Active and Passive Measurements on Campus, Regional and National Network Backbone Paths*

[1] Prasad Calyam, [2] Dima Krymskiy, [2] Mukundan Sridharan, [1] Paul Schopis
[1] OARnet, 1224 Kinnear Road,Columbus, Ohio 43212.
Email:{pcalyam, pschopis}@oar.net
[2] Department of Computer Science and Engineering,
The Ohio State University, Columbus, Ohio 43210.
Email:{krymskiy, sridhara}@cse.ohio-state.edu

***Abstract*— It has become a common practice for Internet Service Providers (ISPs) to instrument their networks with Network Measurement Infrastructures (NMIs). These NMIs support network-wide ”active” and ”passive” measurement data collection and analysis to: 1) identify end-to-end performance bottlenecks in network paths and 2) broadly understand Internet traffic characteristics, on an ongoing basis. In this paper, we present our analysis of the active and passive measurement data collected along network backbone paths within typical campus, regional and national networks which carry traffic of cutting-edge Internet applications such as high-quality voice and video conferencing, multimedia streaming and distributed file sharing. The active measurement data has been obtained by using ”ActiveMon” software, which we have developed and deployed along the above network backbone paths. The passive measurement data has been obtained using SNMP, Syslog and NetFlow data available at the intermediate routers located at strategic points along the same network backbone paths. Our analysis of the measurement data includes studying notable trends, network events and relative performance issues of the network backbone paths which are reflected in the active and passive measurement data collected regularly over several months. Our results thus provide valuable insights regarding traffic dynamics in the different academic network backbones and can be used for better design and control of networks and also to develop traffic source models based on empirical data from real-networks.**

I. INTRODUCTION

Network Measurement Infrastructures (NMIs) [1] [2] [3] [4] in today’s networks support various network-wide ”active” and ”passive” measurement data collection and analysis techniques. Active measurements require injecting test packets into the network. Traditionally, active measurement tools such as Ping and Traceroute were used to determine round-trip delays and network topologies using ICMP packets. Recently, active measurement tools such as H.323 Beacon [5] and Multicast Beacon [6] have been developed that emulate application-specific traffic and use the obtained results of performance of the emulated traffic in the network to estimate the end-user perceived application-quality. Many other active measurement tools such as OWAMP [7], Iperf [8], Pathchar [9] and Pathload [10] that use sophisticated packet probing techniques have also become popular. All these tools are being used in NMIs to routinely monitor network topology, available bandwidth, packet one-way delay, round-trip delay, loss, jitter and Mean Opinion Scores.

In comparison to the active measurements, the passive measurements do not inject test packets into the network. They require capturing of packets and their corresponding timestamps transmitted by applications running on network-attached devices (e.g. switches and routers) over various network paths. Some of the popular passive measurement techniques include collecting Simple Network Management Protocol (SNMP) data, Syslog data and NetFlow data from switches and routers in the network. SNMP data provides switch-level and router-level information such as availability, utilization, packet errors and discards. Syslog data provides details of activities and failures of network switches and routers. NetFlow data provides bandwidth/link utilization information between network backbone routers. This information could be used to determine the flow-level type, duration and amount of application traffic traversing the network.

In this paper, we present our analysis of the active and passive measurement data we have collected over several months on three hierarchically different network backbone paths: campus, regional and national paths. The measurement data presented and analyzed in this paper has been collected from the Third Frontier Network (TFN) Beacon Infrastructure testbed (TBI) [11] we built as part of our TFN Measurement Project. The primary goals in building this testbed are:

- Goal-1: To study end-to-end network performance measurement data reported by various tools to empirically correlate network events and measurement data anomalies in a routine monitoring infrastructure,
- Goal-2: To analyze long-term network performance trends via statistical analysis of active and passive measurement data collected at strategic points on an ongoing basis, and
- Goal-3: To use findings obtained from fulfilling the above Goals 1 and 2, to comprehensively compare performance at campus, regional and national network backbone levels and hence to quantify end-to-end network performance stability in typical academic network backbones.

