

Pitfalls for ISP-friendly P2P design

Michael Piatek* Harsha V. Madhyastha† John P. John* Arvind Krishnamurthy* Thomas Anderson*

ABSTRACT

Peer-to-peer file sharing applications have become enormously popular over the past few years, coming to represent a large fraction of wide-area Internet traffic. A side effect of this explosive growth has been an emerging tussle between users, who want fast downloads, and ISPs, whose flat-rate pricing business model is threatened by the extreme volume of P2P traffic. Because ISP costs scale with usage while their prices do not, many ISPs have attempted to throttle or shut down P2P systems. Recently, several researchers have proposed that this tussle is unnecessary, that small changes in client and/or protocol behavior can lead to a “win-win” solution of better performance for end-users with less wide-area traffic for ISPs. Using a very large scale trace measurement of BitTorrent usage, we find evidence that such a win-win outcome is unlikely for at least one very popular P2P protocol.

1. INTRODUCTION

Peer-to-peer (P2P) networks have emerged as a powerful tool for building robust and scalable systems. Many P2P systems are unmanaged, yet achieve high levels of performance, scalability, and robustness through the use of *randomness* in their communication pattern. Randomization increases path diversity and avoids bottlenecks, but this robustness is not without cost. For example, the most popular use of P2P today is for file sharing in networks such as BitTorrent and eDonkey. Because files in these networks are shared with a random selection of peers spread across the globe, data being transferred often traverses multiple ISPs. The combination of the popularity of P2P services and their inefficiency means that file sharing accounts for a large portion of backbone ISP traffic [11, 21].

Network oblivious sharing increases costs for ISPs. As demand increases on transit links, backbone ISPs are forced to invest in increasing capacity. These costs are passed on to edge ISPs that pay transit costs proportional to the amount of interdomain traffic they generate. Many customer facing ISPs, however, offer flat-rate pricing, forcing them to absorb the costs externalized by P2P file sharing. In response, some ISPs have elected to rate-limit or simply block these “problem” protocols [4, 14], in turn leading developers into an arms race to evade restrictions.

Recent research has suggested that this emerging tussle between ISPs and their users can largely be avoided through small changes in client and/or protocol behavior. We call this class of techniques *ISP-friendly* in that they reduce the burden of P2P applications on providers.

For example, the Ono system [3] proposes that the BitTorrent client be modified to prefer local peers, and the P4P project [20] provides a way for ISPs to notify clients which peers are preferred. Both claim that their designs are “win-win” for both users and ISPs: download speeds will improve and interdomain traffic will be reduced.

To validate these claims and quantitatively compare ISP-friendly design strategies, we conducted a very large scale measurement of BitTorrent usage. Our goal is to develop data to guide system and protocol designers in shaping the tussle space between users and their ISPs. Our measurements span almost twenty thousand swarms (groups of peers downloading the same file) encompassing nearly fifteen million unique IP addresses in total. By using simultaneous measurements of the Internet topology at scale, we can determine how often data in each swarm would transit the boundaries between ISPs. And indeed, we find that, in an idealized setting, most of the interdomain P2P traffic in our trace is unnecessary.

Our principal result, however, is that this benefit is difficult to achieve in practice. Specifically, we find the following pitfalls to adapting BitTorrent to be ISP-friendly:

- **Limited impact:** Contrary to the published literature, client-only optimizations to BitTorrent yield neither better performance nor less interdomain traffic in the common case. Our traces show that BitTorrent clients usually have too few peers to find many in the same ISP. To confirm this in the wild, we show that Ono reduces interdomain traffic by less than 1% when connecting to live swarms through a large residential ISP.
- **Reduced performance and robustness:** Optimizing for locality alone degrades the structural robustness of overlay topologies and, for many users, performance. In contrast to random topologies, those optimized for locality rely on fewer interdomain links to connect clusters of local peers. Also, local peers are not always faster, particularly for users in regions where asymmetric bandwidth capacities are typical, e.g., the US.
- **Conflicting interests:** Reducing interdomain traffic reduces costs for some ISPs, while it reduces revenue for others. We present trace data demonstrating that the set of tier-1 ISPs have a strong incentive to strategically manipulate BitTorrent peering relationships, creating longer paths than necessary and potentially setting up an arms race between ISPs.

