Analysis of Bursty Packet Loss Characteristics on Underutilized Links Using SNMP

Seung-Hwa Chung[1], Deepali Agrawal[1], Myung-Sup Kim[1], James W. Hong[1], and Kihong Park[2]
[1]DPNM Lab., Dept. of Computer Science and Engineering, POSTECH, Korea
[2] Dept. of Computer Sciences, Purdue University, USA
{mannam, deepali, mount, jwkhong}@postech.ac.kr, park@cs.purdue.edu

Abstract – ISPs typically provide sufficient bandwidth according to increasing traffic requirements. However, in the case of QoS-sensitive applications such as VoIP, service quality may not be up to expectations even in underutilized links because of sporadic but non-negligible losses due to traffic bursts. This study aims to detect and analyze packet loss characteristics on underutilized links in an enterprise IP intranet environment. We collected packet loss data from various routers and switches deployed on POSTECH's campus network. To obtain the packet loss information, we fetched data from private and standard SNMP MIB variables of the monitored routers and switches. We analyzed the data and identified parts that are representative of packet loss across three time scales: 5min, 5sec and 1sec. Although preliminary, our multi-resolution analysis shows that links that appear underutilized at coarse time granularity, often the case in production network monitoring, reveal burstiness and resulting losses at fine time granularity relevant for VoIP and other QoS-sensitive traffic.

Keywords – Multi-resolution packet loss detection and analysis, bursty traffic, QoS, SNMP

I. INTRODUCTION

Today, the number of Internet users is continually increasing, along with the number of networked applications. These newly emerging network-based applications including VoIP, teleconferencing, streaming media, peer-to-peer, and games generate a significant amount of traffic. This trend is accelerating because of broadband connections and improved PC performance. Increased traffic load underlies many performance and security related problems in both Internet and enterprise networks.

Most real networks, including backbone, enterprise, and some access networks, employ overprovisioning to mitigate performance problems such as packet loss, delay, and jitter for QoS-sensitive applications. To protect against quality degradation at bottleneck segments in large, heterogeneous IP internets, priority scheduling is used to shield QoS-sensitive traffic from best-effort traffic. For example, Cisco routers use LLQ (low latency queue) in IOS as a basic building block to support VoIP. Unlike telephony, however, where uniform standards enable the allocation of end-to-end bandwidth across multiple providers and equipment vendors via TDM channels, IP internetworks crucially rely on overprovisioning to facilitate end-to-end QoS.

An important component of traffic management for resource provisioning and network planning is traffic monitoring. Unfortunately, traffic monitoring systems such as NetFlow [1], MRTG [2] and NG-MON [3] do not detect packet losses. They monitor only the traffic that passes successfully through the router. Moreover, these systems cannot detect bursty traffic at the second time scale; their coarse time granularity is limited to 5-minute and 1-minute aggregates in the case of MRTG and NG-MON, respectively. Due to averaging over large time intervals we may see links being underutilized (e.g., 20-30%) when, in fact, at the second time granularity, critical for assuring toll quality VoIP, traffic spikes and packet losses are present leading to unacceptable service violations. This is especially relevant given the self-similar nature of Internet traffic [4].

We summarize studies related to packet loss relevant to the present study. Papagiannaki et al. [5] presented a characterization of congestion in the Sprint IP backbone network. They analyzed link utilization at various time scales to measure the frequency and duration of micro congestion. While they detected traffic bursts and their duration, they did not measure the actual packet loss that occurred during these times. This work did not provide byte loss or packet loss counts. Hence, further work is needed on this topic.

Hall et al. [6] analyze the effect of early packet loss on web traffic and its download time. They discovered that the TCP SYN packet loss causes higher latency in web page download than other packet losses. This work concentrates on web traffic. Mochalski et al. [7] studied changes in traffic pattern relative to different points of observation in the network and investigated the contributing factors to the changes observed. They measured the delay across the router and firewall and tried to relate the delay to packet loss, concentrating on analyzing packet loss using delay. We try to analyze packet loss using a range of parameters, such as the number of incoming and outgoing packets, bandwidth, and router performance.

This work was supported in part by the Electrical and Computer Engineering Division at POSTECH under the BK21 program of the Ministry of Education, Korea.

