Modelling Internet Paths Through a Minimum RTT Delay Analysis

Manoel Eduardo Mascarenhas da V. Alves[1], Sergey Nesterov[2] & Reginald P Coutts[2]

[1]Research and Development Center for Information Processing and Automation, Praça das Orquídeas 44/52, Alphaville, Barueri, SP, 06453-000, Brazil
edumasc@cpdia.org.br

[2]Centre for Telecommunications Information Networking, The University of Adelaide, Level 5, Capita Building, Adelaide, South Australia 5000, Australia
{snesterov, rcoutts}@ctin.adelaide.edu.au

Abstract -This paper describes an experimental methodology for previewing path characteristics of interconnecting networks between two given sites on the Internet. The methodology is based on the analysis of Traceroute path samples collected in a regular time basis through a period of time. Paths are classified on either typical paths or minimum typical paths in terms of Traceroute parameters such as RTT, IP/DNS addresses and total number of echo-replies per node. Such a classification is used with a Client-Server Internet connectivity model and a minimum RTT delay model for previewing path characteristics.

1. INTRODUCTION

The Internet has evolved dramatically in the past few years as result of developments in internetworking & telecommunications technologies and increasing market demand for interactive services. While increasing service diversity is noticeable through steady competition for developing and delivering multimedia content, the Internet infrastructure as a whole has evolved in a rapid and almost "organic" fashion, resulting in an enormous mesh of hosts, networks and network peering points – a complex and fault susceptible environment. In addition, the lack of Quality of Service (QoS) standards -and the lack of consensus on what is an adequate QoS level- has led to an atypical situation where customers are serviced in a best-effort basis, without plain guarantees of effective performance.

In the face of these performance issues and the assortment of application service level requirements, our research describes a new methodology to investigate end-to-end Internet delays and to preview path characteristics of interconnecting networks. Internet packet delay is a known problem for applications that require a higher interactivity level *either* between users and applications (e.g., Video on Demand (VoD)) *or* among users (Video-conference and Voice over IP).

In addition, it is important to observe that TCP performance is also affected by long-delay links as in satellite networks. In long delay-links, the slow start and congestion control mechanisms can result in low efficiency of available channel bandwidth due to senders being idle for long periods of time waiting for acknowledgements [1,2].

In the case of Long Fat Pipes (LFN) –i.e., networks where the Bandwidth-Delay Product (BDP) is large-, some applications might have low performance due to small TCP window size configuration [3, 4]. While for short-duration data transfers the TCP window size does not limit performance [5], for bulk file transfers TCP affects negatively end-to-end throughput. The sender stays idle for a long period of time waiting for acknowledgements – which results in low link occupancy. While a TCP window scale option might improve performance, it is important to note that a larger window size increases the probability of more than one packet per window being lost, which results in low throughput performance due to inefficiencies of Fast Retransmit and Fast Recovery algorithms. For a further discussion about LFNs and their performance issues, read [6,7].

2. OBJECTIVES

This paper describes a method for identifying path characteristics between two sites on the Internet without previous knowledge of interconnecting networks. The experimental methodology is achieved by repeated use of *Traceroute* to different Client-Server connections[1] and subsequent study of collected traces to identify path characteristics. Software utilities were developed for automating the *Traceroute* utility and registering *Traceroute* path sample parameters (Node, RTT, IP/DNS addresses and Packet loss) in log files. These

---

[1] Client-Server connections are also called *virtual paths*.

parameters were used for classifying *typical paths* and, later, *minimum typical paths*.

This paper is organised as follows. We introduce a Client-Server Internet connectivity model in Section 3, which is the basis for the experimental analysis. In Section 4, a minimum RTT Delay model is presented. In Section 5, we point a recent research where the software supporting the experiment is described. Section 6 describes experimental assumptions and data analysis procedures. In Section 7, a discussion about a number of findings is provided. Finally, in Section 8, some concluding remarks are made.