In sum, the tussle over P2P traffic between users and ISPs, and between ISPs themselves, is likely to continue. We hope that these challenges will serve to motivate the community to revisit the issue of ISP-friendly P2P design

*Univ. of Washington

†Univ. of California, San Diego

from a holistic perspective, taking into account the interests of content providers, network operators, and users.

2. MEASUREMENT OVERVIEW

We are interested in understanding the interaction between file sharing and ISPs from an *Internet wide* perspective. An Internet wide perspective should include a characterization of the many objects and the many networks involved in distribution. Thus, we use large-scale measurements of both 1) membership in file sharing networks, i.e., which IPs are participating in which swarms, and 2) the AS paths between those peers. For the latter, we use measurements from the iPlane project [12], which refreshes its atlas of the Internet topology daily. We next describe our measurements of BitTorrent.

BitTorrent distributes large files by splitting them into blocks and distributing blocks out-of-order from the data source. Peers bootstrap into the overlay by contacting one or more central servers, called trackers, which maintain a list of active peers and provide a random subset of these upon client request. We focus our attention on BitTorrent because it is among the most popular P2P networks today and represents a significant volume of Internet traffic.

To collect a sample of BitTorrent peers, we crawled a set of well-known websites that aggregate .torrent metadata, downloading the set of 18,370 target swarms hosted by those websites. Crawled swarms range in popularity from new, popular swarms with tens of thousands of users to those with few participants. Swarm metadata includes the total size of the set of files to be downloaded, providing us with the demand (in bytes) of each user in each swarm. From a cluster of ten machines at the UW, we contacted the trackers associated with each swarm repeatedly over a one month period and requested membership information. Because many BitTorrent trackers return only a small subset of total available peers (~ 50), we made multiple requests per swarm, one from each of our measurement nodes, and repeated this query every 15 minutes (membership in each swarm was queried every ninety seconds, on average). Over the course of a month, we observed 14,380,622 distinct IPs, with many occurring in multiple swarms.

3. POTENTIAL REDUCTION IN INTER-DOMAIN TRAFFIC

We first use our traces to analyze locality optimization with cooperative users and complete topology information. This is the best case for reducing interdomain traffic; peers select paths to reduce traffic without regard to performance, and each client has complete knowledge of *all* other peers and AS-level paths to these peers. Of course, individual network operators may have more complex traffic engineering priorities, but reducing interdomain traffic is a broadly shared goal. We consider four approaches for peer selection at clients:

- Random: Matching with random existing users.

Method	Lifetimes	Percentage reduction		
		Small	Large	Overall
Shortest path	8 Hours	11.2	35.6	27.3
	6 Days	30.9	51.6	43.6
Same-AS	8 Hours	4.1	23.6	17.0
	6 Days	17.8	41.9	32.6
Latency	8 Hours	4.7	22.1	16.2
	6 Days	18.2	36.9	29.7

Table 1: Percentage reduction in interdomain traffic relative to random peer matching.

- Latency: Matching users with least delays.
- Same-AS: Matching with users from the same AS when possible. Otherwise, random selection.
- Shortest path: Matching with peers in order of ascending AS path length.

Random peering is the default behavior of BitTorrent trackers, and it serves as our basis for comparison. Some BitTorrent clients have been extended to use latency-based heuristics for choosing among the random set of peers available locally (e.g., [10]). We evaluate the application of a simple latency-based policy globally. Same-AS reflects ISP self-interest with respect to minimizing interdomain traffic, but without distinguishing between short interdomain paths and those that are lengthy. Shortest path attempts to minimize the use of interdomain links overall, representing—in some sense—the common good.

We apply each heuristic during playback of our trace. For each peer join event, we use our measurements to predict AS paths and latencies between the new user and existing peers. Candidate peers are rank ordered according to the given metric and selected in order until the new peer has either satisfied a connection requirement of 30 peers or has exhausted existing peers. To make our analysis tractable, we restrict our consideration to 1,000 randomly sampled swarms from our overall trace for a week long period. Subsequent trace analysis refers to this sample.