Towards achieving Goal-1, in [11] we compiled a few case-studies from the measurement data collected over a 2-month period between the sites in the TBI testbed. The case-studies addressed identifying network measurement anomalies in routine ISP operations due to route changes, device mis-configurations and erroneous data from active measurement tools. In this paper, we extend our analysis and present our work towards fulfilling Goals 2 and 3; i.e. we use both active and passive measurement data collected on a regular-basis over a long-term (six month period for active measurements and four month period for passive measurements) between the TBI testbed sites. During this long-term monitoring period, data was collected for various measurements that were initiated several times each day along the network backbone paths. On this collected measurement data, we use statistical methods to comprehensively analyze the trends and relative end-to-end

---

*This work has been supported in part by The Ohio Board of Regents

network performance. The active measurement data analyzed in this paper has been generated using our ”ActiveMon” software [16] which possesses an extensible and customizable framework for generation and analysis of active measurements that can be used for routine network health monitoring. The passive measurement data analyzed is in this paper includes the data collected by OARnet and the Indiana University Abilene Network Operations Center (NOC) [1] at the various network backbone routers present along the same paths for which ActiveMon was used to obtain active measurement data.

The remainder of the paper is organized as follows: Section II presents the measurement tools and methodologies used to collect the active and passive measurement data presented in this paper. Section III describes the campus, regional and national network backbone paths. Section IV presents our illustrations of anomalies due to network events and our statistical analysis of the trends in the active and passive measurement data collected along the network backbone paths. Section V concludes the paper.

II. MEASUREMENT TOOLS AND METHODOLOGY

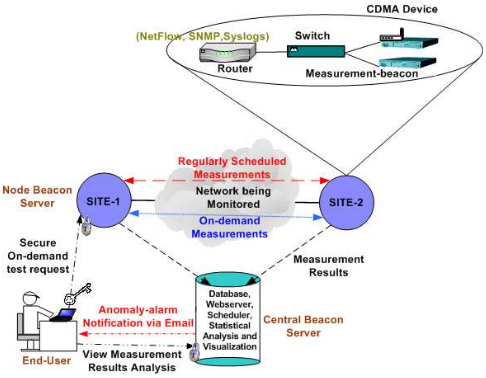

Fig. 1. Active and Passive Measurements in an NMI

TABLE I

ACTIVEMON MEASUREMENT TOOLKIT

Measured Characteristics	**Tool**
Round-trip delay	Ping
High-precision one-way delay	OWAMP
Topology and route changes	Traceroute
Bandwidth capacity: Per-hop	Pathchar
Available bandwidth	Pathload
Bottleneck bandwidth	Pathrate
Transfer bandwidth	Iperf
Performance of voice and video streams	H.323 Beacon

A. *Active Measurements in NMIs*

Fig. 1 illustrates the basic architecture of our ActiveMon software, which uses an active measurement toolkit shown in Table I to perform active measurements between a set of measurement servers (Site Beacon Servers) located at strategic points in the network being monitored. The initiated active measurements across the multiple measurement servers are orchestrated using an efficient measurements scheduler called ”OnTimeMeasure” [17] which regulates the amount of active measurement traffic injected into the network and also prevents conflicts in ongoing active measurements between measurement servers. A central database (Central Beacon Server) is used by ActiveMon to collect and analyze the measurement data collected on a regular and on-demand basis by the measurement servers along multiple network paths. The analysis component of ActiveMon includes a statistical analysis package that processes measurement data for both visualization and alarm reporting functionalities used for alerting appropriate network operations personnel.

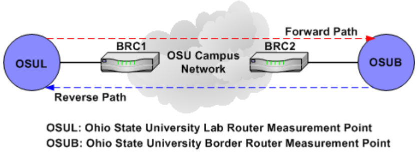

Fig. 2. Campus Network Backbone Path showing end-to-end paths and the edge routers involved in the active and passive measurements

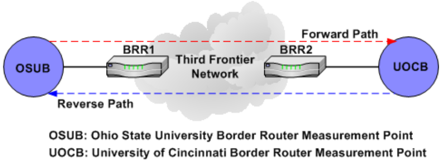

Fig. 3. Regional Network Backbone Path showing end-to-end paths and the edge routers involved in the active and passive measurements

B. *Passive Measurements in NMIs*

As shown in Fig.1, collecting passive measurement data such as Syslogs, NetFlow Records and SNMP Management Information Bases (MIBs) involves regularly polling network devices where the data is collected. Tools shown in Table II are used in the TBI to regularly poll various passive measurement data from critical network devices and the polled data is collected at a central database. This data is then processed for both visualization and alarm reporting functionalities for alerting appropriate network operations personnel, similar to the ActiveMon functionalities.

