Observations on Round-Trip Times of TCP Connections

Phillipa Sessini and Anirban Mahanti
University of Calgary
2500 University Drive NW
Calgary, AB, Canada
{sessinip, mahanti}@cpsc.ucalgary.ca

***Abstract*— Knowledge about properties of network traffic can be beneficial when studying network protocols. It enables realistic models of network traffic to be created and evaluations of current protocols to take place. This study examines trends in round-trip times (RTTs) at a university Web server. Round-trip time is a particularly important characteristic of transport layer Internet traffic to measure because it impacts the throughput of TCP. In addition to examining trends of RTTs, this paper examines the relationship between RTT and the time between the SYNACK and ACK packets. The relationship between this heuristic for estimating RTT and actual measured RTTs is relevant for protocols such as TCP Vegas that need to estimate RTT early in a connection.**

Keywords: Round-Trip Times, TCP, Congestion Control

I. Introduction

In the field of networking, it is important to have accurate information about network conditions. This information can be used for evaluation of current network protocols, construction of realistic networking simulations, and development of new distributed applications. In particular, characteristics of Internet traffic are of interest because the Internet is primarily where networking applications are deployed.

In addition to evaluating current protocols, accurate information about Internet traffic can aid in the development and simulation of new protocols. By carefully analyzing the behavior of the Internet, researchers are able to create traffic models of conditions found on the Internet. Accurate models of network traffic are invaluable, without them there would be little way to know if a system would be successful without deploying a full scale version of it. Deploying a system when it may fail can be expensive, in terms of both time and money.

Although there has been a significant amount of work done in the field of network measurement, it is imperative that studies of network behavior are still undertaken because of the dynamic nature of Internet traffic. In recent years, the popularity of multimedia networking applications such as online games and peer-to-peer programs has changed the characteristics of Internet traffic. In particular, these applications require the transmission of large amounts of data. TCP models of throughput (e.g., [11], [13], [15]) are especially applicable to these larger transfers that spend the majority of their time in the congestion avoidance phase of TCP congestion control. These models show that TCP throughput exhibits dependency on round-trip times (RTTs). For example, the commonly deployed TCP Reno congestion control variants have throughput that is inversely related to RTT. Thus, knowledge of RTT trends can provide insight into the performance of these applications.

This paper examines trends in the round-trip times of TCP connections. Traces analyzed in this paper were obtained from the University of Calgary Web server. These traces include traffic between external hosts and the university Web server only, no internal traffic is captured. The server is connected to the Internet via a 100Mb/s full-duplex Ethernet link. The Web server hosts several files which are of interest to students of the university as well as prospective students.

Our study examines variability of round-trip times at the packet level and at the connection level. The effects of time of day on RTT are also measured. RTTs of traffic flows on the Internet have implications for TCP throughput as well as active queue management techniques deployed at links within the network. The relationships between various characteristics of per connection RTTs and the time between the SYNACK and ACK packets in the TCP handshake are also explored in this paper. The characteristics considered are minimum, median, mean and maximum RTTs seen by a connection. This may have implications for protocols such as TCP

Vegas [7] which uses an estimate of the minimum RTT in its congestion avoidance algorithm [15].

Our main observations are as follows:

- Both packet level and connection level round-trip times exhibit high variability;
- The time between SYNACK and ACK appears to be a reasonable predictor of the average RTT, but is a poor predictor of the minimum, maximum, or median RTT; and,
- RTT measurements show time of day dependence, with higher RTTs being measured when it is day time in Asia and Europe.

The remainder of this paper is structured as follows. Section II describes the measurement framework, trace files, and measurement techniques used. Results are described in Section III. Section IV describes related measurement studies and compares results seen in this study with previous studies. Conclusions are presented in Section V.

II. Methodology

A. Measurement Infrastructure

The primary source of data for this paper is the main router of the University of Calgary campus which is located near the University's Web server. Data is collected using *tcpdump*[1] which runs on a dual-processor Dell server, with 2GB of RAM and 140GB of disk space. This machine is located at the main router on campus, which is connected to the Internet by a 100Mb/s full-duplex Ethernet link. Using port mirroring, traffic from the link is forwarded to the monitor via a 1Gb/s half-duplex Ethernet link. Once collected the data is filtered to include only TCP traffic at the Web server.

