Measuring IP Network Performance
– The SingAREN's "Approach"

Heng Seng Cheng, Lek Heng Ngoh, Jonathan Ong, Chan Keong Yip, Soon Meng Ow and Sharon Lim

***Abstract*—As a network operator, one of the key factors to the success of SingAREN is its ability to maintain a high standard of network performance which can only be achieved by adopting the appropriate QOS metrics and measurement tools. This paper discusses some of the challenges in the field of network measurement which SingAREN is facing as its network expands and as its network traffic increases. To reduce the cost in the international links, SingAREN is likely to employ satellite links in the near future. The different characteristics of satellite links compared to optical fibre links raises some issues for network measurement and these issues are discussed in the paper.**

***Index terms*—Delays, Ethernet, geostationary satellite, packet loss, QOS, TCP.**

I. INTRODUCTION

Singapore Advanced Research and Education Network (SingAREN) is a national initiative to create a high-speed broadband network platform to support R&D and advanced technology development in Singapore, serving users from academia, research organisations and industry [1]. Within Singapore, this comprises a broadband network linking the two local universities and several research organisations [1].

Outside the country, SingAREN has a 14 Mbps link to the United States, a 2 Mbps link to Japan and another 2 Mbps link to Korea [1]. In addition, works are being carried out (using a 2Mbps geostationary satellite link) to prepare the launch of IP services via satellite in the near future.

To support the large number of users that run various types of applications (e.g. file transfer, web browsing and video-conferencing), SingAREN's network has to maintain a high standard of performance. To work towards that goal, SingAREN has adopted a set of quality of service (QOS) metrics, some of which were proposed by the IP Performance Metrics (IPPM) Working Group in the Internet Engineering Task Force (IETF). In this paper, we discuss the experience of SingAREN in the use of various QOS metrics and network measurement tools.

In SingAREN, network performance is measured for several reasons:

(i) to determine if the network performance achieve a minimum level (similar to a GO/NO GO test),

(ii) to diagnose the cause of performance degradation or failure in a network,

(iii) to detect any trend of performance degradation in the network so that preventive measures may be taken before the performance of the network falls below the minimum level and

(iv) to monitor the utilisation of the network for the purposes of upgrading network capacity and cost recovery (i.e. billing) in the future.

In Section II, we look at one-way measurement with particular focus on the measurement results of the Surveyor machine which SingAREN has deployed in its network in collaboration with Advanced Network and Services. Section III is a brief discussion on round-trip measurement with particular focus on its strengths and weaknesses. Section IV covers SingAREN's use of TCP measurement metric and introduces SingAREN's proprietary TCP measurement tool which provides compensation for processing delay. In Section V, SingAREN's work on passive measurement of network traffic is discussed. Passive measurement can be used to facilitate cost recovery (i.e. billing the users for network usage) and to provide critical information on network utilisation prior to an upgrade of network capacity. In Section VI, the issues involved in the measurement of satellite-based network is discussed.

II. ONE WAY MEASUREMENT

Two QOS metrics that are related to one-way measurements have been defined by Almes, Kalidindi and Zekauskas in RFC 2679 and RFC 2680. They are one-way packet loss [2] and one-way delay [3].

Heng Seng Cheng, Lek Heng Ngoh, Jonathan Ong and Chan Keong Yip respectively hold the positions of Associate Research Staff and Operations Manager at SingAREN, Ayer Rajah Crescent #09-11/12, Ayer Rajah Industrial Estate. Heng Seng Cheng, can be contacted at email address: hscheng@singaren.net.sg. Soon Meng Ow and Sharon Yee Pin Lim are undergraduates from Nanyang Technological University, Nanyang Avenue, Singapore 639798.

One-way measurement is useful because the path from a source to a destination in the Internet may be different from the path from the destination back to the source ("asymmetric paths"), such that different sequences of routers are used for the forward and reverse paths [2, 3]. Therefore, round-trip measurements actually measure the performance of two distinct paths together [2, 3]. Measuring each path independently highlights the performance difference between the two paths which may traverse different Internet service providers, and even radically different types of networks [2, 3].

