Detection and Analysis of Packet Loss on Underutilized Enterprise Network Links

Seung-Hwa Chung, Young J. Won, Deepali Agrawal, Seong-Cheol Hong, and James Won-Ki Hong
Dept. of Computer Science and Engineering
POSTECH
Pohang, Korea
{mannam, yjwon, deepali, pluto80, jwkhong}@postech.ac.kr

Hong-Taek Ju
Dept. of Computer Engineering
Keimyung University
Daegu, Korea
juht@kmu.ac.kr

Kihong Park
Dept. of Computer Sciences
Purdue University
West Lafayette, IN, USA
park@cs.purdue.edu

Abstract

ISPs and enterprises usually overprovision their networks as a means of supporting QoS. In spite of that, the service quality of QoS-sensitive applications such as VoIP, video conferencing and streaming media may not be up to expectations. We believe this is due to sporadic but non-negligible packet losses due to traffic bursts even in underutilized links. Our earlier work attempted to detect and analyze packet losses in underutilized links in an enterprise network environment using SNMP. This paper presents an extension of our earlier work by attempting to detect and analyze packet loss at finer granularity of time scale. We have developed a passive traffic capturing system, which can provide smaller time scale analysis of packet loss. We have analyzed the data and identified parts that are representative of packet loss across various time scales: 10 milliseconds, one second, 10 seconds and one minute. Analysis reveals that packet losses on underutilized link do occur due to bursty traffic packets in a small time scale. We also present analysis of other traffic properties such as packet size distribution and flows for the packet loss.

Keywords

Packet loss detection and analysis, packet loss characteristics, underutilized links, bursty traffic, QoS, SNMP

1. Introduction

An important component of traffic management for resource provisioning and network planning is traffic monitoring. Most ISPs and enterprises have chosen to overprovision their network link bandwidth based on the data that is obtained from traffic monitoring systems, such as NetFlow [1], MRTG [2] and NG-MON [3].

However, these systems monitor traffic with large time intervals and cannot provide data for precise analysis results. They are incapable of detecting bursty traffic on a small time scale because their coarse time granularity is limited to five minute and one minute aggregates in the case of MRTG and NG-MON, respectively [4]. Due to averaging over large time intervals, we may see links being underutilized (e.g., 20% or less) when, in fact, at time granularity smaller than a second, critical for assuring toll quality VoIP, traffic spikes and packet losses are present leading to unacceptable service qualities. This is especially relevant given the self-similar nature of Internet traffic [5].

To our knowledge, there are not many studies available in the area of packet loss on underutilized links. However, the following research are partially relevant to our study, focusing on packet loss. Papagiannaki et al. [6] presented a characterization of congestion in the Sprint IP backbone network. They analyzed link utilization at various time scales (millisecond level) to measure the frequency and duration of micro congestion. While they detected traffic bursts, they did not mention the packet loss that occurred during these times. Their study did not provide various traffic parameters except burst, and no information of the packet loss. In this paper, we provide the packet loss characteristics with various traffic parameters.

Hall et al. [7] analyzed 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 downloads than other types of packet losses. This work concentrated on web traffic, and showed that a small amount of packet loss can contribute to serious delays. Mochalski et al. [8] studied changes in traffic pattern relative to different points of observation using a passive network tap [9] 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.

Our earlier work attempted to detect and analyze packet losses in underutilized links in an enterprise network environment using SNMP [4]. We have discovered that using SNMP MIB variables was sufficient for detecting packet loss at coarse time granularity (e.g., one minute). However, it has limitations in detecting packet loss at finer timer granularity, which will provide more accurate data for analysis.

This paper presents an extension of our earlier work by attempting to detect and analyze packet loss at finer granularity of time scale. We have developed a passive traffic monitoring system, which can capture all packets going through a network link and provide traffic data of smaller time (e.g., millisecond level). We have analyzed the data and identified parts that are representative of packet loss across various time scales: 10 milliseconds, one second, 10 seconds and one minute. Analysis reveals that packet losses on underutilized link do occur due to bursty traffic packets in a small time scale. We also present analysis of other traffic properties such as packet size distribution and the number and lifetime of flows [10].

