

Understanding Delay Variations on Internet Paths

Zhenhai Duan, Kuai Xu, and Zhi-Li Zhang

Abstract—In this paper we investigate the *network* factors that may affect the user perceived end-to-end delay jitter. In particular, we identify the following three major factors: *per hop* queueing delay variations along an Internet path; intra-domain multi-path routing; and inter-domain route (i.e., AS path) alterations. By studying traceroute data collected on Internet paths, we find that 1) larger queueing delay variances are likely to be experienced at routers residing at boundaries of (or rather, links between) two (AS) network domains as well as the edge of the Internet; 2) intra-domain multi-path routing may have significant impact on end-to-end delay jitter; and 3) inter-domain route alteration adversely affects user perceived delay jitter. In this paper, we report on these results and discuss their implications in network engineering.

Index Terms—Internet Delay Measurements

I. INTRODUCTION

End-to-end delay variations play a key role in the design and operations of reliable transport protocols and real-time adaptive applications. To a large extent, a user (or rather, an application) perceived end-to-end delay jitter influences how an application behaves, and how resources in end-hosts are allocated. For example, to determine the time to wait for an acknowledgement before retransmitting a packet, TCP flow control needs to estimate both the average and deviation of the packet round-trip times of a connection. In addition, adaptive applications such as video streaming need to properly resize the playback buffer to hold early-arrived packets based on observed end-to-end delay variations. Moreover, large end-to-end delay variations make supporting delay-jitter-sensitive real-time traffic such as voice over the Internet extremely hard.

These protocols and applications in general prefer stable end-to-end delays with little or no variation. However, given the diversity of network entities, complexity of control mechanisms, and size of the Internet, it is very difficult to control the network behavior to achieve this goal. Indeed, even the origins of end-to-end delay variations are not well understood. For example,

advances in optical technology have introduced a glut of bandwidth in the Internet backbone. Consequently, packets traversing the core networks of major (tier-1) ISPs rarely experience queueing delays or losses [5], [6]. On the other hand, we, as end users, still observe large end-to-end delay variations from time to time, which render our experiences with delay-sensitive interactive applications (e.g., VoIP) often unsatisfactory. So, where does the problem (namely, large end-to-end delay variations) come from? It is towards answering this question that motivates us to carry out the measurement study reported in this paper.

Clearly, a variety of factors may contribute to the user perceived end-to-end delay variations, among which are improperly realized TCP/IP stack on end hosts, short-term traffic load fluctuations, and routing policy changes. While it is important to understand the factors pertinent to end-hosts, in this paper we will focus only on *network* factors that contribute to the user perceived delay variations. To gain a better understanding of various network factors, we adopt a “bottom-up” approach and examine the delay variations at three levels: from the link (or “per-hop”) level, to intra-domain (IP router) path level, and then to inter-domain (i.e., AS) path level. At the link level, we examine “per-hop” delay variations experienced by packets along an Internet path. In particular, we are interested in finding out where packets normally experience (more or less consistently) large queueing delay variations: on a link within an access network, or within a backbone network, or between two AS’es? At the intra-domain (IP router) path level, multi-path routing has been employed by many network domains (especially backbone networks) for traffic load balancing and reliability. Depending on how it is implemented, multi-path routing can introduce undesirable delay variations perceivable by end users. Hence we also investigate how the end-to-end delay variations can be affected by this multi-path routing practice. Due to various reasons (e.g., routing policy changes, link failure, BGP misconfiguration), AS level paths to a destination network may be altered from time to time. At the inter-domain level, we therefore examine the effect of inter-domain route alterations on the end-to-end delay variations.