3. CLIENT-SERVER CONNECTIVITY MODEL

For ease of analysis, the Internet environment can be summarised as a collection of Autonomous Systems[1] (AS), which vary in size, geographical coverage and function. A Client-Server model based on AS connections provides an overview of virtual paths through the Internet and this model can be used for better understanding the influence of different AS on overall path performance. A Client-Server Internet connectivity model is shown in Fig. 1 and its main elements are described below:

- **A. Client Machine.** This is a computer running a standard Operating System and having dial-up or dedicated line access to the Internet. For our study, the client ran NT 4.0 Operating System, had a Pentium class hardware configuration and was connected to the University of Adelaide Internet gateway via a 100 BaseT Local Area Network (LAN). The Client machine had the software utilities for automating the *Traceroute* utility and registering *Traceroute* path parameters.
- **B. Local Network**. This is the AS that provides Internet access to the client machine. For a dial-up user such a network is a traditional ISP, for a dedicated line user a broadband ISP, research/commercial organisation or backbone provides the connectivity. For our study, the local access is via the University of Adelaide Internet gateway comprising the later.
- **C. Local/Remote Regional Network.** This is the AS that provides connectivity to local networks. Depending on the geographic location of the regional network (serving the Client or the Application Server), it can be classified either as local or remote.
- **D. Backbone.** A backbone is a transit AS, providing *regional, natio*nal or *international* coverage. For this study, three types of backbone interconnections are defined: the *Australian backbone, the US backbone* and *the remote server backbone*.
- **E. Remote LAN.** This network provides Internet connectivity to an Application Server.
- **F. Application Server.** The Application Server is used for storing digital content such as Web pages and program files. File application servers can have mirror sites closer to users, resulting in better response time and higher throughput [8]. In our case, Application Servers were used as destination targets for the *Traceroute* utility running at the Client machine.

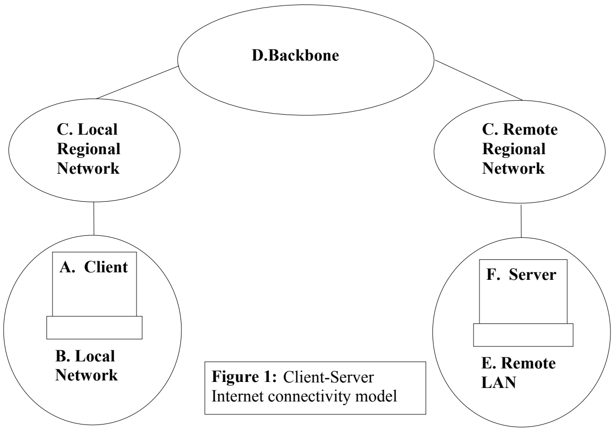

Figure 1: Client-Server Internet connectivity model

This model will be used as the basis of the experiment to assess Internet performance of those terms discussed in Section 2.

4. MINIMUM RTT DELAY MODEL:

A RTT delay equation for a certain virtual path is defined as

$$RTT_{Total} = RTT_{\text{Propagation}} + RTT_{Switching} + RTT_{Queuing}\ , \qquad \textbf{(1)}$$

where RTT_{Total} is the RTT delay from a Client to an Application Server, $RTT_{\text{Propagation}}$ is the RTT delay due to propagation of electrons and photons on a particular medium, $RTT_{Switching}$[2] is the RTT delay for transmitting/switching a packet over a network interface and $RTT_{Queuing}$ is the RTT delay due to packets being delayed by congestion in routers.

A minimum RTT delay model provides an estimate of a best end-to-end RTT performance along a virtual path. Packets experience a minimum RTT delay when there *is no queuing delay* on routers along a virtual path, i.e., there is *no network congestion*. In this case, Equation (1) reduces to

$$RTT_{\text{Minimum}} = RTT_{\text{Propagation}} + RTT_{Switching}\,. \qquad \textbf{(2)}$$

Similarly to the minimum RTT delay model, it is possible to define a minimum One-Way (OW) delay model for end-to-end analysis of a virtual path. In this case, it is said that

---

[1] Autonomous Systems are a group of subnetworks administered by a Single Administrative Authority (SAA) with a set of Interior Gateway Protocols (IGPs). The SAA is normally a Network Service Provider or large organisational network (e.g., campuses and corporate networks).

