

A Measurement-Based Analysis of Multihoming

Aditya Akella
Bruce Maggs*
Srinivasan Seshan
Carnegie Mellon University
{aditya, srini+, bmm}@cs.cmu.edu

Anees Shaikh
IBM T.J. Watson
Research Center
aashaikh@watson.ibm.com

Ramesh Sitaraman*
Univ. of Massachusetts
ramesh@akamai.com

ABSTRACT

Multihoming has traditionally been employed by stub networks to enhance the reliability of their network connectivity. With the advent of commercial “intelligent route control” products, stubs now leverage multihoming to improve performance. Although multihoming is widely used for reliability and, increasingly for performance, not much is known about the tangible benefits that multihoming can offer, or how these benefits can be fully exploited. In this paper, we aim to quantify the extent to which multihomed networks can leverage performance and reliability benefits from connections to multiple providers. We use data collected from servers belonging to the Akamai content distribution network to evaluate performance benefits from two distinct perspectives of multihoming: high-volume content-providers which transmit large volumes of data to many distributed clients, and enterprises which primarily receive data from the network. In both cases, we find that multihoming can improve performance significantly and that not choosing the right set of providers could result in a performance penalty as high as 40%. We also find evidence of diminishing returns in performance when more than four providers are considered for multihoming. In addition, using a large collection of measurements, we provide an analysis of the reliability benefits of multihoming. Finally, we provide guidelines on how multihomed networks can choose ISPs, and discuss practical strategies of using multiple upstream connections to achieve optimal performance benefits.

Categories and Subject Descriptors

C.2 [Computer Systems Organization]: Computer-Communication Networks; C.2.1 [Computer-Communication Networks]: Network Architecture and Design

General Terms

Measurement, Performance

1. INTRODUCTION

Large enterprises and content providers, who depend on the Internet to operate their businesses, require a high level of reliability

*Bruce Maggs and Ramesh Sitaraman are also with Akamai Technologies. This work was supported by the Army Research Office under grant number DAAD19-02-1-0389. Additional support was provided by IBM.

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.

SIGCOMM'03, August 25–29, 2003, Karlsruhe, Germany.
Copyright 2003 ACM 1-58113-735-4/03/0008 ...\$5.00.

from their network connections. Increasingly, these large consumers and producers of network data are turning to *multihoming* as a technique to achieve resilience to service interruptions [4]. Multihoming is defined simply as a customer (or ISP) network having more than one external link, either to a single ISP, or to different providers [12]. The customer typically has its own public AS number, and advertises its address prefixes via all of its upstream providers using BGP [14].

While multihoming to multiple providers is motivated primarily by a need for link-level and provider-level fault tolerance, the advent and expected growth of “intelligent route control” devices and services promises to allow subscribers to leverage multihoming for more than just increased resilience [9, 13]. For example, performance to different parts of the network may vary depending on which upstream provider is used. In such situations, careful route selection can significantly improve performance. Even availability can be managed to some extent by choosing ISPs that have sufficiently diverse connectivity to destinations of interest. In this paper, our primary goal is to quantify the extent to which subscribers can leverage connections to multiple network providers to improve performance. We also provide a study of the reliability benefits.

We characterize performance in terms of the speed and efficiency with which wide-area transfers occur. Conceptually, our approach is to consider a network subscriber in a major metropolitan area, and evaluate the relative benefits of choosing upstream providers from several available options. We are interested in the perspective of both high-volume Web sites and data centers, which are interested in attaining good performance to many parts of the network, and also enterprise subscribers, who are more interested in receiving data from various content providers. We focus on the common case in which the subscriber has little or no control over end-to-end paths, but rather only which ISPs provide first-hop connectivity to the Internet.

Our study draws empirical observations from measurement data sets collected at servers and monitoring nodes deployed by Akamai, a large content distribution service provider. These servers and monitors are attached to a diverse set of ISPs (most nodes connected to a single provider), with multiple Akamai servers located in each of the major metropolitan areas that we analyze. The network performance data collected at these Akamai nodes allows us to compare performance across providers from the perspectives of enterprises or content providers in different metropolitan areas. We analyze performance primarily in terms of observed network latency as the Akamai servers and monitors fetch objects from customer Web servers or other Akamai servers.

Our analysis is based largely on the notion of *k-multihoming* in which we quantify the best performance achieved when a subscriber is multihomed to *k* available providers in a given city. We establish a baseline in which we assume that it is possible for a sub-

scriber to employ all k providers and switch to the best performing provider at each instant. By evaluating the performance as k is increased, we provide some insight into the incremental performance benefit when adding providers. To assess the impact of the chosen set of providers, we also compare the performance of the optimal multihoming solution to random (and worst-case) selections of ISPs. In addition, we quantify the usage of each ISP in the optimal k -multihoming solutions to understand how traffic should be distributed among the k upstream providers to achieve the best performance. We show that, on average, performance can be improved considerably by multihoming for both the enterprise and content provider perspectives. For example, even in a 2-multihoming solution, average performance was improved by 25% for 3 out of 4 metro areas we study. We also find strong evidence of diminishing incremental performance benefits as more providers are added. We observe that increasing beyond $k = 4$ provides little added performance. Comparing the optimal multihoming solution to a random choice of k providers (for $k \leq 4$), we find that random selection degrades performance 15% on average, and as much as 40%. This suggests that a careful choice of providers is key to achieving the full performance benefits of multihoming.

Although our main focus is on performance benefits, we also evaluate the availability improvements due to multihoming. Using a large set of traceroute measurements, we perform an analysis of how much benefit multihoming can provide in terms of path diversity across candidate providers in a given location. One important aspect we do not consider is how to use multiple providers in such a way as to optimize bandwidth costs. Each ISP contract typically has its own pricing structure and bandwidth commitments, which may result in one provider link being more expensive depending on the time-of-day or traffic level. These differences can be exploited to reduce overall bandwidth costs, however we leave the problem of understanding the cost-performance trade-offs of multihoming as future work, and instead confine our study to performance and reliability benefits.

In the next section, we further motivate our work with two case studies that demonstrate temporal differences in performance across network providers in specific situations. These studies are chosen to represent both enterprise and data center perspectives. Sections 3 and 4 present the performance benefit analysis in detail and Section 5 follows with a description and initial results from the reliability analysis. Section 6 evaluates some practical multihoming scenarios and also discusses strategies for choosing providers when multihoming. We relate our work to previous work in Section 7 and summarize the paper in Section 8.

2. MULTIHOMING FOR PERFORMANCE: TWO CASE STUDIES

In this section we present two relatively small sets of measurement data that illustrate the potential for performance improvements due to multihoming. These empirical results provide some evidence that performance differences between ISPs exist and that relative performance changes over time.

2.1 Data Center Multihoming

We first consider an example of three commercial data centers each multihomed to two tier-1 ISPs. These sites host the Web site of IBM Corporation, and receive HTTP requests from clients distributed all over the world. The hosting center is configured with a data collection module that passively collects delay estimates over each of the center's provider links as clients fetch a designated embedded object from the Web site. The estimates are then aggregated by IP address prefix, according to the BGP tables at the site.

Delay estimates are based on *handshake round-trip time* (hrrt),




Figure 1: Relative performance for 2-multihomed data centers

defined as the time between receiving the first TCP SYN packet from the client to open a connection to the object server, and the time at which the final TCP ACK is received to complete the three-way handshake. The data collection module keeps a weighted moving average of hrrt on a per-prefix basis, and reports this average hourly for the 500 most active prefixes (as determined by request rate for the objects used for measurements). The Web site is mirrored across the three data centers, located in the U.S. in the East coast, the Midwest, and West regions. Clients are directed to each location using an IP-level load balancing mechanism, such that there is little overlap in the client prefixes appearing at the sites. All three data centers are multihomed to the same two network providers. In Figure 1, we show the average relative delay observed over the links for a 7-day portion of a trace taken in January 2003. The graphs plot the ratio of the hrrt observed on ISP 1 to the hrrt on ISP 2, averaged over all client prefixes. Hence, when the curve is above 1, ISP 2 provides better average delay. Also, since we have an estimate of the request rate from each client prefix, we weight the average by the request rate in order to emphasize performance differences for those prefixes that generate more traffic. The graphs show the weighted average ratio computed for each hourly sample from the collection module.

As illustrated in Figure 1, each data center has an opportunity to capitalize on performance differences if it is able to dynamically direct traffic over its provider links. For example, both the East and West data centers show multi-hour periods where one of the ISPs provides better average delay than the other. The Midwest trace shows that ISP 2 is better most of the time, though there are some periods when the average performance is very similar.

