

On the Feasibility of Switching ISPs in Residential Multihoming

Ahsan Habib
Siemens TTB Center
Berkeley, CA 94704
ahsan.habib@siemens.com

Nicolas Christin
Information Networking Institute
and CyLab Japan
Carnegie Mellon University
nicolasc@cmu.edu

John Chuang
School of Information
University of California,
Berkeley, CA 94720
chuang@ischool.berkeley.edu

Abstract—In theory, multihomed Internet hosts, that is, hosts simultaneously connected to multiple Internet service providers (ISP) should see increased access capacity, be able to circumvent possible last-mile congestion problems, and experience improved end-to-end quality of service (QoS). In practice however, the advantages one can gain from multihoming are highly dependent on the path switching mechanism used, that is, on dynamically deciding which ISP should be used as a first-hop. This paper is a first step toward understanding the trade-off between performance improvements multihoming can help achieve and the complexity of the decisions that must be made. We measure changes in end-to-end network layer metrics (loss, latency, jitter) over the different paths available from a multihomed host to a large population of Internet hosts. Our measurements indicate that 1) in over 60% of the cases, one only needs to reevaluate the service provided by each ISP every minute to improve the performance of a specific metric, and that 2) in approximately 85% of the cases, decisions to switch from one ISP to another can be treated independently of the service metric of interest. We conclude that multihoming could in practice result in noticeable performance improvements.

I. INTRODUCTION

Recent measurement studies, such as [1], [2], indicate that connecting an Internet host to multiple service providers (ISPs), or *multihoming*, might considerably improve end-to-end response times experienced by the host. With the availability of broadband connections in most households and the increasing number of wireless networks, it is not hard for a user to simultaneously connect to multiple networks. Quite intuitively, by allowing the multihomed host to choose between two or more possible “first hops,” multihoming indeed provides a way to circumvent most of the potential last-mile congestion problems.

Maybe less obvious is the fact that, in addition to having a choice of first hops, multihomed hosts could indirectly benefit from the differences in peering relationships among various service providers, to use significantly disjoint routes to a destination (path diversity). Simply stated, multihomed hosts might be able to dynamically avoid most points of congestion in the network, by taking advantage of the choice among the different ISPs available. We note that not all applications need to choose the best among the available paths. For example,

a file sharing or streaming application might prefer to use all available paths to increase its throughput. However, an application can constantly monitor all available paths and migrate its session from a congested path to a non-congested one, if necessary [3].

This paper explores from a practical standpoint whether, by utilizing path diversity, multihoming is a potentially attractive solution for hosts that require high quality service. We first verify by measurements that multihoming indeed provides more than first hop path redundancy. We then investigate two questions: 1) to which extent end-to-end service can be improved by taking advantage of multihoming? and 2) how fast should a multihomed host be able to detect changes in the network to use multihoming efficiently?

The answers to both questions appear closely tied to the path switching mechanism used, that is, to the mechanism in charge of deciding, possibly for every single packet, which ISP should be used as a first hop. In this paper, rather than designing and evaluating a specific path switching mechanism, our strategy is to measure changes in end-to-end network layer metrics (loss, latency, jitter) over the different paths available from a multihomed host to a large population of Internet hosts. The measurements gathered allow us to quantify the improvement in its end-to-end service a multihomed host can experience compared to a single-homed host, when using an ideal path switching mechanism that always picks the right ISP.

Such an ideal mechanism needs to know in advance which path offers a better service, which, because network conditions constantly change, may be difficult. Thus, to have a better grasp of the practical benefits multihoming can yield, we also assess the trade-off between the performance improvements achievable with multihoming, and the reaction times a practical mechanism might need to infer and adapt to changes in network conditions. We complement our study by measuring how the choice of a given path to optimize for a specific service metric (e.g., latency) impacts other service metrics (e.g., loss, jitter).

There are several measurement based studies related to path diversity and benefits of multihoming. Akella et al. [1], [2] primarily focus on enterprise multihoming for Web access by providing measurements from the Akamai network. Nayak [4] measures path diversity from four different providers (Exodus,

UUNet, Sprint, AT&T). Teixeira et al. [5] study diversity within the Sprint network by measuring paths between points of presence, and Tao et al. [6] show that path diversity is effective to reduce end-to-end losses, in both multihomed and overlay networks. Different from these related works, we consider *residential multihoming* in the context of peer-to-peer connections, and discuss the frequency at which a multihomed host has to switch between ISPs to observe noticeable improvements in service, as well as the potential correlations between different service metrics. We have studied the benefits of simultaneously utilizing all available paths of a multihomed host in [3] especially for connections with long duration such as streaming and file sharing.

In this paper, we study the feasibility of switching among ISPs to take the benefits of the better ISP at any given time. This approach is suitable for connections with long duration such as file transfer or streaming applications that need to switch from a congested ISP to a non-congested one, if available. We show that 1) in over 60% of the cases, one only needs to reevaluate the service provided by each ISP every minute to improve the performance of a specific metric, and that 2) in approximately 85% of the cases, decisions to switch from one ISP to another can be treated independently of the service metric of interest. We conclude that multihoming could in practice result in noticeable performance improvements. Thompson et al. [7] designs a scheduler that analyzes the end-users' networking behaviors to achieve better performance at flow level. Our work is complementary to the work by Thomson et al. By providing insights into the time granularity at which one needs to make switching decisions as well as the appropriate metrics to use for scheduling, our measurement analysis can inform the design of a scheduler of the type proposed by Thomson et al.

While our study presents the main limitation of gathering measurements from just a few multihomed residential hosts, we believe that the experimental setup chosen (DSL and cable, major metropolitan area) is characteristic of the connectivity typically available to a majority of residential users in the United States and Europe, and can therefore provide valuable insights.