TABLE II

PASSIVE MEASUREMENT TOOLKIT

Measured Characteristics	**Tool**
Description of traffic flows	NetFlow [12]
Availability	Nagios [13]
Bandwidth Utilization	MRTG [14]
Errors and Discards	Statscout [15]

III. DESCRIPTION OF THE MEASURED NETWORK BACKBONE PATHS

Fig. 2 shows a campus network backbone path between an Ohio State University Lab router (OSUL) and the Ohio State University border router owned by OARnet (OSUB). Given the large geographical area of OSU, the intermediate path characteristics can demonstrate the network performance that can be expected in a typical campus backbone network path. BRC1 and BRC2 shown in Fig. 2 correspond to the edge routers for which we present passive measurement data in Section IV of this paper.

Fig. 3 shows a regional network backbone path between the Ohio State University border router (OSUB) and the University of Cincinnati border router (UOCB), both owned by OARnet. The intermediate path between these measurement points is covered only by the OARnet network backbone routers, i.e. regional backbone network routers. The intermediate path characteristics can thus illustrate the network performance that can be expected in a typical regional backbone network path.

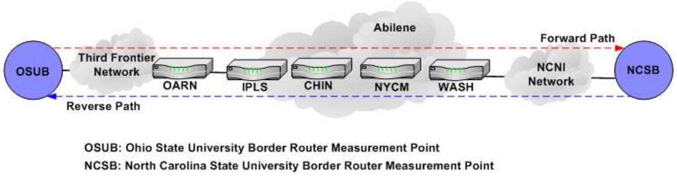

Fig. 4. National Network Backbone Path showing end-to-end paths and the intermediate and edge routers involved in the active and passive measurements

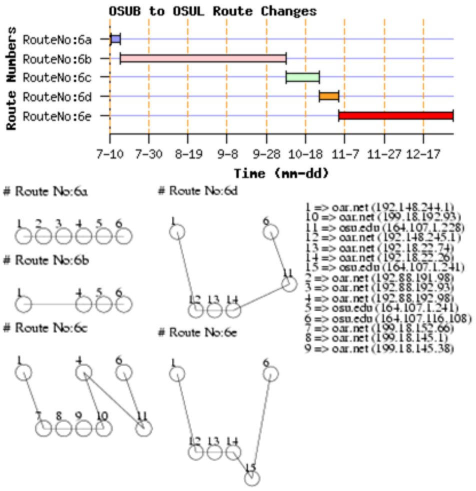

Fig. 5. Route Changes along Campus Network Backbone Path

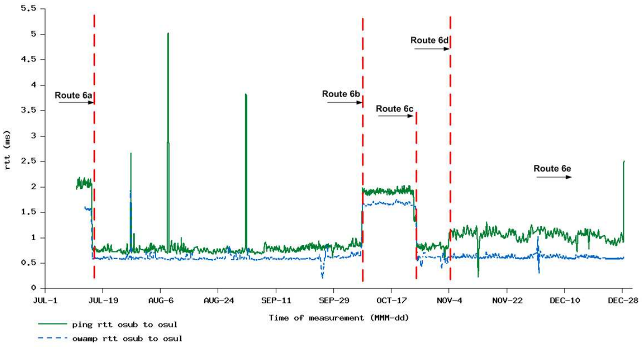

Fig. 6. Delay measurements along Campus Network Backbone Path

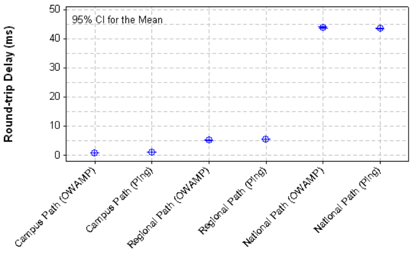

Fig. 7. Mean and 95% CI Plot for Delay measurements

BRR1 and BRR2 shown in Fig. 3 correspond to the edge routers for which we present passive measurement data in Section IV of this paper.