The organization of this paper is as follows. Our packet loss detection and traffic monitoring methods are described in Section 2. Section 3 describes the traffic data collection and experimental environment. In Section 4, we give an analysis of packet loss and IP traffic. Finally, concluding remarks are given and possible future work is discussed in Section 5.

2. Packet Loss Detection and Traffic Monitoring Method

In our work, we have implemented a traffic monitoring system that can detect traffic burst accurately in a small time granularity (e.g., 10 milliseconds). We basically monitored traffic on the link using a passive network tap and at the same time we monitored packet losses by polling Cisco private MIBs using SNMP [11].

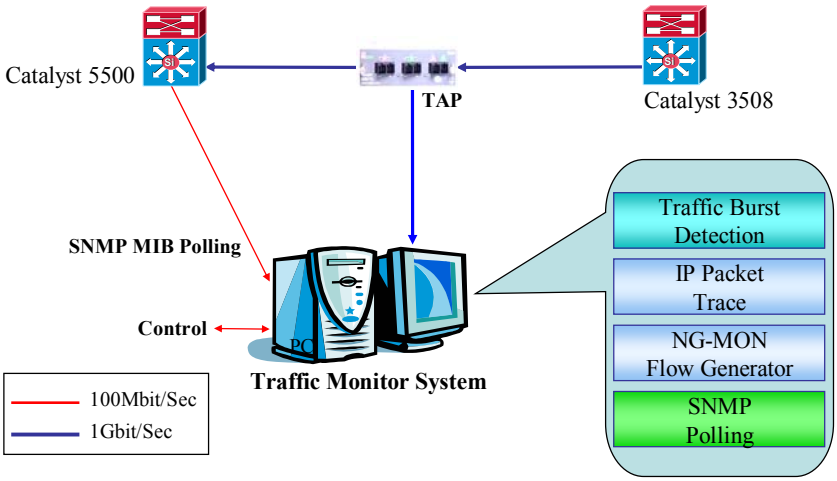

Figure 1: Overview of Traffic Monitoring Modules

There are four independent modules operating in the traffic monitoring system: the SNMP Polling Module, Traffic Burst Detection Module, IP Packet Trace Module, and NG-MON [3] Flow Generator Module. The modules in the monitoring system are shown in Figure 1 and their functionalities are described in Table 1. Although each module serves different purposes, their outputs are correlated to each other for detecting packet loss and analyzing traffic characteristics.

The SNMP Polling Module is responsible for an initial detection of packet loss at the router. We set the SNMP polling interval as 10 seconds, which we found was a reasonable time granularity for the router to update its MIB counters [4]. While analyzing combinations of Cisco private MIB variables [12, 13], we choose rather a simple solution which relies on the increase of Inqueue Drop counter variable to indicate packet loss. In addition, this module polls other MIB variables as well for later analysis, such as ingress packet and byte counts.

On the other hand, inputs of the rest of three modules are from the traffic packets being copied from the link using a passive network tap. The Traffic Burst Detection Module counts the number of packets it receives in small time scale as well as the total bytes, and reports a sudden increase of these values within the accuracy of 10 millisecond time granularity. This was possible through a minor modification to the 3Com gigabit NIC driver [17]. The IP Packet Trace Module operates at the same time to collect packet headers by using the libpcap library [18]. At last, the NG-MON Flow

Generator Module generates a 5-tuple (Src/Dst IP addresses, Src/Dst port numbers, and Protocol) based flows for later use to determine the relationships between packet loss and flow characteristics.

Thus, the SNMP Polling Module detects the packet loss by detecting an increase of the Inqueue Drop count and records its time. For the same time period, we conduct in-depth analysis with the data from the rest of three modules.

Module Name	Module Description
Traffic Burst Detection Module	This module detects traffic burst (number of packet, size of packet) at 10 millisecond time granularity.
SNMP Polling Module	This module polls Cisco private MIB and Standard MIB II values (number of dropped packet, ingress packets and bytes) at 10 second time granularity.
IP Packet Trace Module	This module captures IP packets passively from tap and save IP header information with captured time stamp.
NG-MON Flow Generator Module	This module generates packets into 5-tuple based flows (Src/Dst IP Address, Src/Dst Port and Protocol) and used for detail flow analysis.