The remainder of this paper is organized as follows. We introduce our measurement methodology and tools in Section II. In Section III, we show that, by providing path diversity, multihoming can enhance the end-to-end service hosts can experience. We then, in Section IV, discuss the constraints practical mechanisms need to satisfy to take advantage of multihoming, before drawing brief conclusions in Section V.

II. MEASUREMENT METHODOLOGY

In an effort to complement related work more axed on *enterprise multihoming* for content distribution, our focus in this study is on *residential multihoming*, and on end-to-end measurements between similar Internet hosts. Our main goal is to describe general trends, and avoid *misconceptions* [8]: we want to be certain that what we are trying to measure is indeed what we do measure. Our choice of test-bed, dataset,

and measurement tools tries to help us reduce the risk of misconceptions, as we discuss next.

Testbed. We use two different multihoming testbeds in our experiments. Each testbed consists of a set of two hosts connected to two ISPs via an EDIMAX BR-6524 [9] broadband router. The advantage of the EDIMAX router is that it merely provides connectivity to multiple ISPs, without attempting to perform any advanced function such as load balancing or traffic shaping. Thus, contrary to other choices of hardware, e.g., load balancing systems designed for enterprise networks [10], the router we use should not impact the measured data. The ISPs chosen are the three largest providers in the San Francisco (East) Bay Area: two DSL service providers (SBC and Earthlink) and one cable service provider (Comcast). In the first testbed, the hosts are multihomed via SBC and Comcast. SBC and Earthlink DSL providers are used in our second testbed.

While our chosen experimental setup may appear relatively simplistic, we contend that it is quite typical of the connectivity available to a vast majority of residential users in the United States and Europe. We use two machines behind the router (instead of a single host) to be able to run parallel measurements.

Dataset. We measure end-to-end loss, latency, and jitter between our residential testbed and a set of Internet hosts (destinations), consisting of a set of 35,868 KaZaA clients, a set of 49,742 Gnutella clients, and a set of 109,915 Overnet clients.¹ These peer-to-peer clients are distributed all over the world, so that the measurements we gather should not be impacted by the specifics of local or regional networks.

The motivation for using IP addresses bound to peer-to-peer clients lies in the growing interest from residential users in peer-to-peer applications such as file sharing, voice-over-IP, or peer-to-peer media streaming [12]. The user-perceived experience for Web measurements in the context of residential multihoming is studied in [3]. As the Web sessions are usually very short in duration, it might not be feasible to switch ISPs in the middle of a Web session. A recent history based performance data can be used in selecting paths during Web access.

On the other hand, inexistent IP addresses used for poisoning of file sharing networks [11] might be included in the dataset. In addition, some clients are not traceable, and some clients sporadically go offline. Such unreachable hosts do not alter our measurements, since they do not produce any results.

Tools. Similar to [4], we use *traceroute* data to analyze topological path diversity due to residential multihoming. We first trace the end-to-end paths from our test-bed to each destination through both ISPs. We note that network interfaces in routers are often associated with multiple IP addresses (IP aliasing), which can introduce errors in topology generation from *traceroute* data. We resolve IP aliases using *sr-ally* [13]. *sr-ally* executes an IP identifier-based pairwise alias test to

¹The IP addresses of these clients were obtained through the measurement apparatus described in [11].

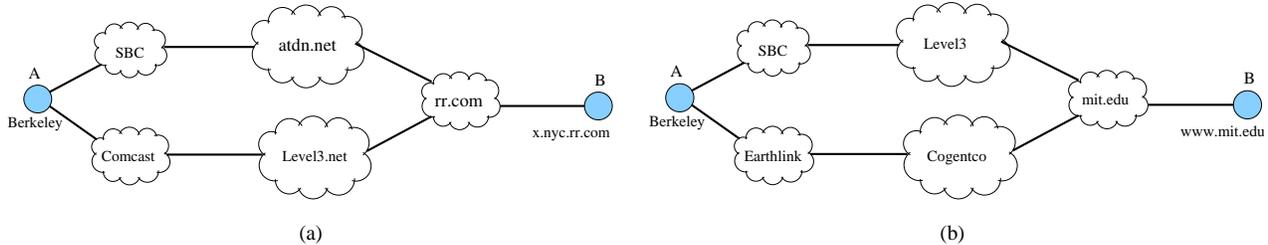


Fig. 1. **Path diversity.** (a) User A has two separate paths to reach user B $x.nyc.rr.com$. On one path, packets traverse through sbc and $atdn.net$ to reach $rr.com$ before reaching the destination. On the other path, packets traverse through $comcast$ and $level3.net$ to reach $rr.com$ before reaching B . (b) A similar example of path diversity for SBC and Earthlink providers.

discover whether two IP addresses belong to interfaces on same machine. Essentially, IP-identifier based alias resolution seeks evidence that the two IP addresses share a single IP-id counter. If packets generated by two different IP addresses have in-order IP identifiers, those IP addresses are likely aliases [13].

To measure latency, loss, and jitter on both ISPs, we send probe packets simultaneously through two network interfaces so that the probe packets travel through both ISPs at the same time. We use *ping* to measure the round-trip time (latency) and packet loss ratio. We fork two processes to run two instances of *ping* simultaneously, which are synchronized with the system clock. The packets sending times are synchronized at a few milliseconds granularity. As suggested in [14], we compute the jitter as the Inter Quartile Range (IQR) of the frequency distribution of the round-trip time estimated by *ping*.

Time window. The measurements are collected over an eight-month time period (December 2004 through May 2005 and December 2005 through February 2006). The relatively large measurement window allows us to limit the impact of time-of-the-day or day-of-the-week effects. The role of seasonal patterns might play in the collected data appears negligible, thanks to the geographical dispersion of the hosts used for measurements.