We perform Internet path delay measurements using

The authors are with the Department of Computer Science & Engineering, University of Minnesota. E-mail: {duan, zhzhang, kxu}@cs.umn.edu.

the well developed and tested tool `traceroute`. To capture both the path characteristics in short-time scales and changes in longer time scales, we conduct *consecutive* `traceroute` measurements for a duration of either one hour or one and half hour several times a day on different days from the University of Minnesota (UMN) to a number of destinations (see the next section for a detailed description of the experiments). For the analysis of *per hop* packet delay variations, it is ideal to *directly* measure packet delays over each link on a given path. However, in practice such measurements are even difficult to perform for the ISPs who own the networks; needless to say, such an approach is impossible for us. To circumvent this problem, we employ an indirect method which uses `traceroute` to measure the round-trip times (RTTs) *from a source to all the intermediate routers as well as the destination*, and based on such measurements to estimate the “per-hop” delay variations. For the analysis of delay variations at intra-domain path level and inter-domain AS path level, we also rely on BGP information in addition to the `traceroute` measurements. To obtain BGP information, we set up a passive E-BGP peer (i.e., an E-BGP listener) with the University of Minnesota’s Gigapop BGP router and obtain all the BGP updates. Our findings can be summarized as follows: 1) Large delay variations are likely to be experienced at the hops (routers or links) near/between boundaries of two (AS) network domains as well as the edge of the Internet. This observation is consistent with the belief or fact that today’s Internet backbone networks are generally well-provisioned; 2) intra-domain multi-path routing may have a significant impact on end-to-end delay jitter, especially when *traffic load balancing is performed at packet level*. For example, in one case we have found that for a specific pair of source and destination, the mean RTT on one path is 116ms , while on another path, it is 81ms , measured over a time period of one and half hour. In this case it seems that packets to the same destination are routed, in a round-robin fashion, among several paths within one particular network domain; And 3) inter-domain route alteration adversely affects user perceived delay jitter. For example, our measurement data show that packets to one destination from UMN may take six different AS level paths (and eight IP router level paths) within a time span of one hour. In addition, on one AS level path, the mean RTT is 122ms , while it is 78ms on another one. In this paper we will report on these results and discuss their implications in networking engineering.

There is rich literature on the study of the Internet path characteristics such as available bandwidth, throughput, and end-to-end delay. However, to our best knowl-

edge, the analysis of *per hop* queueing delay variations reported in this paper is new. In the comprehensive Internet measurement study conducted by Paxson [7], [8], [9], the author identified and discussed the impact of route alteration on end-to-end delays. However, the `traceroute` data in this study was collected in a relatively sparse manner, typically once every one or two days. It is not immediately clear what its implications are for applications, which typically last only for seconds or minutes (and in some rare cases, hours). In our study we are more interested in the time scales that are comparable to the lifetime of applications. In [12], based on simulation study Varadham et al found that small changes in network topology can cause significant packet re-orderings, thereby greatly degrading the performance of transport protocols such as TCP. This study however focused only on the consequence of network topology changes and packet re-ordering behavior. In a more recent work by Zhang et al [13], by analyzing a large dataset collected over the Internet, the authors found that end-to-end delays on Internet paths were not mathematically or operationally “stationary,” but nonetheless are still highly predictable. In this study the authors viewed the Internet as a blackbox for the measurements and did not discuss the causes of the end-to-end delay variations. Labovitz et al [4] studied the Internet routing instability behavior and found that 99% of routing update information may not reflect the real Internet topological changes. They discussed the potential impact of routing instability on the Internet infrastructure, but did not study how it affects the end-to-end delay variations. Savage et al [11] compared the performance on the “default” path of a connection to that on the potential alternate routes, and found that for 30 – 80% of the current paths, there is an alternate path with much better performance. The objective of their study was to examine how good the current Internet routing is, which is clearly different from ours. Other related work includes [1], [2], [3], [10], to name a few.

The remainder of the paper is structured as follows. In Section II, we present the analysis methodology and describe the `traceroute` experiments. A detailed analysis of the measurement data is presented in Section III and the implications of these results on network engineering are discussed therein. Finally we conclude the paper in Section IV and discuss further improvements.