[2] $RTT_{Switching}$ is also known in the literature as $RTT_{Transmission}$

$$OW_{Minimum} = OW_{\text{Propagation}} + OW_{Switching}, \quad \textbf{(3)}$$

where

$$OW_{\text{Propagation}} = \sum_{\text{First link}}^{\text{Last link}} OW_{\text{Inter-city propagation delay}}, \quad \textbf{(4)}$$

$$\text{OW}_{\text{Switching}} = \text{Typical Switching Delay} \times \frac{(2N-1)}{2} \quad \textbf{.(5)}$$

N is the number of nodes within a client-server virtual path. For Equation (4), the One-Way inter-city propagation delay is calculated based on what kind of transmission medium is used for interconnecting two cities. Fibre or GEO satellite links are the typical transmission mediums for backbones and regional networks. While for fibre the propagation speed is 200,000 Km/s, propagation delays on GEO links vary between 250 to 300 ms [9]. In our study, a One-Way GEO satellite propagation delay between two satellite-connected cities is asserted as 250 ms, independently of where geographically the connection is originated and terminated. For a fibre link between two cities, the propagation delay is calculated as

$$OW_{\text{Inter-city propagation delay}} = \frac{\text{Distance between cities(Km)}}{200}, \quad \textbf{(6)}$$

where $OW_{\text{Inter-city propagation delay}}$ is expressed in *ms*.

The *Typical Switching Delay* (shown in Equation 5) is associated to a minimum typical path definition, which is introduced in Section 6.3. The Typical Switching delay is defined in Section 6.4.2.

Moreover, for *symmetric* upstream and downstream paths it can be assumed that

$$RTT_{\text{Minimum}} \approx 2 \times OW_{\text{Minimum}} \quad \textbf{(7)}$$

Furthermore, in the case where Client and Application Server are located at the same geographic region, Equation (3) reduces to

$$OW_{Minimum} = OW_{Switching}. \quad \textbf{(8)}$$

5. SOFTWARE SUPPORTING THE EXPERIMENT

Software utilities for collecting RTT response delays and for studying the typical virtual paths to any given Application Server were developed in Visual Basic code. Due to space limitations, the software description is not provided. For further details about this software, read [10].

6. DATA ANALYSIS PROCEDURES & ASSUMPTIONS

Data was collected to study RTT delays from our Client to remote servers and from this information, identify typical paths and estimate interconnecting networks characteristics. A number of assumptions were made, which are discussed below:

6.1 The Application Servers (See Section 3, Item F) are common file repositories used by Internet users. The popular TUCOWS[1] Web site was selected, which has a number of mirror servers on all continents. The selected mirror sites were:

- *Australia:* South Australia & Victoria;
- *Asia:* Honk Kong;
- *USA:* West Coast & East Coast;
- *Europe:* England & Germany;
- *Middle-East*: Israel;
- *South America:* Brazil & Argentina;
- *Africa:* Zimbabwe & South Africa.

Regional classification was based on expected delays for different geographic regions visible from Adelaide, Australia. One computer was used as a Client from Nov 3 until Nov 30 1999, running a software utility for automating *Traceroute* and collecting path samples to all Servers. This utility ran in parallel to all sites and sampled respective paths every five minutes.

For the purpose of confirming the geographical location of each server, a shareware network analysis tool called NEOTRACE[2] was used for obtaining the server's latitude/longitude coordinates. By checking the coordinates on a map, it was possible to verify the veracity of the 'NEOTRACE coordinate lookup' and the most appropriate city location for each server. When the coordinate lookup would generate a mismatch in relation to the Web map viewer, the TUCOWS server administrator would be contacted for confirming its geographical location.

6.2 *Traceroute* provides *upstream* performance information (RTT, packet loss and IP/DNS addresses) for every node within a virtual path. Each node within the path is probed with three echo-request datagrams. For each node within a path, it is possible to obtain estimated RTTs and IP/DNS address details. In case echo-request datagrams are lost for a particular node, a star (*) is registered to indicate packet loss.

Our observation shows that normally the total number of *Traceroute* echo-replies (RTT replies) for each node

---

[1] Tucows servers are configured based on a suggested hardware/OS configuration. During this study the suggested characteristics were: Pentium 133 or greater, UNIX OS, 32Mb of RAM or greater, 8Gig hard disk space, T-1 or greater bandwidth.

[2] http://www.neoworx.com/

within a virtual path has a similar occurrence distribution. As an example, if there are 1000 RTT replies for node 1 during an experimental period of time *t*, it can be expected that all nodes within the same *virtual path* will also have approximately 1000 RTT replies.

