ASes are brought to the competition, which further distributes the profit.

Finally, Min-AS-Hop reduces the highest profit to its minimum, causing it the best strategy to even out the gain/cost among all ASes. Also compared to Max-AS-Revenue, dropping the degree constraint plays the opposite role, causing fewer ASes to be involved than the case when DC$=4$.

C. *Peer Load*

We finally evaluate the load imposed on each peer by all P2P streaming strategies. CS strategy is excluded since all of its workload is on the server.

Fig. 9 shows the sorted curve for each peer in terms of the number of blocks served. Min-AS-Hop incurs the most uneven load distribution, where more than 100 peers are free riders when there is no degree constraint. Random strategy performs the best with this regard, while the performance of Max-AS-Revenue is in between. We also note that

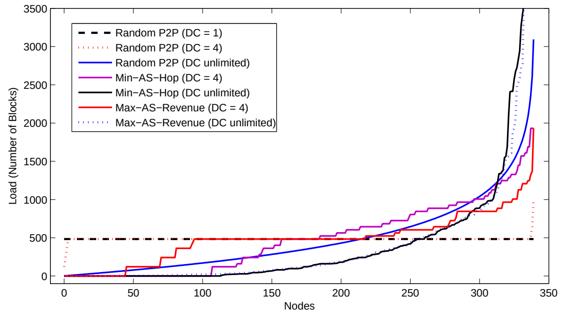

Fig. 9. Load of Each Node

enforcing degree constraint can effectively improve the uneven distribution for all strategies, while the order of their performances stay unchanged.

From our evaluation, we have the following observations. First, compared to the random strategy, AS-aware P2P streaming strategies (Min-AS-Hop and Max-AS-Revenue) can effectively reduce the network cost in terms of minimizing average IP hops and AS hops, and maximizing intra-AS traffic. Second, AS-aware strategies also help even out the financial gain/cost distribution among ASes, where Min-AS-Hop performs the best. Finally, random strategy outperforms the AS-aware strategies by introducing fairer workload among peers, indicating the necessity to tradeoff between the goals of achieving fairness among peers and optimizing AS-related performance such as intra-AS traffic and financial fairness.

V. Conclusions

In this paper, we conduct a comprehensive evaluation study on a set of P2P streaming solutions based on real world Internet topology. We study AS-aware solutions with different objectives, i.e., minimizing AS hops and maximizing AS revenue, as well as the random P2P streaming solution which has no AS-awareness built in. While it clearly shows the advantage of AS-aware solutions at reducing inter-AS traffic and achieving better financial fairness among ASes, they also demonstrate deficiency at evenly distributing peer workload as done by the traditional random solution. As such, our study suggests the necessity to consider, in the design of future P2P streaming solutions, the tradeoff between the goals of achieving fairness among peers and optimizing AS-related performance such as intra-AS traffic and financial fairness.

References

[1] R. Bindal, P. Cao, W. Chan, J. Medved, G. Suwala, T. Bates and A. Zhang , *Improving Traffic Locality in BitTorrent via Biased Neighbor Selection*. 26th IEEE International Conference on Distributed Computing Systems(ICDCS), 2006.

[2] T. Karagiannis, P. Rodriguez, and K. Papagiannaki , *Should Internet service providers fear peer-assisted content distribution*. In Proceedings of the Internet Measurement Conference 2005, 2005.

[3] S. Ren, L. Guo, X. Zhang , *ASAP: an AS-Aware Peer-Relay Protocol for High Quality VoIP* . 26th IEEE International Conference on Distributed Computing Systems(ICDCS), 2006.

[4] C. Huang, J. Li, K. W. Ross, *Can internet video-on-demand be profitable?*. SIGCOMM'07: Proceedings of the 2007 conference on Applications, technologies, architectures, and protocols for computer communications, pp. 133-144, 2007.

[5] N. Spring, R. Mahajan, David. Wetherall and T. Anderson, *Measuring ISP topologies with rocketfuel*. Vol. 12, No. 1, pp. 2-16, IEEE/ACM Transaction on Networking, 2004.

[6] L. Gao, *On inferring autonomous system relationships in the internet*. Vol. 9, No. 6, pp. 733-745, IEEE/ACM Transaction on Networking, 2001.

[7] X. Dimitropoulos, D. Krioukov, M. Fomenkov, B. Huffaker, Y. Hyun, k. claffy and G. Riley, *AS relationships: inference and validation*. Vol. 37, Issue. 1, pp. 29-40, ACM SIGCOMM Computer Communication Review, 2007.

[8] http://www.planet-lab.org. May 2008.

[9] http://www.cymru.com. May 2008.

[10] http://www.caida.org. May 2008.

[11] http://www.youtube.com. May 2008.

[12] http://www.pplive.com. May 2008.

[13] http://www.bittorrent.com. May 2008.

[14] http://www.uusee.com. May 2008.

[15] http://www.skype.com. May 2008.