The connection that is monitored by the packet sniffer is one of two main Internet connections at the University of Calgary. The two connections are used for commercial traffic and educational traffic, with the sniffer being located on the commercial connection. This has impacts on the analysis of data that is collected from the monitoring infrastructure. For example, occasionally only one direction of a connection will use the link being measured while the other direction uses the other connection. Connections that display this property are not considered for analysis in this paper.

Since the monitoring infrastructure is primarily utilized for long term data collection, the majority of the data collection process is automated. The monitor machine is configured to automatically start *tcpdump* whenever the machine is rebooted as well as restarting *tcpdump* every hour. Restarting *tcpdump* every hour ensures that in the event that *tcpdump* exited due to an error it will be restarted within an hour.

[1]http://www.tcpdump.org

TABLE I

Summary Statistics of Trace Data

Trace	Data	Packets	Connections
Aug. 26, 2004 (Total)	7.1GB	12,150,900	328,808
Aug. 31, 2004 (Total)	8.7GB	14,664,458	385,397
Aug. 26, 2004 (Filtered)	2.3GB	3,534,111	54,748
Aug. 31, 2004 (Filtered)	2.9GB	4,371,723	64,127

More complex data processing is not automated. For example, once the disk at the monitor is full the data must be transferred to the file server. This process uses scripts to copy the traces from the monitor to the file server, verify that the data was copied correctly and remove the traces from the monitor. Each of these steps is initialized manually to ensure that any errors are detected before the traces are deleted from the monitor. Before the data is provided to researchers, the trace files are also repaired. This involves the deletion of any incomplete packets at the end of trace files.

B. Traces

The traces analyzed in this study consist of two 24 hour traces collected on Thursday, August 26, 2004 and Tuesday, August 31, 2004, respectively. Since these traces are collected on weekdays they show typical patterns of workday activity. They were collected when school was not in session so the load seen by the server in these traces may be lighter than at other times of the year. In the time of day analysis, these traces are broken into four time periods: late night, corresponding to the hours between 0:00 and 5:59; morning, corresponding to hours between 6:00 and 11:59; afternoon, between 12:00 and 17:59 and evening, the hours between 18:00 and 23:59.

C. Scripting

A fundamental challenge when performing Internet measurement studies is the volume of data being analyzed. Thus it is crucial to select appropriate tools to perform analysis on data sets. For this study, the *Bro Intrusion Detection System*[2] [14] was used to measure properties of the collected traces. Although its primary purpose is for network intrusion detection, *Bro* comes with powerful scripting capabilities which make analyzing large volumes of data more manageable.

[2]www.bro-ids.org

D. Measurement Techniques

Round-trip time is defined as the time from when a packet is sent to when its acknowledgement is received. This concept may seem quite simple, however, in practice there are several factors which can complicate the measurement of RTTs. For example, the client may only acknowledge every other packet in a window if the delayed ACK mechanism is in use. Also, when packets are lost this can result in an over estimation of round-trip time.

Round-trip time is computed as, the difference between the time a packet is sent and the time an ACK for that particular packet is received. To reduce the effects of lost packets on this calculation, connections with lost or retransmitted packets are not included in the analysis. Also, to increase the statistical validity of the connections analyzed, connections with fewer than 8 RTT samples are excluded from the analysis. As a result of buffering at the packet measurement device, it was sometimes the case that the device would report impossibly small RTTs. To decrease the effects of this malfunction on observations made, RTTs of less than 3 $msec$ are excluded from the analysis. Statistics of the traces analyzed in this study before and after the filtering process are summarized in Table I.

Often in Internet measurement studies, complete bidirectional traces are not available, as a result of this many studies use a single RTT measured during the TCP handshake to approximate RTT (e.g. [3], [9], [10]). In this study, the time between when the SYNACK packet is seen by the monitor and when the ACK packet is seen by the monitor is observed. The time between the SYNACK and ACK packets in the TCP handshake is the first round-trip time observed by a server and may be used by protocols as a metric for estimating RTT early in a TCP connection. For example, the TCP Vegas congestion control algorithm requires an estimate of the minimum RTT in order to compute queueing delays. In Section III, we evaluate whether or not the RTT measured by the SYNACK and ACK exchange is a useful indicator of the minimum RTT (`baseRTT`) a connection.

