

Understanding Network Delay Changes Caused by Routing Events

Himabindu Pucha¹, Ying Zhang², Z. Morley Mao², and Y. Charlie Hu¹

¹School of ECE, Purdue University, West Lafayette, IN 47907

²Department of EECS, University of Michigan, Ann Arbor, MI 48109

ABSTRACT

Network delays and delay variations are two of the most important network performance metrics directly impacting real-time applications such as voice over IP and time-critical financial transactions. This importance is illustrated by past work on understanding the delay constancy of Internet paths and recent work on predicting network delays using virtual coordinate systems. Merely understanding currently observed delays is insufficient, as network performance can degrade not only due to traffic variability but also as a result of routing changes. Unfortunately this latter effect so far has been ignored in understanding and predicting delay related performance metrics of Internet paths. Our work is the first to address this shortcoming by systematically analyzing changes in network delays and jitter of a diverse and comprehensive set of Internet paths. Using empirical measurements, we illustrate that routing changes can result in roundtrip delay increase of converged paths by more than 1 second. Surprisingly, intradomain routing changes can also cause such large delay increase.

Given these observations, we develop a framework to analyze in detail the impact of routing changes on network delays between end-hosts. Using topology information and properties associated with routing changes, we explain the causes for observed delay fluctuations and more importantly identify routing changes that lead to predictable effects on delay-related metrics. Using our framework, we study the predictability of delay and jitter changes in response to both passively observed interdomain and actively measured intradomain routing changes.

Categories and Subject Descriptors: C.2.3 Computer Communication Networks: Network Operations

General Terms: Measurement, Performance

Keywords: Network delay changes, Network jitter changes, Routing dynamics, Routing events.

1. INTRODUCTION

Network delays and delay variations are two of the most important network performance metrics directly impacting several wide-area network applications ranging from real-time applications such

as voice over IP [42] and time-critical financial transactions, to multicast streaming applications [13, 36, 5, 6], locality-aware systems for redirection and server selection [41], proximity-aware DHTs [26, 40], positioning systems [17, 8, 11], and overlay routing systems [3, 38]. The wide array of applications sensitive to network delays and its variations underscores the importance of understanding when, by how much, and why network delays vary.

Network performance between a given source and destination host on the Internet can transform drastically over time. As an example, Figure 1 depicts the measured latency (from trace 1, Section 2.1) on three wide area Internet paths over a period of 2-3 days. The figure shows that apart from short-term fluctuations in latency, significant long term latency changes can also occur and persist for some time.

There are two primary factors that contribute to the above significant performance changes: *network topology changes* and *traffic fluctuations*. Network topology modifications such as link failures or traffic engineering are manifested as routing changes affecting the path reaching the given destination. Traffic fluctuations are caused by behavioral modifications of traffic sources, *e.g.*, flash crowd events. Given these two fundamental causes for network performance change, it is critical to understand their impact on the application performance and the predictability of their effect to achieve better network performance guarantees. User behavior contributes significantly to traffic fluctuations and is challenging to model given the presence of unexpected behavior such as DDoS attacks. In contrast, routing changes, in particular at the interdomain level, can be passively observed and directly used for predicting network performance of the resulting network path. The ability to perform such prediction enables useful applications such as host-based proactive mitigation against performance degradation and network-based performance-sensitive route selections.

In this work, we focus on understanding the change in network delay and jitter properties of the stable network path after the routing event has converged relative to the delay performance prior to the routing change. Previous work [22, 31, 32] has studied the performance degradation during routing changes or mostly the transient effects of routing convergence on application performance. Complementary to this work, we perform a study to characterize the change in delay and jitter caused by adopting a new stable path after routing convergence, *i.e.*, the forwarding path has stabilized. Such knowledge is helpful in determining the necessary response to expected network delay changes and can help route selection decisions by taking into consideration expected delay degradation. Intuitively the path chosen after the routing change can be significantly different in its network properties compared to that before the routing change, accounting for the performance deviations.

Existing work [39] in analyzing the constancy of Internet path

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SIGMETRICS'07, June 12–16, 2007, San Diego, California, USA.

Copyright 2007 ACM 978-1-59593-639-4/07/0006 ...\$5.00.




Figure 1: Network delay changes over time for sample source-destination pairs.

properties and predicting network delays based on virtual coordinate systems [18, 9] has so far ignored the effect of routing changes on network performance. However, routing changes are inevitable due to the continuous state of flux of the entire Internet caused by physical and configuration modifications. In this work, we address this shortcoming by systematically examining a comprehensive set of routing events occurring across diverse network locations and their effect on delay-related metrics of the communication between a given source and destination host.

Our work is also motivated by recent research proposals on revealing path diversity through protocol changes [33, 34] or new routing services [15, 10] and commercial products [7, 1] on edge-based load balancing for multihomed networks. Given the inherent path diversity [30] on the Internet, our work provides guidelines for path selection taking into consideration the impact on delay behavior. Recent work [27] on understanding Internet path inflation focused on identifying delay increase caused by topology and routing policies. Between a given source and destination node, there usually exists multiple paths depending on the state of the network. Our work helps identify the delay difference across such paths.

We summarize our key experimental findings here. We identified that more than 40% of paths studied experience delay changes of more than 50ms, confirming the conjecture that network delays of Internet paths are subject to fluctuations caused by routing changes. Further analysis shows that the variability of delay changes is actually small for most path transitions measured, allowing applications to make use of such stability. Thus, we analyze the predictability of network delay and jitter changes caused by routing events and identify network and route properties that lead to predictable delay and jitter fluctuations. Routing changes restricted to within a single network are expected to have minimal impact on network delay changes due to limited size of a single network and typically rich connectivity within large networks. Thus, we find to our surprise that intradomain routing change caused by certain networks can account for significant delay increase, sometimes up to 1 second, due to sparse internal network connectivity. The degree of similarity between the newly advertised path and the current path provides an indication of the expected delay changes. We also identify a few ASes responsible for a large number of delay increases following path changes.

The main contributions of our paper are the following:

- We demonstrate the significant impact of routing changes on network delay and jitter at the Internet path level, focusing on stable paths before and after routing changes. This is the first comprehensive study on the effect of routing events on end-to-end network delay behavior after routing convergence.
- We examine the network properties of routing changes to identify causes for delay changes. By identifying dominant

AS hop contributing to most delay changes, we account for delay variations caused by queuing delays due to persistent congestion as well as increased propagation delays. We also note a few ASes involved in many path changes contributing to significant latency and jitter variations.