Object	OID	Description
ifInUcastPkts	1.3.6.1.2.1.2.2.1.11	The number of subnetwork-unicast packet delivered to a higher-layer protocol
ifOutUcastPkts	1.3.6.1.2.1.2.2.1.17	The total number of packets that higher-level protocols requested be transmitted to a subnetwork-unicast address, including those that were discarded or not sent
ifInOctets	1.3.6.1.2.1.2.2.1.10	The total number of octets receive on the interface, including framing characters
ifOutOctets	1.3.6.1.2.1.2.2.1.16	The total number of octets transmitted out of the interface, including framing characters

Table 1. SNMP MIB II variables used in measurements.

Object	OID	Description
locIfInputQueueDrops	1.3.6.1.4.1.9.2.2.1.1.26	The number of packets dropped because the input queue was full
locIfOutputQueueDrops	1.3.6.1.4.1.9.2.2.1.1.27	The number of packets dropped because the output queue was full
locIfInIgnored	1.3.6.1.4.1.9.2.2.1.1.15	Provides the number of input packets that were ignored because the internal buffers were full
cpuLoad	1.3.6.1.4.1.9.2.1.56	CPU Utilization (5 sec avg.)

Table 2. Cisco Enterprise MIB variables used in measurement.

We monitored key routers and switches from POSTECH’s campus network to collect traffic data every 5 minutes, 5 seconds and 1 second using SNMP [8]. We analyzed the data to detect the bursty traffic occurrences and packet loss characteristics stemming from the bursts.

The organization of this paper is as follows. Our packet loss detection method is described in Section II. Section III describes the traffic data collection. In Section IV, we give an analysis of traffic bursts and packet loss. Finally, concluding remarks are given and possible future work is discussed in Section V.

II. Packet Loss Detection Method

SNMP agent is used to obtain traffic and loss data from the routers and switches on POSTECH’s campus network, most of which support SNMP. From the supported SNMP MIB II, we selected and fetched data for four MIB variables, ifInUcastPkts, ifOutUcastPkts, ifInOctets, and ifOutOctets. The selected SNMP MIB II variables are described in Table 1.

Next, we tried to detect packet loss by comparing the incoming and outgoing packet counters. However, the loss information obtained is inaccurate for the following reasons.

- Some packets are destined to the router. Therefore, there is no outgoing packet for such packets.
- Some packets are generated by the router. Therefore, there is no incoming packet for such packets.
- Some packets are broadcasted by the router. This causes a large difference between the incoming and outgoing packet counters and confuses the loss count.

The above problems cannot be avoided by using SNMP MIB II variables. Hence, we decided to use Cisco enterprise MIB variables [9] that provide packet loss and router performance related information. The four enterprise MIB variables used are: locIfInputQueueDrops, locIfOutputQueueDrops, locIfInIgnored and cpuLoad. The selected enterprise MIB variables are described in Table 2.

The formula for packet loss calculated from Cisco enterprise MIB variables, at a given time granularity, is as follows:

$$\textit{Packet Loss} = \text{locIfInputQueueDrops} + \text{locIfOutputQueueDrops} + \text{locIfInIgnored}$$

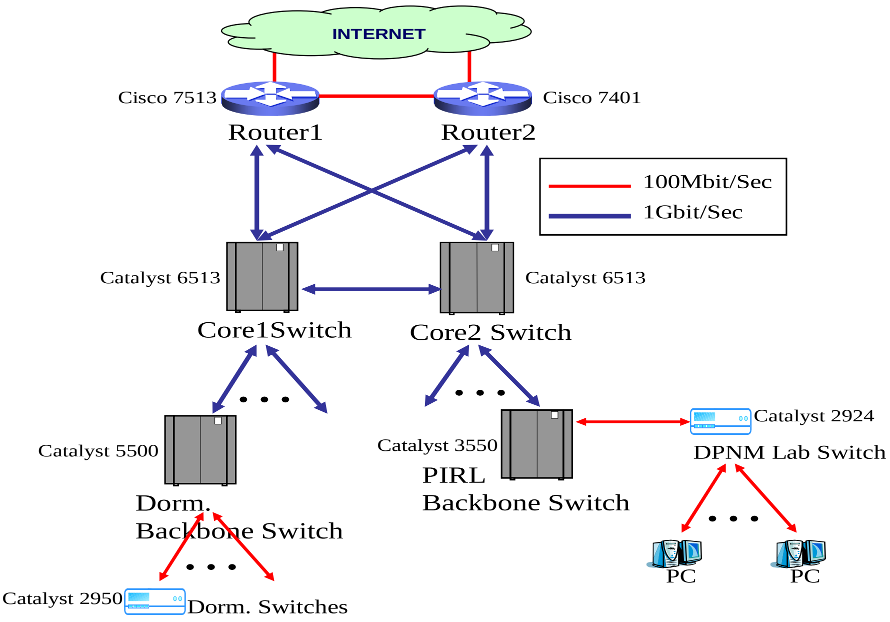

Figure 1. POSTECH Internet and Intranet Infrastructure.