III. MULTIHOMING AND PATH DIVERSITY

The key insight into the potential benefits of multihoming is that not only it provides first hop path redundancy, but more generally it offers highly diverse end-to-end paths both in topology and network layer metrics such as latency, loss, and jitter, as we show in this section. The quality of service of an application might depend on one or a combination of these metrics. When an application needs to choose one path among the all available ones due to multihoming, it can rank the paths based on the metric of its preference.

A. Topological diversity

We discuss the topological path diversity, by analyzing the end-to-end paths over multiple ISPs and the path segments that are shared among all paths. Clearly, it is desirable to have no (or low) overlap among the alternative paths provided by multihoming.

Figure 1 shows specific instances of path diversity due to multihoming. User A at Berkeley has two separate paths to reach user B at New York. On one path, packets traverse through sbc and $atdn.net$ to reach $rr.com$ before reaching the destination. On the other path, packets traverse through $comcast$ and $level3.net$ to reach $rr.com$ before reaching B . In our measurements, the paths have 4-7 hops overlap (with an average value of 5.99) for most of the hosts, whereas the average end-to-end path length is 19.89 in SBC and 19.15 in Comcast. Thus, on average, one third of the hops are overlapped among the alternate paths between a source-destination pair. We now quantify the path overlap for each individual host.

We define the metric Single Source Path Overlap (SSPO) to express the path overlap between a multihomed user and any host in the network. As shown in Figure 1, the path overlap occurs for a multihomed host at the edge network with which the source node is connected to. SSPO is an estimation of the expected fraction of hop overlap, which is the ratio of the shared hops to the total non-shared hops of all paths. Let H_i be the total number of hops through ISP i to reach a destination, and E be the total number of edges of the tree which is constructed from a host to a multihomed user. A general definition of single source path overlap $SSPO$ for k multihoming (connected to k isps) is as follows:

$$SSPO = \frac{\sum_{i=1}^k H_i - E}{E}. \quad (1)$$

The value of SSPO varies in the range from 0 to 1, where 0 represents no overlap and 1 is 100% overlap. A similar metric is used for measuring path diversity in enterprise multihoming [1]. Figure 2 shows the cumulative distribution function (CDF) of the expected fraction of hops that are overlapped in end-to-end paths from our two-homed test-bed to the destination hosts. The SSPO is less than 0.30 for 80% of the hosts and less than 0.5 for 99% of the hosts. The average value of SSPO is only 0.20. This experiment confirms that multihoming is not one hop path redundancy (90-95% overlap), instead, two “almost” non-overlapped paths exist to reach a large number of destinations for a multihomed user. When one path is congested, an application can still reach a host through the other path, provided that the shared path is not

congested. SSPO is a very useful metric in selecting available suppliers in a file transfer or streaming session, which may last for minutes or even hours and any ISP (probably not all) may experience congestion during this time. It is desirable to select suppliers with low SSPO with the receiver host because the file transfer or streaming application would be able to switch from a congested ISP to a non-congested ISP when necessary. If all ISPs experience congestion, it is necessary for a multihomed host to replace the supplier with a different one.

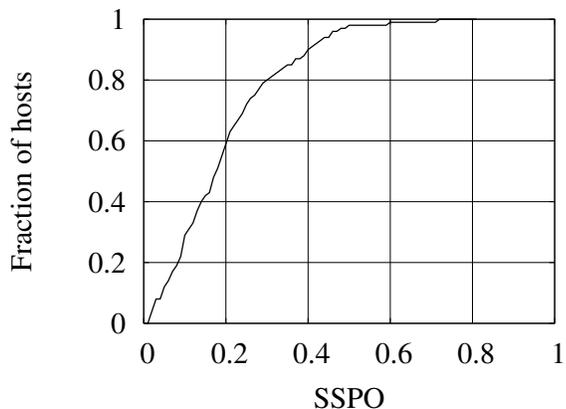


Fig. 2. **Cumulative distribution function of SSPO to the destination hosts.** The expected fraction of hop overlap is less than 0.30 for 80% of the hosts. The average SSPO is 0.20.

B. Latency, loss, and jitter

Next, we quantify the benefits of multihoming by measuring differences in latency, loss ratio, and jitter between SBC and Comcast. As discussed in Section II, all measurements are done concurrently on both ISPs. We refer to SBC as ISP_1 and Comcast as ISP_2 . The metrics measured via SBC and Earthlink have similar properties to the metrics measured in SBC and Comcast. Therefore, we present only one set of data in this section. One exception is that latency difference is higher in SBC and Earthlink comparing to SBC and Comcast. The switching decision in Section IV uses data from our both measurement test-bed.

Latency. The average RTTs of end-to-end paths are 251.08 ms and 264.74 ms for ISP_1 and ISP_2 respectively to reach our large set of KaZaA, Gnutella, and Overnet hosts. Each sample point is an average over 10 ping packets. The 50th percentile and 90th percentile are 181.1 ms and 319.7 ms for ISP_1 and 187.9 ms 350.3 ms for ISP_2 respectively and the maximum RTT goes as high as several seconds for both ISPs. The variation of latency captures the heterogeneity of the hosts that reside all over the world.

Even though, ISP_1 offers a slightly shorter path, on average, than ISP_2 , one ISP does not provide low end-to-end latency for all hosts. Roughly half of the hosts are better off with ISP_1 and one third of them are better off with ISP_2 in terms of providing low end-to-end RTT. An application can

select the right ISP that reduces the end-to-end latency to reach a destination.

To quantify how much latency a user can reduce using multihoming, we plot the CDF of latency difference $|RTT_1 - RTT_2|$ in Figure 3, where RTT_1 and RTT_2 represent round trip time via ISP_1 and ISP_2 respectively. For 30% of the cases, the benefit is not significant (≤ 5) ms. However, we can reduce end-to-end latency by at least 20 ms to reach more than 20% of the destinations. The improvement is at least 40 ms for 10% of the hosts. We compute latency reduction as a percentage of the end-to-end latency. It shows that, a user can reduce at least 20% of the original end-to-end latency to reach 15% of the destinations. In short, multihoming can effectively be used to reduce the latency to reach a large number of hosts.

