An abridged version of this paper appears in *IEEE Network*, November/December 1997.

Wide-Area Internet Traffic Patterns and Characteristics (Extended Version)

Kevin Thompson, Gregory J. Miller, and Rick Wilder

MCI Telecommunications Corporation[†]
vBNS Engineering
2100 Reston Parkway
Reston, Virginia 20191
{kthomp, gmiller, wilder}@mci.net

Abstract *– The Internet is rapidly growing in number of users, traffic levels, and topological complexity. At the same time it is increasingly driven by economic competition. These developments render the characterization of network usage and workloads more difficult, and yet more critical. Few recent studies have been published reporting Internet backbone traffic usage and characteristics. At MCI, we have implemented a high-performance, low-cost monitoring system that can capture traffic and perform analyses. We have deployed this monitoring tool on OC-3 trunks within internetMCI's backbone and also within the NSF-sponsored vBNS. This paper presents observations on the patterns and characteristics of wide-area Internet traffic, as recorded by MCI's OC-3 traffic monitors. We report on measurements from two OC-3 trunks in MCI's commercial Internet backbone over two time ranges (24-hour and 7-day) in the presence of up to 240,000 flows. We reveal the characteristics of the traffic in terms of packet sizes, flow duration, volume, and percentage composition by protocol and application, as well as patterns seen over the two time scales.*

1 Introduction

Sustained, rapid growth, increased economic competition, and proliferation of new applications have combined to change the character of the Internet in recent years. The sheer volume of the traffic and the high capacity of the trunks have rendered traffic monitoring and analysis a more challenging endeavor. In its role as the network service provider for the National Science Foundation's (NSF's) very-high-speed Backbone Network Service (vBNS) project, MCI has developed an OC3-based traffic monitor, known as OC3MON [ACTW97]. This publicly-available tool facilitates measurement and analysis of high-speed OC3 trunks that carry hundreds of thousands of simultaneous flows.

In this paper, we report on traffic measurements taken from two locations on internetMCI's commercial backbone. We characterize the traffic over two time scales, 24 hours and 7 days, in terms of traffic volume, flow volume, flow duration, and traffic composition in terms of IP protocols, TCP and UDP applications, and packet sizes. We also present an analysis of the trunk capacity consumed by ATM protocol overhead. Finally, we perform a flow optimization analysis that examines the measured traffic from the perspective of a router that is capable of creating optimized switching paths to improve the performance of long-lived flows. We show the percentage of packets that would benefit from optimized paths depending upon the number of packets chosen to identify a sufficiently long-lived flow. We find that under our assumptions, optimizing the switching of 20% of the flows benefits about 50% of the packets.

The paper begins with an outline of related previous work in section 2. Section 3 describes the OC3 monitor platform and section 4 provides a description of the data collection points within the backbone. In section 5, we present the detailed traffic data. Section 5 is divided into separate subsections for general traffic characteristics, IP protocol data, TCP and UDP application data, protocol overhead analysis, and flow optimization analysis. In section 6 we provide a summary of our major observations, in section 7 we provide directions for future work, and in section 8 we give pointers for those interested in the availability of vBNS OC3 monitor data or the monitor software.

[†] This work was supported by the National Science Foundation under Cooperative Agreement Number NCR-9321047.

2 Previous Studies

Internet traffic measurement has been a subject of interest as long as there has been Internet traffic. A number of comprehensive studies have been published over the years as the Internet has evolved from the original ARPANET [Kleinrock76] to today's multi-provider commercial environment. Under its centralized administration, the NSFNET, from 1988-1995, provided a single Internet backbone with accessible measurement points. Traffic studies from this era include [CPB93], which investigated detailed NNstat and SNMP NSFNET data in May 1992 and earlier, presenting a range of traffic characteristics on what was at the time a T1 backbone. Application usage was reported as being dominated by FTP, SMTP, NNTP, DNS, Telnet, and a growing amount of "other" traffic. Packet sizes exhibited a bimodal distribution indicating a mixture of bulk data transfers and interactive applications, with a mean packet size of 186 bytes.