II. ANALYSIS AND EXPERIMENTAL METHODOLOGIES

In this section, we first present the analysis methodology for understanding the delay variations on Internet

TABLE I
EXPERIMENT DESTINATIONS AND DURATIONS

Set	Name	Destination	Location	Duration
S_1	compaq	204.123.2.48	California, USA	3/31/2002–4/2/2002
	utexas	128.83.40.144	Texas, USA	3/31/2002–4/2/2002
	yahoo	64.58.76.222	Washington, DC, USA	3/31/2002–4/3/2002
	zolar	202.2.78.242	Hong Kong	4/13/2002
	fh-friedberg	212.201.24.18	Germany	4/14/2002
	info-x	213.161.85.10	United Kingdom	4/15/2002
	kyoto-u	192.50.8.47	Japan	4/16/2002
S_2	info-x	213.161.85.10	United Kingdom	4/20/2002
	kyoto-u	192.50.8.47	Japan	4/21/2002
	zolar	202.2.78.242	Hong Kong	4/22/2002
	fh-friedberg	212.201.24.18	Germany	4/23/2002
S_3	rpionline	206.114.32.10	Michigan, USA	4/24/2002
S_4	connectiva	216.207.67.189	Indiana, USA	4/25,4/27/2002
	teleglobe	195.219.32.214	France	4/26,4/28/2002
	psi	154.13.2.48	Illinois, USA	4/29/2002




Fig. 1. Geographic locations of traceroute servers used in the experiments (the number in the circle indicates the number of traceroute servers used in the state).

paths and then describe the experiment settings. the assumptions for interpreting the data. We conclude this section by discussing the three network factors affecting user perceived end-to-end delay jitter.

A. Analysis Methodology

As mentioned in Section I, we adopt a “bottom-up” approach and examine the delay variations at three levels: from the link level, to intra-domain path level, and then to inter-domain path level.

1) *Per Hop Queueing Delay Variations Along A Path:* Recall that, by traceroute, we are only able to obtain the round-trip times from a source to all the intermediate routers as well as the destination, but not the *per hop* round-trip times. In this subsection, we present a way to estimate the *per hop round-trip queueing delay variances* (or simply *per hop queueing delay variances* by only relying on traceroute measurements. Before




Fig. 2. An illustration of an Internet path

we proceed, we need to define the round-trip *queueing* delay from a source to an intermediate hop (say the i th hop) first. Consider a collection of traceroute measurements. Let d_i^{min} denote the minimum measured round-trip delay from the source to the i th hop in the collection, then the round-trip delays from the source to the i th hop *subtracted* by d_i^{min} are regarded as the *round-trip queueing delays between the source and the hop*.

Consider an arbitrary path and assume that there are N hops on the path. Let h_i denote the i th hop along the path, for $i = 1, 2, \dots, N$. Let random variable X_i denote the round-trip queueing delay from the source to the i th hop h_i , for $i = 1, 2, \dots, N$. Denote by a random variable Y_i the per-hop round-trip queueing delay on hop h_i . In the following, we show how to estimate the per hop queueing delay variances Y_i by X_i 's (Fig 2). Define $X_0 = 0$. It is easy to see that, in a statistical sense, $X_i = X_{i-1} + Y_i$. Therefore

$$\begin{aligned} \text{Var}(Y_i) &= \text{Var}(X_i - X_{i-1}) \\ &= \text{Var}(X_i) + \text{Var}(X_{i-1}) - 2\text{Cov}(X_{i-1}, X_i). \end{aligned} \quad (1)$$

2) *Intra-Domain Multi-Path Routing and Inter-Domain Route Alteration:* In this paper, we employ a simple rule to distinguish intra-domain multi-path routing and inter-domain route alterations. If during the course of a traceroute experiment (a collection of traceroute measurements lasted for certain time. See the next subsection for the accurate definition), we

only see route changes within AS domains, but the AS level path is fixed between a pair of end-hosts, we say that there is intra-domain multi-path routing in the `traceroute` experiment. On the other hand, if the AS level path has also been changed, we would say that there are inter-domain route alterations in the experiment. Note that, it is possible that an experiment containing inter-domain route alterations will also involve intra-domain multi-path routing, but not vice versa.