2.2 Enterprise Multihoming

Our next case study considers the perspective of an enterprise wishing to optimize its multihomed connectivity to application service providers or Web-based services. At first glance, it would appear difficult for an enterprise to control the routing or delivery of data from a content-provider to the enterprise site. One possibility is for the enterprise to control routing announcements to its different providers. This approach has a number of drawbacks, however, including coarse control and slow response by routing to changes. Another possibility is to use address ranges allocated from each provider. Each transfer to the enterprise could use a destination address from the best-performing ISP for the corresponding data source. This address use could be controlled through the use of network address translation (NAT) techniques and clever use of

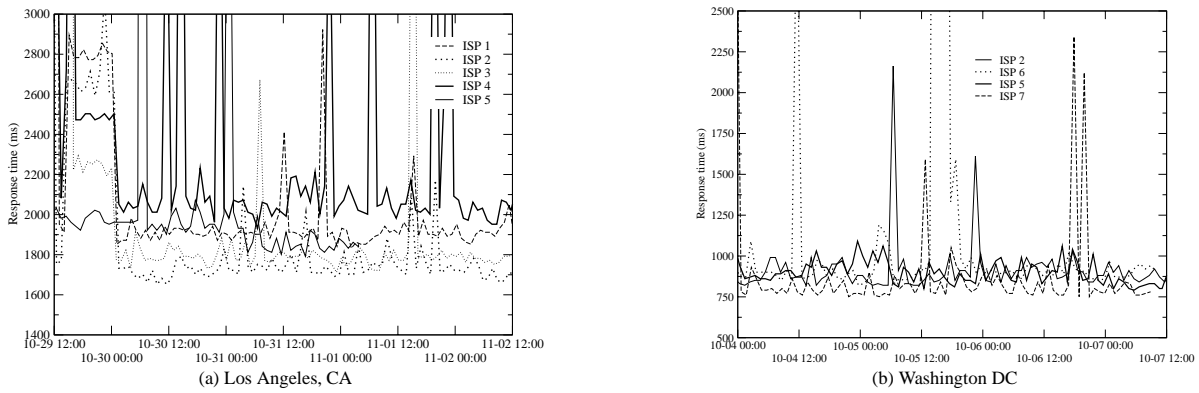


Figure 2: Performance across client ISPs

DNS. For example, for flows initiated from the enterprise, a NAT box could dynamically choose the best ISP and alter the contents of packets appropriately. Similarly, for flows established to the enterprise, the DNS could return the address that provides the best performance to the connecting host. We consider an example of this scenario with response time measurements from Keynote Systems agents which are widely distributed and connected primarily to Tier-1 ISPs (typically with a 10 Mbps link in the U.S.) [6]. By choosing major cities in which there are several Keynote agents attached to different ISPs, we can observe ISP performance differences from the perspective of clients located in each city.

The agents were configured to measure the time to retrieve a complete Web page from a dedicated off-the-shelf cache appliance located in a commercial hosting center in the Midwest. The page is an instance of the index page from a production sports Web site with 29 embedded objects and a total size of 104 KB. The cache was used only for these measurements – it did not serve any production traffic, and hence was unloaded. The content had sufficiently long expiration times such that the cache rarely needed to fetch any object from the origin server. The Keynote agents did not use HTTP/1.1 persistent connections, which likely inflates the absolute response times. Comparisons across agents are still useful, however, as they behave uniformly and are configured identically. Since the clients are homogeneous and well-connected, and the cache is similarly well-connected and unloaded, we argue that observed differences are primarily due to network effects.

Each agent requested the page every hour during a period starting in late August through mid-December 2001. In Figure 2, we show snapshots of continuous traces (*i.e.*, with no agent errors) from Washington D.C. and Los Angeles, CA. Each city has several deployed agents, with some overlap between the specific ISPs. The graphs show that relying on one ISP is prone to prolonged performance degradation. For example, near the beginning of the trace in Figure 2(a), the performance across each ISP is significantly different (by approximately 200 ms) with ISP 5 the best and ISP 1 the worst. After October 30, however, the performance shifts dramatically and a new performance ordering is established. In Figure 2(b) ISP 6 provides similar or slightly better performance than ISP 5 for most of the trace. For several hours between noon and midnight on October 5, however, response time on ISP 6 suffers a severe and persistent degradation. Similarly, while ISP 7 appears to provide the best performance over most of the trace, for a few samples between noon and midnight on October 6, it shows the worst performance among all ISPs.

2.3 Motivation for a Broader Study

These case studies are not necessarily representative of what performance gains can be generally achieved by multihoming to mul-

iple providers. Nevertheless, both the data center and enterprise scenarios suggest that judiciously choosing and using multiple network providers can provide improved network performance. These observations provide additional motivation for the broader study presented in the remainder of the paper.

In our study, we use a large set of measurements to perform a more comprehensive analysis of multihoming benefits. Three of the data sets were collected from servers and performance monitors belonging to the Akamai content distribution network (CDN). We employ these data sets to evaluate the performance benefits of multihoming. The fourth data set was collected using the Keynote infrastructure (described above). We use this data set to perform an analysis of the reliability benefits.

In Sections 3 and 4, we describe the Akamai data used for analyzing performance benefits along with results from our analysis. Section 5 describes the data and metrics used for analyzing reliability benefits of multihoming and the results of our evaluation.

3. ENTERPRISE PERSPECTIVE

The performance benefit of multihoming at an enterprise is reflected in the download performance of requests from the enterprise to destinations of interest, for example important or popular Web sites. Intuitively, the best-case scenario for an enterprise is to be able to use its multiple network connections to achieve *nearly optimal* performance for a large fraction of its Web requests.

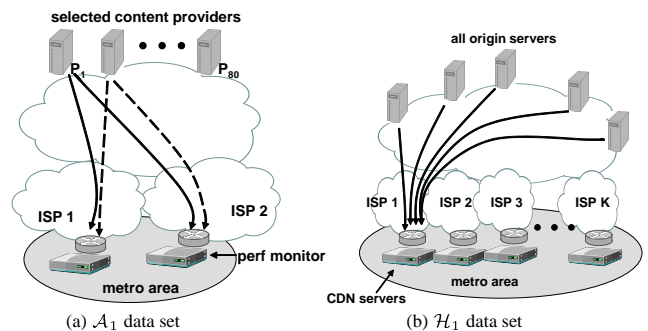


Figure 3: Measurement data sets: In (a) Akamai performance monitors in a given city are connected to different ISPs and download 10KB objects at 6-minute intervals from servers belonging to 80 content providers. In (b) Akamai servers connected to different ISPs in the same city download objects from all customer origin servers in order to serve them to clients. For this data set, turnaround times are averaged over each hour across retrievals from all origin servers.

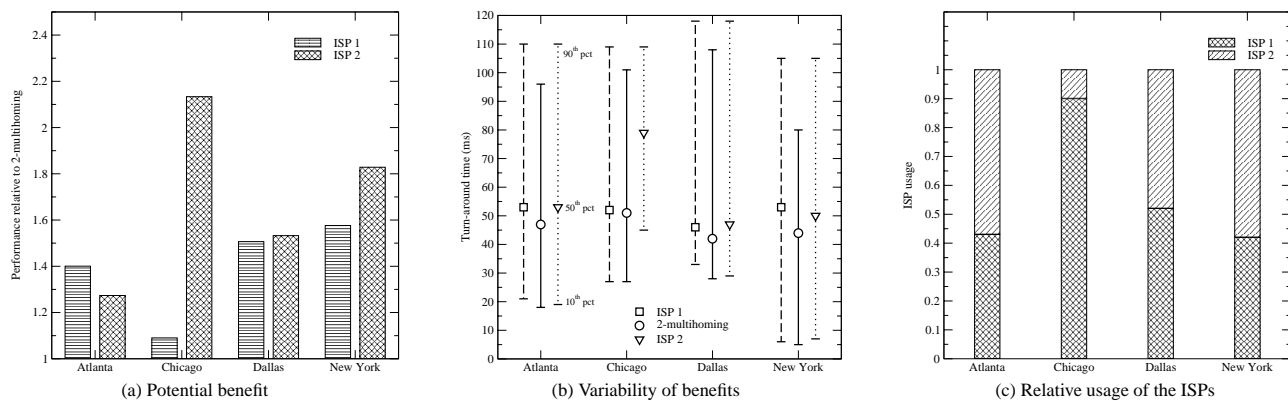


Figure 4: 2-multihoming evaluation: The average benefits are shown in (a). Graph (b) shows the median, 10th and 90th percentile turnaround times for each ISP and for 2-multihoming. The relative usage of the two ISPs in the optimal schedule is shown in (c).