By simultaneously analysing the number of RTT replies for each node in a particular client-server path and the node's IP and DNS addresses, *typical virtual paths* could be identified for *each* Application Server in relation to our Client. It was only considered typical virtual paths as the virtual paths that were persistent [11] and not short-lived.

The main value of characterising a typical virtual path is that such an approach allows a better classification of nodes in accordance to the client-server Internet connectivity model introduced in Section 3. This provides a better understanding of each AS contribution on overall delays. Since AS are administered by different organisations with both dissimilar geographical reach and QoS settings, RTT, packet loss and other network parameters might vary considerably across an Internet virtual path.

For ease of analysis, nodes within a typical path are classified based on the client-server Internet connectivity model described in Section 3 and with the following assumptions:

- The first and last nodes are classified as Local Network nodes. This is done because the first node is expected to be an Internet gateway router either for the Client or the Application Server;
- The second and second last nodes are classified as Local/Remote Regional Network nodes. This is a hierarchical assumption based on the fact that these networks provide connectivity to Local Networks.
- Intermediary nodes are classified as one of the following types of backbone nodes: Regional, remote server, Australian, or US. For this classification, it was observed DNS information and RTT delays between nodes. In some cases, RTT delays indicate a satellite link interconnecting two nodes and therefore, it is possible to have a better idea of a node's geographical location.

In this study, the classification of *typical virtual paths* is used for introducing the *minimum RTT delay model* (See Sections 4 & 6.4) and for hypothesising path characteristics within a typical virtual path. Since this model is based on the calculation of end-to-end RTT delays for a congestion-free path, it does not focus on specific AS contributions within a virtual path on overall performance.

6.3 A *minimum typical path* for a client-server connection is defined as a typical path with the *lowest end-to-end* RTT delay measured during the experimental period. Kalidindi & Zekauskas [12] observed that at least one measurement packet is expected not to experience congestion within a congested Internet virtual path. Because their probing sample rate (2 packets per second) was much higher than our *Traceroute* implementation sample rate (1 *Traceroute* sample for all Application Servers, every 5 minutes), it is assumed that at least one packet is going to experience either *no* congestion *or* a low level of congestion during our experimental period.

6.4 For a *minimum typical path*, some assumptions were also made for estimating the minimum One-Way delay for a particular client-server connection. The minimum One-Way delay calculation is based on Equation (3).

6.4.1 It is assumed that DNS information obtained via the *Traceroute* sample inspection (See Section 6.2) of a minimum typical path provides a good estimate of geographic location[1] for each node.

For estimating the One-Way propagation delay for a minimum typical path, it is assumed that Internet transmission links *between continental cities* are *fibre* links while links between *intercontinental cities might be satellite or fibre links*. This assumption is used for calculating propagation delays between different cities within a *minimum typical path*. The calculation is based on Equation (4) and by considering two facts:

- GEO satellites links have a typical One-Way propagation delay of 250ms;
- Light propagation speed within fibre is 200,000 Km/s.

6.4.2 A *typical switching delay* is defined as the average time for routers in a *minimum typical path* for processing probing packets. For our research, it is used the fact that the South Australian Tucows server is in the same region of our Client for calculating the typical switching delay via Equation (8). It is also assumed that for this client-server connection there is symmetry between upstream and downstream paths (See Equation 7). Therefore, combining Equations (7) & (8), it is found that

$$OW_{Switching} \approx \frac{RTT_{Total}}{2} . \qquad \textbf{(9)}$$

Based on our assumption in Section 6.3 that at least one packet will experience *no* delay or *a low* delay level within a minimum typical path, it was found that the minimum *measured* RTT delay for the South Australian server is 10 ms. Taking into account that the *minimum typical path* for the South Australian Application Server has 7 nodes and using Equations (5) & (7), the *Typical Switching Delay* was calculated as

$$OW_{Switching} = 10 / 2 = 5ms \overset{7\text{ nodes}}{\Rightarrow}$$

---

[1] It was observed from *Traceroute* sample analysis that normally DNS addresses have information about airport codes *or* city initials.

$$\text{Typical Switching Delay} = \frac{5}{6.5} = 0.76\ ms\ . \qquad \textbf{(10)}$$

It is assumed that this typical switching delay is going to be a typical value for router switching in *any* node within *any* virtual path. This value is used in Equation (5) for estimating the $OW_{Switching}$ within *any* minimum typical path.