B. Experimental Settings

We conduct all our experiments using the well-known tool `traceroute`. For ease of exposition, we will refer to both the tool and a measurement conducted using the tool as `traceroute`¹. We will also refer to the packets sent by `traceroute` as *probes*, and a probe with TTL set to n as “hop n ” probe [9]. For all the `traceroute` measurements, we set the maximum TTL (hop limit) to be 30; the time to wait for a response to a probe 5 seconds; and the UDP probe packet size 38 bytes. To eliminate the effects of DNS name lookup on the observed delays, during the measurements we have suppressed the mapping from an IP address to its corresponding DNS name. Moreover, in the hope that a hop n probe will see a similar network condition as the hop $n - 1$ probe, only *one* probe is sent at each TTL setting.

All the `traceroute` source hosts are located at the University of Minnesota. We will simply refer to them as *umn* given that the specific names of the hosts are not of interest. Table I presents the destination hosts used in `traceroute` measurements. Because it is possible that multiple IP addresses are associated a host name, we use destination IP addresses instead of host names. As shown in the table, we group the measurements into four sets, denoted by S_1 to S_4 . They are different in the way we perform the `traceroute` measurements. For the destinations in set S_1 , we conduct `traceroute experiments` eight times a day during the time interval listed in the column “Duration” in the table, starting from 0:00AM. Each experiment last for one and half hour, and we start the next experiment one and half hour later. During each experiment, `traceroute` measurements are conducted consecutively in the sense that after we finish the current `traceroute` measurement, e.g., by receiving the ICMP `PORT_UNREACHABLE` message from the destination, we *immediately* start the next `traceroute` measurement.

¹we may use the term *measurement* to refer to a single `traceroute`.

Because of the concerns of the processing overhead the experiments impose on the intermediate routers and the destination hosts, we change the way to conduct the experiments in set S_2 . For any destination in set S_2 , we conduct `traceroute` experiments four times a day, at 0:00AM, 9:00AM, 3:00PM, and 9:00PM, respectively. Each experiment lasts one hour. Within each experiment, the source, after finishing the current `traceroute` measurement, will wait for an exponentially distributed time interval with a mean of 50 *ms* before starting the next `traceroute` measurement. For the same reason, we further increase the mean time interval between two adjacent measurements to 1 *s* and 2 *s* for the experiments in sets S_3 and S_4 , respectively.

For analyzing the delay properties of these paths, we only consider the *informative* measurements. By informative, we mean the measurements that reach the destinations by visiting the intermediate routers once and only once (no route loop), and contain no timeouts at any hop (annotated by “*” at a hop).

III. RESULTS AND IMPLICATIONS

In this section, we present the measurement results and discuss their implications in networking engineering. We start by examining how per hop queueing delay variations change link by link along Internet paths. We then move on to the effects of intra-domain multi-path routing. Towards the end of this section, we investigate how AS level route alterations affect user perceived end-to-end delay variations.

Before we proceed, we find it is convenient to make a distinction between a *path* and a *route*. For the sake of exposition, we define a *path* to be *any connection* between a source and a destination. It could change from time to time by traversing different intermediate routers. On the other hand, a *route* is a particular *realization* of a path; given a route, the intermediate routers to be traversed are *fixed*. Note, unfortunately, that this definition of path conflicts with the usage of “path” in the term “multi-path routing”. Since multi-path routing is a conventional term and has been used widely, we continue using it to mean “multi-route routing” in our definition.

A. Per Hop Delay Variations

Consider an arbitrary path with N hops. One of the challenges in analyzing per hop delay variations along the path using Eq. (1) is to verify that the hop n probe follows the same route of the hop $n - 1$ probe, for $1 < n \leq N$. We adopt the following simple rule to verify this. Recall that a `traceroute` experiment is a collection of