Table 1: Functional Description of Modules in the Traffic Monitor System

3. Traffic Data Collection

The POSTECH campus network is comprised of a gigabit Ethernet backbone, which 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, as shown in Figure 2. We tried to find a link that was continuously underutilized and has traffic that is used by various Internet applications including QoS-sensitive applications. Our campus Internet access links are heavily utilized and we tried to find the most suitable place to observe packet losses on an underutilized link. We found the link that satisfied the above conditions within the campus dormitory network. This link is on the edge of the POSTECH campus network and is mostly underutilized.

We decided to monitor packet loss from the dormitory backbone switch (Cisco Catalyst 5500 [14]) that is connected to the Core switches (Cisco Catalyst 6513). This dormitory backbone switch is connected with many sub-dormitory switches as shown in Figure 2. We monitored all links between the dormitory backbone switch and sub-dormitory switches and found the link between the sub-dormitory switch (Cisco Catalyst 3508) for “Nakwon APT” and the dormitory backbone switch was the most suitable for our study. This 1-Gbps link is on the edge of the POSTECH campus network and the link bandwidth was continuously underutilized. Most importantly the link showed the frequent occurrence of packet losses.

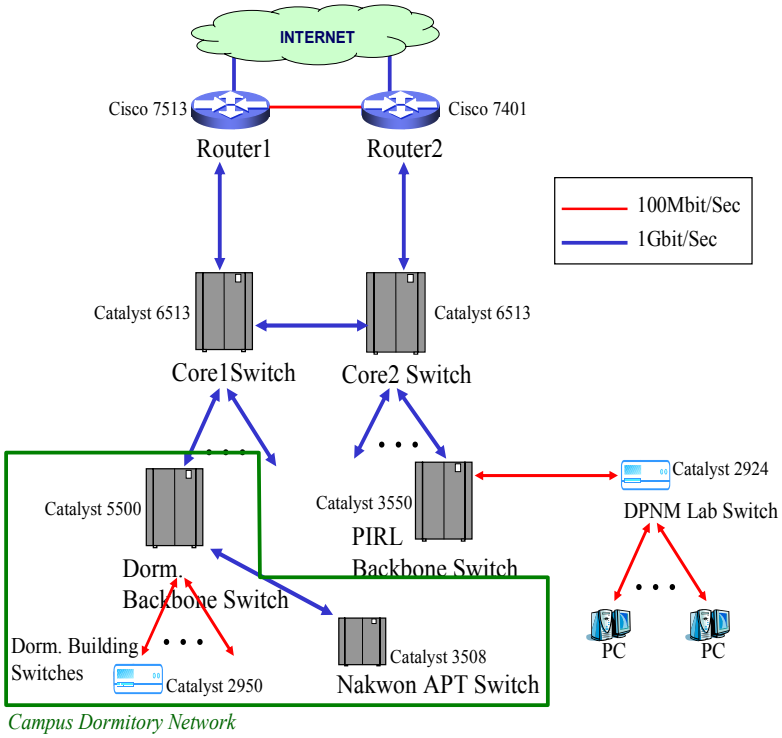

Figure 2: POSTECH's Campus Network

We installed a passive network tap on the link that connects the dormitory backbone switch (Cisco Catalyst 5500) and sub-dormitory switch (Cisco Catalyst 3508). The dormitory backbone switch offers Cisco private MIBs, so we could obtain packet loss data using SNMP.

4. Analysis of Packet Loss and IP Traffic

This section presents an analysis of IP traffic and packet loss. We collected traffic data including the number of packet losses, all packets' IP header information and the number of flows for one week, 2004.11.23 2:00 pm to 2004.11.30 3:00 pm from the specified switch and link. The traffic data is collected passively using a passive network tap and polling MIB values at a 10 second granularity and aggregated to yield one minute and five minute data. We focused on the portions of data that clearly illustrated packet loss.

4.1. Bursty Traffic Analysis

Figure 3 illustrates the distribution of incoming bytes (i.e., from Nakwon APT switch to Dorm backbone switch) observed using SNMP over the five day period at 10 second, one minute and five minute measurement granularity. This data is

collected from the dormitory backbone switch's 1 Gbps port, which is connected to sub-dormitory switch for "Nakwon APT." Overall average utilization of the link was about 5%. We observed that the curve of incoming packets and incoming bytes with five minute average values is fairly smooth. The curves with one minute average values clearly show bursty traffic and a wider distribution than the five minute measurement curves. Whereas, the curves with 10 second measurement values show even more bursty traffic with high peaks and wider distribution than the one minute 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 10 second. In Figure 4, the graph shows that packet loss occurred on a port that is connected to the sub-dormitory switch for "Nakwon APT" when the link on this port was underutilized.