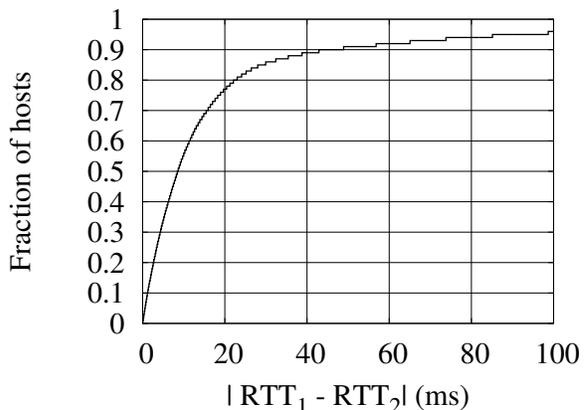


Fig. 3. **Cumulative distribution function of latency difference in both ISPs.** The end-to-end latency can be reduced by at least 20 ms by selecting the proper ISP in 20% of the cases. The improvement is at least 40 ms for 10% of the destinations.

Loss. We estimate the number of packets lost before reaching the P2P hosts. For each host, we send 100 ping packets and count how many of them are lost. We express loss ratio as a fraction of packets that are lost out of 100 packets. In our experiments, 77% hosts experience 0 packet loss on both ISPs and 81% hosts experience identical packet loss (both zero and non-zero loss) on both ISPs. However, a significant number of the hosts experience high losses. The average, 50th percentile and 90th percentile of non-zero loss ratios are 11.9%, 2%, and 33.4% for ISP_1 and 10.8%, 2%, and 27% for ISP_2 respectively. In other words, our measurements confirm that, while losses are an infrequent event in the Internet, their magnitude can be problematic, which can significantly affect the service experienced by applications.

In the context of multihoming, we are interested to see whether the *non-zero* loss events on both ISPs are correlated, i.e., if, when one ISP experiences high loss the other ISP also experiences high loss. The scatter plot in Figure 4(a) shows that loss ratios between ISPs present a low correlation. Figure 4(b) is the CDF of the difference of non-zero loss ratio, and shows that more than 90% hosts that experience loss on one ISP can reduce the loss ratio by 10% or less on the other

ISP. Therefore, loss sensitive applications can reduce overall loss ratio by using multihoming.

Jitter. Jitter captures the variation of latency over time, and is a crucial service metric for delay sensitive applications such as VoIP and video streaming. In our experiments, the average jitter is 20.87 ms on ISP_1 and 24.37 ms on ISP_2 . Each sample point is obtained as the Interquartile Range (IQR) of frequency distribution of 100 RTT samples from each ISP. The 50th percentile and 90th percentile are 5 ms and 34.9 ms for ISP_1 and 10 ms 36.45 ms for ISP_2 respectively. Therefore, the users experience, on average, lower jitter on ISP_1 than on ISP_2 . Like other metrics, one ISP cannot consistently provide low jitter for all hosts. Figure 5 shows the CDF of jitter difference to the destinations. For 50% of the hosts, the jitter improvement is 5 ms or less. The improvement is significant for the rest of the hosts (5-15 ms for 40% of the hosts, and more than 15 ms for 10% of the hosts). Thus, if an application prefers low jitter, multihoming can select the proper ISP to reduce the jitter. In the next section, we discuss how each of the network layer metric can be used in path switching decision in a multihoming environment.

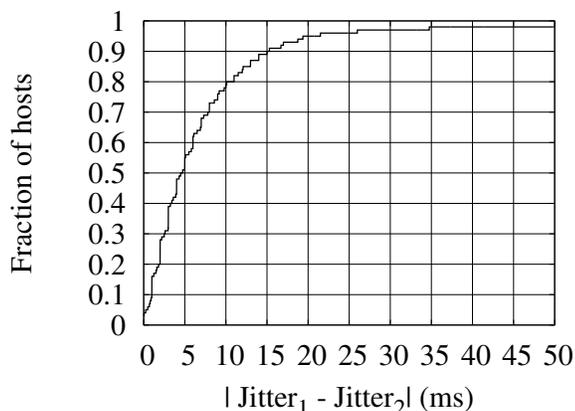


Fig. 5. Cumulative distribution function of jitter difference in both ISPs. More than 40% hosts can reduce jitter by 5-15 ms and 10% hosts can reduce jitter by more than 15 ms.

IV. PATH SWITCHING

We have shown that by selecting the right ISP a multihomed user can reduce latency, loss, and jitter. Therefore, an application needs to decide which ISP (first hop) to select to reach a destination, and how often to switch among ISPs to capture the benefits of path diversity. If an application selects an ISP based on network metric x and x changes too frequently, the switching overhead will be high. On the other hand, if x does not change too often then a simple measurement-based switching algorithm will be able to effectively utilize the benefits of multihoming without incurring high overhead.

In this section, we first show that an application does not need to frequently change ISPs to receive the best possible performance. Then, we show that the decision to switch can be made independently on any network metric, and that most

of the time a decision to switch ISPs based on one metric is consistent with the decision that would be made by considering the other metrics.

A. Path switching frequency

We assess how often an application should switch ISPs to improve the experienced latency, loss, and jitter. For example, any time if $RTT_1 - RTT_2 > \epsilon_R$, it is better to select ISP_2 , where ϵ_R is a threshold that determines a tangible gain of latency. Similarly, we define switching thresholds ϵ_L and ϵ_J for loss and jitter respectively. For each host, we send 5 ping packets every second for an hour. Switching decision is made every second based on the average value of the samples for each ISP. Therefore, for each source-destination pair, 3600 switching decisions are made. We compute the intervals (“switching time”) when switching from one ISP to another improves each of the network layer metrics for a given destination host. We compute the number of switching events and switching intervals for each source-destination pair. If a host switches from ISP_1 to ISP_2 and stays at ISP_2 for several consecutive seconds, it counts as one switching event. Unless, it switches back to ISP_1 again, the number of switching decisions remain constant. For each destination and each network layer metric, we compute the minimum, average, maximum, 95th percentile, and 99th percentile of switching intervals.