The time between SYNACK and ACK packets is observed, rather than the time between the SYN and SYNACK packets because the traces collected in this study are collected closer to the server side of the connection. Thus, in this study, the time between the SYN and SYNACK packets would be an underestimation of the RTT observed by connections between the server and a client. However, if the monitor device was located closer to the client the time between SYN and SYNACK may provide a more accurate impression of RTTs experienced. Also, as suggested in previous work, if the monitor is located centrally in the network the time between the SYN and ACK packets may be a better predictor of RTTs seen in the network [10].

III. RESULTS

A. Packet Level Analysis

The RTTs of packets in the August 31, 2004 trace are plotted in Figure 1. The average RTT observed in this trace was 168.9 $msec$ with 50% of the RTTs falling below 76.1 $msec$ and 75% of the values falling below 151.4 $msec$. The maximum RTT observed is 19.6 sec. This indicates that the minimum RTT observed, even after excluding RTTs of less than 3 $msec$, differs from the maximum RTT by almost 4 orders of magnitude. This indicates that the range of RTTs observed by connections on the Internet is quite large and is consistent with observations made in [1]. In most cases the larger RTTs are observed by connections located geographically very far away from the Web server, for example in Africa.

The standard deviation of the RTT observed by packets is a 402.2 $msec$ and the coefficient of variation is 2.4. This indicates that the RTTs observed by packets at this Web server have high variance. This is to be expected since these packets are from different connections and traverse different network paths.

B. Variability of RTTs Within Connections

Standard deviation can be used to evaluate the variability of RTTs within a connection. Figure 2 shows the cumulative distribution function (CDF) of per connection standard deviation of both days of traces. The average standard deviation is 142.9 $msec$ and half of the connections have a standard deviation of less than 49.4 $msec$. The majority of the connections experience a low standard deviation with 75% experiencing a standard deviation of less than 98.7 $msec$. The connections that experience higher standard deviations in their RTTs likely experienced a routing change during the course of the connection, causing a sudden shift in the RTT seen by the connection. In contrast, connections that saw lower standard deviations of RTTs likely did not change routes during the connection but experienced congestion which caused minor fluctuations in observed RTTs.

C. Effectiveness of SYNACK/ACK for Estimating RTT

The time between the SYNACK and ACK packets in the TCP handshake is the first measure of RTT