Time in Second	In Packets	Out Packets	Lost Packets	In Bytes	Out Bytes	Lost Bytes	CPU %
3:23:20 pm	0	0	0	64	0	64	61
3:23:21 pm	0	0	0	0	0	0	61
3:23:22 pm	400	522	0	79669	454646	0	62
3:23:23 pm	0	0	0	64	609046	0	62
3:23:24 pm	437	611	0	67670	0	67670	62
3:23:25 pm	0	0	0	64	0	64	62

Table 3. Data from DPNM Lab switch

III. TRAFFIC DATA COLLECTION

SNMP agents are running in various network devices deployed in the campus network. POSTECH's campus network is comprised of a gigabit Ethernet backbone, which, in turn, is composed of two Cisco IP routers, two core backbone switches, dozens of gigabit building backbone switches, and hundreds of 100Mbps switches and hubs that are deployed inside the buildings.

We decided to monitor two Cisco IP routers (Cisco 7513 and 7401), two core switches (Catalyst 6513), one dormitory backbone switch (Catalyst 5500) that connects a number of dormitory switches, and a DPNM laboratory switch (Catalyst 2924). The routers and switches exhibit varying utilization and traffic across different ports. We selected links from different routers and switches to represent low (less than 20%), moderate and high (more than 70%) link utilization. From the Core1 switch we selected a 1Gbps link that is connected to Internet gateway router1 which is lightly utilized. From border router 2 we chose a 100Mbps egress link that is heavily utilized. From the dormitory switch we selected a 1Gbps link that is connected to another backbone switch and is underutilized. Both Cisco routers that in the Internet access path are connected to Core 1 and 2 switches. The Core 2 switch is connected to a backbone switch that is connected to the DPNM laboratory switch, and the Core 1 switch is connected to a backbone switch that is connected to the dormitory switch. The physical topology of the POSTECH network used in the traffic measurement is shown in Figure 1.

Time Interval	Min.	Max.	Mean	Standard Dev.	Standard Dev. divided by Avg.
5 minute	1452	18628	5598.1	2155.8	0.39
5 second	24	51256	5607.9	6378.6	1.14
1 second	0	256282	5608.1	18135.9	3.23

Table 4. Incoming packet statistics for Core1 switch 1Gbps link.

Time Interval	Min.	Max.	Mean	Standard Dev.	Standard Dev. divided by Avg.
5 minute	170460	8087676	2986637.8	1426388.7	0.48
5 second	2841	18475514	2990452.1	3626093.3	1.21
1 second	0	92377571	2990452.3	10072900.6	3.37

Table 5. Incoming bytes statistics for Core1 switch 1Gbps link

The columns of the table show MIB polling time, incoming packets, outgoing packets, number of packet loss, incoming bytes, outgoing bytes, difference between incoming and outgoing bytes (if outgoing bytes are more than incoming bytes then it writes zero), and CPU utilization, respectively. The data indicate the presence of events when the number of incoming packets was zero for a period of 1 second, which we know is not true because the DPNM Lab switch continually received packets during the specified time interval. The discrepancy is caused by MIB counters not being updated immediately after a packet arrives. A router/switch's main priority is forwarding packets, which can cause counter updates to be delayed, especially when traffic load is heavy. To find the minimum time interval that can be used to reliably poll the selected devices using SNMP, we increased the polling interval from 1-second to 2-second and so on. We found that with the polling interval of 5-second, the data obtained is sufficiently accurate.

IV. Analysis of IP Traffic and Packet Loss

This section presents an analysis of IP traffic and packet loss. We collected data for three days, 2004.5.27 3:22 pm to 2004.5.30 9:32 pm, from the specified routers and switches. We collected data at 1-second granularity and aggregated to yield 5-second and 5-minute data. We focused on the parts that clearly illustrated packet loss.