SingAREN has deployed a one-way measurement called Surveyor which was built and remotely controlled by Advanced Network & Services. The Surveyor machine performs one-way packet loss measurement and one-way delay measurement with more than thirty sites outside Singapore. Most of these sites are in the United States while the rest are in Japan, Korea and Europe.

The Surveyor machine has demonstrated to us the usefulness of one-way measurement on quite a number of occasions. For example, on 10 September 1999, the Surveyor machine detected that SingAREN's overseas link to Asia-Pacific Advanced Network Consortium (APAN) in Japan was experiencing a high percentage of packet loss and large delay variations in the forward path (i.e. in the direction from SingAREN to APAN, Japan). In the return path (from APAN, JAPAN to SingAREN), very low percentage of packet loss was encountered which is a very good example of asymmetricity within the Internet. Results of the measurements are on display at the web site http://www.advanced.org/surveyor/.

The Surveyor machine also provides traceroute information in both forward and return paths this information have been proven to be extremely useful in isolating the cause of a network problem. For example, we were able to pin-point a router (which implemented rate control mechansim to regulate traffic into SingAREN's international link with the United States) in the Abilene network as the most likely cause of delay variations detected by the Surveyor machines on 10 September 1999 between SingAREN and two US universities (University of Chicago and Stanford University).

In that instance, the one-way delay metric from Stanford University to SingAREN and the one-way delay measurement from University of Chicago to SingAREN exhibited similar pattern of delay variations. On the other hand, the one-way delay metric from Stanford University to University of Chicago exhibited no delay variations. After comparing the traceroute information from Stanford University to SingAREN with the traceroute information from Stanford University to University of Chicago, it was discovered that both paths were identical to each other except on the last few routes (see Figure1):

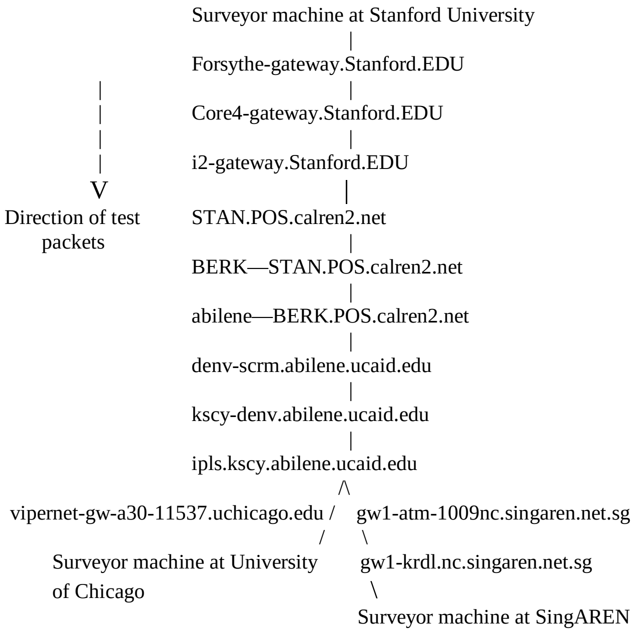

Figure 1: The domain names (or routers) which some of the test packets from the Surveyor machine in Stanford University passed through on 10 September 1999 as they were routed to SingAREN and University of Chicago.

Since the path from Stanford University to University of Chicago had insignicant delay variation, the possible causes of the delay variations are one (or more) of the routers at following domains:

(a) ipls.kscy.abilene.ucaid.edu

(b) gw1-atm-1009nc.singaren.net.sg

(c) gw1-krdl.nc.singaren.net.sg

We found out that the routers at gw1-atm-1009nc.singaren.net.sg and gw1-krdl.nc.singaren.net.sg did not implement rate control to regulate traffic coming from Abilene network. We also found out that other paths from Abilene network to the SingAREN's surveyor machine did not exhibit delay variations. This allows us to eliminate the routers at gw1-atm-1009nc.singaren.net.sg and gw1-krdl.nc.singaren.net.sg as the cause for the delay variations experienced.