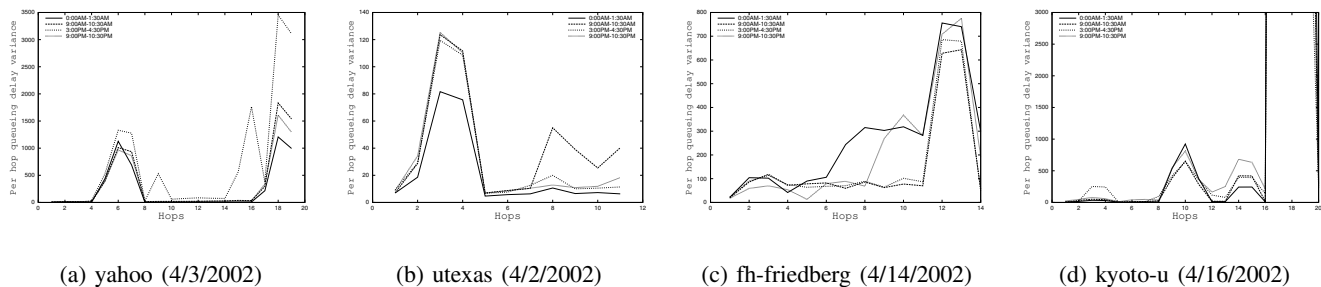


Fig. 3. Per hop queueing delay variances (S_1).

measurements conducted within one and half hour (data set S_1) or one hour (in other data sets). We say that all the probes of a traceroute measurement follow the same route if any of the following conditions holds true.

- Within a traceroute experiment, all the traceroute measurements expose the same single route.
- There is a dominant route used by the traceroute measurements in an experiment. All other routes only appear occasionally, and when they do, they are used *continually* by traceroute measurements (i.e., not interleaved with the dominant route).

For all the measurement data used in this subsection, we have verified that one of the above conditions holds. In the case the second condition holds, the per-hop delay variations are computed using only the dominant route.

Fig. 3(a) presents the per hop delay variances along the path from *umn* to *yahoo* on April 3, 2002, for four experiments. The first experiment (0:00AM-1:30AM²) and the second one (9:00AM-10:30AM) share the same route (Table II), while the third (3:00PM-4:30PM) and the last (9:00PM-10:30PM) share another route (it differs from the first route only at the 17th hop, with 216.33.98.19 replaced by 216.33.98.3). In Table II, we also include the host (or rather the interface) name and the AS that a host belongs to.

Note first that at hops 5, 6, and 17 to 19, we see significantly larger per-hop delay variations across all the experiments. Hops 5 and 6 belong to AS 217, while 17 to 19 belong to 3967, both can be considered as at the *edges of the Internet*. From the figure we also see that for all the experiments, hop 7 also incurs fairly high delay variations, which together with hop 6, connect AS 217 and AS 5006. The results from the experiment from 3:00PM to 4:30PM (during which time period the Internet tends to be more heavily used) deserve a closer examination. In this experiment, packets also experience

²All times are in Central Standard Time.

TABLE II
A PATH FROM *umn* TO *yahoo*

Hop	IP address	AS	Domain name
1	128.101.32.253	217	ecscsi-2-rsm.cs.umn.edu
2	192.168.99.30	217	
3	160.94.26.70	217	tc3x.router.umn.edu
4	160.94.26.98	217	tc2x.router.umn.edu
5	192.42.152.134	217	otr-tc2.northernlights.gigapop.net
6	192.42.152.14	217	
7	137.192.3.254	5006	core1-ge1-1-0.msc.mr.net
8	137.192.5.9	5006	core1-sol-0-1.ply.mr.net
9	63.237.33.53	209	63-237-33-53.cust.qwest.net
10	205.171.20.33	209	chi-core-01.inet.qwest.net
11	205.171.20.174	209	chi-core-03.inet.qwest.net
12	205.171.8.161	209	dca-core-03.inet.qwest.net
13	205.171.9.9	209	dca-core-01.inet.qwest.net
14	205.171.9.14	209	dca-brdr-01.inet.qwest.net
15	216.32.173.249	3967	ibr01-p5-1.stng01.exodus.net
16	216.33.99.83	3967	dcr03-g6-0.stng01.exodus.net
17	216.33.98.19	3967	csr22-ve241.stng01.exodus.net
18	216.35.210.126	3967	
19	64.58.76.222	3967	w1.dcx.yahoo.com

larger than normal delay variations at hops 9, 15, and 16, of which hops 9 and 15 are border routers between ASes. From these results, we observe that larger delay variations are likely to be experienced at routers residing at boundaries of (or rather, links between) two (AS) network domains as well as the edges of the Internet. Similar observations hold on other paths, some examples of which are shown in Figs. 3(b), 3(c), and 3(d). Due to space limitations, we do not discuss them in detail here.