A. Detection of Bursty Traffic in Small Time Scale

Figures 2 and 3 illustrate the distribution of incoming packets and incoming bytes over the 3-day period at 1-second, 5-second and 5-minute measurement granularity. This data is collected from the Core1 switch 1 Gbps link, which is connected to Internet router1 and is lightly utilized. Overall average utilization of the link is 2.4 percent. We can observe that the curve of incoming packets and incoming bytes with 5-minute average values is fairly smooth and underutilized.

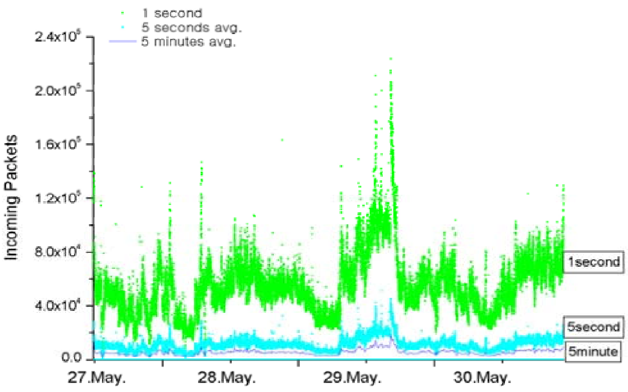

Figure 2. Incoming packet distribution.

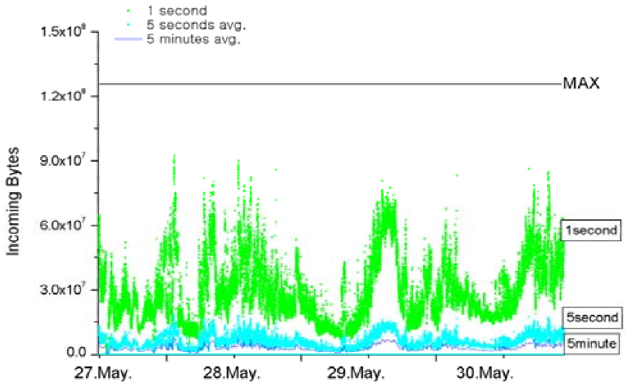

Figure 3. Incoming byte distribution.

The curves with 5-second average values clearly show bursty traffic and wider distribution than 5-minute measurement curve. Whereas, the curves with 1-second measurement values show even more bursty traffic with large peaks and wider distribution than 5-second measurement curve due to these traffic bursts. These graphs indicate that burstiness in traffic exist in the underutilized link at the small time scale of 1 and 5 seconds.

Table 4 and 5 respectively show the statistics for incoming packets and bytes for Core 1 switch 1 Gbps link. From the tables we can observe that the mean value of incoming packets and bytes is similar for all time (5-minute, 5-second and 1-second) scales. On the other hand, standard deviation shows that the distribution of packets and bytes gets wider as the time granularity shortens.