In our analysis of enterprise multihoming, we use two distinct sets of data, \mathcal{A}_1 and \mathcal{H}_1 (described below), collected from servers and monitoring nodes deployed by Akamai. An important feature of this data is that the collection points are connected to a large variety of ISPs. Moreover, there are multiple metropolitan areas in which a number of collection points are located, each connected to a different ISP. We use monitoring nodes and Akamai servers in a single metro area connected to different ISPs as stand-ins for a multihomed enterprise.

Data Set \mathcal{A}_1 : This data set comprises statistics collected by 27 geographically distributed Akamai monitoring nodes. One or two nodes are located in major cities in the U.S., with multiple nodes in the same city attached to different upstream provider networks, as shown in Figure 3(a). Every 6 minutes, on average, these nodes download designated objects directly from a large number of content providers that are Akamai customers. For each attempted download, the performance monitor logs a number of statistics, including the HTTP response code, turnaround time for the request (if successful), the size of the object downloaded, the total response time, and any errors (if unsuccessful). We focus, in particular, on the *turnaround time*, which is defined as the time between the transfer of the last byte of the request from the Akamai node and the receipt of the first byte of the response from the origin server. Hence, the turnaround time offers a reasonable estimate of network delay. We collected these statistics at all 27 performance monitors for downloads made from about 80 customer content providers. The data was collected between Thursday, 23rd January, 2003 and Sunday, 26th January, 2003 (inclusive). Of the 80 content providers, 20 are the top customers of Akamai; that is, those for which the Akamai network serves the largest number of bytes.

Data Set \mathcal{H}_1 : For each Akamai server in a given city, this data set contains the average turnaround times for requests made by Akamai servers back to the origin content provider servers (Figure 3(b)). These requests are typically initiated when an Akamai server does not have a valid object cached and has to retrieve it from the origin server. These turnaround times are averaged every hour across *all* the requests sent to *every* origin content provider. We collected this data for each hour over two five-day periods: Monday, 6th January 2003 to Friday, 10th January 2003 and Monday, 13th January 2003 to Friday, 17th January, 2003 (both inclusive).

As mentioned above, our primary performance metric is the turnaround time, which indicates roughly the delay on the underlying path to the Web server. Since the customer content providers of the CDN are large Web servers, we expect their servers to be well-provisioned, and therefore the observed turnaround time should be constituted mainly of network delay with almost no delay due to the

Web server itself. Note that this delay, and its variation, is one of the crucial factors determining the performance of downloads from the content provider (since the TCP throughput is dependent on the observed round-trip time of the underlying path). A more complete metric would have been the absolute throughput for the transfer, or its combination with the turnaround time. Although we did have the download times for objects in data set \mathcal{A}_1 , the objects were typically on the order of 10KB and hence the download times may not be indicative of the long-term TCP throughput (or typical download speed) on the path. Nevertheless, the turnaround time metric accurately captures the performance of small downloads (< 10KB) and also captures the key component determining the performance of larger downloads.

3.1 Performance Benefits: 2-Multihoming

To quantify the performance benefits of enterprise 2-multihoming, we use the data set \mathcal{A}_1 . We compare the performance achieved by using the best provider link for each download, relative to that of using a single provider for all downloads. We average this ratio over downloads from all of the content providers and report this normalized performance metric. We also must be careful to compare only those transactions for which both performance monitors successfully downloaded the object at roughly the same time. We select cities in the U.S. with 2 performance monitors, giving us four locations: Atlanta, Chicago, Dallas and New York. The rest of the cities have only one performance monitor. The monitor nodes, each connected to different upstream providers can be used to measure the benefits of 2-multihoming employing the respective providers. More formally, the computation may be expressed as:

$$N_X = \frac{\sum_{i,t} (M_X(P_i, t) / M_{best}(P_i, t))}{Numvalid(P_i, t)}$$

where N_X is the performance of using ISP X , relative to 2-multihoming. $M_X(P_i, t)$ denotes the value of the turnaround time for the transfer initiated at time t by the monitor node attached to ISP X to retrieve an object from content provider P_i . Similarly, $M_{best}(P_i, t)$ is the best value (across both ISPs) of the the turnaround time for a transfer to the same city at time t from content provider P_i . The sum in the numerator is over all P_i, t pairs such that there was a transfer logged at time t to content provider P_i via *both* the providers A and B . $Numvalid(P_i, t)$ is a function that simply counts the total number of such P_i, t pairs. Notice that the optimal value of N_X is 1 and this occurs whenever one of the two ISPs is consistently better than the other. If $N_X > 1$, then $N_X - 1$ denotes the maximum

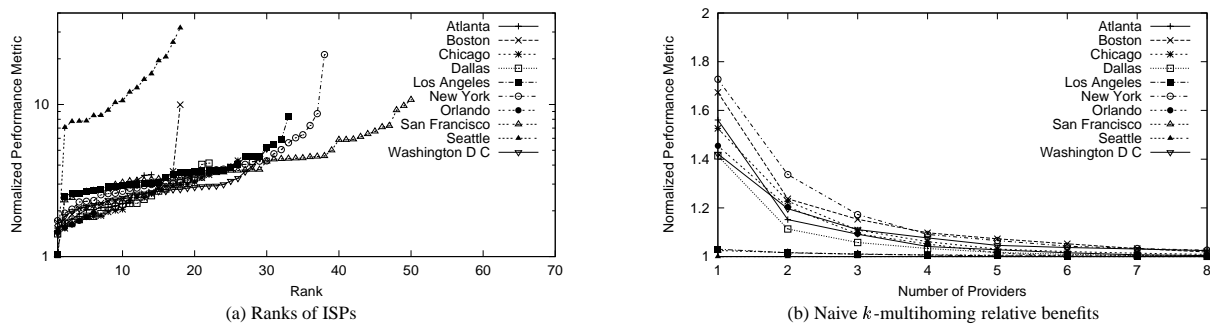


Figure 5: Naive k -multihoming: Figure (a) shows the 1-multihoming performance of the ISPs in each city, with ISPs ranked according to their performance. Figure (b) shows the diminishing returns from k -multihoming in each city.




Figure 6: Naive k -multihoming absolute benefits: The graphs show the 10th percentile, median and 90th percentile turnaround times for k -multihoming solutions. Note that the y-axes are on different scales.

improvement in performance possible from multihoming to both the ISPs (*i.e.*, from 2-multihoming), compared to the performance seen while using ISP X alone.

We quantize the time stamps on each download in \mathcal{A}_1 to integers corresponding to the number of minutes elapsed from a fixed point. Since the monitors download objects at roughly 6 minute intervals, we round the time stamp to a multiple of 6 that is at most 3 (minutes) away from the true time stamp. If there are two or more downloads (in cases where this happened, there were at most 2 downloads) from the same monitor to the same content provider mapping to the same rounded time-stamp, we pick one randomly.

The results for the performance benefits from 2-multihoming at each of the four cities are shown in Figure 4(a), which indicates the value of N_X for each ISP X . Each of the two ISPs in the four cities were tier-1 providers (*i.e.*, very large national carriers) [15]. In all four cities, 2-multihoming clearly offers performance benefits, albeit to varying degrees. For example, Chicago’s ISP1 provides nearly optimal performance by itself ($N_{ISP1} = 1.09$). However, in each of the other three cities, the minimum performance benefit from 2-multihoming is at least 25% on average. Figure 4(b) illustrates the absolute performance improvement for the median, 10th, and 90th percentile turnaround time. Note that 2-multihoming uniformly improves the maximum turnaround times, but has less effect on the median and minimum performance. Also, the extent of the absolute improvement varies across cities. Figure 4(c) shows the fraction of time when one of the two ISPs provides better performance than the other. Except in Chicago where $ISP1$ is used almost 90% of the time, both the ISPs in the other cities are put to use for roughly equal amounts of time in the optimal schedule.

3.2 Enterprise k -Multihoming, $k > 2$

So far, we have only considered multihoming to two upstream

providers. The data set \mathcal{A}_1 does not permit us to analyze k -multihoming for $k > 2$, since there are at most two monitor nodes per city. However, we can find a lower-bound on the maximum performance benefit from k -multihoming for $k > 2$ from the data set \mathcal{H}_1 as described below.