The observed (relatively) large delay variations at the hops near/between boundaries of two (AS) network domains as well as the edges of the Internet are consistent with the belief or fact that today’s Internet backbone networks are generally well-provisioned. It indicates that, to provide a better service quality to end users, we need to pay special attention to both access networks and network boundaries. In order to provide satisfactory *end-to-end* quality of services, it is *not* sufficient to only “over-provision” the network cores.

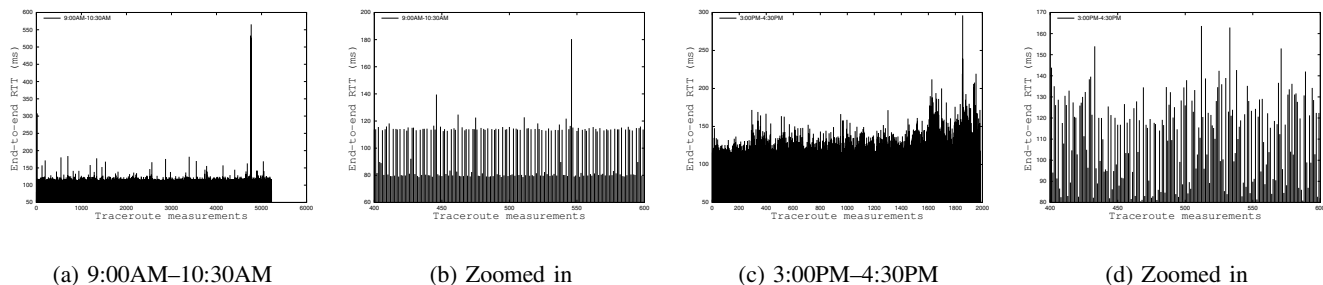


Fig. 4. End-to-end RTT (compaq, April 2, 2002)

B. Intra-Domain Multiple Path Routing

We now investigate the potential impact of multi-path routing on user perceived path characteristics. In Fig. 4, we plot the RTT on the path between *umn* and *compaq* for two experiments, The first one (a) from 9:00AM to 10:30AM, while the second one (c) from 3:00PM to 4:30PM. For both of them, we also show a “zoomed-in” portion of the experiments ((b) and (d), respectively) to illustrate the end-to-end delay behavior more clearly. From both Figs. 4(b) and 4(d), we can clearly see that the end-to-end delays on the path are bi-modal. A closer examination of the *traceroute* data reveals that eight routes are used within the domain *alter.net*, where four of them share one link, and the another four share another link. Fig.5 sketches the “routers” involved in the multi-path routing within *alter.net* and their connectivities exposed from the *traceroute* experiments (we also marked the eight routes at the 10th hop). Table III gives the mean RTT to the 9th and 10th hops on the eight routes respectively. From the table we see that the eight routes have a similar RTT till the 9th hop. At the 10th hop, we see that there are two distinct groups, where routes 1 to 4 have RTTs on the order of $40ms$, while routes 5 to 8 have RTTs on the order of $70ms$. we surmise that routes 1 to 4 actually traverse the same set of routers but different interfaces at the 10th hop. The same thing applies to routes 5 to 8. So essentially, we suspect that there are two IP *router* level routes on the path between *umn* and *compaq*, which leads to the observation of bi-modal RTTs on the path.