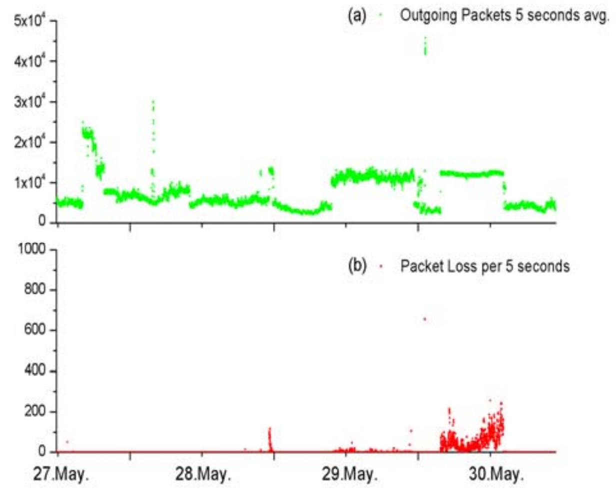

Figure 4. Outgoing Packets & Loss Distribution

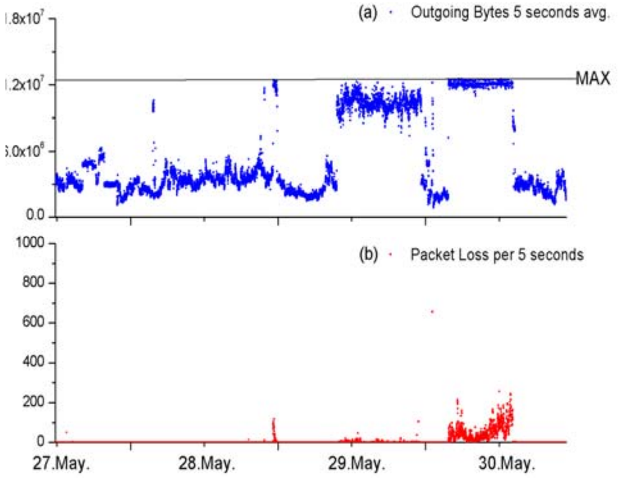

Figure 5. Outgoing Bytes & Loss Distribution

B. Packet Loss in Highly Utilized Links

Figures 4 (a) and (b) illustrate the distribution of outgoing packets 5-second average values and packet loss per 5-second over the period of 3-day. Figures 5 (a) and (b) illustrate the distribution of outgoing bytes 5-second average values and packet loss per 5 second over the period of 3 days. This data is collected from Internet router 2 100Mbps Internet link.

The curve of outgoing bytes, Figure 5 (a), shows bursty traffic with some large peaks reaching the maximum link utilization point. These bursts cause the link to be highly (77%) utilized over the period of time from 2004.5.29 9am to 2004.5.30 1pm that cause many packet losses in the link. These high outgoing packets are caused by the two 1 Gbps links of the same router that sent high traffic outside our campus network through this 100Mbps link. We can observe the cluster of points showing many lost packets corresponding to the high peak of outgoing bytes and outgoing packets. From analysis of these plots we can conclude that when the traffic burst is so high that it reaches the maximum link utilization point, packet losses are very high. When the traffic burst is a little smaller so that it reaches close to the maximum link utilization point but remains below it, the losses decrease dramatically. Thus the 5-second time scale can provide useful information on the packet loss process from the aggregate traffic process.

C. Packet Loss in Underutilized Links

Figures 6 (a) and (b) illustrate distribution of incoming packets and outgoing packets for 5-second average values over 3-day time interval. Figures 7 (a) and (b) illustrate distribution of incoming bytes and outgoing bytes for 5-second average values over 3-day time interval. This data is obtained by measurement at Core 1 switch 1 Gbps link that is connected to Internet router 1 and is lightly utilized. Overall average link utilization is 2.4 percent. Figures 6 (a) and (b) show that the number of incoming and outgoing packets is not high and Figures 7 (a) and (b) show that the link is well underutilized for this period of time. No high traffic bursts that reach maximum link utilization can be seen even with 5-second measurement values.

Figure 8 illustrates distribution of packet loss over time. We can observe that packet loss exist even when the link is well underutilized. Because this link is very lightly utilized the packet losses are few. Hence, we present below the zoomed in plots for another link that is 5.5% utilized and has more losses than this link.