The main factor that controls locality in the BitTorrent sharing workload is choice. Only when users have many available sources from which to request data can they take advantage of locality-aware peer matching. When there are few available peers, users will exchange with anyone that can provide needed data, regardless of locality. Choice in a given swarm is determined by two properties: 1) the swarm’s *fundamental* popularity and 2) peer lifetimes. Fundamental popularity refers to the total number of users interested in a data object over its lifetime. If these users are also persistent, i.e., they continue to share after a download completes, many choices will be available.

We distinguish among these four cases in our analysis as each has different traffic and locality properties. Specifically, we perform the same trace playback using each different peer matching strategy and with peers having either a 6 day or an 8 hour lifetime. Note that our trace method-

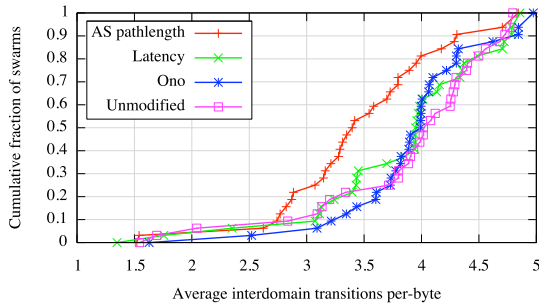


Figure 1: The cumulative fraction of swarms with a given average of interdomain transitions per-byte.

ology is unable to measure the actual peer lifetime, but other measurements indicated that most BitTorrent clients disconnect within a few hours after completing a download. Hence the 6 day lifetime is designed to illustrate what would happen if clients were incentivized to continue sharing well past download completion. We further separate these results into those for large, popular swarms (the top 10%) and remaining smaller swarms.

Results for each of these trials are summarized in Table 1. These results show across-the-board reductions in interdomain traffic. The greatest benefit is achieved when many choices are available; i.e., for popular swarms with long-lived peers. Moreover, locality optimization is most effective when using information about the underlying network topology; latency is a poor predictor of the number of interdomain crossings.

4. REDUCING INTERDOMAIN TRAFFIC IN PRACTICE

The encouraging results of our trace replay are consistent with published literature attesting to the benefits of locality-aware peer selection. In practice, however, we found these benefits difficult to achieve. In this section, we report results of a real-world comparison of three methods of reducing interdomain traffic.

From a measurement node connected via Comcast, a popular US residential cable ISP, we joined a set of 32 popular, recently created candidate swarms drawn from a popular BitTorrent aggregation website and performed back-to-back downloads with instrumented BitTorrent clients. Each download was performed four times with four different peer selection strategies: 1) shortest AS path length, 2) minimum latency, 3) the Ono client plugin, and 4) unmodified BitTorrent. To select peers with minimal AS path length or latency, we use path predictions from iPlane [12]. Ono estimates whether two peers are local by measuring the overlap in CDN replicas to which each node is directed [3], and peers with high CDN replica overlap are considered local. Because Ono is a plugin for the Azureus BitTorrent client that we use for all trials, we provide a direct, apples-to-apples comparison. To limit the time between the first and last downloads of a given swarm, we downloaded only the first 30 MB of each file.

Figure 1 compares the average interdomain transits per-byte obtained from downloading candidate swarms with each peer selection strategy. These results show the limited real world reduction in interdomain traffic realized by a single locality-aware client today. The median value is reduced from the unmodified baseline of 4.01 to 3.39 for shortest AS path, a reduction of just 15% with a negligible difference in performance (not shown in graph). The median ratio of download times between a client using shortest AS path and unmodified peer selection is 0.98. This is by design; we swap distant for local peers only when the switch is expected to maintain or improve performance. Our assumption is that users are unlikely to adopt a locality-aware client that reduces performance, which is a risk when optimizing for locality alone, a topic we return to in the next section.