Therefore, it may be concluded that the rate control mechanism at the router in Abilene network (which regulated traffic into SingAREN's international link with the United States) is the most likely cause of the delay variations. To make diagnostic even faster, a program can be written to automatically identify the "trouble-making" node based on the algorithms that have just been described.

The main disadvantage of one-way measurement is the need for each measurement machine to install a global positioning system (GPS) device that is used to synchronise the clocks at each machine with the clocks at another machine. As a PCI-based GPS device is quite expensive, deployment of one-way measurement machine can be an expensive undertaking. The alternate solution to a GPS device is to employ *Network Time Protocol (NTP)* which was is used to synchronize the time of a computer client to another server or reference time source [4]. The disadvantage of *NTP* is its inaccuracies can be as high as a few tens of milliseconds whereas the inaccuracies of GPS can be as low as tens of microseconds.

After several months of experience with the Surveyor machine, SingAREN are fully convinced that one-way measurement is an essential measurement methodology for its network.

III. ROUND-TRIP MEASUREMENT

In comparison to one-way measurement, round-trip measurement is much easier to implement as the same measurement machine is used to transmit and receive the test packets. This means that there is no need to perform clock synchronisation between two separate measurement machines.

To have a better understanding of the capabilities of round-trip measurement, SingAREN has deployed a Skitter measurement machine on its network. The Skitter measurement machine which was built and controlled remotely by the Cooperative Association for Internet Data Analysis (CAIDA) in the United States can be used to measure the round-trip time between itself and hosts around the world.

The disadvantage of a round-trip measurement machine is it actually measures the forward and the return paths together. Because of this, if the measurements indicate poor performance between the measurement machine and a host, it is impossible (or at least extremely difficult) to determine whether the forward or the return path is having a problem. Because of this, a round-trip measurement can only be used to perform a GO/NO GO test on a network.

However, because of its simplicity and relatively low cost (since no GPS is required), SingAREN will continue to employ round-trip measurements for checking its network performance.

IV. MEASUREMENT OF TCP PERFORMANCE ACROSS THE INTERNET

It is estimated that 90 to 95 percent of all Internet traffic are Transmission Control Protocol (TCP) packets [5]. Because of this, TCP performance is an important QOS indicator. TCP performance will provide network operators with the insights into the network performance from the end users' perspective.

Besides, one-way measurement and round-trip measurement may not be able to detect performance problems that is associated with TCP protocols. This is because there is evidence that most TCP implementations exhibit non-linear performance over some portion of their operating region to the extent that an increase in capacity may actual reduce the perceived quality of the network [6].

An Internet-draft on the measurement of TCP performance across a network was written by Mathis and Allman [6]. In this internet-draft, the QOS metric for TCP performance is called Bulk Transport Capacity (BTC) which is defined as a measure of a network's ability to transfer significant quantities of data with a single congestion-aware transport connection (e.g. TCP) [6]. The intuitive definition of BTC is the expected long term average data rate (bits per second) of a single ideal TCP implementation over the path in question [6].

A major challenge to specifying the BTC metric is the existence of many congestion avoidance algorithms permitted by IETF [6]. The allowed diversity is sufficient to lead to situations where different implementations will yield non-comparable measures and potentially fail the formal tests for being a metric [6]. The solution to this problem is to specify the BTC metric more tightly than the typical IETF protocol [6].

As congestion avoidance algorithms of TCP are still evolving [6], the authors believe the standardisation of the BTC metric may take some time to be completed and it is possible that we may end up with a BTC metric that itself needs to evolve with the TCP standards. This may make widespread acceptance of BTC metric as a QOS metric rather difficult. Hence, SingAREN has decided to use a popular TCP implementation for its BTC metric in the meantime and migrate to the standard BTC metric when it becomes available.