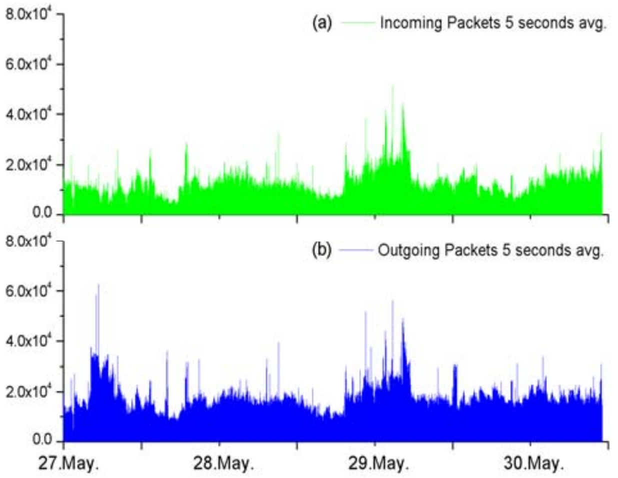

Figure 6. Incoming & Outgoing Packet Distribution

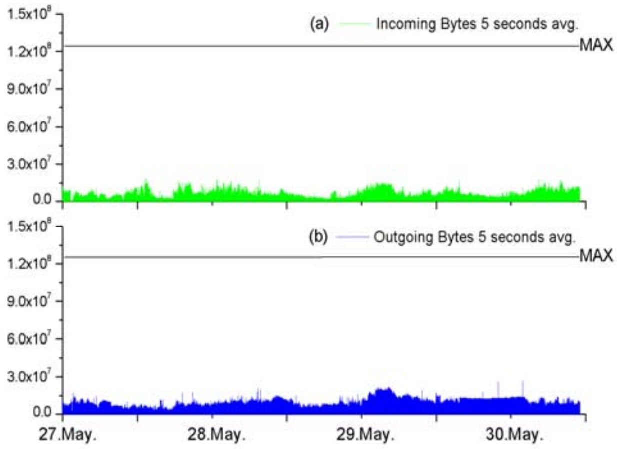

Figure 7. Incoming & Outgoing Byte Distribution

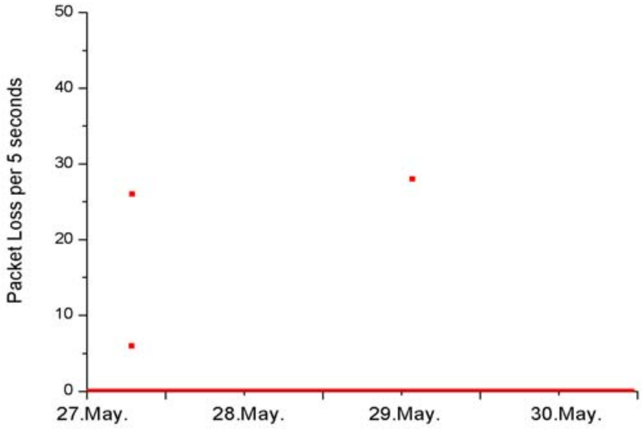

Figure 8. Packet Loss Distribution

Figures 9 (a) and (b) illustrate distribution of incoming packets and outgoing packets for 5-second average values over 12 hours time interval. Figures 10 (a) and (b) illustrate distribution of incoming bytes and outgoing bytes for 5-second average values over 12 hours time interval. This data is obtained by measurement at dormitory backbone switch 1 Gbps link that is connected to another backbone switch and is lightly utilized. Overall average link utilization is 5.5 percent. We presented the zoomed in graphs for 12 hours to show the packet losses more clearly.

Figure 11 illustrates distribution of packet loss over time. We can observe that some peaks of the packet loss exist even when the link is well underutilized. We believe that these losses occurred due to the traffic bursts that occurred in the time granularity that is smaller than 5-second and hence could not be seen with 5-second measurements.