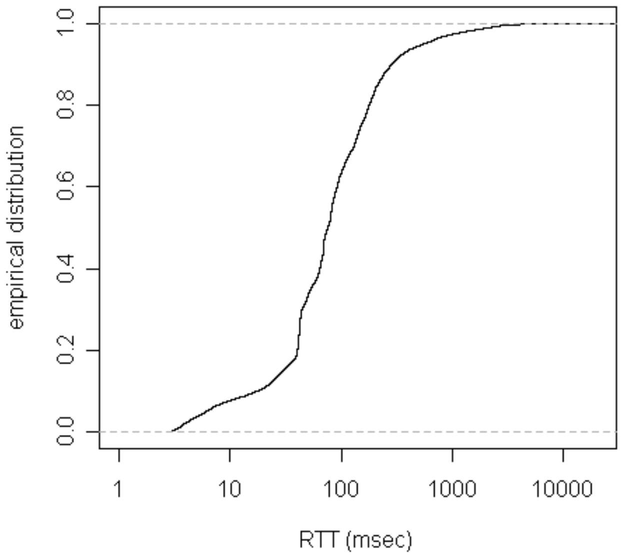

Fig. 1. CDF of per packet RTT

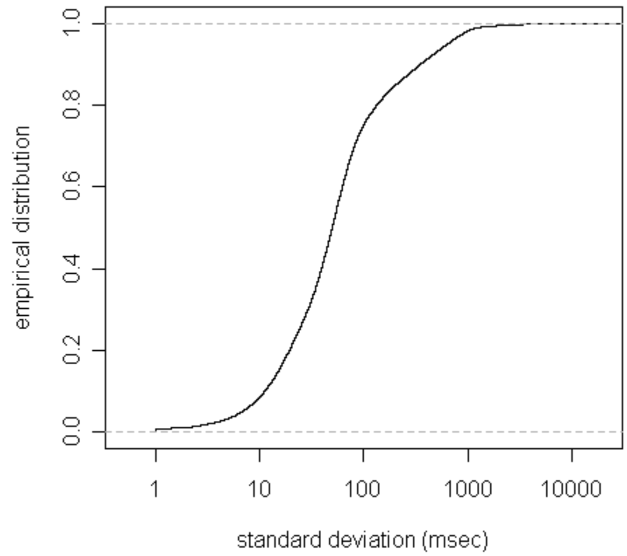

Fig. 2. CDF of per connection standard deviation

observed by a Web server. Thus, for servers that require an estimate of RTT time, this may be an appropriate mechanism for approximating RTT early in a connection. The distribution of the time between the SYNACK and ACK packets is compared with the distribution of minimum, median, mean and maximum RTTs observed by connections in Figure 3. The distribution of time between SYNACK and ACK packets closely follows the distribution of the mean and median RTT of a connection. The minimum and maximum RTTs observed by connections show noticeably different trends than the time between SYNACK and ACK packets. The average minimum, median, mean, max RTT seen by connections is 14.0 $msec$, 181.8 $msec$, 182.1 $msec$ and 629.0 $msec$, respectively. The average time between the SYNACK and ACK packets is 145.7 $msec$.

To gain a better understanding of the relationship between the time between SYNACK and ACK packets and various characteristics of RTTs, the absolute difference between the SYNACK/ACK and minimum, median, mean and maximum values is compared in Figure 4. The average absolute difference between the SYNACK/ACK time and the mean RTT is the lowest at 80.9 $msec$. The average absolute differences for minimum, median and maximum are higher at 132.0 $msec$, 199.2 $msec$ and 489.3 $msec$ respectively. The 75th percentile shows a similar trend as the mean time between the SYNACK/ACK time and the mean RTT. 75% of the connections have an absolute difference between the SYNACK/ACK and mean RTT that is less than 51.7 $msec$. The 75th percentiles for the absolute difference between the SYNACK/ACK time and minimum, median and maximum RTTs are 146.4 $msec$, 195.7 $msec$ and 207.5 $msec$, respectively.

This indicates that the time between the SYNACK and ACK packets may be a possible heuristic for the mean RTT observed by a connection. However, it is not a good indicator of the maximum RTT since it is incorrect by over one third of a second on average. Its performance for estimating minimum and median RTTs is marginal and there may be more effective heuristics for estimating these characteristics of RTTs.

D. The Effects of Time of Day on RTT

Time of day can also influence trends in RTTs. This is as a result of late night in North America corresponding to peak business hours in Asia and Europe. Often RTTs observed late at night will be larger than those seen during the day. This is examined in Figure 5. These figures show that the highest RTTs are observed between midnight and 6 am MST. The lowest RTTs are observed between noon and 6 pm. Evening and morning RTTs fall between these values. The average RTTs for late night, morning, afternoon and evening are 244.7 $msec$, 199.6 $msec$, 144.9 $msec$ and 178.0 $msec$ respectively.

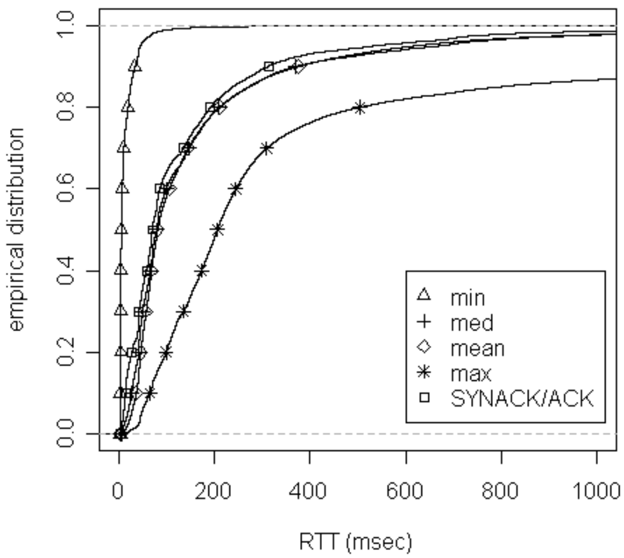

Fig. 3. CDF of per connection SYNACK/ACK time, minimum, median, mean and maximum RTTs