Ideally, a TCP measurement machine should only measure the performance of a network and not let its processing capability (of the TCP protocols) affect the measured results. In practice, it is acceptable if the BTC of the underlying network dominates the performance as perceived by a user [6]. However, this is not always the case especially for a network spanning a small geographical area such as Singapore where the processing delay of TCP measurement machines (end-hosts) will contribute significantly towards the measurement results.

Figure 2(a) is a simple illustration of an ideal TCP session where the end-hosts do not contribute to any delay and Figure 2(b) is an illustration of a typical TCP session where the processing delay of the end-hosts contribute (degrade) to the overall performance of the TCP session.

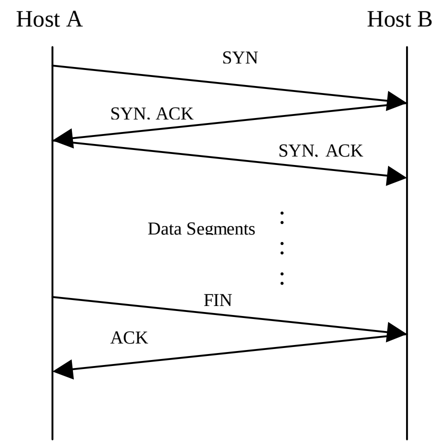

Figure 2(a): of an ideal TCP session where the end-hosts do not contribute to any delay

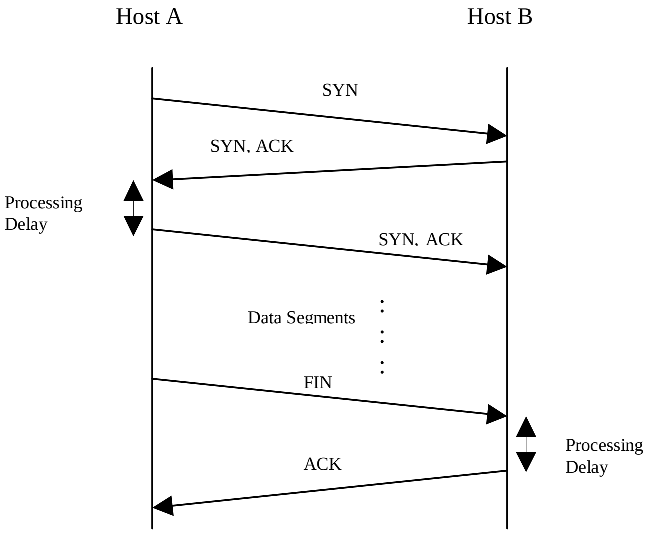

Figure 2(b): of a typical TCP session where the processing delay of the end-hosts contribute (degrade) to the overall performance of the TCP session.

To investigate the impact of the processing delay of the end-hosts on the performance of a TCP session, we performed a TCP data transfer between two end-hosts. During the TCP data transfer, a HP Internet Advisor was used to record the timestamps of TCP packets arriving and leaving one of the end-hosts (a SUN Sparc5 workstation). We also measured the round-trip delay between the two end-hosts.

The experiment revealed that while the round-trip delay between the two end-hosts is approximately 3 ms, the longest processing delay incurred by the SUN Sparc5 workstation is 1.77 ms which is proportionally large. This shows that if a TCP measurement machine is implemented using a SUN Sparc5 or any other general purpose computers with similar computing power, then the processing delay cannot be ignored.

To reduce the processing delay of a TCP measurement machine, one may implement the TCP protocols in a programmable Ethernet card where the CPU is sufficiently powerful. This may significantly reduce the processing delay as there is no passing of TCP packets between the network card and the general purpose computer via the PCI bus. Further improvement in performance can be achieved if there is little or no interrupts to the CPU as it executes the software for the TCP protocols. Such TCP measurement machines are commercially available but they are likely to be quite expensive which makes them economically unviable for deployment on a large scale.