7. PATH CHARACTERISTICS BASED ON THE MINIMUM RTT ANALYSIS

In this section we apply our *Minimum RTT Delay Model* for calculating the *minimum typical path* characteristics for a number of Application Servers. For a thorough discussion about all servers and their typical paths on a step-by-step basis, read [10].

After classifying the *typical paths* to each Server and identifying among the typical paths which one had the *measured minimum RTT* delay, a *minimum typical path* is chosen for further analysis (See Sections 6.3 & 6.4) . It is assumed that upstream and downstream have a certain level of symmetry. Based on the *measured minimum RTT*, the type of propagation mediums that interconnects our Client to the Remote Servers is hypothesised. This is done by calculation as described in Section 4 and by assuming that Internet transmission links between continental cities are *fibre* links while links between intercontinental cities might be *satellite* or *fibre* links. For ease understanding, each calculation is discussed in a server-by-server basis. As a result of space limitations, *Traceroute* typical paths are not shown in this Section.

7.1 SOUTH AUSTRALIA

The South Australian server only had one typical path with 7 nodes. Therefore, this path is the *minimum typical path*. As discussed in Section 6.4.2, any propagation delay for this Server can be disregarded and therefore *the minimum measured RTT* delay is used for finding a *typical switching router* delay. This value is used for *all* virtual paths for calculating the total switching delay within a path.

7.2 VICTORIA, AUSTRALIA

Since this Server only had one *typical path* (6 nodes), this path is the *minimum typical path*. The *minimum measured RTT* was 13.33 ms.

This Application Server is approximately 662 Km from our Client; therefore, the two-way propagation delay is calculated as 6.62 ms (Equation 6). By adding the switching delay contributions of all routers within this path, the minimum RTT delay is calculated as 14.98 ms (6.62+ 0.76 x 11)[1]. This is found to be a close value to our *measured minimum RTT* delay. For this case, deviations might be result of overestimating the router switching delays, measurement inaccuracy [13] or path asymmetries.

7.3 USA WEST COAST

There were 3 *typical paths* for this Server. It was found that the *minimum typical path* for this server has 14 nodes and the *minimum measured RTT* is 300.33 ms. The minimum typical path follows the sequence: Adelaide, Sydney, San Francisco, Los Angeles, San Francisco. While upstream and downstream paths are assumed symmetric, inter-continental path characteristics might differ though (e.g., upstream via satellite, downstream via fibre).

While estimating the *minimum typical path* characteristics, it was found three possible path configurations. Since it is assumed that all *intra-continental* links are fibre links, path differences are due to dissimilar medium considerations for the *inter-Pacific* links, which can be *all-fibre, all-satellite* or *hybrid fibre-satellite*. By calculating all possible *minimum RTTs through different medium combinations, we found* that the closer value (352 ms) to the minimum measured RTT suggested a hybrid Trans-Pacific link. The hybrid Trans-Pacific link hypothesis was confirmed by OPTUS, which is the Trans-Pacific link carrier for traffic requests originated within AARNET network sites. All *upstream traffic* to the US is sent via fibre links while *downstream traffic* flows via either cable or satellite links. The downstream link selection will depend on *both* the type of traffic as router-map filters are applied *and* on the capacity of the links (particularly, OPTUS has *much higher* satellite link capacity *than* fibre link capacity). For example, traffic coming from Asia-Pacific links has higher preference within OPTUS network and thus follows a downstream fibre link. This hybrid Trans-Pacific link implementation is used for supporting our discussion in the following sections.

7.4 HONG KONG

Since this Application Server only had a *typical path* during the experimental period and the upstream and downstream paths were found to be *asymmetric*[2], it is only possible to say that the typical path measured is the *minimum upstream typical path* to the Server. The *Traceroute* shows an upstream path following: Adelaide, Sydney & Hong Kong. The downstream path, which was identified after contacting OPTUS, the carrier providing backbone facilities for AARNET[3], follows: Hong Kong, California, Sydney & Adelaide. The *minimum measured RTT* was found as 170 ms.

By considering this path asymmetry, different path topologies were calculated. By comparing the calculated

---

[1] For not being repetitive, the switching delay calculation is omitted to other servers. One can easily calculate it by using the Typical Switching Delay (0.76 ms) in Equation (10).