Fig. 4 shows a national network backbone path between the Ohio State University border router (OSUB) owned by OARnet and the North Carolina State University border router (NCSB) owned by NCNI. The intermediate path between these measurement points is covered only by OARnet, Abilene and NCNI network backbone routers, i.e. two regional backbone networks connected via a national backbone network. The intermediate path characteristics can thus illustrate the network performance that can be expected in a national backbone network path. OARN router shown in Fig. 4 is the edge-router located at the peering point of TFN and Abilene. The routers- IPLS, CHIN, NYCM and WASH shown in Fig. 4 correspond to the intermediate routers for which we present passive measurement data in Section IV of this paper.

IV. MEASUREMENT DATA AND ANALYSIS

A. *Description and Analysis of Active Measurement Data*

In this section, we describe our active measurement data collected along the campus, regional and national network backbone paths. The active measurement data presentation focuses on a period of approximately six months (between July 2004 and Dec 2004).

1) Route Changes: The most common anomalies observed while monitoring network paths are those caused by route changes. Route changes are attributed to ”route flaps” caused by suboptimal routing protocol behavior, network infrastructure failures, re-configuration of networks or load-balancing strategies used by ISPs to improve network performance.

Fig. 5 shows our visualization method used in ActiveMon to represent route changes indicated by Traceroute along the campus network backbone path. We can observe from Fig. 5 that in the six month period, route changes have occurred four times. We also observed that during the same time period, only two route changes occurred in the regional network backbone path and none occured in the national network backbone path. The reason for the high frequency of route changes in the campus network backbone path was due to various network management activities and also due to the different phases involved in transitioning of the campus traffic from an old ATM network to our new Third Frontier Network. The route numbers and hop numbers shown in Fig. 5 correspond to distinct route signatures and host IP addresses stored in our TBI’s central database tables.

2) Delay : Delay is the time taken for a packet to traverse from a sender end-point to a receiver end-point. Commonly ”round-trip delay” is used to characterize network delay. Fig. 6 shows the round-trip delay measurements provided by Ping and OWAMP during the six month monitoring period for the campus network backbone path. For OWAMP which measures high-precision one-way delay for a path, we calculate the round-trip delay using the sum of the bidirectional one-way delay measurements along the path. We can notice that the variations in trends of the Ping and OWAMP measurements respond in a similar fashion and the variations directly correspond to the route changes along the same path as shown in Fig. 5. We can also observe sudden short-lived dips and peaks in the measurements due to miscellaneous temporal network dynamics in the network path.

Fig. 7 compares the delay measurements along the campus, regional and national network backbone paths using an ”interval plot”. The comparison shows the mean and the 95th percentile confidence interval for the six month data sets. The confidence interval provides the range (lower and upper limits) in which the delay values fall for the given set of predictor values. As expected, the delay ranges are least in the campus path with least number of intermediate hops and the corresponding values for the national path are largest due to the relatively higher number of intermediate hops.

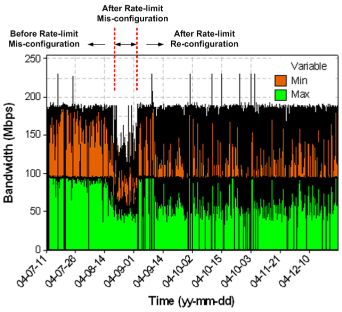

Fig. 8. Bandwidth measurements along Regional Network Backbone Path

3) Bandwidth: Bandwidth measurements reflect the congestion levels caused by the network dynamics and the link provisioning in the paths. Fig. 8 shows the area plot of the available bandwidth measurements provided by Pathload during the six month monitoring period for the regional network backbone path. In particular, we can notice three distinct trends. The trends as indicated in the Fig. 8 are noticeable before, during and after a measurement anomaly in the Pathload measurements that involved an actual rate-limit mis-configuration and a later re-configuration that corrected the problem.