- We analyze at a fine-grained level the stability of path latency and jitter as well as the predictability of changes in delay and jitter due to routing events. We identify network and route properties to help predict delay degradation caused by routing changes.
- Additionally, we develop a new measurement and analysis methodology for studying the effects of routing changes on latencies of converged network paths, enabling a more comprehensive study of the delay constancy property of Internet paths.

The paper is organized as follows. Section 2 describes our experimental methodology for data collection and analysis. Section 3 summarizes high-level observations of routing changes: their extent and impact on network delay changes. Section 4 then performs macroanalysis of the observations by separating interdomain from intradomain path changes. Section 5 further performs microanalysis by studying the network properties of routing changes to identify causes for delay changes. Section 6 summarizes the findings and their impact on latency predictability as well as application performance. Finally, Section 7 discusses related work and Section 8 draws conclusions.

2. EXPERIMENT METHODOLOGY

To effectively capture the impact of routing events on network delays of Internet paths after convergence, we need to construct a monitoring system to record the delay values of converged network paths and the network route information to identify routing changes. Our goal is to achieve sufficiently representative coverage of the Internet by selecting a wide range of network paths without incurring significant measurement overhead. Routing changes at the interdomain level can be *passively* identified using a monitoring BGP session with the local BGP router, which is easily set up. However, this does not capture intradomain routing changes that may not be visible in BGP. Passively identifying routing changes internal to an AS (Autonomous System or network) requires access to intradomain routing protocol data such as OSPF and IS-IS, making it infeasible to obtain such information from all networks of interest. Therefore, tracking fine-grained IP-level routing changes requires continuous monitoring using active probing. We developed two complementary approaches in collecting data, balancing the trade-off between overhead and coverage as well as data granularity.

2.1 Data sets

Trace 1: In the first setup, we perform continuous monitoring of 60,000 Internet paths from 200 vantage points to 300 destination hosts. Both the vantage points and destination hosts are chosen from the PlanetLab testbed [23] to achieve both geographic and network diversity. Vantage points and destination hosts are from unique PlanetLab *sites* and have significant overlap. These IPs cover 258 distinct /24 network prefixes and 186 distinct origin Autonomous Systems or ASes. The NANOG traceroute tool is used to perform traceroute approximately every 20 minutes to each of the destination host from each source node, using a time-out value of 4 seconds per traceroute run and three probes for each router hop. Network roundtrip delay or RTT values are directly obtained from traceroute results. ICMP-based delay measurements are quite accurate, as Govindan and Paxson have shown that ICMP generation times are negligible for estimating RTT values [12]. RTT values capture both the forward and reverse delays and are used to analyze the impact of routing changes of the forwarding path on the one-way network delay. We later discuss the implication of this approximation and verify the effect of reverse path delay changes. The data collection using this active probing based setup lasted 20 days in October 2006 and we term this first data set *Trace 1* in subsequent discussions.

Trace 2: In the second experiment setup, active probing is triggered from passively monitored real-time local BGP routing updates at five network locations of the RON testbed [25]. These five hosts with four distinct upstream providers respectively reside at Global Crossing network in Chicago, MIT, an edge network in Seattle, Global Crossing network in New York, and University of Michigan. Each BGP update indicates that the local route to a given destination prefix has changed from BGP’s perspective. As soon as an update is received, one traceroute run is executed to an identified live IP of the prefix. Traceroute is repeated as long as the consecutive IP level paths differ, indicating that convergence is still in progress. Ten ping probes are sent as soon as the path is converged to obtain RTT values.

Several precautions are undertaken to limit the probing frequency preventing excessively probing a host or a network. We restrict the maximum continuous probe duration for each prefix to be within 10 minutes as routing events usually converges within several minutes [14]. At most one probing process is permitted for each prefix, and the minimum probing interval for the same prefix is restricted to 10 minutes.

Live IPs responding to ICMP probes are needed to probe delay changes. Combining active probing [35] with traffic logs such as DNS data, we collected live IPs covering 61% of all announced prefixes and 68% of all ASes on the Internet. We find the set covers a large percentage of ASes in different tiers of the Internet hierarchy. These live IPs cover 58% of all prefixes and 62% of ASes in routing updates monitored during our study.

The data set for which results are presented in this paper was collected for 12 days starting from October 19th and termed as *Trace 2*. RTT of the stable path is measured by taking the average of 10 ping responses. We treat two paths as different if any of the IP hops differ. If the number of missing hops is above 5, the path is not assumed to have converged.

In summary, both measurement approaches are complementary by providing a comprehensive view of how both interdomain and intradomain routing events influence end-to-end delay and jitter between a source and a destination host.

2.2 Justifications

We discuss limitations with our two experimental setups and

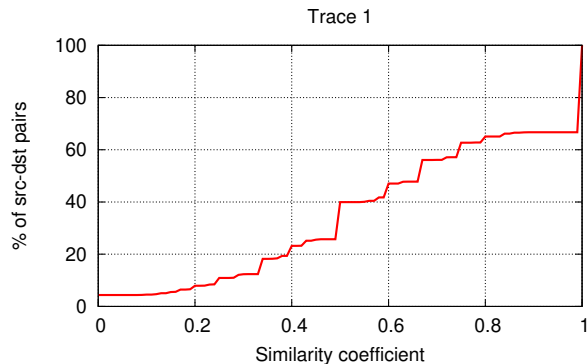


Figure 2: Distribution of similarity coefficient for all node-pairs.

argue why they do not affect our analysis. First, classic traceroute tools are susceptible to loop, cycle, and diamond problems [4] caused by load balancing routers. In Trace 1, we purposely ignore paths containing obvious anomalies such as loops and cycles, but may still overestimate the number of unique paths. This limitation may cause us to associate RTT values of the same IP-level path with different paths; however, for the remaining true routing changes, our analysis of delay performance changes induced by routing events is unaffected. As future work, we will use Paris Traceroute [4] to more accurately identify IP-level routing changes.

We intentionally focus on delay behavior after routing is stabilized, thus traceroute measurements ideally should be made during time periods with stable routing. As each hop is probed three times, we can detect potential load-balancing or routing changes through disagreeing IP addresses returned for the three probes at each hop. In both traces, such paths are ignored to reduce the transient effect of routing changes.