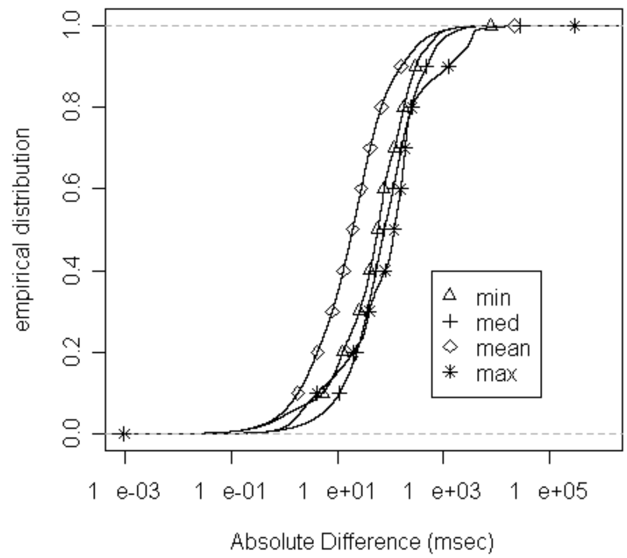

Fig. 4. CDF of absolute difference between SYNACK/ACK measurement and various RTT measures

In contrast to the mean RTT, the standard deviation of RTTs observed by connections does not vary depending on the time of day. This is illustrated in Figure 6. The distributions of standard deviations observed by connections across various times of day follow a consistent trend. The average standard deviations of RTTs within a connection for late night, morning, afternoon and evening are 174.2 $msec$, 152.3 $msec$, 123.0 $msec$ and 142.2 $msec$ respectively. Connections between the hours of 12:00 and 17:59 display slightly smaller average, standard deviations of observed RTTs, however, their overall trend is similar to that of other times of day.

In the case of maximum, mean, median, and minimum RTTs, the late night and afternoons display the most extreme values with morning and evening falling in between. This suggests that during the morning and evening hours there may be some transition time between the work day in Asia and Europe and the work day in North America. However, time of day appears to have little effect on the variability of RTTs observed by connections.

IV. Related Work

Several studies have been done to characterize TCP flows on the Internet (e.g. [1]–[6], [8]–[12], [16]). However, few have focused on RTTs [1], [3], [6], [10]. Even fewer of these studies have had access to bidirectional traffic traces. A lack of bidirectional data has resulted in the need for estimation techniques to measure RTTs. In contrast, this study analyzes bidirectional traffic traces from a Web server and is thus able to compare a heuristic for estimating RTT to measured values for RTTs.

Measurements of a campus network were studied by Aikat *et al.* [1]. In particular, their focus was on the variability of RTTs in TCP connections. The distribution of RTTs seen in [1] is similar to the distribution seen in this paper. Aikat *et al.* also examine the distribution of standard deviation of RTTs seen in TCP connections, the distribution is consistent between the two papers. In contrast to [1], this paper studies the time between the SYNACK packet and the ACK packet, rather than the time between the SYN and SYNACK packet which is studied by Aikat *et al*.

Jiang and Dovrolis [10] discuss methods of passively estimating RTTs. They suggest using the time between when the SYN packet is received and the ACK packet is received as a metric for measuring RTTs within a network. When the monitor is centrally located this metric can provide an accurate estimation of RTTs. They also examine the effect of time of day on RTTs and they find that RTTs are greater when it is day time in Asia as compared with when it is day time in North America. Their result is similar to the results seen in the analysis of time of day and its effects on RTT in this paper.