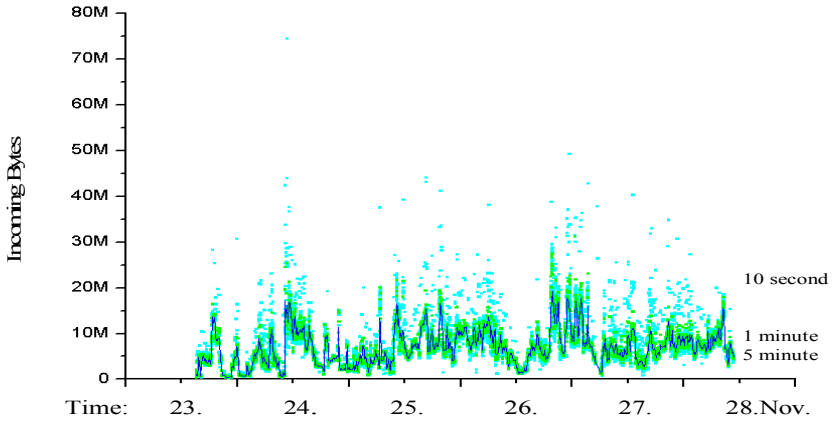

Figure 3: Incoming Byte Distribution

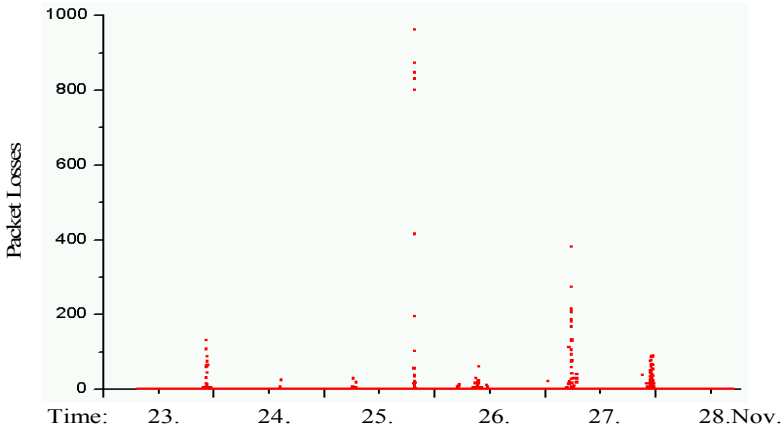

Figure 4: Packet Loss Distribution

In Figure 5, the graph (a) shows the total packet loss that occurred on the dormitory backbone switch except for the port that is connected to the sub-dormitory switch for "Nakwon APT." We tried to analyze traffic data when the port in focus alone had a packet loss because we wanted to analyze packet loss characteristics unaffected by traffic of other ports. Figure 5 (b) shows the packet loss occurrences that occurred only on the "Nakwon APT" port. We specified the data points at 2004.11.23 23:43:40 and analyzed traffic properties in this time to discover packet loss characteristics. The rest of collected traffic data during packet loss periods shared similar characteristics with the chosen one; however, the data collected during this time period revealed more clear analysis results than the others.