Other studies have reported on traffic growth trends of the NSFNET [Heimlich89, Frazer95]. Merit's final NSFNET report summarized statistics on the backbone from 1988-1995, showing exponential traffic growth through 1994 [Frazer95]. Information storage and retrieval applications such as Gopher and Web were noted as beginning to overtake mail and file transfers in their share of network traffic. Traffic composition by packet counts in April 1995 showed this distribution: other (27%), WWW (21%), FTP-data (14%), NNTP (8%), Telnet (8%), SMTP (6%), IP (6%), Domain (5%), IRC (2%), Gopher (2%), and FTP-control (1%).

Wide-area traffic studies have been conducted from the perspective of one to a few specific sites, revealing the traffic breakdown in terms of protocols and applications [Caceres89, CDJM91, Paxson94]. Caceres showed that TCP offered 90% of the byte count and 80% of the packet count at Bell Labs in 1989. SMTP was responsible for more than 50% of all TCP packets [Caceres89]. Paxson's long-term site study showed growing usage of applications like SMTP, FTP, and X11 in the pre-Web days [Paxson94].

Since the disappearance of the NSFNET backbone in 1995, multiple, commercial-carrier administered backbones have emerged, connected via Network Access Points (NAPs) and private interconnects. Large-scale traffic measurement studies have been published less frequently in this more competitive environment. Today, the Web provides on-line access to current and historical SNMP statistics at public inter-exchange points (see [stats]). The National Laboratory for Applied Network Research (NLANR) makes available IP protocol and application usage data from an NNstat collector at the Federal inter-exchange point called FIX-West (see [FIXWEST]). Some commercial Internet providers are now taking advantage of the ability to take measurements on their own backbones with a similar high level of detail using Cisco Systems routers. MCI, ANS and BBN use an experimental system, called Cflowd, to collect and process data produced by the flow export feature in Cisco's Internetwork Operating System (IOS) [MH97].

Other recent Internet measurement and analysis studies have examined network performance issues, such as route stability and end-to-end performance. [Paxson97] notes the difficulties in obtaining end-to-end measurements on scales larger than a single site. This comprehensive report analyzes 20,000 TCP transfers of 100 kilobytes each among 35 sites over about 1,000 Internet paths. His findings cut across areas such as routing pathologies, loss characteristics, specific TCP implementation behaviors, and packet queueing and delay variations.

In this paper, we present detailed measurements of wide-area trunks on MCI's Internet backbone, focusing on the daily and weekly patterns exhibited by the traffic. We characterize the traffic in terms of the IP protocols and applications present, the packet sizes seen, and the properties of the various flows observed. Our measurements confirm the continuation of the trend observed by earlier studies [BC94] whereby the Web has grown to be the largest single application on the Internet.

3 Description of the OC3 Monitor

In this section, we briefly describe OC3MON, our special-purpose OC3-rate monitoring and analysis platform, developed in a collaboration between MCI's vBNS Engineering and NLANR. We used this tool to gather the data presented in the remainder of the paper. The goal of the OC3MON project was to address three incompatible trends: Current, widely-used statistics gathering tools, which are largely FDDI and Ethernet based, have difficulty scaling to higher speeds. ATM circuits at OC3 rate are increasingly used for high-volume backbone trunks and interconnects. And finally, detailed, flow-based analysis is important to understanding usage patterns and growth trends, but such analysis is not generally possible with the data that can be obtained directly from today's routers and switches.

The current OC3MON implementation satisfies the need for a high-speed monitoring and flow analysis tool while meeting the project's two driving design constraints of flexibility and low cost. OC3MON is a programmable data collection and analysis tool that can be easily modified as we codify and refine our understanding of the desired statistics. Furthermore, it is inexpensive to build, which facilitates widespread deployment. Both the flow analysis code and monitor architecture are in the public domain.

3.1 Description of the OC3MON Hardware

The OC3MON platform is an IBM personal computer clone with 128 MB of main memory, a 166 MHz Intel Pentium processor, an Ethernet interface, two ATM interface cards, and a 33 MHz 32-bit-wide PCI bus. The ATM interface card used in the current OC3MON implementation is the Fore Systems ATM network interface card (NIC) for the PCI bus. The Intel i960 processor on this interface card allows us to optimize OC3MON operation with custom firmware.