The impact of Ono on both performance and locality is negligible. With respect to performance, the median ratio of download times between Ono and an unmodified Azureus BitTorrent client is 1.02. Figure 1 summarizes the impact on interdomain traffic. The median weighted AS path length for Ono is 3.99, versus 4.02 for the unmodified client. Although these results might seem contradictory given previously published measurements of Ono, the difference is simply one of presentation. While 33% of *Ono-recommended peers* are within a single AS and download rates increase by 31% for *recommended paths*, Ono’s end-to-end benefit is limited by the vanishingly small fraction of peers it recommends as “local”, even when applied to new, popular swarms.

These results expose many pitfalls in adapting BitTorrent to be ISP-friendly.

- Client-only ISP-friendly designs suffer from the lack of complete information regarding concurrent downloaders. Maximizing efficiency depends on peer matching with a global perspective, e.g., at the tracker.
- Download-and-depart behavior limits the potential for reducing interdomain traffic. Increasing exploitable locality depends on users continuing to share even after downloads complete. But, BitTorrent includes no incentives to do so.
- Even if local replicas exist, a client is likely to prefer non-local peers that provide higher download rates. In the next section, we consider the performance implications of optimizing for locality in isolation.

5. DOWNLOAD PERFORMANCE

Locality optimization exposes a performance tradeoff. Without any effort to remain performance neutral, preferential exchange with local peers may reduce performance for some users. The bandwidth capacity of BitTorrent users is highly skewed, and the majority of total capacity comes from a small minority of high capacity peers [7]. But, these peers are not uniformly distributed; clustering peers globally on the basis of locality also tends to clus-




Figure 2: The ratio of the capacity of each user’s set of local peers when matched to minimize AS path length and randomly.

ter them by capacity.¹ Although the total amount of capacity remains unchanged, clustering high capacity peers increases download rates for the high capacity minority while reducing the overall average download rate per-user.

To make these issues concrete, we consider the potential change in download rates resulting from an idealized shift to shortest AS path matching. We consider the 100 most popular swarms from our trace. For the set of observed peers from each of these swarms, we simulated a tracker that selected either 50 users at random or selected 50 users based on minimal AS path length. We assign capacities on a prefix level, using bandwidth measurements of more than 100,000 BitTorrent users collected from popular swarms in 2006 [7].

To express the change in per-user performance, we compute the ratio of the average capacity of each client’s peers when matched to minimize AS path length and when matched randomly. A ratio greater than 1 implies that the average capacity of peers per-user increases when using shortest AS path matching and a ratio less than 1 implies that average capacity decreases. The distribution of these ratios is shown in Figure 2. These results show that for the majority of peers in the majority of swarms, the total capacity of their peers is greater under random matching than under shortest AS path matching. One might expect the median ratio of average download rate to be 1; i.e., for each peer in a swarm, some nearby peers will be slower, but others will be faster. Instead, the median ratio is 0.15. This is because most BitTorrent *peers* from popular swarms in our trace come from the United States, while most *capacity* comes from comparatively high bandwidth peers in Europe.

6. STRUCTURAL ROBUSTNESS

To what extent does locality-aware peer selection degrade the resilience of the overlay graph? Unfortunately, there is no standard metric for quantifying the robustness of a network topology. We apply the following heuristic. For each peer in a swarm, we compute the shortest path in the overlay graph to all other nodes. Next, edges are

¹BitTorrent’s tit-for-tat policy attempts to achieve bandwidth matching *locally* by choosing among a subset of peers, but its effectiveness is limited by a partial view of potential peers, slow convergence, and churn [15].

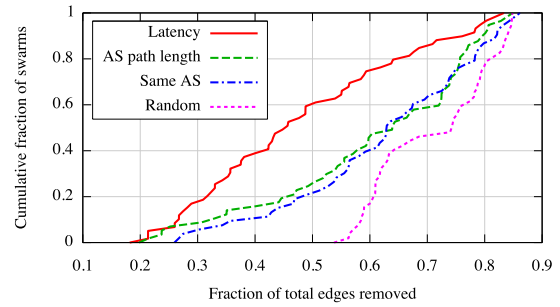


Figure 3: The cumulative fraction of swarms as a function of the fraction of edges removed to disconnect at least half of peers in the overlay topology.

ordered by their popularity among the set of shortest paths from all nodes. We then remove the most popular 1% of these edges and repeat this process until at least half the peers are disconnected from the largest connected component. This metric measures the extent to which the robustness of the overlay topology depends on a small minority of crucial connections that have relatively low redundancy.