Throughout this paper, we use RTTs to capture the effect of a routing change from a source node to a destination node on the network delay and jitter of the directed path. In the event that the reverse routing change is not correlated with the forward change, the change in RTT reflects the contribution from the forward path change alone. However, if a reverse routing change coincides with the forward change and is in fact a result of the forward change or vice-versa, the change in RTT captures the effect of both these events. Note that most applications tend to care about round-trip delays instead of just purely one-way delays due to bidirectional communication. Thus, understanding how routing changes on the forward path can influence RTT is important for delay or jitter sensitive applications.

To understand the likelihood of forward and reverse changes coinciding, we examine Trace 1 which consists of a large amount of symmetrical probes. We first quantify the amount of sharing between two AS-level paths (the forward and reverse path in this case) by defining a similarity coefficient γ between two paths P_i and P_j . Given that the set of ASes in P_i and P_j are A and B respectively, γ_{P_i, P_j} is calculated as

$$\gamma_{P_i, P_j} = \frac{|A \cap B|}{|A \cup B|}$$

Figure 2 shows that the amount of sharing at AS-level between the forward and reverse path in Trace 1. It shows that 35% of the paths have similarity coefficients larger than 0.8. Hence reverse path changes can be potentially correlated with forward path changes.

The analysis of Trace 2 is unlikely to be affected by the inac-




Figure 3: Network delay changes over time annotated with routing changes for sample source-destination pairs.

curacy introduced by delay changes on the reverse paths. This is mainly due to the fact that most routes in commercial Internet tend to be asymmetric; thus, routing changes on the forward paths do not usually correlate with those on the reverse paths. Unlike Trace 1, Trace 2 consists mainly of paths traversing the commercial Internet. In addition, analysis of Trace 2 is further unlikely to be affected by the reverse path changes as the networks associated with source nodes in Trace 2 are well connected and have good routing stability as verified by examining updates associated with these prefixes from RouteViews BGP data [19].

3. HIGH-LEVEL OBSERVATIONS OF PATH CHANGES

In this section, we show that (1) Path latency changes happen, (2) When they do happen, they can be significant enough that applications need to care about them, and (3) The underlying availability of multiple network level paths is one cause of latency changes.

Consider the latency fluctuations shown in Figure 1. We identified the different network-level paths associated with the three source destination pairs in the figure and correlated them with the latency changes. The annotated latency changes with the routes used are shown in Figure 3. The figure clearly shows that routing changes can cause large increases in latency and such changes can persist. For example, on the path planet3.berkeley.intel-research.net \rightarrow planetlab1.cs.uoi.gr when route R1 changed to route R3 the latency increased by 60% and then remained high for 8184 seconds before again reducing by 60% when R3 changed back to R1. Similarly, the latency on the path from planetlab-4.cs.princeton.edu \rightarrow planetlab2.cs-ipv6.lancs.ac.uk increased by around 20ms and stayed high for days when the IP level route changed from R1 to R2.

Given that routing changes caused latency changes in the sample scenarios described above, we are interested in knowing the bigger picture of how often routing changes occur. Figure 4 answers this question by depicting the distribution of the fraction of times routing changes were observed when monitoring a source-destination pair. In fact, Figure 4 shows that for 60% of source-destination pairs in Trace 1, one single routing change was observed every 100 times the source-destination path was consecutively probed while for 20% of source-destination pairs 10 routing changes were observed for every 100 probes!

While this is a significant rate of change, it is possible that the routing change doesn't really matter, i.e. it did not have any discernible impact on the network delay properties. To ascertain this aspect, we also extracted the distributions of maximum latency differences observed in our measurements from Trace 1. The results in Figure 5 show that the routing changes *do* cause network delay properties to change. We observe that 50% of source-destination pairs had a maximum latency change of more than 30ms while




Figure 4: Fraction of path transitions per source-destination pair.




Figure 5: Distribution of maximum difference between median latency when IP paths switch per source-destination pair.

30% of the source-destination pairs had a latency change of more than 100ms whenever routing changes occurred. This gives an idea about the maximum impact possible on applications due to latency fluctuations. We found that latency fluctuations can result in a change of up to one second.

In summary, routes change often result in significant changes to network delay properties. Hence, we now study more closely the characteristics of routing changes and the impact of routing events on latency changes.

4. MACROANALYSIS: INTERDOMAIN VS. INTRADOMAIN ROUTING CHANGES

In this section, our goal is to understand the interplay between routing changes and network delay and jitter properties. To this end, we characterize routing events and the causal effect of routing

events on network delay and jitter properties by analyzing the two traces. We focus on the difference in the impact of the two types of possible routing changes: intradomain (within an AS) and interdomain (across ASes). This macroanalysis serves to answer how often path changes occur, how long they last, and how network delay properties change due to routing events.

4.1 How often do path changes occur?

Before correlating path events with their impact on the network delay and jitter properties, we first characterize the extent of path changes for the 60,000 paths proactively monitored for 20 days in Trace 1 and the live IPs covering 61% of all announced prefixes on the Internet reactively monitored for 12 days in Trace 2.

Figure 6 provides information on how fast routing changes are observed in both traces for both interdomain and intradomain changes. Essentially it plots a distribution of the fraction of times a routing change is observed while continuously monitoring the path between a source-destination pair. The first figure shows that for 60% of source-destination pairs, at least one *intradomain* routing change is observed every time the path is sampled 100 times consecutively. Additionally, 20% of source-destination pairs see an even higher rate of *intradomain* routing changes ranging from 10 every 100 samples to 70 every 100 samples. We also observe that the *interdomain* routing changes occur less frequently, e.g., almost all source-destination pairs observe less than 10 changes every 100 samples. Thus, interdomain changes are far less frequent than intradomain changes. While intradomain changes appear to occur frequently, it remains to see *how many unique paths* exist given these frequent routing changes.

Thus, the next logical step is to characterize the actual amount of path-diversity, i.e., the number of unique paths that a given source-destination path can potentially use. Figure 7 shows the CDF of the number of unique AS-level and IP-level paths seen by each source-destination pair, for Trace 1 and Trace 2. The results show that for Trace 1, 50% of source-destination pairs witnessed 6 or more unique IP paths, and 20% of source-destination pairs witnessed 12 or more unique IP paths, while only about 6% source-destination pairs did not experience any path changes. In addition, about 20% of source-destination pairs had more than three unique AS paths. Interestingly, around 12% of source-destination pairs had more than 20 unique IP-level paths. Most of these source-destination pairs had a large number of paths with small changes likely due to load balancing etc. In addition, certain paths such as those from an AT&T PlanetLab node to nodes in Germany and Switzerland had a large number of unique paths to reach each other and did in fact frequently switch between them.