Because of this reason, SingAREN has considered the option of developing her own TCP measurement machine where the TCP protocols are implemented on a programmable Ethernet card installed with a powerful CPU. However, the commercially available programmable Ethernet card that we encountered so far uses a CPU that is several times less powerful than those found on a normal PC which is not helpful in reducing the processing delay. Although it is possible to design and develop our own programmable Ethernet card that uses a powerful CPU, the development time is likely to be relatively long which makes this an unattractive option.

Due to the need to deploy TCP measurement machines very soon, SingAREN has taken the option of building TCP measurement machines with commercially avaliable general purpose computers and network cards. A proprietary algorithm will also be implemented to compensate the measured TCP throughput for the processing delay inccurred by the general purpose computers. The following subsection describes SingAREN's TCP measurement machine in greater details.

A. A TCP Measurement Machine With Compensation For Processing Delay

In this approach, the TCP measurement machine consists of

(i)an Intel-based personal computer (PC) which implements a popular TCP standards and

(ii) an embedded PCI board with Ethernet/Fast Ethernet that can be programmed to timestamp the arrival and departure of Ethernet frames. Figure 3 shows the block diagram of SingAREN's TCP measurement machines and the setup configuration.

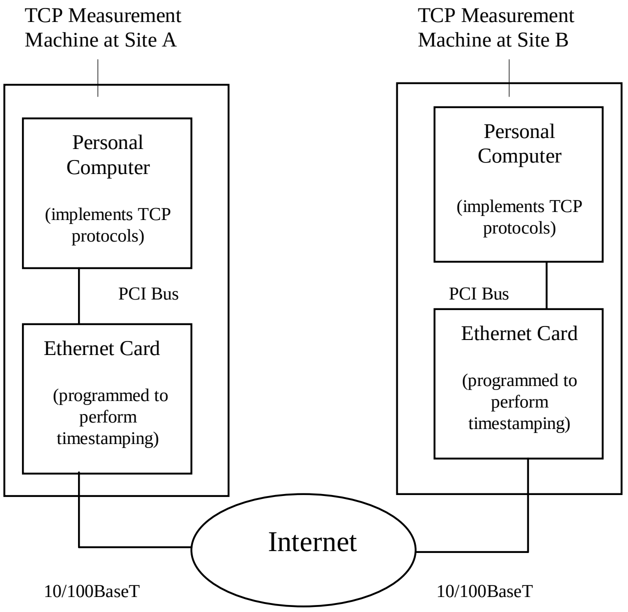

Figure 3: Block Diagram of SingAREN's TCP measurement machines and the setup configuration.

Figure 4 shows the time line of a TCP session where the arrival and departure of Ethernet frames are accurately timestamped by the on-board microprocessor on the Ethernet card. With the knowledge of the timestamping information, we will then be able to compensate the TCP session time for the processing delay of the PC:

TCP session time (with compensation for the processing delay)

= t_a8 - t_a1 – (Host A's processing delay) – (Host B processing delay)

= t_a8 - t_a1 – [(ta3 – ta2) + (ta4 - ta3) + (ta5 – ta4) + (ta6 – ta5) + (ta8 – ta7)] – [(tb2 – tb1) + (tb7 – tb6) + (tb9 – tb8)] (1)

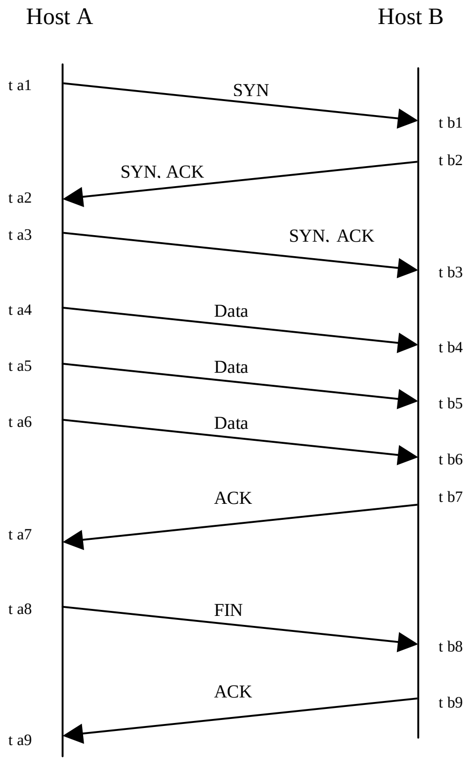

Figure 4: A TCP session where data is transferred from Host A to Host B.