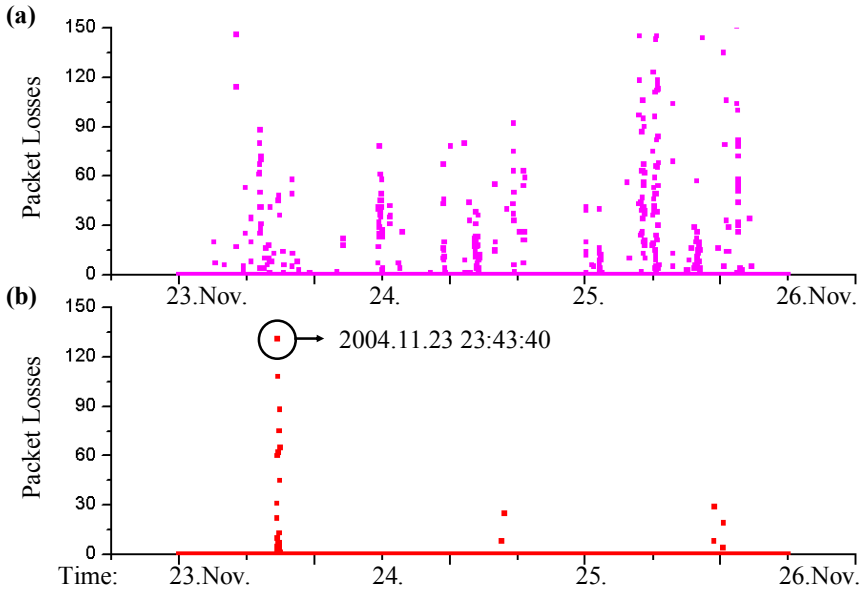

Figure 5: Packet Loss Distribution

During the five day period we monitored, the link bandwidth was continuously underutilized. As we can see in Figure 4, there were packet losses even on the underutilized link. Yet, we do not know the causes of such phenomenon. The following analysis will illustrate possible causes of the packet loss.

4.2. Bursty Traffic in Small Time Scale

The traffic monitoring system collected bursty traffic data from the link (between the dormitory backbone switch and sub-dormitory switch for "Nakwon APT") using a passive network tap in a time scale of 10 milliseconds. Figure 6 illustrates the bursty traffic when there was no packet loss detected (Figure 6 (a)) and when packet losses were detected (Figure 6 (b)). We can see that when there were lost packets the distribution shows more burstiness than when there is no packet loss.

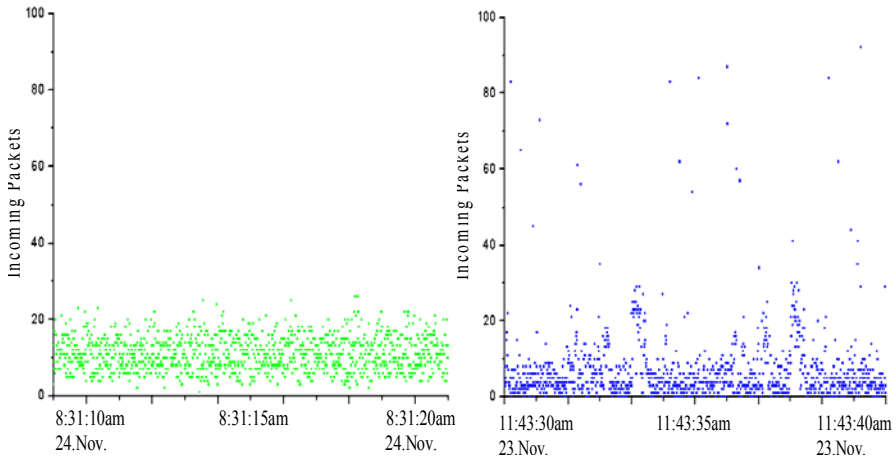

(a) without Packet Loss (b) with Packet Loss

Figure 6: Incoming Packet Distribution

Table 2 shows the statistics for incoming bursty packets for both cases. From the table we can observe that the mean value of incoming packets without packet loss and with packet loss is similar for both (the mean value of incoming bursty packets without packet loss is little higher). On the other hand, the standard deviation shows that the distribution of bursty packets gets highly variable when there is packet loss.

Case	Min.	Max.	Mean	Standard Dev.
(a)	1	32	11	4
(b)	0	196	8	15

Table 2: Incoming Packet Statistics

Figure 7 shows the incoming byte distribution (link bandwidth usage in granularity of 10 milliseconds) when there was no packet loss detected (Figure 7 (a)) and when packet losses were detected (Figure 7 (b)). Figure 7 shows an interesting distribution. We can see that when there was no packet loss the packet distribution shows more burstiness than when there were packet losses. From the result, we can observe that bursty packets are strongly related to packet loss but bursty bytes are not an important factor for packet loss in underutilized links.

Table 3 shows the statistics for incoming bytes for both cases. From the table we can observe that the mean value of incoming bytes without packet loss and with packet loss is similar for both (the mean value of incoming bursty packets without packet loss is little higher). On the other hand, the standard deviation even shows that the distribution of bursty bytes gets narrower when there is packet loss.