We attach the two OC3MON ATM NICs to an OC3 optical fiber pair carrying IP-over-ATM traffic. We connect the receive port of each ATM card to the monitor port of an optical splitter. The splitter carries a fraction of the light from each fiber to the receive port of one NIC. Attached to an OC3 trunk that terminates on a switching device (e.g., an ATM switch or a router), one of the OC3MON NICs sees all traffic received by the switching device and the other NIC sees all traffic transmitted by the switching device. The OC3MON NICs capture traffic on the two directions of an OC3 link independently.

3.2 Description of the OC3MON Software

The custom-developed OC3MON firmware is implemented in C++ and assembly code and provides full flexibility in terms of collection and analysis capability. OC3MON is capable of three types of data collection: raw cell trace capture, active flow reporting, and expired flow analysis. In raw trace mode, OC3MON captures either every cell, or the first cell of every packet (AAL5 frame), that appears on the link. In this mode, OC3MON does not perform any analysis on the captured data, but simply produces a time-stamped raw cell trace. The maximum length of a raw cell trace collected by OC3MON is limited to the amount of RAM in the monitor. In the two flow analysis modes, OC3MON collects statistics regarding *flows*, either *active* or *expired*, where the definition of a *flow* is configurable.

In this paper, we define a flow as a uni-directional traffic stream with a unique <*source-IP-address, source-port, destination-IP-address, destination-port, IP-protocol*> tuple. Any flow for which the monitor has not seen a packet within the last 64 seconds is considered to be an *expired flow*. This flow definition was proposed in [Claffy96]. A flow for which a packet was seen within the last second is considered to be an *active flow*. The monitor reports statistics based on expired flows whenever polled by a collector. After reporting the expired flows statistics to the collector, the monitor clears its stored state on all expired flows. It continues to maintain state information on flows that have not yet expired. These continuing flows, which have not yet been reported on, are referred to as *known flows.* The data presented in this paper is derived from expired flows statistics as reported by the monitors when being polled around the clock on 5-minute intervals. Flows that have been in progress (*known*) for one hour are artificially expired by the monitor so they can be reported. This mechanism for delimiting long-lived flows can affect the data that is reported by producing artificial traffic spikes.

The monitor does not maintain statistics on expired flows individually, but instead aggregates them on a per-protocol basis. For this reason, flow statistics such as duration, byte volume, and length in packets are reported in terms of averages. We intend to modify the monitor in the future to allow collection of more detailed distribution information, as we recognize that averages can be of limited value in describing Internet traffic characteristics because of their wide variation.

4 Description of the Monitoring Points

MCI has deployed the OC3MON monitor at all vBNS supercomputer sites and switching points as well as at two locations on the internetMCI backbone as of September 1997. In this paper we present data from the commercial Internet backbone only. Traffic on the vBNS, which has been to date used primarily by supercomputer centers for research and experimentation, is not representative of typical commodity Internet traffic. The most prevalent application on the vBNS presently is inter-cache communication among Web caches deployed at the edges of the vBNS backbone (see [cache]).

Two OC3MONs are installed on OC3 links within nodes on the internetMCI backbone. Both OC3MONs monitor traffic on a fiber pair between a core router and a backbone ATM switch. The first monitoring point (see Figure 1) is within a node that serves as a junction for several backbone trunks as well as an access point for local customer traffic near a major U.S. East Coast city. The measurements from this point were taken the week of August 24, 1997, beginning 00:00 Eastern Daylight Time (EDT) Sunday, August 24, 1997 and ending 24:00 EDT Saturday, August 30, 1997. Our one-day graphs from this site focus on Wednesday, August 27, 1997. The second monitoring point (see Figure 2) is within a node that includes a U.S.-U.K. trans-Atlantic DS-3 trunk. In this case, the monitor is installed on an OC3 fiber pair between the router to which the international DS-3 trunk is homed and the backbone ATM switch. The measurements at this point were conducted from 00:00 EDT Saturday, May 17, 1997 through 24:00 EDT Friday, May 23, 1997. In this case, the one-day graphs focus on Friday, May 23, 1997. Because we wanted to isolate the international traffic as much as possible, we were unable to use more recent data from this monitor. In late May 1997, MCI began homing new customer access links to the router where the international trunk terminates, making it more difficult to measure exclusively the international traffic from our measurement point.