Figure 3 summarizes the results, showing the cumulative distribution of the fraction of edge removals required to disconnect half the overlay peers from the largest connected component. The median fraction of removals required decreases from the random baseline of 0.75 to 0.45 for a latency-based overlay. Using either same-AS or shortest path preferences results in more resilience but still falls short of the robustness of a randomly constructed overlay. This data reflects the varying impact of locality-aware selection depending on the type of swarm. Each selection strategy sometimes constructs overlays much more easily disconnected than random pairing ($y\text{-axis} \leq 0.1$). These generally correspond to very popular swarms with significant exploitable locality amenable to a particular selection strategy. Also, for very unpopular swarms ($y\text{-axis} \geq 0.9$), the selection strategy has little influence on robustness since choice among peers is so limited.

7. STRATEGIC ISPS

So far, we have considered ISPs as cooperating to achieve a common goal: reducing interdomain traffic. In practice, ISP-friendliness is not well-defined. Individual ISPs may have specific traffic engineering goals within their network that are not taken into account by our analysis. And, while minimizing interdomain traffic may represent the common good, individual ISPs derive little benefit from minimizing AS path length once traffic exits their network.

More fundamentally, what is friendly to one ISP may be unfriendly to another, as ISPs themselves have commercial relationships with each other that are strongly affected by user choices. Tier-1 ASes would prefer *more* interdomain traffic, not less, since their customers pay for transit traffic. For ASes with customers that pay based on peak usage or per-byte, an increase in locality means a likely reduction in revenue. This raises the question: can an ISP *increase* its revenue by influencing the file sharing choices




Figure 4: Strategic peer matching. Circles denote P2P users and dashed lines represent monetized paths.

made by its clients, and how much of an impact would that have on global efficiency?

We next examine the following questions: 1) What mechanisms can a strategic ISP use to influence sharing among BitTorrent users? and 2) what is the impact of strategic ISP behavior on efficiency? We find that opportunities for subverting locality to increase revenue are frequent. Strategic behavior increases average path lengths relative to shortest AS path matching by 72.6%.

Matching strategy: Generally, a strategic AS would prefer that its users connect to other P2P users in other ASes according to its default BGP policy: with its 1) customers, 2) other users in its AS or peer ASes, or 3) users reachable through its provider(s), in decreasing order of preference. If the strategic AS is not a tier-1, clearly any match which requires transit on its provider links (case 3) should be avoided if possible. Alternatively, if the strategic AS can induce P2P users in its customer ASes to send and receive additional traffic through other customers (case 1), such matches are preferable as they may generate revenue. If such profitable matches cannot be found, matches that are made should prefer P2P users within the strategic AS or its peer ASes (case 2) and avoid the provider link. How might a strategic ISP implement this policy?

- In the case of BitTorrent, if the ISP hosts the tracker or can mandate the use of a particular client, it can implement a strategic policy directly.
- If a P2P network supports ISP-provided hints about which paths are ISP-friendly, as suggested by P4P [20] or the recently proposed ALTO/BGP IETF draft [16], a strategic ISP can explicitly mark peers on revenue generating paths as ISP-friendly regardless of locality.
- If the ISP can shape traffic on its network, it can indirectly induce profitable sharing by making revenue generating paths fast and expensive paths slow. By design, BitTorrent attempts to discover (and use) fast paths.

Although the strategic policy that we apply will never cause a strategic AS to lose revenue, it does not guarantee that each connection with a customer results in revenue gain. This depends on the circumstances of the customer. Figure 4 shows two typical cases: (a) when the strategic AS is the customer’s only provider and default route, and (b) when the customer AS is multi-homed.