[2] This finding was result from a throughput performance research that we have carried out. For further information, read [14]

[3] Australian Academic Research NETwork

values to the *minimum delay measured*, it is possible to say that an all-fibre intercontinental topology (i.e., Australia to Hong Kong, Hong Kong to California and California to Australia) is the best choice. The calculated RTT delay for an all-fibre path is 178.34 ms. This hypothesis was later confirmed by OPTUS. Since Asia-Pacific links have higher local preference within OPTUS, traffic coming from these regions via the US will flow via fibre links.

7.5 ISRAEL

It was found that the minimum typical path for this server has 17 nodes and follows: Adelaide-Sydney-Los Angeles-New York-Telaviv. The minimum measured RTT was 471 ms.

While hypothesising about the medium configurations for this minimum typical path, two path configurations were found to have a RTT close to the *minimum measured* RTT. The calculated RTT for these paths are 469 ms and 484 ms. For ease understanding, we call these paths as A and B, respectively. While continental links are fibre links (by assumption), path configuration differences between paths A & B are due to a hybrid Trans-Pacific link (path A) or a hybrid Trans-Atlantic link (path B). In addition, Path A has an all-fibre Trans-Atlantic link and path B has an all-fibre Trans-Pacific link. Since most of the Client-Server connections cross the standard OPTUS US gateway in California (See Section 7.3) and the Trans-Pacific backbone provider has confirmed the hybrid fibre-satellite implementation within its link, it is possible to argue that Path A is the most adequate medium combination.

7.6 ENGLAND

It was found that the *minimum typical path* for this server has 15 nodes and a *minimum measured RTT* of 450.66 ms.

Similarly to the Israeli Server, while calculating the minimum typical path, it was noticed two path configurations having a RTT close to the minimum measured RTT (438.9 ms & 469.8 ms). Likewise, it is asserted that the best path is the one that has the hybrid Trans-Pacific link and the all-fibre Trans-Atlantic link (438.9 ms). Deviations might be result of measurement inaccuracies *or* path asymmetries.

7.7 ARGENTINA

It was found that the *minimum typical path* for this server has 17 nodes and follows: Adelaide-Sydney-Los Angeles-Atlanta-Orlando-Buenos Aires. The *minimum measured RTT* was 788 ms and it was noticed that three different paths have RTT close to this value.

Since there was just one path with a hybrid Trans-Pacific link, this was the selected path. This path has a hybrid Trans-Atlantic link and deviations from the observed value might be due to path asymmetries & measurement inaccuracies.

7.8 BRAZIL

It was found that the *minimum typical path* for this server has 17 nodes and follows: Adelaide-Sydney-Los Angeles-New Jersey-São Paulo. The *minimum measured RTT* was 691 ms and similarly to the Server in Argentina, it was noticed three different paths that have RTT close to this value.

Since there is only one path with a hybrid Trans-Pacific link, this is the selected option (calculated RTT of 673.44 and hybrid intercontinental links). Deviations from the minimum measured RTT are result of path asymmetries and measurement inaccuracies.

7.9 ZIMBABWE

It was found that the *minimum typical path* for this server has 16 nodes and the *minimum measured RTT* was 851.33 ms. It was observed that two paths have values around this minimum RTT. By observing that only one of the paths has a Hybrid Trans-Pacific link, it is asserted that this path is the best option.

8. CONCLUSIONS & FURTHER WORK

Our research proposes a methodology to estimate Internet path characteristics based on the analysis of minimum RTT delays for a typical Internet path and some experimental assumptions. The methodology was tested and a number of results confirm with great accuracy the experimental assumptions.

Deviations are due to overestimating router switching delays, measurement inaccuracy or path asymmetries. In the general case, it was assumed that upstream and downstream paths are likely the same. Nevertheless, in particular cases such as Hong Kong, the paths are different because of peering agreements between ISPs and their backbone providers. For providing more accurate results, *Traceroute* could be deployed in both client and server machines, which would provide more accurate results. Synchronisation could be done by GPS [15] or a hybrid GPS-Network Time Protocol (NTP) [16].

Another improvement that would result in better analysis is separating the packet generation process from the measurement process. As discussed by Cleary *et al* [13], RTT measurements might differ up to 30 ms when the probing packet generation and data measuring processes are carried out in the same machine.