For Trace 2, Figure 7 shows that 50% of source-destination pairs witnessed 3 or more unique paths and 20% of source-destination pairs witnessed 5 or more unique paths. Note that the number of IP-level paths seen in Trace 2 are expected to be smaller since the trace collection was reactive (driven by BGP updates) and did not sample all the paths possible between a source-destination pair.

Given that there are many unique paths available to a source-destination pair to route packets to each other, it is interesting to see whether the use of all these paths are equally likely or do some paths get preferential treatment, i.e., do some unique paths dominate the path selection for a source-destination pair? In other words, it is important to note whether most path changes are transient and most of the time the pair uses a small set of unique paths, each of which cumulatively lasts a long time and hence are more stable than other, transient paths.

To confirm this, for a given source-destination pair, we define the *dominant* paths as those that occupy a significant fraction P of

the total duration, and plot in Figure 8 the CDF of the number of dominant paths for the source-destination pairs for the two traces, varying P between 10% or 30%.

Figure 8 shows that for the P value of 10%, about 50% of the source-destination pairs in Trace 1 have a single dominant path and around 17% of the source-destination pairs have more than two dominant paths. For the P value of 30%, about 82% of the pairs have a single dominant path and the remaining have two dominant paths. The AS-paths show even a higher degree of dominance, i.e., less than 20% of source-destination pairs have more than one AS level path. Similarly, Figure 8 also shows that for Trace 2 as well, less than 20% of source-destination pairs have more than one dominant AS-level path. These results suggest that despite the potentially large number of path changes experienced by the source-destination pairs in the two traces, the vast majority of them have only one or two dominant paths.

Together, both traces show that path changes indeed occur frequently. The proactive monitoring in Trace 1 samples the network more frequently than in Trace 2 and hence discovers additional unique IP-level paths. In addition, both traces show that there exist a large number of unique paths between node-pairs in the Internet which increases the chance of switching between these paths and possible consequent network delay variations. However, most source-destination pairs have one or two dominant paths which are used preferentially over the other available unique paths.

4.2 How long do paths last?

The large number of path changes and small number of dominant paths between a source-destination pair suggest that the non-dominant paths generally have short absolute durations while dominant paths have long absolute durations. Thus, it is important to understand what exactly the distribution of absolute path durations is on a large scale.

Figure 9 depicts the absolute duration of any path occurrence per source destination pair for Traces 1 and 2. Thus, each time a unique path is seen we note how long it persists.

The results for Trace 1 show that only 20% of IP-level paths last for more than 55 minutes. In fact, 60% of the IP-level paths have low absolute duration of less than 25 minutes. On the other hand, AS-level paths have longer absolute durations: 40% of them last for more than 100 minutes. Thus, while paths are highly transient at the IP-level, AS-level paths have significantly higher absolute duration when they occur. The relatively low absolute duration of the majority of path changes have implications on how different applications should react to path changes. For example, relatively short-running jitter-sensitive applications such as VoIP should be more aggressive in adjusting path selections, e.g., via overlay routing, in response to frequent path changes.

The results for Trace 2 actually show a larger absolute duration for AS level paths. In that trace, 60% of paths last for more than 1000 minutes. This is likely due to the fact that the vantage points for Trace 2 are well connected while the vantage points for Trace 1 are from the widely distributed PlanetLab testbed contributing to many large intercontinental paths that are prone to change.

While these results show how long paths last *when they occur*, a related and interesting question to ask is what the aggregate contribution of this path is, i.e. how much aggregate time it exists for, counting over all its occurrences. This question is answered in Figure 10 for both Trace 1 and Trace 2. To understand and quantify the exact aggregate duration of paths for the source-destination pairs, we plot the CDF of aggregate duration of the unique paths of all the source-destination pairs, normalized to the total duration of the trace collection.




Figure 6: Fraction of path transitions per source-destination pair.

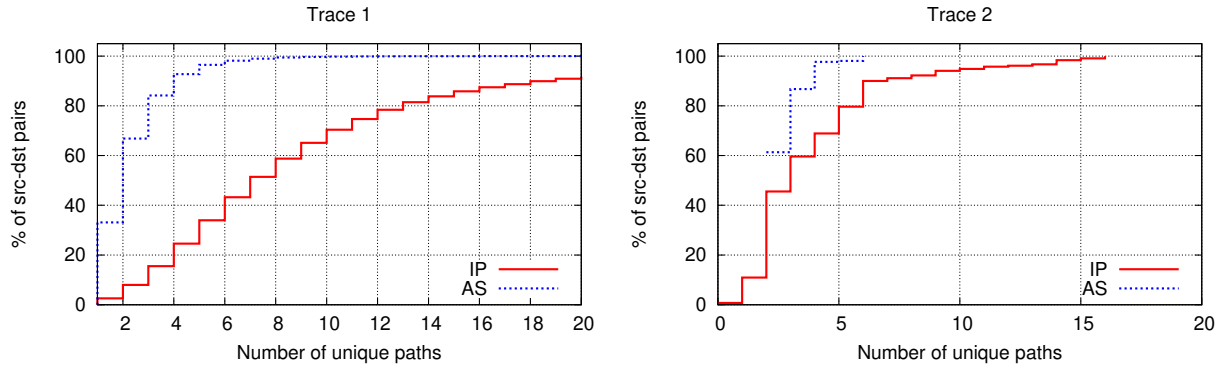


Figure 7: Number of unique paths per source-destination pair.

For Trace 1, Figure 10 shows that the vast majority of the IP-level paths have relatively short aggregate durations. However, the aggregate duration of AS level paths is higher: 50% of the paths last for more than 20% of the total duration of the trace collection and 35% of AS-level paths did not change during the trace collection. A similar trend emerges in Trace 2 as well: about 50% of the AS-level paths did not change during the course of the trace collection.

In summary, AS-level paths last longer. Thus, interarrival between interdomain routing events is longer than that between intradomain routing events. Furthermore, Section 4.1 shows that intradomain routing events are more frequent than interdomain routing events. Next, we characterize the impact of these routing events on network delay properties.

4.3 What is the impact on network delay properties?

To understand the impact of interdomain and intradomain routing events on the network delay properties between source-destination pairs, we calculate various measures that characterize the latency difference during path changes. Latency difference for a path transition from path P1 to P2 is the difference of the median RTT values of P1 and P2 observed during their corresponding occurrence.