5 Detailed Traffic Analysis

We group the data contained in this section into a number of distinct categories for the purposes of presentation. Initially, we examine traffic on the two measured links separately, presenting general traffic characteristics for each link. Here, we depict the overall traffic volume in terms of bytes, packets, and flows as measured over both 24-hour and 7-day time ranges. We also show the average packet size on the link as it varies over these time periods and we present packet size distribution data. The initial section concludes with an analysis of the composition over time of the traffic on each link in terms of the IP protocols and applications present. Next, we examine the traffic patterns exhibited by the individual IP protocols measured. We focus on TCP and UDP, which are the two most prevalent protocols on the measured links. Finally, we study the breakdown of the more prevalent TCP and UDP applications on the two measured links. For TCP, we present WWW client and server data, and for UDP we present data on DNS and RealPlayer (i.e., RealAudio and RealVideo) traffic.

5.1 General Characteristics

In this section, we present general traffic patterns for the two measured links. For each link, we present six graphs that illustrate traffic volume. The graphs in right-hand column cover a 7-day time period while the graphs in the left column focus on a 24-hour period. The first row of graphs, (a) and (b), shows traffic volume in Megabits per second, the second row, (c) and (d), shows volume measured in thousands of packets per second, and the final row, (e) and (f), shows thousands of known flows, both average and maximum seen per collection interval. Following the traffic volume graphs are figures showing packet size data. The first two packet-size graphs show the average packet size as a function of time over the 24-hour and 7-day time ranges. The second two packet-size graphs are histograms of a single 5-minute sample, one in linear scale and the other in logarithmic scale. The final packet size graph is a cumulative distribution plot. Finally, we show "region plots" that graphically depict the percentage composition of the traffic on one of the links over a 24-hour period, both in terms of IP protocols and TCP and UDP applications.

5.1.1 Domestic Link Traffic Patterns

The graphs in Figure 3 show traffic volumes in Megabits per second, kilopackets per second, and known flows at the U.S. East Coast domestic link measurement point. The Megabit per second and kilopacket per second graphs show the volume on each direction of the link, while the flow plots depict the average and maximum number of known flows per second aggregated for both directions. These graphs reveal a traffic pattern that is reflected in the majority of the plots presented in this paper. Namely, traffic volume follows a clear, predictable 24-hour pattern that repeats daily. The graphs in Figures 3 illustrate this pattern over both daily and weekly intervals.

We see traffic nearly triple in volume on both directions of the link from 6:00 through about 12:00 noon EDT. Weekend days show a similar pattern, although the weekend traffic volume decreases by about 25% in bytes and packets, and 20% in flows. We see a disparity in volume on the two directions on the link. The south direction consistently carries a significantly higher packet rate than the north direction, often as much as 50% higher. However, the direction carrying the higher bit-rate changes daily at about 8:00 and 19:00, where the two curves intersect on Figure 3a.

Figures 3e and 3f show two statistics related to known flows: the average number of simultaneous known flows per collection period, as computed from samples taken once per second; and the maximum number of simultaneous known flows seen during each 5-minute collection interval. We observe that during peak hours the monitor keeps state on as many as 240,000 known flows. The flow-volume pattern (Figures 3e and 3f) follows the byte-volume pattern (Figures 3a and 3b) fairly closely. Also, the average number of known flows during any given collection interval is usually within 15% of the maximum for that interval, indicating that there is little variation in the number of known flows on a 5-minute time scale.

Figure 4 shows the average packet size versus time on each direction of the U.S. regional trunk over 24-hour and 7-day time periods. We see that, like traffic volume, average packet size varies over the course of the day, with the same 24-hour pattern repeated daily throughout the week. The packet size patterns on the two directions of the link differ, with the averages for the two directions being roughly inversely proportional to one another. That is, the larger the average packet size in one direction on the link, the smaller the average packet size in the other. This relationship makes sense if we consider that as the average packet size in one direction is driven up by large TCP data packets, most often from HTTP transfers, the average packet size in the reverse direction will be driven down by the corresponding small TCP acknowledgement segments. The packet-size discrepancy between the two directions is responsible for the bit-rate discrepancy shown in Figures 3a and 3b. Recall that while the packet-rate in the south direction is always higher than the packet-rate in the north direction, the bit-rate in the north direction is higher than the south's from 8:00 until 19:00 each day. This observation is explained by the change in the direction carrying the larger average packet size daily at approximately 8:00 and 19:00, shown in Figure 4a.