Figure 6 shows average and 95th percentile of switching intervals based on latency, loss, and jitter respectively. The X-axis is the average or 95th percentile of switching intervals in second and the Y-axis the fraction of total hosts that experience the switching interval. The graphs give several insights about path switching. First, switching based on latency or jitter happens more frequently than switching based on loss because all ISPs experience low loss. Second, switching greedily result in considerable overhead compared to switching based on a threshold ($\epsilon > 0$), even for very small thresholds. Therefore, switching between ISPs should take place only when a tangible gain can result. We find that $\epsilon_R \geq 20$ ms, $\epsilon_L \geq 0$, and $\epsilon_J \geq 5$ ms can be a good choice of thresholds to provide good performance at a reasonable overhead. For example, $\epsilon_R \geq 20$ ms results on average in switching ISPs 12 times per hour, whereas $\epsilon_R = 0$ causes in the order of 300 changes of ISPs. Similarly, $\epsilon_J \geq 5$ ms yields 50 changes per hour, whereas $\epsilon_J = 0$ causes 155 switches in ISPs per hour. With positive thresholds, the average switching interval based on latency, loss, or jitter is at least one minute for 60% of the hosts for SBC and Comcast, as shown in Figure 6(a), Figure 6(c), and Figure 6(e). For 80% of the hosts, the average switching interval is 23 seconds or more regardless of the metric considered. The switching interval is even longer for a host that is multihomed via SBC and Earthlink (Figure 6(b), Figure 6(d), and Figure 6(f)), i.e., when one a user selects SBC for its performance, it does not need to switch to Earthlink for several minutes or vice versa.

Next, we measure the fraction of destinations to which a decision to switch between ISPs results in end-to-end service

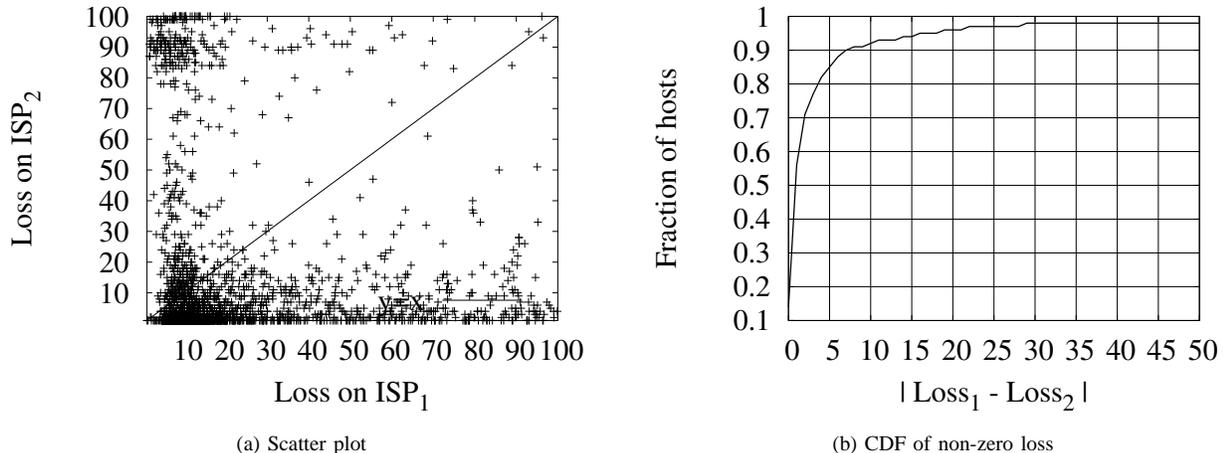


Fig. 4. **Loss measurements.** The plots show (a) a weak correlation between packet loss, and (b) that more than 90% hosts that experience loss on one ISP can reduce the loss ratio by 10% or less on the other ISP.

improvement. Figure 7 shows that multihoming can be effective in improving the service to most destinations even with reasonably large switching intervals. In fact, 95% of the hosts can effectively utilize multihoming to reduce latency as high as 50 ms when switching interval is 20 seconds. In summing, a multihomed user does not need to switch often between ISPs to reap most of the benefits multihoming can provide.

B. Discordance between metrics

Now, we investigate whether the decision to switch among ISPs based on one metric such as latency will strongly conflict with the switching decision based on other metrics such as loss or jitter. To capture this disagreement, we define the *discordance ratio* (D) as follows:

$$D_{x,y} = \frac{\text{time during switching on } x \text{ and } y \text{ disagrees}}{\text{total time}} \quad (2)$$

The discordance ratio takes real values between 0 and 1. A low discordance ratio means that the decision can be made independently of metric x or y and switching to an ISP based on x will not negatively impact metric y or vice versa. On the other hand, a high discordance ratio indicates that if an application switches to an ISP based on x , a host will experience high value of y on that ISP. In such a case, the host should prioritize its metrics of interest in ISP selection. Moreover, a reduction in one metric should not be negligible compared to an increase in other metrics.

The discordance ratio in Eq. (2) is related to Kendall's tau [15], which is a statistical measurement of association between two bivariate variables. To measure the association, both concordant and discordant pairs of two variables are computed in Kendall's tau. We are only interested in discordant pairs because the application will experience poor service if the metrics constantly disagree with each other in switching decision.