Figure 11(a) first shows the impact of interdomain and intradomain routing events on the maximum of the latency difference per source-destination pair to capture the worst case performance degradation. It shows that intradomain path changes cause larger maximum latency differences than interdomain routing events. For example, intradomain changes cause 20% of source-destination pairs to have a latency difference of more than 100ms while interdomain path changes have similar impact on only 10% of source-destination pairs. Interestingly, we found maximum latency differences for routing changes to anycast prefixes to be more pronounced as a result of increased geographic span when routing to such prefixes.

We then improve the granularity of our observations by plotting 1 value for every path transition occurrence (Figure 11(b)). Furthermore, to capture the stability of latency difference from a path transition, we log the latency differences observed for unique path transitions and plot the median and standard deviation of these differences in Figures 11(c),(d). Surprisingly, although intradomain events have larger maximal impact, Figures 11(b),(c) show that the average impact of interdomain events is larger than intradomain events.

An intriguing question here is: why are intradomain events as significant as interdomain events in the impact they have on network delay properties? Typically, one would assume that intradomain changes are due to internal load-balancing and traffic engineering decisions which should not have significant impact on network delay properties. On the other hand, interdomain changes that actually change the AS level path signal a change in traversal of networks with potentially different characteristics. Thus a large latency difference resulting from an interdomain change may be a result of change in geographical distance traversed or persistent congestion experienced in the new path. Thus, intuitively one would expect interdomain events to have further reaching impact than intradomain events. Surprisingly, our results demonstrate otherwise.

We investigated the characteristics of intradomain events that cause large changes in network delay properties in order to shed light on this surprising conclusion. Our first observation was that the latency difference was affected by transient congestion events rather than the underlying intradomain change. For example, Figure 12 demonstrates the range of latency values experienced on a peering link on a path from plab1.eecs.ksu.edu to planet2.att.nodes.planet-lab.org. Intradomain changes occurring on this path showed a high latency difference primarily due to the presence of this *common* congested link which took on RTTs from 5ms to 500ms depending on the time of the intradomain change. In fact, when we




Figure 8: Number of dominant paths per source-destination pair with different thresholds.




Figure 9: Absolute duration of any path occurrence per source-destination pair.




Figure 12: Peering point congestion leads to large intradomain latency changes. Link connects between the Qwest and AT&T networks.

correlated the RTT change from an intradomain event with the resulting hop change, we found that in 14% of cases, the RTT change was not due to the path change but due to congestion.

Second, for the remaining 86% of intradomain changes that showed high latency difference, the latency difference in fact correlated with the point where the path changed. Two reasons explain this high latency difference due to intradomain events: (1) A path change in low tier or edge network ISPs can cause significant differences in the congestion experienced and thereby result in high latency differences. (2) Another surprising commonality that emerged when we logged the ASes involved in the intradomain path change was that a significant fraction of these ASes were sparse networks such as Abilene and GEANT. Since these ASes are not richly connected, an intradomain change can also result in a large latency change due to the increase in physical distance traversed. For example, we observed that when routing within the US, a path from Denver to California was rerouted through Denver-Seattle-California which resulted in an intradomain latency change of 40ms. Thus, all these reasons together make intradomain changes as significant as interdomain changes in terms of latency difference.

Finally to understand how consistent the latency difference is for a given source-destination pair when paths change, we plot in Figure 11(c) the standard deviation of the latency difference values for every unique path transition. We observe that for about 80% of

interdomain path transitions, the standard deviation of latency differences is less than 10 ms. These results suggest that for a given source-destination pair, the latency difference caused by an interdomain path change is fairly predictable. Furthermore, we find that latency differences due to interdomain routing events are more stable than those caused by intradomain routing events.

The corresponding measurements for Trace 2 for these properties are shown in Figure 13. Note that Trace 2 only captures interdomain routing events. Trace 2 demonstrates a larger impact on maximum latency difference than Trace 1 (20% of source-destination pairs have a maximum latency difference of larger than 100ms). Similarly, the median latency difference per unique path transition and the latency difference across all transitions is larger than Trace 1. However, Trace 2 shows severe variability in the latency difference as compared to Trace 1. This is likely attributed to the wide variety of destinations visited in Trace 2 which could include a large number of cable/DSL hosts with long queuing delay.

Another quantity of interest to many real-time applications such as VoIP is jitter. Thus it is important to understand the difference in jitter due to a path transition. In this study, we define jitter to be the standard deviation of RTT for a given path. We now plot the distribution of the difference in the jitter values for a path transition in Figure 14. The results indicate that both intradomain and interdomain path changes result in significant difference in jitter properties. The results show that the network latency for any given path can indeed fluctuate extensively, in particular, some paths have lower jitter than other paths. For example, in Trace 2, 30% of path transitions cause a jitter difference of larger than 20ms. Thus a VoIP stream may be potentially subjected to 20ms increased jitter caused by a path change during a call.

In summary, our measurements indicate that interdomain and intradomain events both have significant impact on the latency and jitter observed by an application.

5. MICROANALYSIS: NETWORK PROPERTIES OF ROUTING CHANGES

Previously Section 4 demonstrated that surprisingly intradomain routing events can have as much impact on roundtrip delay changes




Figure 10: Normalized aggregated duration of all paths from Traces 1 and 2.




Figure 11: Distribution of maximum, median and standard deviation of latency difference when paths switch.

as interdomain routing events. In this section, we examine in more detail network properties of the routing changes to identify causes for delay changes.

Notations and terminology: For the ease of exposition, we use P_{bi} to denote the stable path prior to a routing change i and P_{ai} to denote the converged path after this routing change. Although both paths connect between the same source and destination host pair, they can differ significantly from each other in terms of network characteristics which explain divergent delay behavior.

We define the *dominant hop* as the AS-level hop contributing extensively to the RTT delay values based on a contribution threshold defined later. To capture the *similarity* of two paths, one metric used is based on whether they share the dominant AS hop. The second metric is the percentage of overlapping AS level hops when comparing the AS numbers. This is computed using the ratio between the common number of ASes and the total number of distinct ASes of the two paths. Intuitively, similar paths tend to have similar delay and jitter behavior.

Explaining network delays: In packet switched networks, network delay or latency mostly consists of propagation delay due to the actual physical distance traversed by packets and queueing delay caused by network congestion. Transmission delay is usually

negligible. Thus, in the search of explanations for the differences in delay properties between the paths, we focus on factors influencing propagation and queueing delays. Note that we do not attempt to distinguish the root cause for delay along a given hop due to lack of sufficient information. We plan to use additional metrics such as bandwidth and packet loss to discern the cause in our future work.