Recall that the \mathcal{H}_1 data set includes turnaround times recorded each hour, averaged across requests from a given CDN server to all customer content providers. As a result, we cannot analyze the case where the enterprise chooses the best link for transfers on a per-destination basis. The performance benefit, in this case, is computed based on a coarser form of multihoming, in which a given provider is used for *all* transfers to and from the enterprise, regardless of the destination. We refer to this as *naive k -multihoming*. This is in contrast to the analysis above for 2-multihoming in which the enterprise is able to pick the best provider for each destination at each time instant (we call this *true k -multihoming*). The maximum performance benefit from true k -multihoming is bound from below by naive k -multihoming.

We compute the performance benefits from naive multihoming in a manner similar to the 2-multihoming case:

$$N_{OP_k} = \frac{\sum_t (HT_{OP_k}(t)/HT_{best}(t))}{Numvalid(t)}$$

where N_{OP_k} is the performance of using the k -multihoming option OP_k in a given city, relative to the performance of using all available ISPs. $HT_{OP_k}(t)$ denotes the best average turnaround time performance among the k ISPs in the set OP_k at hour t . $HT_{best}(t)$ is the best average turnaround time performance at hour t over all the available carriers. The sum in the numerator is taken over all hours t for which all the k ISPs have the average turnaround time statistics logged in the data set \mathcal{H}_1 . $Numvalid(t)$ counts the num-

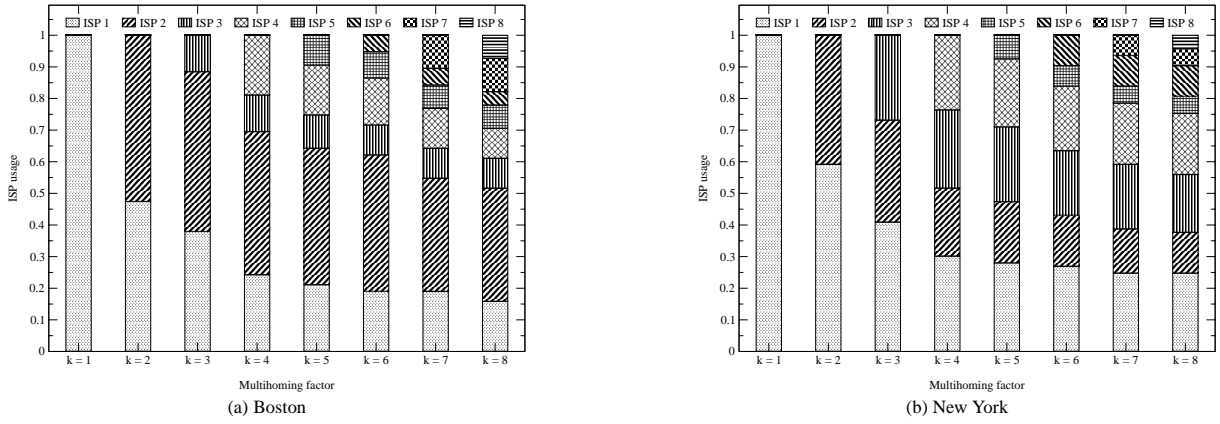


Figure 7: Relative utilization of ISPs: For the cities of Boston and New York, respectively, the graphs show the fraction of time the ISPs in the naive k -multihoming solutions at the city are utilized in the optimal schedule.

ber of such instances t (for a very small fraction of the hours, the average turnaround time data was unavailable for certain networks).

In Figure 5(a) we plot the performance metric NO_{P_1} for each ISP in the city against its rank (The ISP with rank 1 is the best in the city). The graph shows the first week of data; the second week is very similar. Notice that in some of the cities, there are a few ISPs (sometimes just one ISP) that give significantly better performance than the others. For example, the best ISP in Seattle provides at least 7 times better performance as any other ISP. There are also cities in which many ISPs provide similar performance.

From Figure 5(a) it is apparent that, in some cities, there were in excess of 50 providers (e.g., San Francisco). Evaluating all the $\binom{50}{k}$ options for k -multihoming to determine the best naive k -multihoming option is computationally expensive. We reduce the amount of computation by evaluating k -multihoming options against the performance of up to 20 top providers in each city (chosen based on their 1-multihoming performance). This has a negligible impact on our results, as our analysis showed that the performance of the top 20 ISPs is virtually indistinguishable from the performance using all available ISPs in the city (these results are omitted).

In Figure 5(b), we show the maximum performance benefits from naive k -multihoming for the first week of data in \mathcal{H}_1 (again, the second week results are similar). Notice that $k > 1$ provides significantly better performance than 1-multihoming in most locations. For a few cities, however, the performance benefit is not as substantial due to a single ISP providing the best performance almost all the time (e.g., Los Angeles). Also, beyond $k = 4$ the benefit from naive k -multihoming is only marginally better than at smaller values of k for most cities.

ISP	Rank	1-multihoming performance	k -multihoming performance
ISP 1	1	1.72	1.72
ISP 2	2	1.93	1.33
ISP 3	9	2.61	1.17
ISP 4	3	2.05	1.09
ISP 5	4	2.29	1.07
ISP 6	19	3.16	1.04
ISP 7	17	3.03	1.03
ISP 8	13	2.93	1.03

Table 1: Ranks of the ISPs in the k -multihoming solutions at New York, $k \leq 8$, in the order in which they are added, along with the incremental performance improvement.

Figures 6(a)-(c) show the absolute improvements due to multihoming in the 10th percentile, median, and 90th percentile turnaround

times. In most cities, k -multihoming improves the maximum turnaround time performance up to $k \leq 4$. However, most of the lines in (a) and (b) are fairly flat, which indicates that the median and minimum turnaround times are not reduced much as k is increased. To summarize, enterprise multihoming offers the greatest performance benefits to high-latency transfers.

Table 1 shows the order in which ISPs get added to the k -multihoming solution in New York for increasing values of k . For each ISP, we also show its 1-multihoming rank and performance. Notice that the best k -multihoming solution does not necessarily comprise the k best 1-multihoming options (e.g., the third ISP has a rank of 9 based on its 1-multihoming performance). Rather, ISPs are added based on their contribution to the overall k -multihoming performance.

We also consider how often each of the providers is employed in the optimal schedule for enterprise multihoming. In particular, we are interested in whether a provider’s contribution towards performance improvement is proportional to the frequency with which it is used in the optimal schedule. The results for two cities, Boston and New York, are illustrated in Figure 7. The results show clearly that the contribution to performance is not proportional to the usage. For example, the 6th ISP in New York is used for a significant fraction of time in the 6-multihoming solution (Figure 7(b)). However, the marginal benefit of adding this ISP to the 5-multihoming solution was less than 0.02 (Figure 5(b)). It is also possible that an ISP belonging to the best k -multihoming solution is utilized for a very small fraction of time in the optimal schedule, but, whenever used, contributes significantly to improving the overall performance. For example, ISP1 is used for smaller fraction of time than ISP2 for the best naive 2-multihoming solution in Boston (Figure 7(a)). However, the contribution of ISP1 to the overall benefit due to 2-multihoming is clearly larger than that of ISP2.




Figure 8: \mathcal{A}_2 data collection: Number of ISPs per city are in (a). Akamai servers (“clients”) in (b) download objects from designated servers in 5 cities, each connected to a different upstream ISP (“multihomed Web servers”).

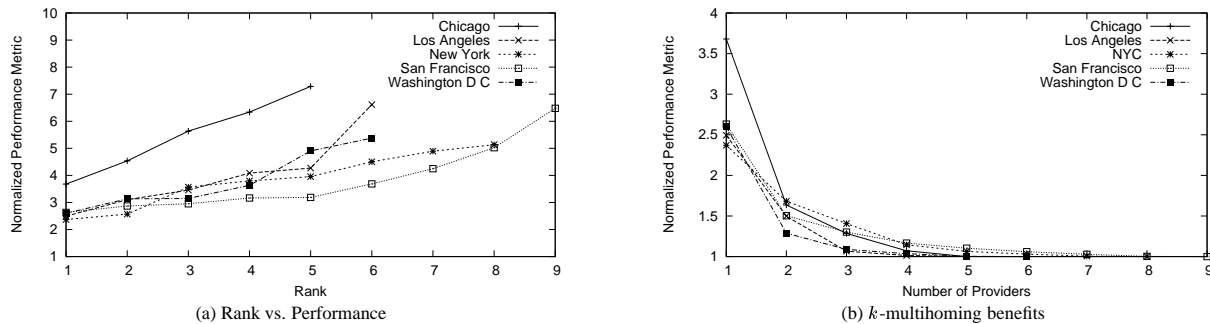


Figure 9: Web server multihoming: Figure (a) plots the 1-multihoming performance of ISPs vs. their rank. Figure (b) shows the diminishing returns from Web server multihoming.