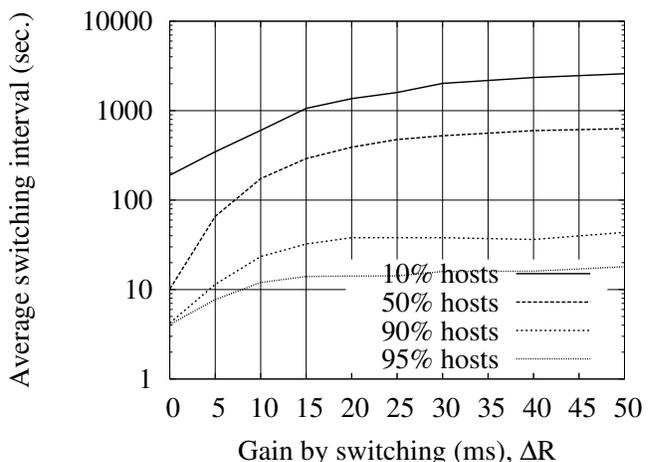
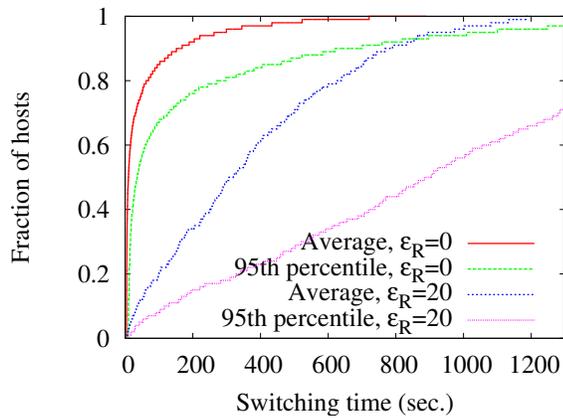


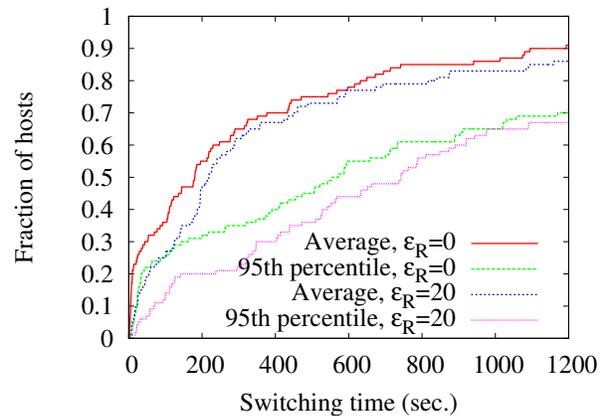
Fig. 7. **Switching interval and latency.** 95% of the hosts will reduce the latency by as high as 50 ms if switching time is 20 seconds. Therefore, it is not necessary to switch frequently between ISPs to effectively utilize the benefits of multihoming.

We compute pairwise discordance ratio for latency-loss, loss-jitter, and jitter-latency and plot the CDF of the discordance ratios in Figure 8. $D_{\text{latency,loss}}$ and $D_{\text{jitter,loss}} < 0.1$ for 90% of the hosts for both ISP-pair (Figure 8(a-d)), i.e., if an application switches to an ISP based on latency or jitter, it is highly likely that the loss experienced with this ISP will be modest as well. This is because loss is an infrequent event in the Internet and both ISPs experience similar loss ratio most of the time. Therefore, decisions based on latency or loss can be made independently of each other. The same property holds for jitter and loss.

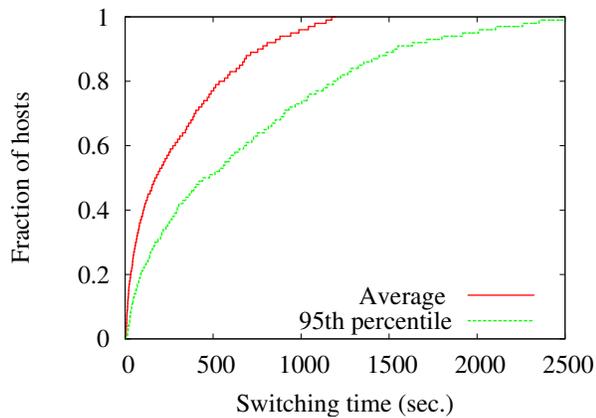
The discordance ratio $D_{\text{latency,jitter}} < 0.1$ for 85% of the hosts when $\epsilon_R=20$ ms and $\epsilon_J=5$ ms (Figure 8(e)(f)). Switching greedily ($\epsilon = 0$) causes a lot of unnecessary switching that



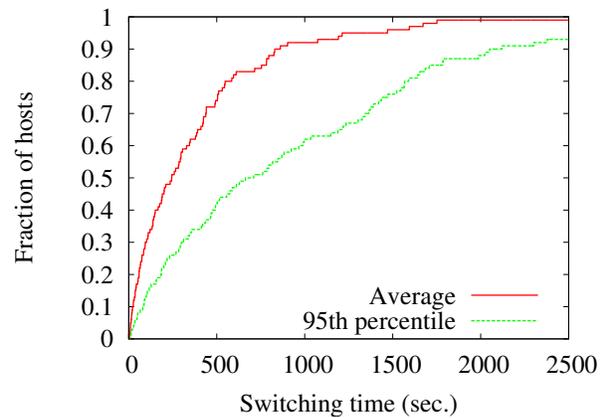
(a) Latency (SBC & Comcast)



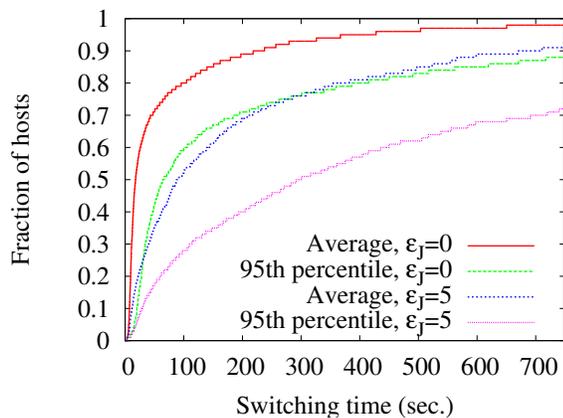
(b) Latency (SBC & Earthlink)