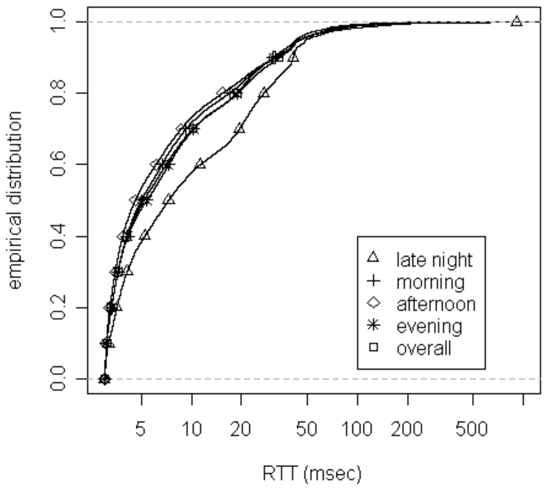

(a) Minimum

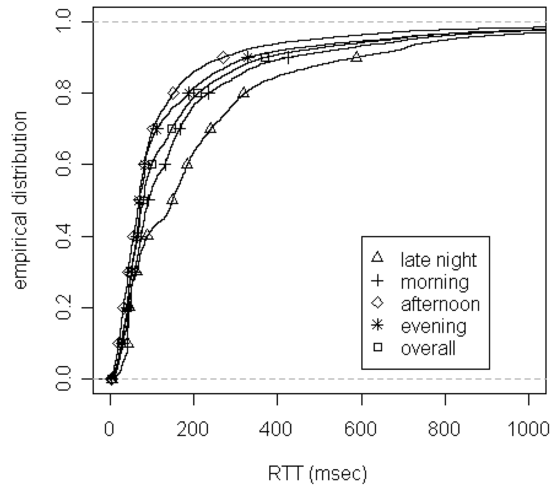

(b) Median

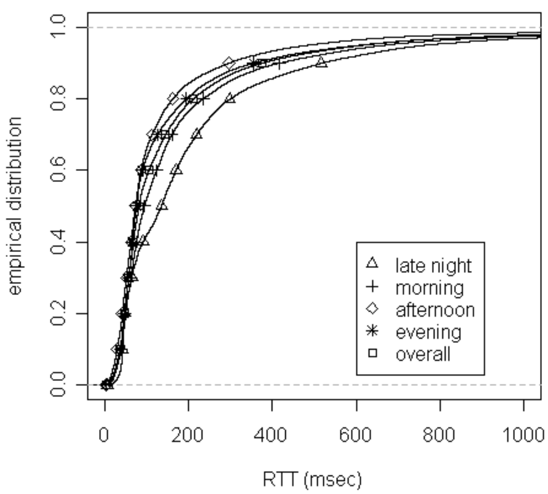

(c) Mean

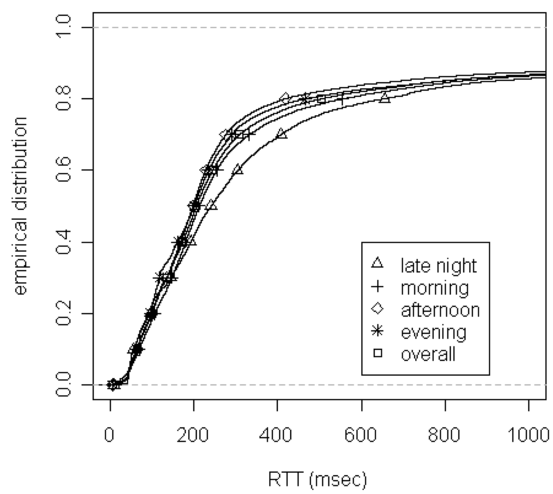

(d) Maximum

Fig. 5. Comparison of the RTT characteristics of connections across different time of day

In [9], Fraleigh *et al.* perform a large scale analysis of bidirectional traffic traces from the Sprint IP backbone. They estimate average RTT based on the time between the SYN packet and the ACK packet similar to [10]. The do not report measures of variability of RTTs in their study.

V. Conclusions

This study analyzed trends in RTTs of TCP traffic at a Web server at both the packet and connection level. At the packet level, it is found that RTTs experienced by packets have high variance. This can be attributed to many connections having very small RTTs and very few having large RTTs, most likely caused by the geographical location of the users.

At the connection level, variability in the standard deviation of RTTs within connections also displays high variance. The variability can be attributed to few connections experiencing path changes during the connection and many connections experiencing congestion causing high and low standard deviations respectively.

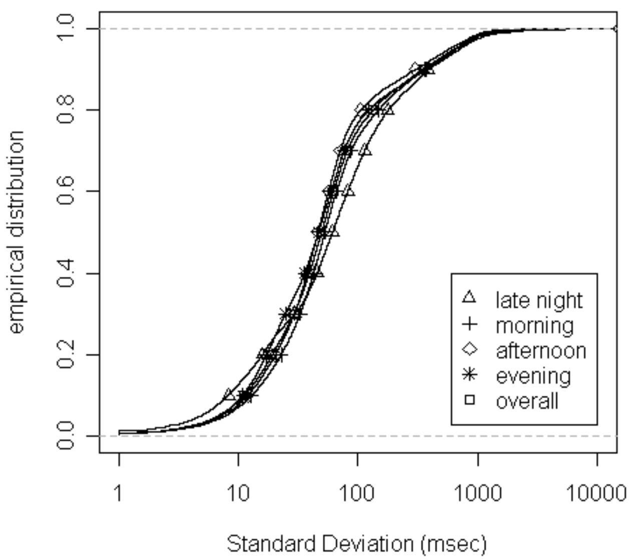

Fig. 6. Comparison of per connection standard deviation across different time of day

When evaluating the time between SYNACK and ACK packets as a predictor of RTTs, it is found that the time between SYNACK and ACK packets best predicts the average RTT. The time between SYNACK and ACK packets is shown to be a poor predictor of the maximum RTT seen by a connection. It is important to bear in mind with this result that the time between SYNACK and ACK packets is a good measurement of RTT when the network monitor is located at the server side of the network. However, if the monitor is located at the client end of the network or in the center of the network the time between the SYN and SYNACK or SYN and ACK packets may provide a better estimate of RTT.