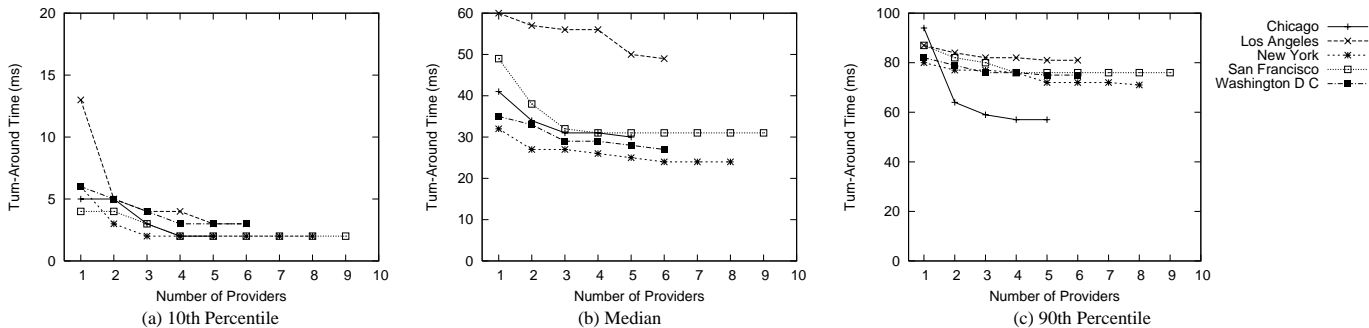


Figure 10: Absolute benefits from Web server multihoming: The three figures plot the 10th percentile, median, and 90th percentile download time for k -multihoming options. Note that the y -axes are on different scales.

4. WEB SERVER PERSPECTIVE

From a Web server’s point-of-view, the performance benefits from multihoming should be reflected in the end-to-end performance of the requests it serves to a large number of widely distributed clients. With this goal in mind, we collect a new data set \mathcal{A}_2 to understand the benefits of multihoming.

Data Set \mathcal{A}_2 : In five cities – Chicago, Los Angeles, New York, San Francisco and Washington D. C. – we select Akamai servers attached to distinct upstream carriers. The servers in each city collectively act as stand-ins for a multihomed Web server operating in that city. We select Akamai servers in various U.S. cities, other than the above five to serve as distributed Web clients. We perform periodic Web transfers from each server in the five cities to each “client,” as illustrated in Figure 8(b).

The number of Akamai servers in the five cities (reflecting the number of ISPs we test for Web server multihoming) is tabulated in Figure 8(a), totaling 34. The number of servers in the other U.S. cities was 40.¹ For each Web server stand-in, all of the remaining servers ($40 + 33 = 73$ in total) act as stand-ins for clients downloading the same 50KB JPEG object from the server (Figure 8(b)). The downloads occur at regular 6 minute intervals. All the servers collect statistics for these downloads, identical to those collected by the monitor nodes in the data set \mathcal{A}_1 (described in Section 3). As before, we focus on the turnaround times for the downloads. Our results are based on data collected between 4th June 2003 and 8th June 2003.

¹Using more servers was precluded by Akamai’s contractual agreements with the respective carriers; since we perform active downloads, we had to be careful not to violate these contractual constraints.

4.1 Performance Benefits

To understand performance benefits of Web server multihoming, we adopt a similar methodology as with enterprise multihoming. For each download, we compare the client-perceived turnaround time achieved by using the best provider among all those available in the city, with that of using the best provider in a candidate multihoming option. We average this ratio over transfers to all clients, and report the minimum normalized performance metric (the minimum is taken over all candidate options). As before, we compare only those transactions for which there was a successful transfer over all ISPs at roughly the same time.

Formally, $M_{best}(A_i, t)$ denotes the best value of the turnaround time for a transfer to client A_i ($i = 1, \dots, 73$) at time t , across all available carriers in a city. For a k -multihoming option OP_k , let $M_{OP_k}(A_i, t)$ be the best turnaround time across just the ISPs in the set OP_k . We compute the performance benefit from the option OP_k as follows:

$$N_{OP_k} = \frac{\sum_{i,t}(M_{OP_k}(A_i, t)/M_{best}(A_i, t))}{Num_{valid}(t)}$$

The sum is over those times t when transfers occur from all the ISPs in the city to client A_i . $Num_{valid}(t)$ is the number of such time instances.

In Figure 9(a), we show the normalized 1-multihoming performance of each ISP as a function of its rank. The performance provided by the ISPs in a given city is quite different (except in San Francisco, where the top five ISPs exhibit virtually identical performance). However, no single ISP provides ideal, or close-to-ideal, performance in any city. In Figure 9(b), we plot the normalized benefits from Web server multihoming as a function of the number of providers. Again, similar to the case of enterprises, multihoming

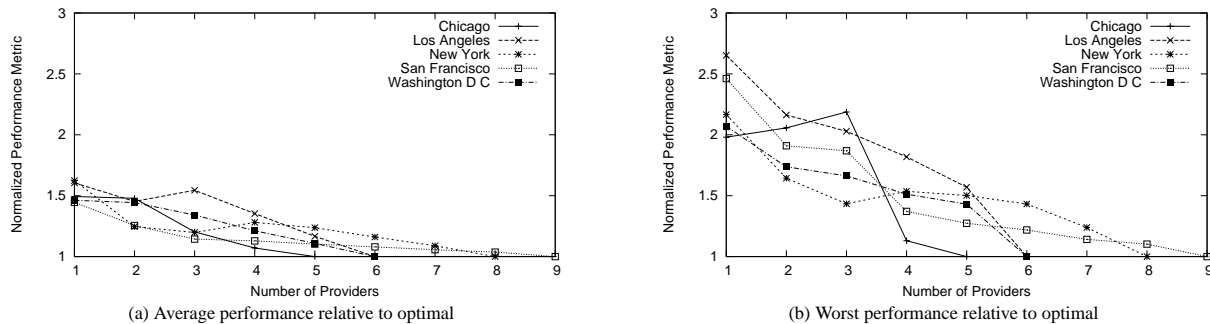


Figure 11: Sub-optimal choices: Graph (a) shows the *average* performance across all k -multihoming options. Graph (b) shows the performance of the *worst* k -multihoming option. In both graphs, the y-axis is relative to the optimal k -multihoming solution.

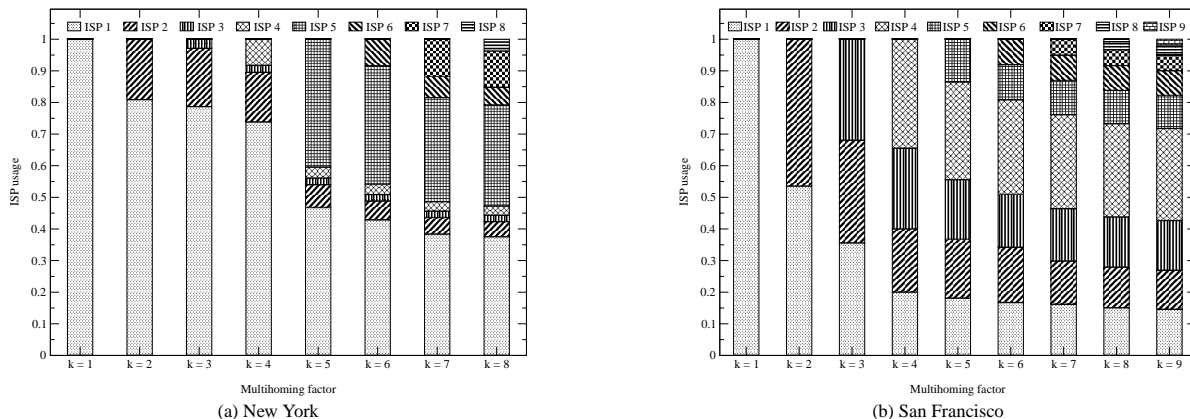


Figure 12: ISP usage in Web server multihoming: The graphs show relative fractions of time the ISPs in the best k -multihoming solutions are utilized in two cities.

significantly improves average performance. We also see that the marginal benefit is small beyond 4 upstream providers.