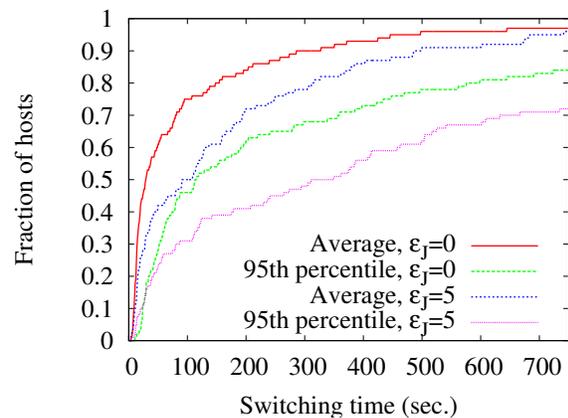
(c) Loss (SBC & Comcast)



(d) Loss (SBC & Earthlink)

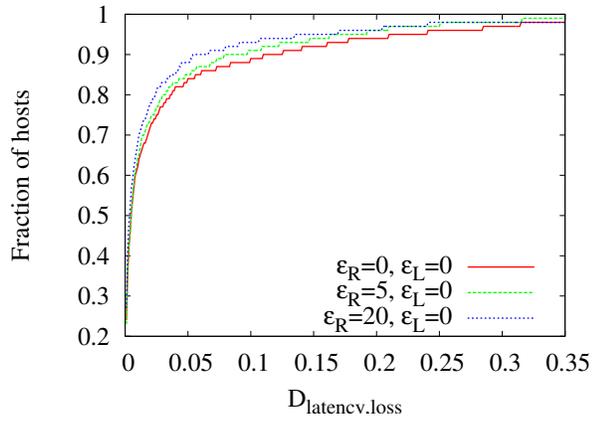


(e) Jitter (SBC & Comcast)

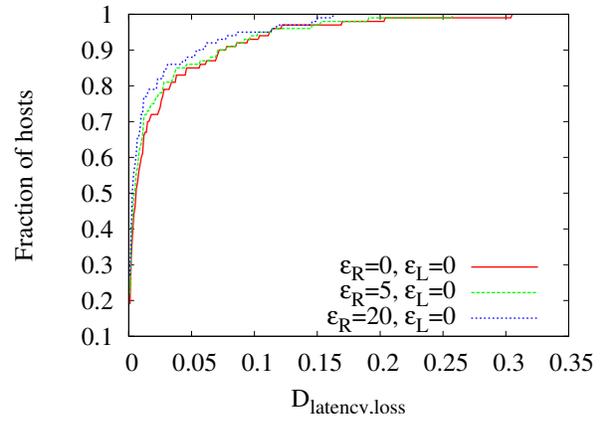


(f) Jitter (SBC & Earthlink)

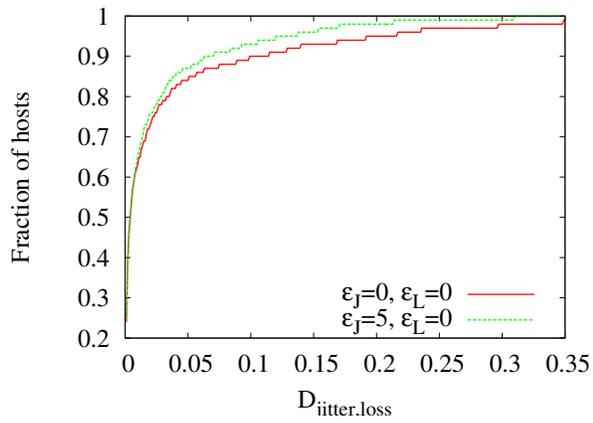
Fig. 6. **Average and 95th percentile of switching intervals based on latency, loss ratio, and jitter.** A positive threshold in switching decision reduces a lot of unnecessary switching among ISPs and provides tangible gain to the applications. Switching interval is even higher for SBC and Earthlink comparing to SBC and Comcast. Therefore, a user switches ISP less often when it is multihomed via SBC and Earthlink comparing to SBC and Comcast.