Time of day is also found to have an impact on RTTs observed in this study. Times corresponding to day time in Asia and Europe show higher RTTs than times corresponding to day time in North America. Time of day does not, however, seem to have an impact on the variability of RTTs observed by connections.

ACKNOWLEDGEMENTS

The authors thank Martin Arlitt for providing the traces used in this paper and also for his help with Bro. This work was supported by the Natural Sciences and Engineering Research Council of Canada.

REFERENCES

[1] J. Aikat, J. Kaur, F. Smith, and K. Jeffay. Variability in TCP Round-trip Times. In *IMC '03: Proceedings of the 3rd ACM SIGCOMM Conference on Internet Measurement*, Miami Beach, USA, October 2003.

[2] M. Allman, W. Eddy, and S. Ostermann. Estimating Loss Rates with TCP. *SIGMETRICS Performance Evaluation Review*, 31(3), December 2003.

[3] M. Arlitt, B. Krishnamurthy, and J. Mogul. Predicting Short-Transfer Latency from TCP Arcana: A Trace-Based Validation. In *IMC '05: Proceedings of the 5rd ACM SIGCOMM Conference on Internet Measurement*, Berkeley, USA, October 2005.

[4] M. Arlitt and C. Williamson. An Analysis of TCP Reset Behaviour on the Internet. *ACM Computer Communication Review*, 35(1), January 2005.

[5] P. Benko and A. Veres. A Passive Method for Estimating End-to-End TCP Packet Loss. In *Proceedings IEEE Globecom 2002*, Taiwan, November 2002.

[6] S. Biaz and N. Vaidya. Is the Round-trip Time Correlated with the Number of Packets in Flight? In *IMC '03: Proceedings of the 3rd ACM SIGCOMM Conference on Internet Measurement*, Miami Beach, USA, October 2003.

[7] L. Brakmo and S. O'Malley. TCP Vegas : New Techniques for Congestion Detection and Avoidance. In *SIGCOMM '94: Proceedings on Communications Architectures, Protocols and Applications*, Portland, USA, October 1994.

[8] M. Crovella and A. Bestavros. Self-similarity in World Wide Web Traffic: Evidence and Possible Causes. In *Proceedings of the 1996 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, Philadelphia, USA, May 1996.

[9] C. Fraleigh, S. Moon, B. Lyles, C. Cotton, M. Khan, D. Moll, R. Rockell, T. Seely, and C. Diot. Packet-Level Traffic Measurements from the Sprint IP Backbone. *IEEE Network*, 2003.

[10] H. Jiang and C. Dovrolis. Passive Estimation of TCP Round-Trip Times. *ACM SIGCOMM Computer Communications Review*, 32(3), July 2002.

[11] J. Padhye, V. Firioiu, D. Towsley, and J. Kurose. Modeling TCP Throughput: A Simple Model and its Empirical Validation. In *Proc. of ACM SIGCOMM 1998*, Vancouver, Canada, September 1998.

[12] J. Padhye and S. Floyd. On inferring TCP behavior. In *SIGCOMM '01: Proceedings of the 2001 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, San Diego, USA, August 2001.

[13] N. Parvez, A. Mahanti, and C. Williamson. TCP NewReno: Slow-but-Steady or Impatient? In *Proceedings of the 2006 IEEE International Conference on Communications (ICC)*, Istanbul, Turkey, June 2006.

[14] V. Paxson. Bro: A System for Detecting Network Intruders in Real-Time. *Computer Networks*, 31(23-24), December 1999.

[15] C. Samios and M. Vernon. Modeling the Throughput of TCP Vegas. In *Proceedings of the 2003 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, San Diego, USA, June 2003.

[16] Y. Zhang, L. Breslau, V. Paxson, and S. Shenker. On the Characteristics and Origins of Internet Flow Rates. In *SIGCOMM '02: Proceedings of the 2002 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, Pittsburgh, USA, August 2002.