Single-homed: In (a), we compare strategic and locality aware matching. At left, P2P users in both provider and customer prefer local peers to minimize path length and interdomain traffic. Although several connections transit the provider, more revenue can potentially be generated by inducing local peers to connect to remote ones, as shown at the right of (a). In this case, P2P users in the provider AS are directed to connect preferentially to users of its customer AS. In this case, the monetized paths result in a net traffic gain. However, because most demand is unlikely to be satisfied by intradomain users, the potential for gain is limited. Most of the P2P traffic will traverse the provider regardless.

Multi-homed: For multi-homed customers (b), a strategic provider benefits by competing for transit P2P traffic from its customer, increasing the chance that a monetized path will generate billable traffic. Each customer connection represents billable data that may have been routed through another provider. In this case, a strategic ISP should try to induce users in a customer AS to communicate with its P2P users rather than those of another provider.

In general, strategic ASes may not know detailed information about the routing policies of their customers, and so should be strategic with respect to all customers whether singly or multiply homed.

Efficiency loss: The strategic policy we apply avoids local peers, leading to longer AS path lengths on average. We quantify the extent of this increase by replaying our BitTorrent trace and choosing a target, strategic AS. For comparison, we record path lengths using the previously described shortest AS path and latency matching methods as well as random matching. When making strategic matches, we apply AS relationship data from CAIDA to predicted paths. While we report results for a single, large AS, similar trials for other ASes yield similar results.

Table 2 gives the overall efficiency loss for each method relative to shortest AS path and the overall efficiency gain relative to random matching. On the whole, strategic behavior results in an increase in interdomain paths by 72.6% relative to shortest path matching, with most of that increase resulting in revenue for the strategic ISP.

8. RELATED WORK

There have been various measurement studies of the traffic generated by P2P systems and evaluations of ways

	% reduction relative to random	% increase relative to shortest AS path
Shortest AS-path	51.4	0.0
Latency	40.7	21.9
Strategic	16.1	72.6

Table 2: Percentage change in AS path lengths relative to random and shortest AS path based peer matching.

to mitigate the resulting load. Saroiu et al. [17] and Gummadi et al. [6] examine the Gnutella and Kazaa workloads, document the increasing popularity of P2P systems, study the impact of caching and the potential for bandwidth savings of a locality aware mechanism. Sen and Wang [18] perform trace analysis of P2P traffic along the border routers of a single ISP and provide data that suggests that application-level traffic engineering might help.

Other researchers have studied the interactions between P2P systems and ISPs. Karagiannis et al. [8] study the impact of peer-assisted content distribution on ISPs. Also related to our work are efforts that examine whether ISPs and P2P systems can work together to perform traffic engineering and propose various solutions to achieve the necessary cooperation. For instance, Keralapura et al. [9] show that P2P systems could have an adverse impact on the stability of traffic engineering techniques currently used by ISPs in the absence of cooperation. Aggarwal et al. [1] and Bindal et al. [2] propose that ISPs use “oracles” to recommend peerings that are locality preserving.

The notion of ISPs manipulating existing protocols to accomplish traffic engineering goals or for strategic benefit has also received attention by researchers. Wang et al. report widespread use of path prepending to influence routing [19]. Mahajan et al. suggest additional protocol mechanisms by which ISPs coordinate their actions to overcome common inefficiencies in interdomain routing [13], but efficient outcomes depend on mutual trust between ISPs. More recently, Goldberg et al. [5] examine the incentives for ISPs to manipulate routing announcements to attract generic revenue-generating traffic and find that ensuring honesty likely requires substantial restriction in policy freedom. We apply similar ideas to the interaction between ISPs and P2P applications and quantify the potential for increasing revenue with measured workloads.

9. CONCLUSION

P2P systems and ISP operators currently have an adversarial relationship. The random matching of senders and receivers typical of current P2P file sharing networks generates significant amounts of interdomain traffic that increases costs for ISPs. In this paper, we have reported measurements of BitTorrent file sharing and network-level paths, examining the potential for locality-awareness to align the interests of users and ISPs. We find that while locality exists, simple heuristics are not sufficient to fully

exploit it and may hamper network robustness. Further, large ISPs that provide transit service derive revenue from today’s P2P file sharing patterns suggesting that systems designed to discover locality should not expect universal cooperation from ISPs.