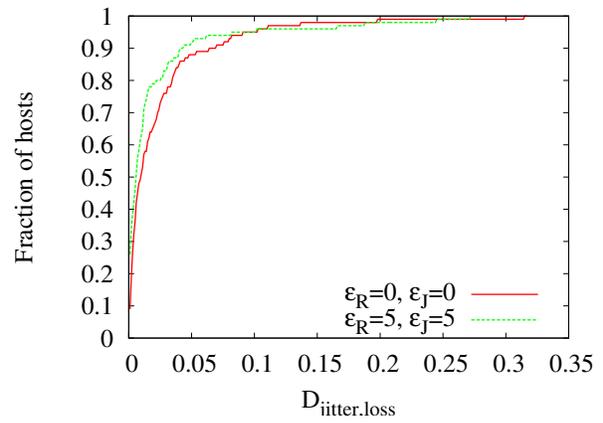
(a) SBC & Comcast



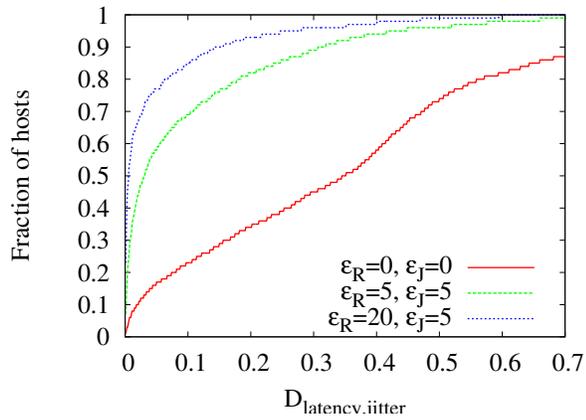
(b) SBC & Earthlink



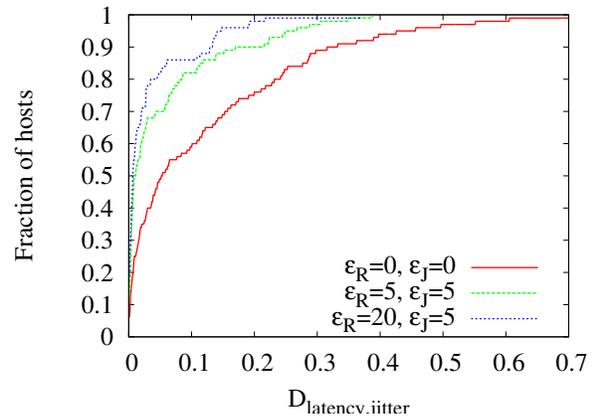
(c) SBC & Comcast



(d) SBC & Earthlink



(e) SBC & Comcast



(f) SBC & Earthlink

Fig. 8. **Discordance ratios.** With threshold $\varepsilon > 0$, Switching based on latency or loss, jitter or loss, and latency or jitter disagrees only 10% of total time for more than 85% of the hosts when the host is multihomed via SBC and Comcast. The disagreement is even lower for SBC and Earthlink.

does not improve the performance significantly and causes frequent disagreement in switching decision. These pairwise discordance ratios show that, using a positive threshold, an application can switch independently on latency, loss, or jitter, and at the same time the application can obtain tangible benefits from multihoming. When there is a conflict, the application has to pick the metric based on its preferences to select the right ISP. For example, a VoIP application or bulk TCP data transfer application might put high priority to loss and a video streaming application might put high priority to jitter.

V. CONCLUSION

We assess the benefits of residential multihoming, and the complexity of the mechanism needed to harvest the advantages provided by multihoming. Our measurements show that multihoming provides highly diverse paths with less than 30% overlap among the different paths over the available ISPs. A host can utilize this path diversity to reduce latency, loss, or jitter substantially, thereby improving application performance. Our measurements also indicate that, in over 60% of the cases a multihomed host can obtain tangible gains only by reevaluating the service provided by each ISP every minute, regardless of the service metric considered. Furthermore, in approximately 85% of the cases, decisions to switch from one ISP to another can be treated independently of the service metric of interest. In summing, one can envision that practical mechanisms could rely on multihoming to provide noticeable end-to-end performance improvements. The design of such mechanisms is an avenue we are currently exploring [16].

REFERENCES

- [1] A. Akella, B. Maggs, S. Seshan, A. Shaikh, and R. Sitaraman, "A measurement-based analysis of multihoming," in *Proceedings ACM SIGCOMM*, Karlsruhe, Germany, Aug. 2003.
- [2] A. Akella, S. Seshan, and A. Shaikh, "Multihoming performance benefits: An experimental evaluation of practical enterprise strategies," in *Proceedings USENIX Annual Technical Conference*, Boston, MA, June 2004.
- [3] A. Habib and J. Chuang, "Improving application QoS with residential multihoming," *Computer Networks*, 2007 (to appear).
- [4] K. Nayak, "Measuring provider path diversity from traceroute data," in *ISMA Winter Workshop*, San Diego, CA, Dec. 2001.
- [5] R. Teixeira, K. Marzullo, S. Savage, and G. M. Voelker, "In search of path diversity in ISP networks," in *In Proceedings of ACM/USENIX IMC*, Miami Beach, Florida, Oct. 2003.
- [6] S. Tao, K. Xu, Y. Xu, T. Fei, L. Gao, R. Guerin, J. Kurose, D. Towsley, and Z. Zhang, "Exploring the performance benefits of end-to-end path switching," in *Proceedings of IEEE ICNP*, Berlin, Germany, Oct. 2004.
- [7] N. Thompson, G. He, and H. Luo, "Flow scheduling for end-host multihoming," in *Proceedings of IEEE INFOCOM*, 2006.
- [8] V. Paxson, "Strategies for sound Internet measurement," in *Proceedings of ACM/USENIX IMC'04*, Taormina, Italy, Oct. 2004, pp. 263–271.
- [9] "Edimax, <http://www.edimax.com/>," 2005.
- [10] "Rether networks inc., <http://www.rether.com/>," 2005.
- [11] N. Christin, A. Weigend, and J. Chuang, "Content availability, pollution and poisoning in peer-to-peer file sharing networks," in *Proceedings of ACM EC'05*, Vancouver, BC, Canada, June 2005.
- [12] M. Hefeeda, A. Habib, B. Botev, D. Xu, and B. Bhargava, "PROMISE: Peer-to-peer media streaming using CollectCast," in *Proceedings ACM Multimedia '03*, Berkeley, California, Nov. 2003.
- [13] N. Spring, R. Mahajan, and D. Wetherall, "Measuring ISP topologies with rocketfuel," in *Proceedings ACM SIGCOMM*, Pittsburgh, Philadelphia, Aug. 2002.
- [14] L. Cottrell, W. Matthews, and C. Logg, "Tutorial on internet monitoring & PingER at SLAC," 2001, <http://www.slac.stanford.edu/comp/net/wan-mon/tutorial.html>.
- [15] M. Hollander and D. A. Wolfe, *Nonparametric statistical methods*. John Wiley & Sons, 1999.
- [16] A. Habib, N. Christin, and J. Chuang, "Taking advantage of multihoming with session layer striping," in *Proceedings Global Internet Symposium (Global Internet '06)*, Barcelona, Spain, April 2006, pp. 102–107.