Figure 4c shows a linear-scale histogram of packet sizes from a 5-minute time period on both directions of the domestic link. Figure 4d shows the same data on a logarithmic scale. These figures show the expected predominance of small packets, with peaks at the common sizes of 44, 552, 576, and 1500. The small packets, 40-44 bytes in length, include TCP acknowledgement segments, TCP control segments such as SYN, FIN, and RST packets, and Telnet packets carrying single characters. Many TCP implementations that do not implement Path MTU Discovery use either 512 or 536 bytes as the default Maximum Segment Size (MSS) for nonlocal IP destinations, yielding a 552-byte or 576-byte packet size [Stevens94]. A Maximum Transmission Unit (MTU) size of 1500 bytes is characteristic of Ethernet-attached hosts.

Figure 4e is a cumulative distribution plot for the same packet size data presented in Figures 4c and 4d. This graph shows that almost 75% of the packets are smaller than the typical TCP MSS of 552 bytes. Nearly half of the packets are 40 to 44 bytes in length. A linear slope on Figure 4e between 572 and 1500 bytes indicates an equal distribution of packets in that size range. Almost 100% of the packets are 1500 bytes or smaller.

5.1.2 International Link Traffic Patterns

The graphs in Figure 5 depict bit-rate, packet-rate, and flow volumes for the international link measurement point. Again, bit-rate and packet-rate volumes are shown for each direction on the link. We see that the traffic levels on the two directions of the link differ, with the U.S.-to-U.K. volumes following the same 24-hour pattern exhibited on the domestic link, except that it is shifted back 5-6 hours. The time shift in the pattern is presumably due to the 5 to 6 hour time difference between the U.S. East Coast and the U.K and Europe. It is interesting to note that while the U.S.-to-U.K. direction has more than twice the bit-rate during peak morning hours than the reverse direction does, its packet rate is consistently lower, by a factor of 1/2 to 2/3, for a 24-hour period. This implies that packet sizes in the U.S.-to-U.K. direction are larger (as confirmed by packet size measurements shown in Figure 6). We hypothesize that the larger packets are generated by Web servers in the U.S. in response to requests from Web clients in the U.K. and Europe.

Figures 5e and 5f show that during peak hours, the monitor witnesses up to 160,000 known flows. Like on the domestic link, the flow-volume pattern follows the byte-volume pattern and the average and maximum flow-volume values are nearly the same.

The graphs in Figure 6 depict packet-size data for traffic on the international link. Average packet size on the international link follows a cyclic pattern daily, as it does on the domestic link. Figure 6a shows, even more so than on the domestic link, a clear inverse proportionality relationship between the average packet sizes in the two directions on the link. Again like the domestic link, the direction carrying the larger average packet size reverses daily.

The packet size histogram and distribution graphs, Figures 6c, 6d, and 6e, show a similar packet size composition on the international link as on the domestic link. We again see that over 50% of the packets are smaller than 45 bytes in length and that nearly all packets are 1500 bytes or smaller. This packet-size data, including the temporal variation of the average and the modality evident in the packet size histograms, calls into question the concept of average packet size as a meaningful statistic, which is, at best, time-of-day and direction dependent.

5.1.3 Traffic Composition

The graphs in Figure 7 reveal the composition of the traffic over a 24-hour period on the domestic link. The plots in the left-hand column show the traffic constituency in terms of IP protocols and the right-hand column shows the breakdown of TCP and UDP applications. Focusing first on IP protocols, we observe that TCP is by far the dominant protocol on the link. Over the course of a day, TCP averages about 95% of the bytes, 90% of the packets, and at least 75% of the flows on the link. UDP has the second highest levels at roughly 5% of the bytes, 10% of the packets, and 20% of the flows on average. The other IP protocols plotted in Figure 7 are IPv6, encapsulated IP (IP-in-IP), ICMP, and an aggregate category for the remaining protocols labeled “other.” These other protocols individually make up a negligible percentage of the overall traffic. ICMP constitutes the third highest packet percentage after TCP and UDP, but still makes up less than 2% of the overall packets and 0.5% of the overall bytes.