10. REFERENCES

- [1] V. Aggarwal, A. Feldmann, and C. Scheideler. Can ISPs and P2P systems cooperate for improved performance? *ACM CCR*, 2007.
- [2] R. Bindal, P. Cao, W. Chan, J. Medval, G. Suwala, T. Bates, and A. Zhang. Improving Traffic Locality in BitTorrent via Biased Neighbor Selection. In *Proc. of IEEE ICDCS*, 2006.
- [3] D. R. Choffnes and F. E. Bustamante. Taming the Torrent: A practical approach to reducing cross-ISP traffic in P2P systems. In *SIGCOMM*, 2008.
- [4] M. Dischinger, A. Mislove, A. Haeberlen, and K. P. Gummadi. Detecting BitTorrent blocking. In *IMC*, 2008.
- [5] S. Goldberg, S. Halevi, A. D. Jaggard, V. Ramachandran, and R. N. Wright. Rationality and traffic attraction: Incentives for honest path announcements in BGP. In *SIGCOMM*, 2008.
- [6] K. P. Gummadi, R. Dunn, S. Saroiu, S. Gribble, H. Levy, and J. Zahorjan. Measurement, Modeling and Analysis of a Peer-to-Peer File-Sharing Workload. In *SOSP*, 2003.
- [7] T. Isdal, M. Piatek, A. Krishnamurthy, and T. Anderson. Leveraging BitTorrent for end host measurements. In *Proc. of PAM*, 2007.
- [8] T. Karagiannis, P. Rodriguez, and K. Papagiannaki. Should Internet Service Providers fear Peer-Assisted Content Distribution? In *IMC*, 2005.
- [9] R. Keralapura, N. Taft, C.-N. Chuah, and G. Iannaccone. Can ISPs take the heat from overlay networks? In *HotNets*, 2004.
- [10] J. Ledlie, P. Gardner, and M. Seltzer. Network Coordinates in the Wild. In *NSDI*, 2007.
- [11] Lightreading.com. P2P plagues service providers. http://www.lightreading.com/document.asp?doc_id=31767, Apr. 2003.
- [12] H. V. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. Anderson, A. Krishnamurthy, and A. Venkataramani. iPlane: An Information Plane for Distributed Services. In *OSDI*, 2006.
- [13] R. Mahajan, D. Wetherall, and T. Anderson. Mutually controlled routing with independent ISPs. In *NSDI*, 2007.
- [14] Packet Forgery by ISPs: A Report on the Comcast Affair. <http://www.eff.org/wip/packet-forgery-isps-report-comcast-affair>.
- [15] M. Piatek, T. Isdal, A. Krishnamurthy, and T. Anderson. Do incentives build robustness in BitTorrent? In *NSDI*, 2007.
- [16] P. Racz and Z. Despotovic. An ALTO service based on BGP routing information. <http://www.ietf.org/id/draft-racz-bgp-based-alto-service-00.txt>.
- [17] S. Saroiu, K. P. Gummadi, R. Dunn, S. Gribble, and H. Levy. An Analysis of Internet Content Delivery Systems. In *OSDI*, 2002.
- [18] S. Sen and J. Wang. Analyzing peer-to-peer traffic across large networks. *IEEE/ACM Trans. on Netw.*, 2004.
- [19] H. Wang, R. K. C. C. Dah, M. Chiu, and J. C. S. Lui. Characterizing the performance and stability issues of the as path prepending method: taxonomy, measurement study and analysis. In *SIGCOMM Asia Workshop*, 2005.
- [20] H. Xie, R. Yang, A. Krishnamurthy, Y. Liu, and A. Silberschatz. P4P: Provider portal for P2P applications. In *SIGCOMM*, 2008.
- [21] ZDNet. ISPs see costs of file sharing rise. http://news.zdnet.com/2100-9584_22-1009456.html, 2003.