Hence, the TCP throughput with compensation for the processing delay is:

TCP throughput (with compensation for the processing delay)

$$= \frac{\text{Total Data Transferred in the TCP Session}}{\text{TCP session time (with compensation for the processing delay)}} \qquad (2)$$

The advantages of SingAREN's TCP measurement machine can be summarised as follows:

(i) The TCP protocols can be implemented on a PC or SUN workstation which are relatively inexpensive, thereby allowing widespread deployment of such TCP measurement machines.

(ii) TCP throughput measurements are compensated for the processing delay of the end-hosts which implement the TCP protocols.

(iii) Adoption of the latest BTC metric can be performed by installing the operating system that implements the updated TCP protocols onto the general purpose computer.

V. PASSIVE MEASUREMENTS

Besides active measurements which are described in the previous sections, SingAREN has also been carrying out passive measurements of its network traffic in order to understand the demand for bandwidth by the various types of network protocols and applications. Figure 5 shows the amount of traffic (in bytes) that were processed by SingAREN's main router on 27th December 1999.

Protocols/Applications	Amount of traffic Processed (bytes)
Multicast	34.31 G
tcp-others	7.31 G
tcp-http	6.65 G
tcp-proxy	3.46 G
tcp-ftp-data	1.18 G
udp-others	896.01 M
Icmp-others	683.77 M
tcp-nntp	359.84 M
tcp-smtp	338.19 M
tcp-shell	216.55 M
tcp-telnet	121.53 M
udp-domain	47.16 M
tcp-ssh	29.27 M
ipv6-others	4.56 M
udp-sunrpc	2.8 M
tcp-whois	2.41 M
tcp-login	2.34 M
tcp-domain	718.22 K
udp-ntp	116 K
udp-tftp	70.55 K
udp-echo	32.64 K

Figure 5: Classification of traffic according to protocols and applications.

The information shown in Figure 5 is useful for determining the adequacy of SingAREN's network resources in meeting the traffic demands of the users. It also provides useful information during the planning stage of upgrading SingAREN's network capacity. If there is a need to provide QOS guarantees to certain types of traffic, the information in Figure 5 can also be very useful.

Another area in which passive measurements will be useful is in the areas of cost recovery (i.e billing the users of the network). Currently, the users do not have to pay for the usage of SingAREN's network as the Singapore government is bearing the cost of operating the network. This is a disadvantage as the users do not see the incentive of using SingAREN's network in an efficient manner. For example, the users may send traffic into the SingAREN's network out of convenience instead of necessity. The consequence will be unwaranted network loadings.

For reasons of practicality, billing of network usage have to be carried in such a way where the organisations (e.g. research institutes, universities and industrial partners) pay SingAREN for the usage of its network. These organisations may then decide if they want to breakdown the bill according to the usage by their individual departments.

To achieve fairness in the billing system, the amount of money which each organisation has to pay should be directly proportional to the number of bytes of data that the organisation sends and receives from SingAREN's network. This means that if an organisation (that uses SingAREN's network) is having a TCP session with a remote site where the TCP packets traverse an Internet Service Provider (ISP) that is experiencing a high percentage of packet loss, the cost of retransmitting the TCP packets should be borned by the organisation instead of SingAREN.

VI. THE IMPACT OF HAVING SATELLITE LINKS WITHIN THE NETWORK

The biggest expenditure in operating SingAREN's network at the moment goes to the leasing of international links which are provided by submarine cables. Currently, there are very few competitors that provide international links using submarine cables and this may be the reasons for their relatively high charges. To reduce the cost in operating the network, SingAREN would like to use satellite links to carry some of its traffic to overseas since the cost of satellite bandwidth is lower than submarine cable bandwidth. Besides, satellite network can be made up of asymmetrical links which can further reduce cost as Internet traffic in the international links are most likely to be asymmetrical.