Along uncongested network paths, the propagation delay becomes the dominant component in determining the network delay. Thus, paths traversing similar network distances without going through any congested links tend to have similar delay values, unless the link speed differs significantly, *i.e.*, a satellite link vs. an optical fiber link. Two paths sharing the same bottleneck link either due to congestion or physical distance, accounting for the main fraction of the network delay, will also have similar delay values. Note that we mostly focus on persistent congestion points in understanding delay value changes. Transient congestion is less predictable and accounts for the delay jitter property. If network congestion occurs within the source or the destination network, or near these networks without much path diversity, any two paths connecting the source and destination host will suffer from the same bottleneck links, thus experiencing similar network latencies. For Trace 1, we do not expect congestion to occur frequently near the




Figure 13: Distribution of maximum, median and standard deviation of latency difference when paths switch.




Figure 14: Distribution of jitter difference.

source or the destination host as the PlanetLab sites tend to have good network connectivity. However, for Trace 2, the bottleneck link may occur near the destination network as the probing targets cover a much wider range of address blocks.

Explaining network delay jitter: Delay jitter or variability of network delays is of strong interest to real-time multimedia applications such as voice over IP and video streaming. Delay variations are largely due to transient congestion caused by traffic variability, routing instability, and network failures. The delay jitter behavior is directly influenced by the queueing discipline implemented in the routers traversed.

Using measurement data, we test several intuitive hypothesis that account for the delay difference between the two paths. We found that the IP or AS hop count difference is not helpful in predicting delay or jitter changes. This is not surprising as the IP or AS hop count may not correlate well with delays due to the presence of dominant hops as well as the fact that most paths have similar AS hop counts.

Hypothesis 1: Paths sharing the dominant hop or with high similarity values tend to have similar delay behavior.

The *dominant hop* is identified by examining the distribution of the latency contribution by the top AS hop normalized by the path

AS hop count shown in Figure 15. Based on the figure, 5 appears to be a good threshold value to define dominant hop balancing between the coverage and the selectivity, as some paths inherently do not have a dominant hop with each hop contributing roughly equal delay.

Figure 16 indicates that with larger similarity (based on percentage of overlapping ASes) between two paths across a given routing change, the delay difference has a larger variance indicated by the error bars, suggesting that this similarity metric does not reflect well similarity in the delay behavior. We next use a more direct metric to compare the path similarity based on whether the dominant AS hop is shared in both paths in the case there exists such an AS using the contribution threshold of 5. Figure 17 indicates as expected latency difference is on average smaller if the dominant AS hop is preserved across routing events.

Hypothesis 2: The effect of transient congestion on the latency difference between P_{bi} and P_{ai} is non-negligible.

In Figure 18, we plot the difference between minimum RTT values of the two paths vs. the difference between the median RTT values to illustrate the effect of transient congestion. For both traces, we observe many data points to be close along the $y = x$ diagonal line, indicating minimal effect of congestion. Points below the line




Figure 18: Effect of congestion on delay difference estimation caused by path changes.




Figure 15: Identifying the dominant AS hop.

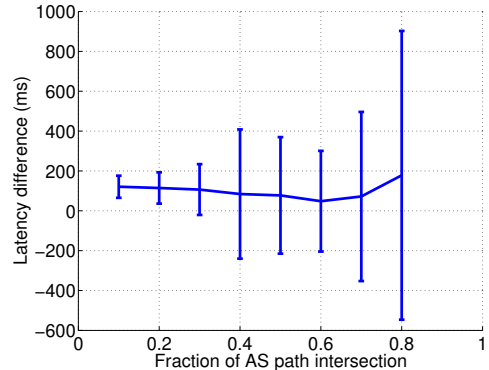


Figure 16: Latency difference vs. similarity of AS path

indicate congestion caused the delay difference to be smaller than estimated using minimal RTT values. There tend to more points above the diagonal line, especially for intradomain routing changes, indicating inflation of delay differences by congestion. Without considering path changes, existing studies on predicting network distance such as GNP will just choose one of the many points on the graph with little chance of identifying the actual difference of minimum RTTs.

Hypothesis 3: There are a few ASes involved in path changes contributing to significant latency and jitter differences.

This hypothesis is verified by Figure 19 indicating a few ASes contributing to most of the routing changes at the interdomain level. The ASes are sorted in reverse order based on their contribution to routing changes based on occurrence. An AS is considered to “contribute” to a routing change if it does not lie within the intersection of the two AS paths excluding the source and destination AS.

6. DISCUSSION

In this section, we discuss how feasible it is to predict latency changes. We further discuss the implications of our measurement study on the performance of various applications.

6.1 Predictability

Table 6.1 summarizes the findings of our paper and how they relate to the predictability of network delay properties. The ability to predict is helpful to construct proactive mitigation responses to possible delay or jitter increases. Note that precisely predicting the exact delay or jitter change may be difficult if not impossible; however, forecasting the possibility of potential degradation in response to a routing change is beneficial as it allows possible application-layer mitigation responses.

6.2 Impact on applications

In this section, we use the insight gained from analyzing network delay dynamics to qualitatively discuss the impact of network delay changes on various applications. We choose delay-sensitive applications and characterize the nature of impact on delay and possible solutions to alleviate potential consequent performance degradation.

Voice over IP: Voice over IP (VoIP) [42] is an increasingly important application that requires paths with low latency, jitter and loss properties. Latency is an important parameter for VoIP applications due to the sensitivity of humans to delays above certain values.

Our study shows that significant path fluctuations do happen in both latency and jitter. Furthermore, the path changes in both intradomain and interdomain cases are significant. Thus, it might be necessary for VoIP applications to use overlay networking [2, 24, 29]. Our methodology can be used to identify the dominant paths in overlay hops and select them according to their characteristics. It might also be necessary to use multiple disjoint paths to react effectively to latency changes. Our study also enables proactive response to interdomain route changes and aids in overlay route selection.