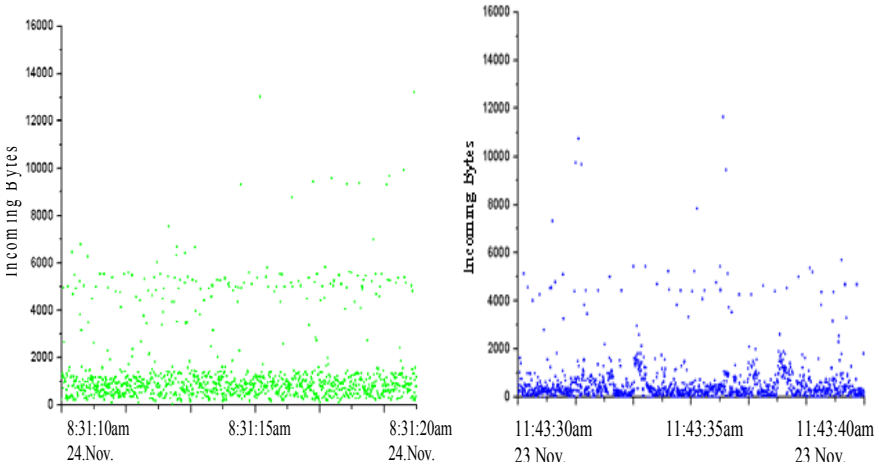

(a) without Packet Loss (b) with Packet Loss

Figure 7: Incoming Byte Distribution

Case	Min.	Max.	Mean	Standard Dev.
(a)	60	13948	1730	1932
(b)	0	12212	680	1187

Table 3: Incoming Byte Statistics

4.3. Packet Size Distribution

The traffic monitoring system collected IP packet headers with timestamp. Figure 8 illustrates packet size distribution over one minute. In the graph, one dot represents one packet and during this period, packet loss occurred in the marked time period. However, no special characteristics are found during the packet loss period. Internet traffic is generated by many different applications and certain packet sizes are more popular than others but it does not seem to be a property for packet loss. There are no traffic characteristics for packet loss related to packet size distribution on an underutilized link.

4.4. Flow Analysis

The traffic monitoring system collected the passing packets and generated flows using five tuple data (Src/Dst IP addresses, Src/Dst ports and Protocol) from the captured headers. Here, we present analysis of lifetime of flows and destination of flows to discover any relationship to packet loss.

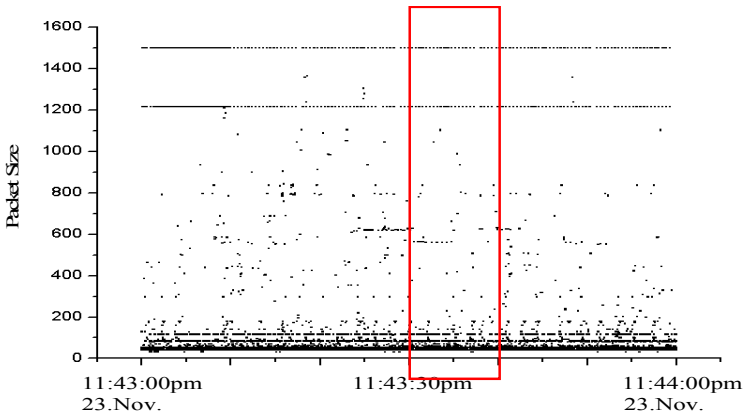

Figure 8: Packet Size Distribution

4.4.1. Flow Life Time

We analyzed flow data in two parts: long-life flow and short-life flow. In our monitoring system, long-life flow means the flow that is alive longer than or equal to one minute, and short-life flow means the flow which is alive for a period shorter than one minute. Long-life flows have the high probability that the flow is generated by lengthy file transfer. Nonetheless, because of TCP's slow start mechanism [15], the hub directly connected to the computers that generate long-life TCP flows may experience many and continuous packet losses.

Figure 9 illustrates the TCP flow distribution over a period of 40 seconds at a 1 second measurement granularity. Packet loss occurred in the marked time period. As shown in the graph, there are no special characteristics when packet loss occurred. Before the traffic from hosts in the edge reaches the dormitory backbone switch, the traffic needs to pass through several hubs and 100 Mbps links. The effect of TCP properties (e.g., TCP slow start) that might be responsible for packet loss is almost negligible on the dormitory backbone switch.