Fig. 9 compares the available minimum and maximum available bandwidth measurements along the campus, regional and national network backbone paths using an interval plot. We can observe that the regional path has the least congested and the most provisioned path. The national path, whose traffic traverses through multiple ISP domains, is the most congested and the least provisioned path owing to the higher possibility of end-to-end congestion at any given time instant. The above bandwidth characteristics of the paths significantly influence the other end-to-end performance metrics observed, as shown in the following subsections.

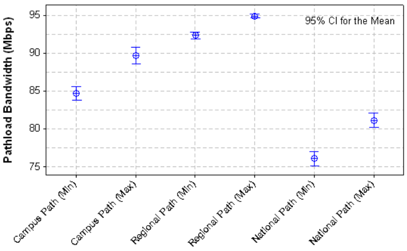

Fig. 9. Mean and 95% CI Plot for Bandwidth measurements

4) Jitter: Jitter, which represents the variations in the network delay, is measured using a moving average computation technique for a given stream of UDP packets, as described in RFC 1889. Fig. 10 shows the jitter measurements provided by Iperf during the six month monitoring period for the campus network backbone path. We can notice that the variations in trends of the jitter measurements occured only for the first route change along the path as shown in Fig. 5. The other route changes did not affect the jitter along the path noticeably. We can also notice that on December 15th, a random network event, which does not correspond to a route change, caused an increased variance in the jitter measurements.

Fig. 11 compares the jitter measurements along the campus, regional and national network backbone paths using an interval plot. We can observe that the regional path jitter measurements, due to the greater available bandwidth as shown in Fig. 9, are much lower in comparison and also exhibit relatively the least spread in the measured values. Similarly as expected due to the path bandwidth limitations, we can observe that both regional and national paths exhibit significant spread in the jitter values; the national path performing the worst in terms of both the mean and spread of the jitter values.

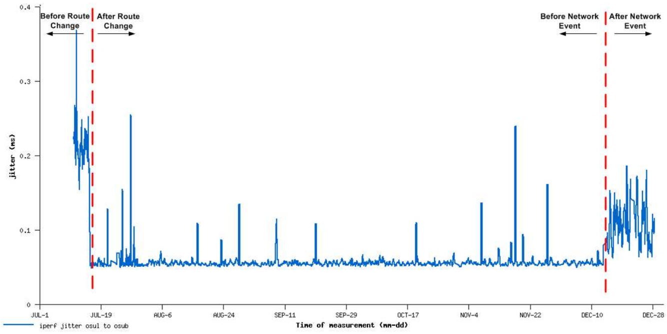

Fig. 10. Jitter measurements along Campus Network Backbone Path

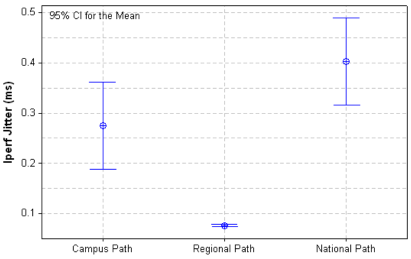

Fig. 11. Mean and 95% CI Plot for Jitter measurements

5) Loss: Loss indicates the percentage of packets lost as observed at the receiver end-point for a given number of packets transmitted at the sender end-point. To compare the loss measurements of UDP packets along the campus, regional and national network backbone paths, we use an interval plot for the loss values reported by Iperf during the six month monitoring period as shown in Fig. 12. We can observe that the regional path experiences almost negligible loss in packets. Interestingly, in comparison to the jitter performance, the loss performance of the national path is seen to be better than the campus path, in terms of both the mean and spread of the loss values. A reasonable explanation for this loss phenomenon seen on the campus path could be due to the last-mile bandwidth sharing limitations amongst departmental networks within university campuses.