TABLE III

AVERAGE RTT TO AN INTERMEDIATE “ROUTER” (MS) (*compaq*, 9:00AM-10:30AM, 4/2/2002)

route	1	2	3	4	5	6	7	8
9th	35.6	35.6	35.5	35.9	35.6	36.1	35.8	36.5
10th	46.4	46.5	46.1	46.3	75.3	75.4	75.2	75.6

Table IV lists the mean end-to-end delays on the eight routes of the path. From the table we see that, during




Fig. 5. Illustration of *alter.net* multi-path routing.

the time period from 9:00AM to 10:30AM, if a packet takes a route from the route set 1 to 4, its expected RTT will be on the order of $80ms$, while from another set (5 to 8), it is on the order of $110ms$. For the time period from 3:00PM to 4:30PM, they are $90ms$ and $130ms$, respectively. We suspect that the bi-modal end-to-end delays on this path are due to the vastly different propagation delays on the two *actual* router-level routes used within the domain *alter.net*.

From the original *traceroute* measurements alone, however we were unable to verify directly whether all the probes of a *traceroute* measurement followed a given route. To further understand how the multi-path routing is performed, we conducted another set of experiments, where three probes were sent back-to-back to each hop. Fig. 5 shows one example of such an experiment, where a truncated snapshot of the *traceroute* measurement is shown. Note that at hop 10 we received ICMP responses from two different addresses, which seems to indicate that multi-path routing within *alter.net* is carried out on packet-by-packet basis. This may explain why packets on the same path experience the observed bi-modal delay behavior with dramatic end-to-end delay variations. Our experiment results suggest that we need to be careful with how multi-path routing is practiced, and to understand its impact on end-to-end delay variations, especially when deploying delay-jitter-sensitive applications such as VoIP over the

TABLE IV
AVERAGE RTT (MS) (*compaq*, 4/2/2002)

route	1	2	3	4	5	6	7	8
9:00-10:30AM	81.9	82	82.5	81.6	116.2	116.2	115.6	116.2
3:00-4:30PM	97.8	96.7	98.3	97.3	130.1	131.7	133.5	133

1.	128.101.34.253	0.511 ms	0.441 ms	0.452 ms	
2.	192.168.99.30	1.779 ms	1.300 ms	1.187 ms	
3.	160.94.26.70	1.528 ms	1.258 ms	1.469 ms	
4.	160.94.26.98	1.476 ms	1.438 ms	1.585 ms	
5.	192.42.152.130	2.711 ms	1.982 ms	2.832 ms	
6.	63.145.65.117	36.936 ms	33.474 ms	33.334 ms	
7.	205.171.16.41	33.738 ms	40.663 ms	39.188 ms	
8.	157.130.172.41	34.613 ms	34.951 ms	34.240 ms	
9.	152.63.93.202	34.649 ms	34.541 ms	35.559 ms	
10.	152.63.94.54	45.197 ms	44.316 ms	152.63.94.50	46.792 ms

.....

Fig. 6. A traceroute measurement from *umn* to *compaq*.

Internet.




Fig. 7. End-to-end RTT (*rpionline*, April 24, 2002)




Fig. 8. AS Paths taken by traceroutes (*rpionline*, April 24, 2002)

C. Inter-Domain AS Level Route Alterations

In this sub-section, we study how the inter-domain route changes affect the user perceived end-to-end delay variations. By analyzing the BGP update information database, we identify a list of subnets that frequently announce BGP updates and choose a host within each such subnet as a destination for the traceroute

experiments. Due to space limitation, here we will only present the results on one path, from *umn* to *rpionline*. Fig. 7 plots the RTT on the path for two different time periods, one from 0:00AM to 1:00AM (referred as the first experiment), another from 9:00PM to 10:00PM (referred as the second experiment). Fig. 8 shows the corresponding AS level routes taken by traceroute measurements over the duration of the two experiments, respectively. The y-axis in this figure ranges from 0–8, where 1–7 represent 7 different AS routes on the path.