Figure 10 shows the values of the 10th percentile, median, and 90th percentile turnaround times for the best k -multihoming options. For most of the cities, we see that the improvement in the 10th percentile and median is more pronounced compared to the 90th percentile (in contrast to the enterprise multihoming scenario in Figure 6 in which the highest turnaround times were improved the most). In summary, Web server multihoming does not necessarily improve only the high latency downloads; low-latency transfers could also benefit.

In Figure 11, we illustrate the impact of choosing a sub-optimal set of providers for a k -multihoming solution. The values on the y-axis in these graphs are relative to the performance of the optimal solution shown in Figure 9. For $k \leq 4$, the *average* performance of k -multihoming is at least 15% worse than that of the optimal choice and could even be as bad as 40% (e.g., $k = 2$ in Chicago). The difference between optimal and random choices of ISPs is small for $k > 4$. In Figure 11(b) we show the performance of the worst k -multihoming option. A poor choice of upstream providers could result in performance that is at least twice as bad as the optimal choice. Therefore, while multihoming offers potential for significant performance benefits, it is crucial to carefully choose the right set of upstream providers.

Finally, Figure 12 shows the frequency at which the ISPs in the best k -multihoming solution are utilized for two sample cities: New York and San Francisco. As with the enterprise perspective, we see that even though an ISP is used for a significant fraction of the time in the optimal solution, it may offer only marginally superior performance (e.g. ISP 5 in New York).

5. RELIABILITY

Multihoming enhances the reliability of stub networks by helping them stay connected to the Internet during wide-area routing failures. However, the extra reliability offered by multihoming depends, to a large extent, on the redundancy or diversity introduced by multihoming in the underlying network paths. For example, if a multihomed network chooses upstream providers which route traffic to distant peers via paths with significant overlap, a failure in the overlapping portions is likely to disconnect the network from many destinations. Our focus in this section is to understand the reliability benefits of multihoming by quantifying the diversity in network paths that multihoming provides.

To analyze the reliability benefits of multihoming, we use a data set called \mathcal{T}_1 , which contains traceroute measurements from a set of 50 geographically diverse nodes deployed by Keynote Systems to select Akamai servers located in three cities: Chicago, New York, and San Francisco. The Keynote nodes are located in 27 different cities, with two Keynote nodes per city, each singly-homed to different tier-1 providers. Thus, they represent the perspective of well-connected endpoints (e.g., large enterprises connected to a major PoP) rather than individual end-users. We choose Akamai servers that are each singly-homed to the twenty top ISPs serving each city (in terms of performance). Therefore, the data set \mathcal{T}_1 is a collection of 3000 traceroutes ($50 \times 3 \times 20$). Note that this data set only provides information about the IP-level connectivity of the network, and hence does not give an indication about lower-level physical redundancy (e.g., cables, fiber trunks, etc.). As such, the reliability measurements from this data describes robustness to IP-level failures (e.g., routing, router configuration, congestion/traffic




Figure 13: Path diversity metrics: Figure (a) illustrates the value of \mathcal{R}_1 for the “dashed” subtree of the tree rooted at the Keynote node K . Figures (b)–(e) illustrate the computation of \mathcal{R}_2 .

flooding, etc.), and not to hardware failures (e.g., power outages, fiber cuts, MPLS failure, etc.).

5.1 Quantifying the Reliability Benefits

Our basic approach in evaluating the reliability benefits of multihoming is to combine paths from each Keynote node to each of the Akamai servers in a given city. The combined paths result in a tree rooted at a Keynote node, with leaves that are Akamai servers in the same city, and connected to different ISPs. The Akamai servers are stand-ins for a multihomed network and the Keynote nodes represent typical destinations with which the multihomed network communicates.

We denote the twenty selected Akamai servers in a city as S_1, \dots, S_{20} , and $OP_k = S_{j_1}, \dots, S_{j_k}$ is a k -multihoming option (where $k \leq 20$) consisting of a subset of these servers. The tree, $T_{i,k}$, is the union of paths from a Keynote node, K_i , to each of the servers S_{j_1}, \dots, S_{j_k} . $E_{i,k}$ is the total number of edges in the tree $T_{i,k}$. $P_{i,k}$ denotes the sum of the hop-counts of the individual k paths that constitute the tree $T_{i,k}$. Thus, from these definitions, $E_{i,20}$ is the number of edges in the tree, $T_{i,20}$, rooted at K_i with all 20 Akamai servers in the city as leaves.

Our analysis of diversity due to multihoming is based on two metrics:

- $\mathcal{R}_1(OP_k) = \frac{1}{50} \sum_i \frac{E_{i,k}}{E_{i,20}}$: $\mathcal{R}_1(OP_k)$ is proportional to the fraction of edges in tree $T_{i,20}$ that also belong to tree $T_{i,k}$, averaged over trees rooted at all Keynote nodes i . So, for example, in Figure 13(a) the tree induced by the 2-multihoming solution with ISPs S_1 and S_2 (dashed lines) shares 4 links with the tree for all 20 ISPs in the city. Intuitively, $\mathcal{R}_1(OP_k)$ estimates how much of the total redundancy provided by all 20 ISPs in a city can be achieved by using just the k providers in the set OP_k .
- $\mathcal{R}_2(OP_k) = \frac{1}{50} \sum_i \frac{P_{i,k} - E_{i,k}}{E_{i,k}}$: $\mathcal{R}_2(OP_k)$ is proportional to the expected fraction of edges in tree $T_{i,k}$ that are shared by two or more paths in the tree. Therefore, $\mathcal{R}_2(OP_k)$ estimates the expected fraction of overlap in end-to-end paths resulting from choosing OP_k . As an example, Figure 13(b) shows two paths from K , each with hop-count 3, thus $P_{i,k} = 6$. The number of edges in the tree is 5, and \mathcal{R}_2 is computed as shown. Figure 13(c) shows the effect of having more shared edges in a similar topology.

Note from the above definitions that higher values of \mathcal{R}_1 and lower values of \mathcal{R}_2 are preferable.

\mathcal{R}_1 is primarily dependent on the initial choice of ISPs (the 20 ISPs in our case). If most of the chosen ISPs have a significant overlap in the paths to and from arbitrary points in the Internet, then

most k -multihoming options will have high values for \mathcal{R}_1 , despite providing only modest path diversity. However, this is effectively captured by \mathcal{R}_2 which would give a value close to 1 if there is a significant overlap in the underlying paths.

\mathcal{R}_2 has an unfavorable bias against greater numbers of ISPs. That is, adding any provider to a k -multihoming solution to give a $(k+1)$ -multihoming solution may result in an inferior value of \mathcal{R}_2 . This is because the $(k+1)$ -st provider will likely have a non-zero intersection in the network paths with the remaining providers (compare Figures 13(b) and (d), for example).

Both the metrics \mathcal{R}_1 and \mathcal{R}_2 have an undesirable bias in favor of long paths. For two k -multihoming options OP_1 and OP_2 with the same number of shared hops in the trees rooted at some K_i , if the paths due to OP_2 are longer than those due to OP_1 , this will result in OP_1 having inferior values for both \mathcal{R}_1 and \mathcal{R}_2 . This is evident for \mathcal{R}_2 in comparing Figures 13(b) and (e). Although we present results for \mathcal{R}_1 and \mathcal{R}_2 independently, they should be evaluated in combination when considering the reliability benefits of a multihoming option.

5.2 Reliability Benefits of Multihoming

We compute the most reliable k -multihoming options in the three cities according to the above two metrics. The results for the three cities are shown in Figure 14. To avoid having to factor in the biases of the metrics into our analysis, we restrict our observations to the comparison of the optimal k -multihoming solution to other sub-optimal solutions, namely ISP choices that give average and worst diversity according to our metrics.

Comparing the optimal k -multihoming solution to a random choice, we see that there is a clear performance benefit for all values of k at each of the three cities. According to either metric, the optimal solution offers at least 25% improvement over a random choice. The difference between the optimal choice of providers and a poor choice of providers is even more pronounced, with the optimal solution being roughly 50% better according to metric \mathcal{R}_2 and 30% better according to \mathcal{R}_1 . The graphs demonstrate a clear advantage from multihoming. However, they also underscore the importance of choosing a good set of upstream providers.

Notice that the above analysis of multihoming applies roughly to both the enterprise and Web server perspectives. In the former case, the Keynote nodes, with their deployment in major cities and connectivity to large providers, serve as stand-ins for popular Web servers from which a typical enterprise might receive most of its data. In the latter perspective, the paths to Akamai servers from Keynote nodes represents a sample of typical paths taken in the Internet to reach popular Web servers. This is because a large number of paths usually traverse through large, national ISPs before reaching the destination. Our traceroutes capture the latter part of these paths, from the large Internet carriers to the ultimate destinations.