6) Mean Opinion Score (MOS): MOS measurements reported by the H.323 Beacon are useful in evaluating network capability to support Voice and Video over IP (VVoIP) applications. The MOS values are reported on a quality scale of 1 to 5; 1-3 range being poor, 3-4 range being acceptable and 4-5 range being good. Additional details pertaining to the

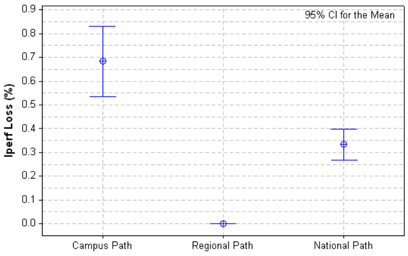

Fig. 12. Mean and 95% CI Plot for Loss measurements

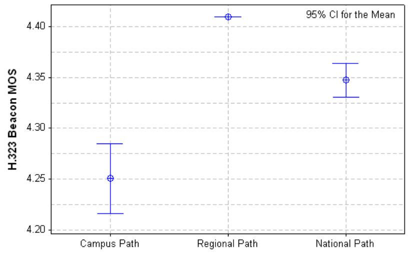

Fig. 13. Mean and 95% CI Plot for MOS measurements

H.323 Beacon MOS measurements can be found in [5]. Fig. 13 compares the MOS measurements along the campus, regional and national network backbone paths using an interval plot. Since MOS is a function of the performance of delay, jitter and loss values, the performance of these metrics as described in the previous subsections, significantly influence the MOS measurements. Due to the higher impact of increased loss in the campus path on the MOS values, we can notice that the campus path performs the worst in comparison to the national path in terms of both the mean and spread of the MOS values. As expected, the regional path MOS measurements perform the best with neglible spread.

7) Stability: Finally, in this subsection, we use the statistical coefficient of variation (ρ) as a measure to evaluate the overall stability metric of the network backbone paths for the active measurements. is calculated as a percentage-

$$\rho = \frac{S}{\bar{X}} * 100$$

where -

$$S = \sqrt{\frac{\sum(x_i - \bar{X})^2}{X}}; \bar{X} = \sum \frac{x_i}{N}$$

Here: x_i is the i^{th} observation, N is the number of non-missing observations, $\bar{X}$ is the Mean and S is the standard deviation.

The ρ values of the active measurements for the campus, regional and national network backbone paths are shown in Table III. Since better stability is implied by a lower value of ρ, we can conclude that regional paths have the overall best network performance stability and cause the least degradation of the end-to-end network performance that affects end-user applications. Whereas, both campus and national paths show higher values of ρ and hence are more probable to cause end-to-end performance bottlenecks, more so in the case of campus paths.

TABLE III
STABILITY ANALYSIS USING CO-EFFICIENT OF VARIATION

Tool Characteristic	Campus	Regional	National
Pathrate Max. Bandwidth (Mbps)	24.23	5.5	6.74
Iperf Jitter (ms)	745.95	45.1	499.89
Iperf Loss (%)	517.63	62.48	127.43
H.323 Beacon MOS	18.48	0.03	9.63
OWAMP Delay (ms)	52.87	10.64	13.47
Ping Delay (ms)	58.52	5.65	24.4

B. Description and Analysis of Passive Measurement Data

In general, it is not common to find notable correlations between active measurements and passive measurements. In [18], a detailed correlation study of loss measurements obtained by active and passive measurement techniques is presented and the study concludes that there is a low degree of agreement between the measures. However, passive measurement data provide a good context to understand the conditions under which the active measurement data was collected. Also, passive measurements provide a different perspective in evaluating the end-to-end performance of a network path.

In this section, we describe our passive measurement data collected along the same campus, regional and national network backbone paths along which the active measurements were collected. The passive measurement data presentation focuses on a period of approximately four months (between July 2004 and October 2004).

1) Availability: Availability is calculated by measuring the uptime or downtime of a network device or service using passive measurements. Scheduled outages (e.g. network devices or services are shutdown for maintenance purposes) are not considered while calculating availability. We collected and evaluated the availability metric as given by the Nagios tool for the routers BRC1, BRC2, BRR1, BRR2, OARN, IPLS, CHIN, NYCM and WASH described in Section 3. We found that all these routers recorded 100% availability. On inspection of the Syslog data of these routers, we further determined that none of the network events observed in our measurement datasets were caused by unexpected router hardware or software failures.