In some cases (e.g., USA East Coast), the analysis methodology could not be employed due to ICMP echo-requests being blocked within a router. In addition, as reported in other research [17], ICMP packets might have low priority and suggest a poorer network performance than the network actually has. For improving the methodology, the RTT measurement could be implemented with a different measurement protocol such as the IP Measurement Protocol (IPMP). For further information about the IPMP, read [18].

REFERENCES

[1] *M. Allman,* "Improving TCP Performance Over Satellite Channels", Master's thesis, Ohio University, June 1997, http://roland.grc.nasa.gov/~mallman/papers/thesis.ps

[2] *H. Kruse,* "Performance of Common Data Communications Protocols Over Long Delay Links – An Experimental Examination", 3rd International Conference on Telecommunication Systems Modeling and Design, pp. 409-415, 1995, http://jarok.cs.ohiou.edu/papers/kruse.ps

[3] *M. Allman & Aaron Falk,* "On the Effective Evaluation of TCP", Computer Communication Review, Vol. 29, No. 5, pp. 59-70, October 1999, http://roland.grc.nasa.gov/~mallman/papers/tcp-evaluation.ps

[4] *J. Semke, J. Mahdavi & M. Mathis,* "Automatic TCP Buffer Tuning", Computer Communication Review, Vol. 28, No. 4, pp. 315-323, October 1998, http://www.psc.edu/networking/ftp/papers/autotune_sigcomm98.ps

[5] *M. Allman,* "A Web Server's View of the Transport Layer", ACM Computer Communication Review, Vol. 30, No. 5, October 2000, http://roland.grc.nasa.gov/~mallman/papers/webobs-ccr.ps

[6] *W.R. Stevens & G. R. Wright,* "TCP/IP Illustrated – Volume 1", 16th printing, Chapter 24, pp. 344-345, Addison-Wesley, 2000.

[7] *V. Jacobson, R. Braden & D. Borman,* "TCP Extensions for High Performance", RFC1323, IETF, 1992, http://www.ietf.org/rfc/rfc1323.txt?number=1323

[8] "Australia's International Bandwidth", Telecommunication Journal of Australia, Vol. 49, No. 1, pp. 3-12, 1999.

[9] *A. Tanenbaum,* "Computer Networks", 3rd edition, Chapter 6, pp. 555-572, Prentice Hall PTR, USA, 1996.

[10] M. E. M. da V.Alves, "Characterisation of End-to-End Perfomance for Web Based File Server Repositories", Master's Thesis, The University of Adelaide, Australia, January 2001.

[11] *V. Paxson,* "End-to-End Routing behaviour in the Internet", IEEE/ACM transactions on Networking, Vol. 5, No. 5, pp. 601-615, Oct. 1997.

[12] *S. Kalidindi & M. J. Zekauskas,* "Surveyor: An Infrastructure for Internet Performance Measurements", Advanced Network & Services, http://telesto.advanced.org/~kalidindi/papers/INET/inet99.html

[13] *J. Cleary, S. Donnelly, I. Graham, A. McGregor & M. Pearson,* "Design Principles for Accurate Passive Measurement", in Proceedings "Passive & Active Measurement Workshop", Hamilton, New Zealand, April 3 & 4, 2000.

[14] M. E. M. da V. Alves, S. Nesterov & R. P. Couts, "Characterisation of End-to-End Performance for Web Based File Repositories", Proceedings of "International Conference on Performance and QoS of Next Generation Networking – P&Q Net 2000", November 27-30, 2000, Nagoya, Japan.

[15] *S. Kalidindi & M. J. Zekauskas,* "Surveyor: An Infrastructure for Internet Performance Measurements", Advanced Network & Services, http://telesto.advanced.org/~kalidindi/papers/INET/inet99.html

[16] "Executive Summary – Computer Network Time Synchronization", Electrical & Computer Engineering Department, University of Delaware, http://www.eecis.udel.edu/~ntp/ntp_spool/html/exec.htm

[17] *W. Mattews & L. Cottrell,* "The Pinger Project: Active Internet Performance Monitoring for the HENP Community", IEEE Communications magazine, Vol. 38, No. 3, pp. 130-136, May 2000.

[18] "IP Measurement Protocol", National Laboratory for Applied Network Research, http://watt.nlanr.net/AMP/IPMP/