As part of our plan to use communication satellite for our international links, SingAREN is currently particpating in the "*ATM Via Satellite*" project together with 1-Net Pte Ltd (which is the Asynhcronous Transfer Mode (ATM) service provider for SingaporeONE) and Temasek Polytechnic's Satellite-Internet Competency Unit (SICU). One of the aims of the project is to build *a Satellite ATM Network* that is interconnected to SingAREN's network. The main components of the *Satellite ATM Network* are a 2.048 Mbps satellite channel, satellite modems, satellite antennas and ATM switching devices.

SingAREN chose ATM as the transport mechanism for carrying traffic over a satellite link for several reasons. The first reason is the emergence of standards for *Satellite ATM Network* as a results of efforts by the International Telecommunication Union (ITU) as well as the ATM Forum. The second reason is the ability of ATM to provide QOS guarantees which is necessary for applications such as video and voice. The third reason is the ease of partitioning bandwidth of a link using ATM's concept of virtual circuits.

The fourth reason is ATM has been adopted as the transport mechansim for Broadband-Integrated Services Digital Network (B-ISDN) which is the technology used by SingAREN as well as telecom providers for their backbone infrastructure. Hence, the use of ATM will provide the ability to interconnect the satellite network to SingAREN's terrestrial network and the public telecommunication network.

The main issues of implementing ATM over a satellite link are [7, 8, 9]:

(i) the impact of satellite errors on the performance of ATM. Due to channel coding, transmission errors introduced by a satellite occur in bursts. When the burst errors corrupt the ATM cell headers, the ATM receiver will have to discard the entire cell resulting in higher cell loss ratio.

(ii) the impact of the inherently long propagation delay (approximately 250 ms) of a geostationary on transport protocols that rely on feedback for congestion control mechanisms and

(iii) the need for rate adaptation as a result of the speed mismatch between the satellite network and the terrestrial network.

Because a satellite network is vastly different from a terrestrial network (fibre links have low bit error rate (BER) and relatively low latency), a separate set of measurement procedures and QOS objectives need to be catered for the satellite network.

Currently, both the ITU and ATM Forum [10] are working on documents that specify the QOS objectives for the ATM over a geostationary satellite link. Figure 6 shows the QOS Class definition and network performance parameters which are found in [10].

	CTD	2-pt CDV	CLR0 +1	CLR	CER	CMR	SECBR
Default Objectives	No	No	No	No	4*10-6	1/day	10-4
QOS Classes							
Stringent Class	400 ms	3 ms	3*10-7	None	Default	default	default
Tolerant Class	U	U	10-5	None	Default	default	Default

CTD - Cell Transfer Delay
CDV - Cell Delay Variation
CLR - Cell Loss Ratio
CER - Cell Error Ratio
CMR - Cell Misinsertion Ratio
SECBR – Severely Errored Cell Block Ratio

Figure 6: Class definition and network performance parameters [10].

It is also recommended that measurement tools such as HP Broadband Series Test Systems (BSTS) should be deployed at the edge of the satellite network to ensure that it meets the QOS objectives shown in Figure 5.

As the *Satellite ATM Network* is meant to be integrated transparently into the terrestrial ATM network, SingAREN can use the same IP QOS metrics and measurement tools to measure the entire network. Because of the long propagation delay of a geostationary satellite link, one has to be aware that a poor TCP throughput performance could be due to the exhaustion of window size (caused by the large bandwidth-delay product of a satellite link) and the combined effects of the flow control algorithms of the TCP protocols and and rather than a lack of bandwidth [11, 12].

VII. CONCLUSIONS

To maintain a high standard of network performance and ensure that its network resources are utilised in an efficient manner, SingAREN will continue to invest in the area of network measurements. The challenge is to achieve our objectives in a cost effective manner.