2) Discards and Errors: Discards is a SNMP metric that indicates the number of packet discarded for a particular network interface. Similarly, Errors is also a SNMP metric that indicates the number of interface errors (e.g., Frame Check Sequence (FCS) errors). Large values of discards and errors are an indication of excessive network congestion at any given point of time.

Table IV shows the discards and errors data as reported by the Statscout tool for the BRC1, BRC2, BRR1 and BRR2 routers, i.e., for the routers in the campus and regional network backbone paths. We can observe that in general, occurrence of errors and discards are not common on network interfaces of backbone routers and their values are normally very low value or close to zero. However, we did record a small amount of errors and discards in the case of the BRR1 router. Further analysis made us realize that these errors and discards were caused in an old ATM network backbone router. After a network upgrade was performed in July, there were no more occurrences of errors and discards.

TABLE IV
SNMP DATA ALONG CAMPUS AND REGIONAL PATH LINKS

Router	In(%)	Out(%)	Errors(Mb)	Discards(Kb)
BRC1	1.4971	2.989	0	0
BRC2	2.124	1.77	0	0
BRR1	6.954	1.806	4.341	63.94
BRR2	12.451	12.451	0	0

3) Utilization: Utilization is a SNMP metric that compares the amount of inbound and outbound traffic versus the bandwidth provisioned on a link in a network path. Table IV shows the inbound and outbound utilization values for the BRR1 and BRR2 router links in the regional network backbone path. Fig. 14 similarly indicates the inbound and outbound utilization values for the IPLS and CHIN link along the national network backbone path.

We can note that the bandwidth utilization averaged over the four months is quite low in all of the above links. Similar observations were recorded for the utilization of the IPLS-CHIN, CHIN-NYCM and NYCM-WASH links. Since Abilene

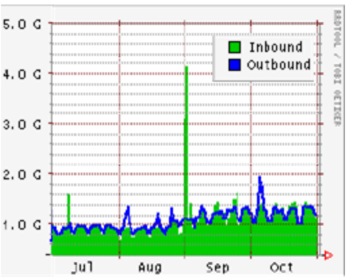

Fig. 14. Utilization on the link between IPLS and CHIN

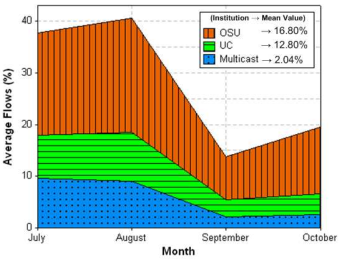

Fig. 15. NetFlow data grouped by Institution at OARN Router

has a capacity of 10Gbps, observing the utilization values expressed in Gbps, we can conclude that the link utilization values have a range that is less than 10%, similar to the cases of BRC1, BRC2, BRR1 and BRR2 shown in Table IV.

4) Flow Information: Fig. 15 shows the NetFlow information for OSU and UC schools generating the traffic flows processed by the OARN router. We selected OSU and UC since they are in the path of our regional and campus network test paths. The Multicast flows in Fig. 15 include the flows that might have originated from OSU or UC to a multicast address. We can observe a dip in the average amount of traffic flows during the months of August and September which correspond to the summer break at these schools. Upon reopening of the schools in the following month, we can observe that this dip is followed by a rise in the average amount of traffic flows. We can also observe from Fig. 15 that the combined traffic from UC and OSU which are two of the largest State Universities in Ohio that have access to Abilene alone constitutes to about 30% of the flows entering Abilene from TFN.

We also analyzed the NetFlow information grouped by the protocol in the flows that have been processed by the OARN, IPLS, CHIN, NYCM and WASH routers during the four month monitoring period. As an example, we show our analysis for the WASH router in Fig. 16. In all the cases, over 80% of the flows corresponded to TCP flows, which include web traffic and other file sharing/transfer traffic. About 10-15% of the flows corresponded to UDP flows which include VVoIP traffic and other streaming media traffic. About 1-3% of the flows constituted ICMP traffic flows, which mainly relate to the network connectivity testing traffic generated by Pings and Traceroutes. We observed very insignificant amount of flows that corresponded to IPv6 traffic over TFN and Abilene, even though both TFN and Abilene use native IPv6 on all the backbone routers. Our above observations regarding the amount of various traffic flows are comparable to observations on commercial network backbone paths as well [19].