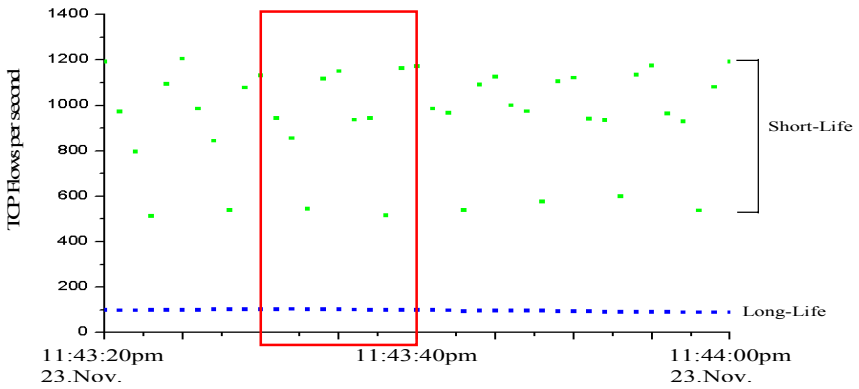

Figure 9: TCP Flow Distribution

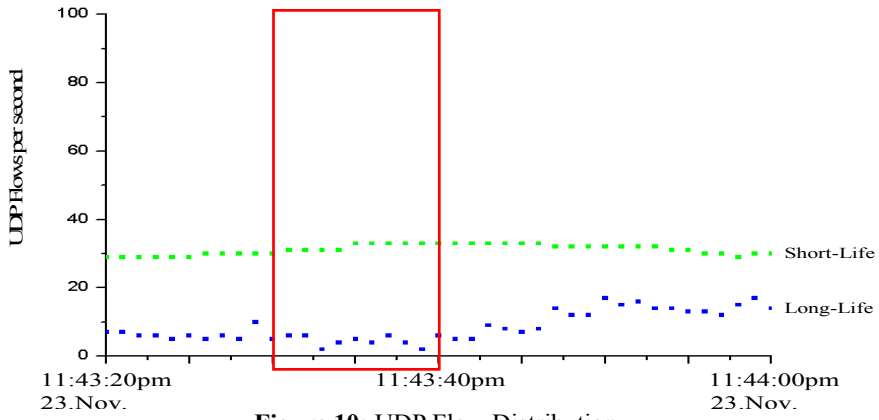

Figure 10: UDP Flow Distribution

Figure 10 illustrates UDP flow distribution over a period of 40 seconds at a 1 second measurement granularity, and the graph shows the same conclusion, that there are no special characteristics when the packet loss occurs, as shown in the TCP flow distribution graph.

4.4.2. One-Tuple Based Flows

Figure 11 illustrates the distribution of flows that are one-tuple based (i.e., packets with the same destination IP address) over a one minute period at 1 second and 10 millisecond measurement granularity. From the previous analysis, we found the five-tuple based flows (Src/Dst IP addresses, Src/Dst ports and Protocol) do not show special characteristics related to the packet loss, so we tried to merge packets into a one-tuple based flows because the switching packets inside router or switch depends on the destination IP address.

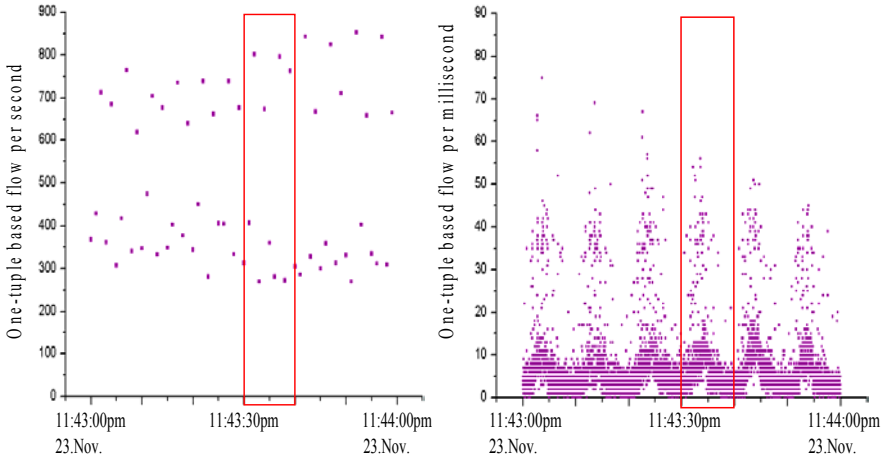

Figure 11: One-tuple based flows at 1 second and 10 millisecond granularities

The marked time on the graphs shows when packet losses have occurred. The graphs show no special flow properties on packet loss time, same as five-tuple based flow characteristics on the packet loss. We can observe that the packet destination does not affect packet loss on underutilized links.