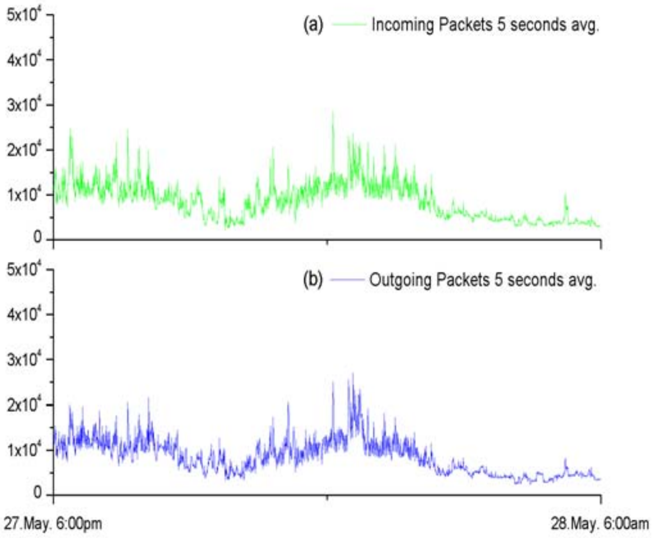

Figure 9. Incoming & Outgoing Packet Distribution

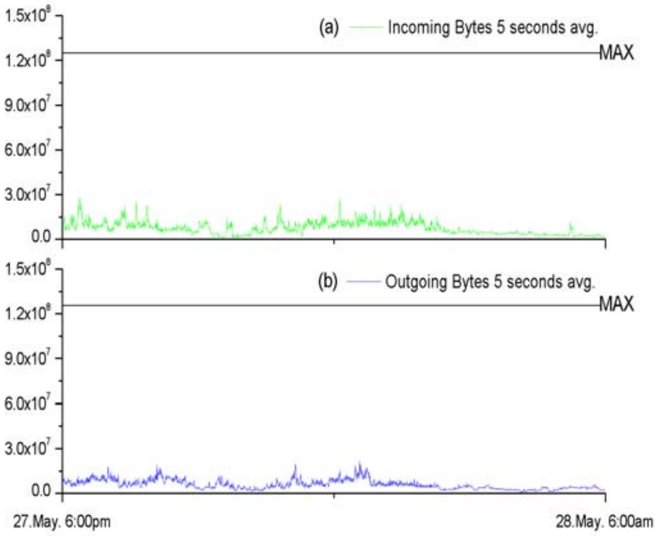

Figure 10. Incoming & Outgoing Byte Distribution

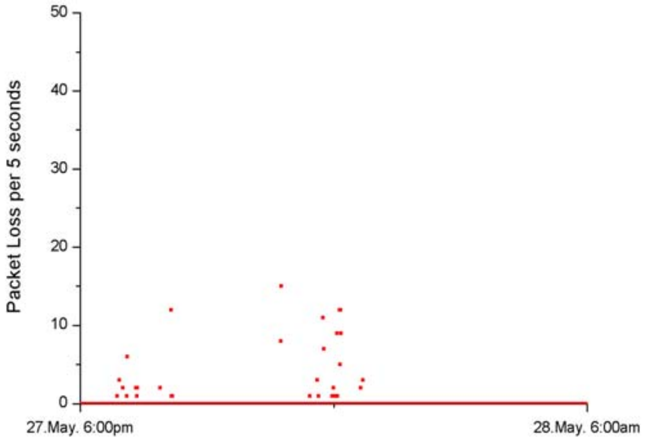

Figure 11. Packet Loss Distribution

V. Concluding Remarks

In our work, we collected data from SNMP MIB variables from POSTECH campus routers and switches and observed IP traffic and packet losses. Analysis of this data shows that the traffic bursts occur at small time granularity such as 1 second and 5 seconds, and we could find the packet losses in underutilized links.

Analysis reveals that some packet losses are caused by traffic bursts that cause a link to be highly utilized for small periods of time. We could also observe the packet losses in highly underutilized links and when no bursts can be seen with 5-second measurement values. We believe this happens due to the traffic burst that occurs at time granularity smaller than 5 seconds.

In this paper, we reduced the measurement time granularity down to 1 second, with 5 second data giving reliable measurement information due to router processing overhead. If it is further reduced then we expect the network to reveal other interesting properties. We are trying to detect packet loss phenomena with even smaller time granularity and study packet loss characteristics according to different applications such as Web, FTP, and P2P.

References

[1] Cisco, “NetFlow Services and Applications,” Cisco White Papers, http://www.cisco.com/warp/public/cc/pd/iosw/ioft/neflct/tech/napps_wp.htm.

[2] MRTG, “Multi Router Traffic Grapher,” http://people.ee.ethz.ch/~oetiker/webtools/mrtg/.

[3] Se-Hee Han, Myung-Sup Kim, Hong-Taek Ju and James W. Hong, "The Architecture of NG-MON: A Passive Network Monitoring System,” Distributed Systems: Operations and Management, Montreal Canada, October 2002, pp. 16-27.

[4] K. Park and W. Willinger. “Self-Simlar Network Traffic and Performance Evaluation,” Wiley-Interscience, 2000.

[5] Konstantina Papagiannaki, Rene Cruz and Christophe Diot, “Network Performance Monitoring at Small Time Scales,” Internet Measurement Conference, Miami, Florida USA, October 2003.

[6] James Hall, Ian Pratt, Ian Leslie and Andrew Moore, “The Effect of Early Packet Loss on Web Page Download Times,” Passive and Active Measurement Workshop, La Jolla, California USA, April 2003.

[7] Klaus Mochalski, Jörg Micheel and Stephen Donnelly, “Packet Delay and Loss at the Auckland Internet Access Path,” Passive and Active Measurement Workshop, Fort Collins, Colorado USA, March 2002.

[8] J.Case, M. Fedor, M. Schoffstall and J. Davin, “A Simple Network Management Protocol,” RFC 1157, May 1990.

[9] Cisco, “MIB Compilers and Loading MIBs,” Cisco Technical Notes, http://www.cisco.com/en/US/tech/tk648/tk362/technologies_tech_note09186a00800b4cee.shtml.