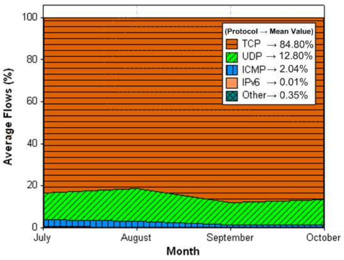

Fig. 16. NetFlow data grouped by Protocol at WASH Router

V. CONCLUSION

In this paper, we presented our analysis of active and passive measurement data collected in a testbed (TBI) over several months on three hierarchically different network backbone paths: campus, regional and national paths. We collected the active measurements using a NMI software called ”ActiveMon” which we have developed and the passive measurements were collected from strategic routers located along the same paths for which we collected the active measurements.

We performed a detailed statistical analysis of the active measurements and described trends and notable anomalies observed in the data due to network events during a long-term monitoring period. We also evaluated the relative performance along the backbone paths and demonstrated that regional network backbone path performed the best and the campus network backbone path performed the worst. Our evaluation involved comparing the mean, spread and stability characteristics of the various active measurements of: delay, bandwidth jitter, loss and MOS. Our analysis of the passive measurement data involved studying various SNMP metrics and analysis of Syslog data followed by a flow-level analysis of the traffic for the selected routers during the monitoring period. In addition to showing notable trends in the passive measurement data, we showed that in academic networks, the utilization levels are lower than 10% as reported by SNMP data. We also showed that these networks generally have 100% availability with negligible or low amounts of router interface-level discards and errors. Finally, using Netflow data, we showed that various kinds of traffic flows in these networks have a predominant amount of TCP traffic and an insignificant amount of IPv6 traffic.

REFERENCES

[1] Abilene Observatory - http://abilene.internet2.edu
[2] SLAC IEPM Project - http://www-iepm.slac.stanford.edu
[3] NLANR AMP Project - http://watt.nlanr.net
[4] NLANR PMA Project - http://pma.nlanr.net
[5] P. Calyam, W. Mandrawa, M. Sridharan, A. Khan, P. Schopis, ”H.323 Beacon: An H.323 application related end-to-end performance troubleshooting tool”, ACM SIGCOMM NetTs, 2004.
[6] NLANR Multicast Beacon - http://dast.nlanr.net/Projects/Beacon
[7] S. Shalunov, B. Teittelbaum, ”One-way Active Measurement Protocol (OWAMP)”, IETF RFC 3763, 2004
[8] NLANR Iperf - http://dast.nlanr.net/Projects/Iperf
[9] Pathchar - http://www.caida.org/tools/utilities/others/pathchar
[10] C. Dovrolis, P. Ramanathan, D. Morre, ”Packet Dispersion Techniques and Capacity Estimation”, IEEE/ACM Transactions on Networking, 2004.
[11] P. Calyam, D. Krymskiy, M. Sridharan, P. Schopis, ”TBI: End-to-end network performance measurement testbed for empirical bottleneck detection”, IEEE TRIDENTCOM, 2005.
[12] M. Fulmer, S. Romig, and R. Luman,”The OSU Flow-tools Package and Cisco NetFlow Logs”, Lisa 2000
[13] Nagios - http://www.nagios.org
[14] MRTG - http://people.ee.ethz.ch/ oetiker/webtools/mrtg
[15] Statscout - http://www.statscout.com
[16] ActiveMon - http://www.itecohio.org/activemon
[17] P. Calyam, C.G. Lee, P.K. Arava, D. Krymskiy, D. Lee, ”OnTimeMeasure: A scalable framework for scheduling active measurements”, IEEE E2EMON, 2005.
[18] P. Barford, J. Sommers, ”Comparing Probe and Router-based Methods for Measuring Packet Loss”, IEEE Internet Computing - Special issue on Measuring the Internet, 2004.
[19] C. Fraleigh, S. Moon, et. al, ”Packet-Level Traffic Measurements from the Sprint IP Backbone”, IEEE Network, 2003.