Because of the diversity in communication technologies employed by SingAREN and the diversity of applications that run on its network, we believe that SingAREN's needs in the area of network measurements can only be met by adopting several different QOS metrics. After several months of studies and evaluation, it has been decided that one-way measurement tool, round-trip measurement tool and SingAREN's own TCP measurement tool will be deployed.

To ensure that the network is not utilised in an ineffcient manner, there may also be a trial to start charging organisations which use SingAREN's network. These measures when put in place will help to ensure that only useful traffic will pass through SingAREN's network and prevent anyone from wasting valuable network resources.

Some of the tools (developed by others) that we have evaluated so far are quite expensive for deployment on a large scale and therefore their use will be restricted to parts of the network that are extremely critical (e.g. SingAREN's international links).

Finally, the employment of satellite links can significantly enhance SingAREN's network capacity at a cost effective manner. However, this new potential also brings with it new challenges as the satellite network needs to be characterised due to its vastly difference from a fibre optical link. This calls for comprehensive testings and experimentations.

Acknowledgment

The authors would like to thank Mr Li Xin at NTRC, NTU for his discussion on the issues of network measurements and Miss Poh Cheng Yee at SingAREN for her help in the the configuration of the network measurement equipment.

References

[1] "About SingAREN", 1998. Available at http://www.singaren.net.sg/html/about_singaren.

[2] Almes, G., Kalidindi S., Zekauskas M., "A One-way Packet Loss Metric for IPPM", RFC 2680, Network Working Group, IETF, Sep 1999. Available at ftp://ftp.isi.edu/in-notes/rfc2680.txt

[3] Almes, G., Kalidindi S., Zekauskas M., "A One-way Delay Metric for IPPM", RFC 2679, Network Working Group, IETF, Sep 1999. Available at ftp://ftp.isi.edu/in-notes/rfc2679.txt

[4] Walt Kelly, "The Network Time Protocol (NTP) Distribution". Available at http://www.eecis.udel.edu/~ntp/ntp_spool/html/index.htm

[5] Newman D, Mandeville R., "Corporate-Class Internet ? Don't Count On It !", Data Comms, Nov 1998. Available at http://www.data.com/issue/981107/isp.html

[6] Mathis M., Allman M., "Empirical Bulk Transfer Capacity", draft-ietf-ippm-btc-framework-02.txt , Network Working Group, IETF, Oct 1999. Avaliable at http://www.ietf.org/internet-drafts/draft-ietf-ippm-btc-framework-02.txt

[7] Chitre D. M., Gokhale D. S., Henderson T., Lunsford J. L., Mathews N., "Asynchronous Transfer Mode (ATM) Operation Via Satellite: Issues, Challenges and Resolutions", International Journal of Satellite Communications, Vol. 12, pp211-222, 1994.

[8] Cuevas E. G., Doshi B., Dravida S., "Performance Models for ATM Applications over 45 Mb/s Satellite Facilities", Tenth International Conference on Digital Satellite Communications, pp228-235, Brighton, 1995.

[9] Akyildiz I. F., Jeong S. H., "Satellite ATM Networks: A Survey", IEEE Comms, Vol. 35, pp30-43, 1997.

[10] Sastri K., Jain R., Deak T., Yegenoghu F., "A Progress Report on the Standards Development for Satellite ATM Networks", ATM Forum/98-0828, Nov 1998.

[11] Allman M., Dawkins S., Glover D., Griner J., Tran D., Henderson T., Heidemann J., Touch J., Kruse H., Ostermann S., Scott K., Semkem J., "Ongoing TCP Research Related to Satellites", RFC2760, Network Working Group, IETF, 2000. Available at http://www.ietf.org/internet-drafts/draft-ietf-tcpsat-res-issues-12.txt

[12] Allman M., Glover D., Sanchez L., "Enhancing TCP Over Satellite Channels using Standard Mechanisms", RFC2488, Network Working Group, IETF, 1999. Available at ftp://ftp.isi.edu/in-notes/rfc2488.txt