Multicast Streaming: Multicasting stored files or live streams to a set of receivers is another important and fast-growing application on the Internet. Most deployed protocols use application-layer implementations. Such multicast protocols include ESM [13], Host-Multicast [36], NICE [5], and SplitStream [6]. Many applications of multicast streaming involve long running broadcasts (*e.g.*, of conference sessions). The longer time period makes such applications susceptible to latency changes from routing events. Streaming requires low delay and jitter properties. Thus, changes in delay

Finding	Impact on predictability
Routes do change often, both at the intradomain and interdomain level	Prediction is important since it is expensive and slow to proactively probe.
A manageable set of unique routes exist. A smaller set of routes dominate.	History-based approach is tractable for prediction.
Intradomain routing changes occur at higher frequency but have shorter duration than interdomain routing events.	Predicting delay differences due to interdomain routing events have longer-term benefit than intradomain events.
Both intradomain and interdomain routing events can cause significant latency differences.	Both events need to be accounted for. Predicting intradomain events is harder than interdomain events.
Average latency difference for interdomain events is larger than that for intradomain events. Also, its maximum and standard deviation are lower.	More variability in intradomain latency differences make them less predictable.

Table 1: Summary of findings and their impact on predictability.

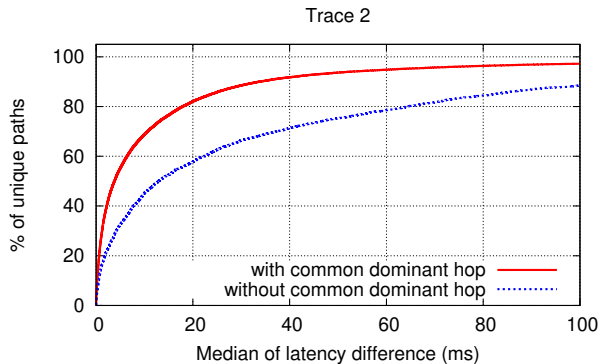


Figure 17: Comparing the latency difference distribution distinguished by the shared dominant AS.

properties over time affect the optimality of the multicast tree constructed.

Locality-based Systems: Many systems take advantage of network locality to reduce delay, *e.g.*, by selecting the closest server or mirror site. Since TCP throughput is influenced by RTT [20], the choice for a nearby and stable server is important. Similarly, other such systems where latency is an important performance metric include locality-based DHTs [26, 40], CDN selection, proxy selection, and Akamai style redirection [41]. Such systems could benefit from stability information of paths to a server to avoid future performance degradation from a change in network delay properties.

Positioning Systems: As mentioned earlier, Internet positioning systems for predicting network distance between two arbitrary nodes using protocols such as IDMAPS [11], GNP [17], Vivaldi [8], IDES, are directly impacted by the stability of identified network latencies, which may be affected by routing events. Such positioning systems may need to trigger callbacks to recalculate coordinates when path changes occur. Apart from physical delay changes over time, path changes may also cause new or worse triangular inequality violations which increase the positioning error for some of these systems. Additionally, our methodology can also help identify stable landmarks for systems such as GNP.

Overlay Routing Systems : Overlay routing systems such as RON [3] and MTCP [38] can improve network performance by finding better overlay paths compared to the direct path to a server. However, if the paths are chosen for downloading large files, it may be useful to select paths whose network delay properties are stable over time.




Figure 19: Small number of ASes contribute to most changes

7. RELATED WORK

Internet network delay is an important property studied extensively in previous work including seminal work by Paxson [21]. For time-critical applications or real-time applications, network delay is an essential performance metric. Network delay is influenced by two main components: propagation delay and queueing delay. Directly related to our study is previous work on predicting [18, 18], explaining [27, 28], or modeling [37] observed network delay focusing on propagation delay only, ignoring effects from transient congestion. For example, some of the existing work on understanding network delay properties focus on predicting mostly *minimum* or *stable* network delay between two arbitrary hosts using a virtual coordinate system either in an infrastructure-based approach [18] or in a decentralized manner [9]. Recently Madhyastha *et al.* adopted a structural approach to predict path latencies using topology and routing policy information [16]. Zhang *et al.* studied the constancy of Internet path properties assuming the path is stable or unaffected by routing changes [39]. Analysis of path inflation causes also examines only one particular stable path for each source destination pair at the time the measurement is performed. This group of work does not take into consideration delay changes caused by transient congestion or routing changes which may result in longer-lived delay changes.

Another class of related work directly focuses on transient effects on network performance caused by routing events during convergence process. BGP's path exploration during routing convergence may last up to several minutes [14]. Thus, studies such as [22, 31, 32] analyze performance degradation during routing changes. Our study investigates a complementary and arguably

equally if not more important problem of identifying the effect of routing change on network delay after routing convergence. As far as we know our study is the first work in this area.

8. CONCLUSIONS

In summary, our study is the first to quantify the network delay variations caused by switching to a new IP-level path caused by routing changes, which are not infrequent events on today's Internet. Intradomain routing changes are found to be more prevalent than interdomain routing changes, which nevertheless still affect more than 58% all prefixes in our measurement study lasting for 12 days. Our findings indicate that there exists stability in the resulting delay differences for the path before and after the routing event and correlation between network properties of route changes and delay impact. Our work presents a first important step towards using routing change information to predict resulting network delay changes of converged paths.

Acknowledgments

We would like to thank the anonymous reviewers for their helpful comments. We are also grateful to Abhinav Pathak for his help with trace collection. This work was supported in part by NSF grants CNS-0338842, CyberTrust-0430204, CAREER-0238379, and CAREER-0643612.