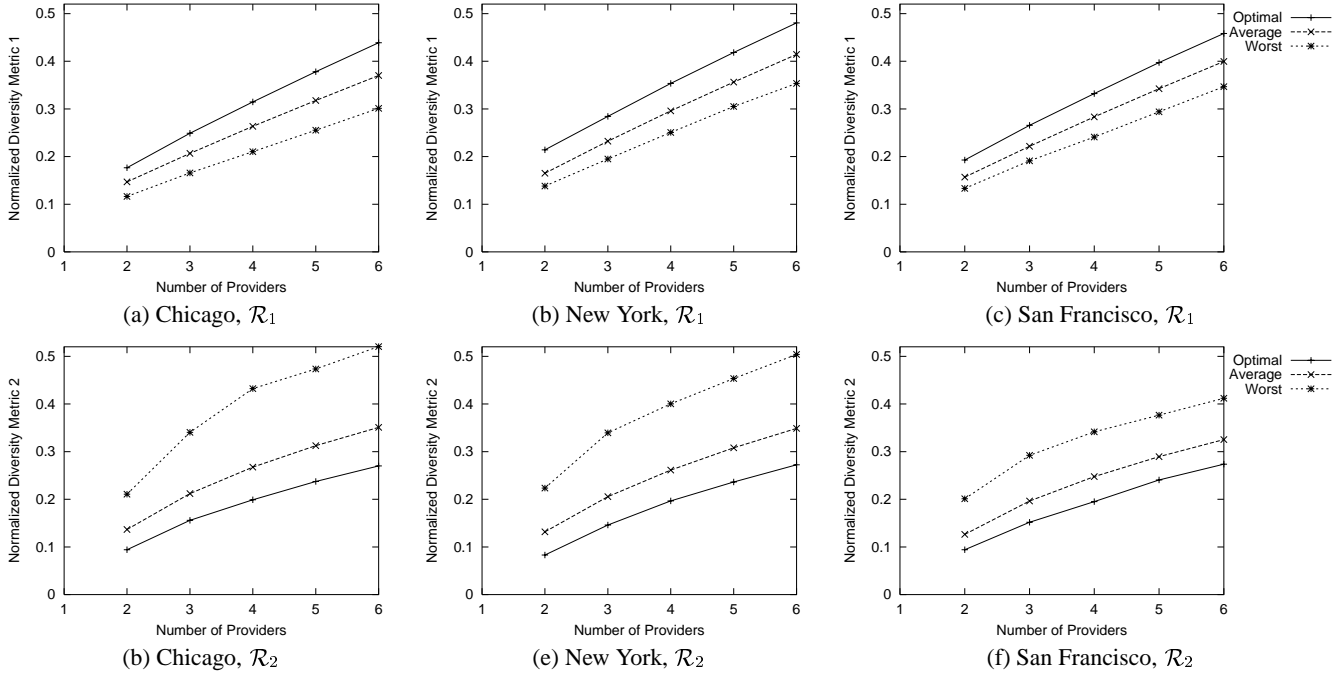


Figure 14: Path diversity benefits: Figures (a), (b) and (c) show the values of the metric \mathcal{R}_1 for the optimal, worst and average multihoming for the three cities. Figures (d), (e) and (f) show the corresponding results for \mathcal{R}_2 .

Our current data set limit our study of network reliability to evaluating one direction of network communication. However, reliable bi-directional connectivity is necessary for useful communications. Collection of the necessary measurement data and analysis of bi-directional connectivity is future work.

6. DISCUSSION

In the previous sections, we highlighted the maximum benefits from multihoming in each metropolitan area that we measured. In this section, we briefly explore some related issues, including how to select an ISP or set of ISPs, and the impact of more practical route selection strategies. We also discuss some of the limitations of our methodology.

6.1 Practical Issues

From a performance perspective, we evaluated an almost optimal form of control over data traffic – networks could decide to change routes frequently and were able to choose the optimal ISP for each transfer. In practice, however, networks must choose some reasonable time granularity on which to make routing changes and use current and past observations to guide their decisions. In this section, we briefly explore the implications of this more realistic scenario.

We assume that a network regularly monitors the performance of transfers using all of its chosen k ISPs. These measurements are used to drive the ISP selection at time intervals of, say, T minutes using an exponentially-weighted average of the performance over each ISP for a given destination². Recent samples of performance are given more weight compared to older samples. At the end of every T minutes, the network chooses the ISP with the best weighted

²If performance at time t_k was s_{t_k} and the previous performance sample was from time t_{k-1} , then the weighted average performance at time t_k is: $Avg_{t_k} = (1 - e^{-(t_k - t_{k-1})/\alpha})s_{t_k} + e^{-(t_k - t_{k-1})/\alpha}Avg_{t_{k-1}}$ where $\alpha > 0$ is a constant.

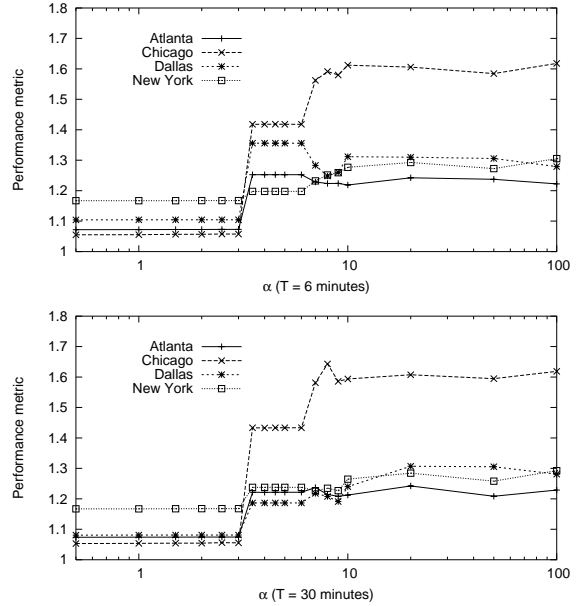


Figure 15: Implementation of 2-multihoming at an enterprise: The graphs correspond to two values of the time interval at which route selection decisions are made – 6 and 30 minutes.

performance for transfers to the destination. Although the network has chosen a single ISP for the given destination for this time period, we assume that it continues to monitor the performance of all k ISPs over the next period. This is in general a challenging problem. Some common approaches include active probing from enterprises to known endpoints of interest (e.g., branch offices, partners, etc.), or using designated objects to measure delays at Web sites (e.g., as described in Section 2.1). Using this more realistic




Figure 16: k -multihoming at a Web server: The three graphs show the impact of attaching increasingly higher importance to stale data for route selection.

route selection algorithm, we revisit the 2-multihoming experiment of Section 3 and the k -multihoming experiment of Section 4. Figure 15 shows the impact of changing the time period length, T , and the weight given to recent samples, α , on the benefit from enterprise multihoming. The three graphs show that, irrespective of the length of T , attaching higher weights to more recent samples (*i.e.*, smaller values of α) offers a distinct performance advantage. In fact, comparing with Figure 4(a), we see that that beyond $\alpha = 3$, the performance is sometimes worse than simply using the better of the two ISPs in the city, indicating that stale information could lead to selecting sub-optimal routes.

In Figure 16, we show the result of using a $T = 30$ -minute time interval for route selection in a k -multihoming Web-server for $\alpha = 1, 5, 10$. The performance improvement is the highest (and significant) when $\alpha = 1$. However, even using all providers, the performance is substantially worse than optimal. At $\alpha = 5$, this algorithm offers only a marginal performance improvement as more ISPs are added. However when α is even higher, the performance actually degrades with the number of providers. In this case, the network often makes incorrect decision on which ISP to use due to increased reliance on older data. These observations show that using timely and accurate performance samples is key to extracting performance benefits.

6.2 Choosing ISPs

One of the goals of this work was to provide guidance to subscribers on how to choose ISPs in a multi-provider multihoming scenario. Like any optimization, satisfying both the reliability and performance goals of a customer simultaneously can be difficult. Also, we have largely ignored economic considerations, which play an important role in ISP selection. In addition, we have not considered how providers might respond to customers multihoming for performance reasons, for example by changing their cost structures. Despite these limitations, we believe that our findings can be used to help guide a reasonable ISP selection strategy as follows.