Note first that, during one hour time span, there are six distinct AS level routes (nine IP level routes) in the first experiment; while in the second experiment, there are six AS routes (eight IP level routes). The changes of AS level routes are quite frequent. Moreover, comparing Fig. 7 and Fig. 8, we see that there is a direct connection between the round-trip times on the path and the AS route taken by the path. For example, we can easily see three segments with largely different RTTs. Indeed, the average RTT of the path is $78ms$ when packets take route 6 (Fig 8). On the other hand, if packets take route 3, it is $122ms$. If an application encounters any such AS level route changes during the course of its lifetime, it will experience significantly large delay variations.

Clearly, frequent AS level route alterations have negative effects on the user perceived path characteristics. However, unlike intra-domain multi-path routing, which can be controlled by an ISP to limit its impact on user-perceived delay variations, AS level route alterations are often outside the control of any ISP alone. Nonetheless, dampening frequent AS level route alterations at a global level is important to reduce user-perceived end-to-end delay variations.

IV. CONCLUSIONS AND ON-GOING WORK

In this paper we investigated the *network* factors that may affect the user perceived end-to-end delay jitter. In particular, we studied the effects of the following three major factors: queueing delay variations at each hop along an Internet path; intra-domain multi-path routing; and inter-domain route (i.e., AS path) alterations. Our findings are: 1) larger queueing delay variances are likely to be experienced at routers residing at boundaries of (or links between) two (AS) network domains as well as the edge of the Internet; 2) intra-domain multi-path routing may have significant impact on end-to-end delay jitters; and 3) inter-domain route alteration adversely affects user perceived delay jitters.

In the current study, all the `traceroute` data are collected at the University of Minnesota (UMN). Therefore, the findings related to the source stub network (UMN) may not be representative. In order to perform a more comprehensive study, we are planning to collaborate with other institutions to conduct similar experiments.

REFERENCES

- [1] M. Allman and V. Paxson. On estimating end-to-end network path properties. In *Proc. ACM SIGCOMM*, Cambridge, MA, September 1999.
- [2] J-C. Bolot. End-to-end packet delay and loss behavior in the internet. In *Proc. ACM SIGCOMM*, pages 289–298, San Francisco, CA, September 1993.
- [3] A. Downey. Using pathchar to estimate internet link characteristics. In *Proc. ACM SIGCOMM*, Cambridge, MA, September 1999.
- [4] C. Labovitz, G. Malan, and F. Jahanian. Internet routing instability. *IEEE/ACM Transactions on Networking*, 6(5):515–528, 1998.
- [5] NSF ITR QoS Workshop. Annapolis, MD, April 2002. <http://www.cs.umn.edu/research/networking/itrworkshop>.
- [6] K. Papagiannaki, S. Moon, C. Fraleigh, P. Thiran, and F. Tobagi and C. Diot. Analysis of measured single-hop delay from an operational backbone network. In *Proc. IEEE INFOCOM*, New York, NY, June 2002.
- [7] V. Paxson. End-to-end routing behavior in the Internet. In *Proc. ACM SIGCOMM*, Stanford, CA, August 1996.
- [8] V. Paxson. End-to-end Internet packet dynamics. In *Proc. ACM SIGCOMM*, Cannes, France, September 1997.
- [9] V. Paxson. *Measurements and Analysis of End-to-End Internet Dynamics*. PhD thesis, University of California, Berkeley, 1997.
- [10] D. Sanghi, A. Agrawala, O. Gudmundsson, and B. Jain. Experimental assessment of end-to-end behavior on internet. In *Proc. IEEE INFOCOM*, pages 867–874, 1993.
- [11] S. Savage, A. Collins, E. Hoffman, J. Snell, and T. Anderson. The end-to-end effects of internet path selection. In *Proc. ACM SIGCOMM*, pages 289–299, 1999.
- [12] K. Varadhan, D. Estrin, and S. Floyd. Impact of network dynamics on end-to-end protocols: Case studies in tcp and reliable multicast. Technical Report USC-CS-TR 98-672, University of Southern California, Information Sciences Institute, April 1998.
- [13] Y. Zhang, N. Duffield, V. Paxson, and S. Shenker. On the constancy of internet path properties. In *Proceedings of ACM SIGCOMM Internet Measurement Workshop*, San Francisco, CA, November 2001.