5. Concluding Remarks and Future Work

In our work, we collected data from a link using a passive network tap and SNMP MIB variables from POSTECH campus dormitory network and observed IP traffic and packet losses. Analysis of this data shows that traffic bursts occur at small time granularity such as 10 seconds and 10 milliseconds, and we could detect packet losses in underutilized links.

Our analysis reveals that packet losses on underutilized links are caused by a high number of bursty packets rather than bursty bytes in a small time scale and time interval between captured packets is an important factor for packet loss. We also observed that packet size distribution does not affect packet losses on underutilized links. We analyzed various types of flows such as long-life and short-life TCP/UDP flows and a one tuple-based flow, and discovered that the characteristics of flows we analyzed during this study do not affect packet loss at all. The Internet traffic collected at our experimental spot consists of various types of traffic by different applications. However, it did not show the special characteristics of the packet size and the flow on the packet loss. Only the bursty packets are affecting the packet loss on underutilized link.

In this paper, we proved the existence of packet losses on underutilized links. We monitored the link and switch using our traffic monitoring modules implemented on a linux system. We analyzed packet loss with various types of traffic properties but the 10 millisecond time granularity can be still too large of a time scale to discover accurately packet loss characteristics for router/switch processes.

For future work, we will monitor packet loss in microsecond units with the help of hardware (e.g., DAG card [16]) and detect bursty CPU loads at a small time scale (Cisco enterprise MIB offers minimum of 5 second avg. value for CPU load). We also plan to study packet loss characteristics according to different applications such as Web, FTP, and P2P on first-contact hub.

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] Seung-Hwa Chung, Deepali Agrawal, Myung-Sup Kim, James W. Hong, and Kihong Park, “Analysis of Bursty Packet Loss Characteristics on Underutilized

Links Using SNMP," 2004 E2EMON, San Diego, California, USA, October 2004, pp. 68-74.
[5] K. Park and W. Willinger. *Self-Simlar Network Traffic and Performance Evaluation*, Wiley-Interscience, 2000.
[6] Konstantina Papagiannaki, Rene Cruz and Christophe Diot, "Network Performance Monitoring at Small Time Scales," Internet Measurement Conference, Miami, Florida USA, October 2003.
[7] 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.
[8] 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.
[9] Net Optics, Network Taps, http://www.netoptics.com/.
[10] Siegfried Lifer, "Using Flows for Analysis and Measurement of Internet Traffic," Diploma Thesis, Institute of Comm. Networks and Computer Engineering, University of Stuttgart, 1997.
[11] J.Case, M. Fedor, M. Schoffstall and J. Davin, "A Simple Network Management Protocol," RFC 1157, May 1990.
[12] Cisco Systems, "MIB Compilers and Loading MIBs," Cisco Technical Notes, http://www.cisco.com/en/US/tech/tk648/tk362/technologies_tech_note09186a00800b4cee.shtml.
[13] Cisco Systems, "Input Queue Overflow on an Interface," Cisco Technical Notes, http://www.cisco.com/en/US/products/hw/modules/ps2643/products_tech_note09186a0080094a8c.shtml.
[14] Cisco Systems, Switches. http://www.cisco.com/en/US/products/hw/switches/index.html/.
[15] W. Richard Stevens, *TCP/IP Illustrated, Volume 1: The Protocols,* Addison-Wesley, 1994.
[16] Endace, DAG Network Monitoring Interface Cards, http://www.endace.com.
[17] 3Com, Gigabit NICs. http://www.3com.com/.
[18] V. Jacobson, C. Leres and S. McCanne, *libpcap*, Lawrence Berkeley Laboratory, http://www.tcpdump.org/.