We have described various metrics to evaluate the path diversity provided by a set of ISPs. Although we have not provided a mapping of this metric to real world reliability, we believe that such a mapping might be possible. First, a customer could use these metrics to roughly identify how many ISPs are likely to be needed to meet their reliability requirements. This would determine k in a k -multihoming scenario. Second, the customer could use the metrics as a simple screening mechanism to eliminate the k -ISP combinations that fall below a particular reliability threshold.

Given the remaining available combinations of ISPs, a customer needs to identify the set that provides the best performance. As shown in Section 3.2, simply choosing the set with the k best individually performing ISPs is not sufficient – the ISPs must be evaluated in how they complement each other.

7. RELATED WORK

BGP’s policy-driven routing, with its inefficiencies and the need for mechanisms to counter these inefficiencies, continue to capture the attention of networking researchers. Several studies have identified and measured the deficiencies of relying on BGP routes in the Internet [10, 16]. These studies have shown that policy-routing largely exacerbates delays and, consequently, end-to-end performance in the Internet.

Taking cue from these studies, researchers have come up with several interesting ways to circumvent these problems. Most of these propose routing packets along an overlay network, effectively bypassing BGP routing [1, 11]. These studies propose techniques to choose overlay routes, and show that these alternate paths offer significant improvements in bandwidth, loss, and delay. In contrast, our focus is on quantifying the performance improvement achievable with an optimized use of existing BGP paths, where stub networks can choose among different ISP links to reach a particular destination but have no further control on the rest of the path.

This issue has been addressed in part by “intelligent route control” solutions marketed by companies such as RouteScience [9] and Sockeye [13]. These commercial products allow enterprises and data centers to dynamically select among their upstream providers for optimal performance. However, it is unclear as to how the set of providers should be chosen. In addition, to the best of our knowledge there has been no study that objectively quantifies the impact that such products can have on observed performance.

The issue of path diversity, and its applications and prevalence in the Internet, has been similarly well-studied. There are several solutions for streaming media that exploit the diversity of wide-area paths [3, 7, 2]. In these schemes, media streams are encoded into multiple complimentary descriptions and transmitted over provably diverse paths, for example leveraging CDN servers to attain diverse paths.

Preliminary thoughts about the specific issue of measuring provider path diversity appear in [8]. Based on measurements from a single PoP in San Jose, and paths through four different providers to about 8500 end-points, the authors show that the observed diversity largely depends on the view point (in this case San Jose) and also on the probed end-points. More recent work studies diversity within the Sprint network by measuring paths between PoPs, and observes that 90% of the PoP pairs have 4 or more link-disjoint paths between them [17]. This work also examines Internet-wide diversity using data from CAIDA [5], and concludes that exploiting the high level of inter-AS path diversity requires that stub networks are multihomed.

Our approach complements these efforts by focusing on the path diversity enabled through multihoming. Our measurements concentrate on capturing the diversity characteristics of significant portions of representative Internet paths originating in tier-1 provider

networks, and destined for multiple networks in several major cities. In addition, we define two intuitive metrics that, we believe, capture key desiderata of path diversity to a reasonable extent.

8. SUMMARY

Multihoming to several network carriers provides a way for large enterprises and content providers to improve the resilience of their network connectivity to failures and outages. Recently, there has been increasing interest in leveraging multihoming to improve network performance as well, by choosing the best performing provider link for transfers to various destinations. In this paper, we quantify the maximum extent to which multihoming to multiple providers improves average performance. We analyze variations of optimal k -multihoming strategies from both enterprise and content provider perspectives, in which the best available ISP is used for data transfers at each instant. Our performance analysis is based on three large data sets consisting of measurements taken at servers and monitoring nodes in the Akamai CDN. These nodes are widely distributed and allow us to evaluate performance across a variety of ISPs in several major metropolitan areas.

From the enterprise perspective, we observe an average performance improvement of 25% or more for 2-multihoming in three of the four large metropolitan areas considered. We also analyze the benefits of multihoming to more than 2 providers, and find little incremental improvement beyond 4 providers. Interestingly, with increasing k (beyond 4 providers), each new ISP added is used for a significant fraction of time to achieve the optimal performance, despite its marginal contribution toward improving performance.

We observe similar trends when considering multihoming benefits from the content provider perspective. Again, the performance benefit of multihoming is significant, but the improvements are very small beyond 4 providers in all 10 cities measured. We also find that the optimal multihoming solution when $k \leq 4$ exceeds the performance of a random selection of providers by an average of 15% and by as much as 40%. Our findings clearly suggest that while multihoming offers substantial performance benefits, a careful choice of upstream providers is crucial.

In an evaluation of reliability benefits due to multihoming, we quantify the IP-level path diversity with measurements from a widely distributed set of well-connected nodes to hosts connected to about 20 different ISPs in each of 3 major cities. We evaluate reliability of a given k -multihoming solution with two metrics: i) the fraction of the total path diversity captured by the solution, and ii) the degree of overlap in the paths. These initial results show that the optimal k -multihoming solution (with respect to each metric) provides a roughly 25% improvement over a random choice. Again, we see that realizing reliability benefits from multihoming requires a careful choice of providers.

Finally, we consider more practical multihoming scenarios in which routing decisions are made on a coarser time-scale, using past observations to guide link selection. We also describe a strategy using our performance and reliability metrics to choose providers in a k -multihoming scenario when there is a trade-off between reliability and performance.

Acknowledgment

We are very grateful for the support and assistance of Roberto De Prisco and Ravi Sundaram of Akamai Technologies, and Herbie Pearthree and John P. Cazares of IBM Global Services. We also thank Geoff Voelker and our anonymous reviewers for their valuable feedback and suggestions.

9. REFERENCES

- [1] D. Andersen, H. Balakrishnan, M. Kaashoek, and R. Morris. Resilient Overlay Networks. In *Proceedings of the Symposium on Operating System Principles*, Banff, Canada, October 2001.
- [2] K. Andreev, B. Maggs, A. Meyerson, and R. Sitaraman. Designing overlay multicast networks for streaming. In *Proceedings of ACM Symposium on Parallel Algorithms and Architectures*, San Diego, CA, June 2003.
- [3] J. Apostolopolous, T. Wong, W. Tan, and S. Wee. On multi-description streaming with content delivery networks. In *Proceedings of IEEE INFOCOM*, New York, NY, June 2002.
- [4] T. Bu, L. Gao, and D. Towsley. On characterizing BGP routing table growth. In *Proceedings of IEEE Global Internet Symposium*, Taipei, Taiwan, November 2002.
- [5] B. Huffaker, D. Plummer, D. Moore, and k claffy. Topology discovery by active probing. In *Proceedings of the Symposium on Applications and the Internet*, Nara, Japan, January 2002.
- [6] Keynote Systems, Inc. <http://www.keynote.com>, June 2003.
- [7] Y. J. Liang, E. Setton, and B. Girod. Channel-adaptive video streaming using path diversity and rate-distortion optimized reference picture selection. In *Proceedings of IEEE Workshop on Multimedia Signal Processing*, St. Thomas, US Virgin Islands, December 2002.
- [8] K. Nayak. Measuring provider path diversity from traceroute data. In *ISMA Winter Workshop*, San Diego, CA, December 2001.
- [9] RouteScience Technologies, Inc. <http://www.routescience.com>, June 2003.
- [10] S. Savage, A. Collins, E. Hoffman, J. Snell, and T. Anderson. The end-to-end effects of internet path selection. In *Proceedings of ACM SIGCOMM*, Boston, MA, September 1999.
- [11] S. Savage et al. Detour: a case for informed Internet routing and transport. *IEEE Micro*, 19(1):50–59, 1999.
- [12] P. Smith. BGP multihoming techniques. NANOG 23, October 2001. <http://www.nanog.org/mtg-0110/smith.html>.
- [13] Sockeye Networks, Inc. <http://www.sockeye.com>.
- [14] J. W. Stewart. *BGP4: Inter-Domain Routing in the Internet*. Addison-Wesley, 1999.
- [15] L. Subramanian, S. Agarwal, J. Rexford, and R. H. Katz. Characterizing the Internet hierarchy from multiple vantage points. In *Proceedings of IEEE INFOCOM*, June 2002.
- [16] H. Tangmunarunkit, R. Govindan, and S. Shenker. Internet path inflation due to policy routing. In *Proceedings of SPIE ITCOM*, Denver, CO, August 2001.
- [17] R. Teixeira, K. Marzullo, S. Savage, and G. M. Voelker. Characterizing and measuring path diversity of Internet topologies. In *Proceedings of ACM SIGMETRICS*, San Diego, CA, June 2003.