9. REFERENCES

- [1] NANOG Panel: Smart Routing Technologies. <http://www.nanog.org/mtg-0206/smart.html>, June 2003.
- [2] Y. Amir, C. Danilov, S. Goose, D. Hedqvist, and A. Terzis. An Overlay Architecture for High Quality VoIP Streams. *IEEE Transactions on Multimedia*, 2006.
- [3] D. G. Andersen, H. Balakrishnan, M. F. Kaashoek, and R. Morris. Resilient Overlay Networks. In *Proc. of ACM SOSP*, 2001.
- [4] B. Augustin, X. Cuvellier, B. Orgogozo, F. Viger, T. Friedman, M. Latapy, C. Magnien, and R. Teixeira. Avoiding Traceroute Anomalies with Paris Traceroute. In *Proc. of IMC*, 2006.
- [5] S. Banerjee, B. Bhattacharjee, and C. Kommareddy. Scalable Application Layer Multicast. In *Proc. of ACM SIGCOMM*, 2002.
- [6] M. Castro, P. Druschel, A.-M. Kermarrec, A. Nandi, A. Rowstron, and A. Singh. SplitStream: High-Bandwidth Multicast in Cooperative Environments. In *Proc. of ACM SOSP*, 2003.
- [7] Cisco Systems, Inc. Cisco Optimized Edge Routing. http://www.cisco.com/en/US/products/ps6628/products_ios_protocol_option_home.html.
- [8] F. Dabek, R. Cox, F. Kaashoek, and R. Morris. Vivaldi: A Decentralized Network Coordinate System. In *Proceedings of ACM SIGCOMM*, August 2004.
- [9] F. Dabek, R. Cox, F. Kaashoek, and R. Morris. Vivaldi: A Decentralized Network Coordinate System. In *Proceedings of ACM SIGCOMM*, 2004.
- [10] N. Feamster, L. Gao, and J. Rexford. How to lease the Internet in your spare time. Technical Report GT-CSS-06-10, Georgia Tech, August 2006.
- [11] P. Francis, S. Jamin, V. Paxson, L. Zhang, D. Gryniewicz, and Y. Jin. An Architecture for a Global Internet Host Distance Estimation Service. In *Proceedings of IEEE INFOCOM*, March 1999.
- [12] R. Govindan and V. Paxson. Estimating Router ICMP Generation Delays. In *Proc of Passive and Active Measurement (PAM)*, 2002.
- [13] Y. hua Chu, S. G. Rao, and H. Zhang. A Case for End System Multicast. In *Proc. of ACM SIGMETRICS*, 2000.
- [14] C. Labovitz, A. Ahuja, A. Bose, and F. Jahanian. Delayed internet routing convergence. In *Proc. ACM SIGCOMM*, 2000.
- [15] K. Lakshminarayana, I. Stoica, S. Shenker, and J. Rexford. Routing as a Service. Technical Report UCB/ECS-2006-19, UC Berkeley, 2006.
- [16] H. Madhyastha, T. Anderson, A. Krishnamurthy, N. Spring, and A. Venkataramani. A Structural Approach to Latency Prediction. In *Proc. of IMC*, 2006.
- [17] T. S. E. Ng and H. Zhang. Predicting Internet Network Distance with Coordinates-Based Approaches. In *Proceedings of IEEE INFOCOM*, June 2002.
- [18] T. S. E. Ng and H. Zhang. Predicting Internet Network Distance with Coordinates-Based Approaches. In *Proc. of IEEE INFOCOM*, 2002.
- [19] U. of Oregon Route Views Archive Project. <http://www.routeviews.org>.
- [20] J. Padhye, V. Firoiu, D. Towsley, and J. Kurose. Modeling TCP Throughput: A Simple Model and its Empirical Validation. In *Proc. of ACM SIGCOMM*, 1998.
- [21] V. Paxson. *Measurements and Analysis of End-to-End Internet Dynamics*. PhD thesis, University of California at Berkeley, 1997.
- [22] D. Pei, L. Wang, D. Massey, S. F. Wu, and L. Zhang. A Study of Packet Delivery Performance during Routing Convergence. In *Proc. of IEEE International Conference on Dependable Systems and Networks (DSN)*, 2003.
- [23] PlanetLab. <http://www.planet-lab.org>.
- [24] S. Ren, L. Guo, and X. Zhang. ASAP: an AS-Aware Peer-relay protocol for high quality VoIP. In *Proc. of the 26th International Conference on Distributed Computing Systems (ICDCS'06)*, 2006.
- [25] RON. <http://nms.csail.mit.edu/ron/sites/>.
- [26] A. Rowstron and P. Druschel. Pastry: Scalable, Distributed Object Location and Routing for Large-Scale Peer-to-peer Systems. In *Proc. of ACM/IFIP/USENIX Middleware*, November 2001.
- [27] N. Spring, R. Mahajan, and T. Anderson. Quantifying the causes of internet path inflation. In *Proc. of SIGCOMM*, 2003.
- [28] L. Subramanian, V. N. Padmanabhan, and R. H. Katz. Geographic Properties of Internet Routing. In *Proc. of the USENIX Annual Technical Conference*, 2002.
- [29] S. Tao, K. Xu, A. Estepa, T. Fei, L. Gao, R. Guerin, J. Kurose, D. Towsley, and Z.-L. Zhang. Improving VoIP quality through path switching. In *Proc. of INFOCOM*, 2005.
- [30] R. Teixeira, K. Marzullo, S. Savage, and G. Voelker. In Search of Path Diversity in ISP Networks. In *Proc. of the USENIX/ACM Internet Measurement Conference*, 2003.
- [31] F. Wang, L. Gao, J. Wang, and J. Qiu. On Understanding of Transient Interdomain Routing Failures. In *Proc. International Conference on Network Protocols*, 2006.
- [32] F. Wang, Z. M. Mao, J. Wang, L. Gao, and R. Bush. A Measurement Study on the Impact of Routing Events on End-to-End Internet Path Performance. In *Proc. ACM SIGCOMM*, 2006.
- [33] W. Xu and J. Rexford. MIRO: Multi-path Interdomain Routing. In *Proc. ACM SIGCOMM*, 2006.
- [34] X. Yang and D. Wetherall. Source Selectable Path Diversity via Routing Deflections. In *Proc. ACM SIGCOMM*, 2006.
- [35] A. Zeitoun and S. Jamin. Rapid Exploration of Internet Live Address Space Using Optimal Discovery Path. In *Proc. Global Communications Conference*, 2003.
- [36] B. Zhang, S. Jamin, and L. Zhang. Host Multicast: A Framework for Delivering Multicast To End Users. In *Proc. of IEEE INFOCOM*, June 2002.
- [37] B. Zhang, T. S. E. Ng, A. Nandi, R. Riedi, P. Druschel, and G. Wang. Measurement-Based Analysis, Modeling, and Synthesis of the Internet Delay Space. In *Proc. of IMC*, 2006.
- [38] M. Zhang, J. Lai, A. Krishnamurthy, L. Peterson, and R. Wang. A transport layer approach for improving end-to-end performance and robustness using redundant paths. In *Proc. of USENIX*, 2004.
- [39] Y. Zhang, N. Duffield, V. Paxson, and S. Shenker. On the Constancy of Internet Path Properties. In *Proc. of the ACM SIGCOMM Internet Measurement Workshop (IMW)*, 2001.
- [40] B. Y. Zhao, L. Huang, J. Stribling, S. C. Rhea, A. D. Joseph, and J. Kubiatowicz. Tapestry: A Resilient Global-Scale Overlay for Service Deployment. *IEEE JSAC*, 2004.
- [41] Akamai. <http://www.akamai.com/>.
- [42] Skype. <http://www.skype.com>.