The right-hand column of Figure 7 shows the percentage composition of the most prevalent TCP and UDP applications measured on the domestic link over a 24-hour period. For each application, we combine client-to-server and server-to-client (and in the case of DNS, server-to-server) traffic into a single category. We see that the Web is the dominant application on the link, comprising up to 75% of the bytes, 70% of the packets, and 75% of the flows when client and server traffic are considered together. In measuring applications, we end up with a larger “other” category than when measuring IP protocols. The “other” category is spread among a wide range of TCP and UDP port numbers, no one of which represents a significant percentage of the traffic by itself. Among the most common port numbers in the “other” category are 81, 443, 3128, 8000, and 8080, which are all Web-related, indicating that the Web may actually be slightly under-represented in our measurements.

In addition to Web traffic, we identify 5 other applications that contribute an appreciable percentage of traffic: DNS, SMTP, FTP (data connections), NNTP, and Telnet. In terms of flows, DNS traffic represents the second largest application at nearly 18% of the overall flows. However, DNS flows are small, accounting for only 3% of the total packets and 1% of the bytes on average. SMTP averages 5% of the bytes, 5% of the packets, and 2% of the flows. FTP data connections, on average, constitute roughly 5% of the bytes, up to 3% of the packets, and less than 1% of the flows. NNTP represents 2% of the bytes, and less than 1% of the packets and flows. Finally, Telnet accounts for about 1% of the packets, and less than 1% of the bytes and flows.

In the following sections we examine more closely the characteristics of the two most prevalent IP protocols, TCP and UDP, and three of the applications: WWW, DNS, and RealPlayer. In those sections, we analyze the daily variation in the traffic composition evident in Figure 7, in addition to the averages reported here.

5.2 IP Protocol Traffic

In this section, we present data on the various IP protocols observed on the two measured trunks. We focus on TCP and UDP, which are by far the two most prevalent protocols, but we also include brief discussion on ICMP, IP-in-IP, and IPv6. For TCP and UDP, we show a total of six graphs for each protocol from each measurement point. The three graphs in the left column of each page depict data plotted on a 24-hour time scale, while the graphs on the right show a 7-day range. For each time period, we present three types of plots for the category in question: aggregate traffic, fractional traffic, and flow characteristics. The aggregate graphs show the volume (per 5-minute interval) of the protocol’s traffic versus time, as measured in bytes, packets, and flows. The fractional graphs show the protocol’s traffic as a percentage of the total traffic, again as measured in bytes, packets, and flows. Finally, the flow characteristics graphs show the average number of bytes, packets, and seconds per flow versus time, for the flows pertaining to the protocol in question. For each protocol, the aggregate graphs (a and b) are presented first, followed by the fractional graphs (c and d), and finally the flow graphs (e and f). All quantities in this section are based on data from expired flows, and all graphs show aggregated data for both directions on a particular link.

5.2.1 TCP

The graphs in Figure 8 show Transmission Control Protocol (TCP) traffic on the international link. The traffic volume graphs (Figures 8a and 8b) show packets, bytes, and flows each more than doubling from midnight levels into the 10:00 hour. Evening hours show a corresponding decrease toward midnight. The fractional traffic graphs (Figures 8c and 8d) show TCP's dominance in the traffic mix as noted in the previous section for the domestic link: 95% and more of total bytes, and 85%-95% of total packets. TCP flows oscillate between 75% and 85% of the total. We see that TCP is a higher percentage of overall traffic during business hours than in the evening or overnight. As Figures 8e and 8f show, the average number of packets per TCP flow is as high as 22 during overnight hours and about 17 during the day. Average kilobytes follow a similar time-of-day pattern, ranging from 5-8 kilobytes per flow. From this data, an approximation of the average TCP packet size is 300 bytes. The average number of packets per flow is 16-20, and the average flow duration is 12-19 seconds, with time-of-day variation.

The graphs in Figure 9 show TCP traffic on the U.S. regional trunk. Both the 24-hour and 7-day aggregate TCP traffic graphs are consistent in pattern with the overall traffic patterns for the trunk. This pattern is not surprising because of TCP's dominance on the link, in this case about 95% of the overall bytes, 85-90% of the overall packets, and 70-75% of the overall flows are TCP. Fluctuations downward in the packet fractions (in Figures 9c and 9d) are due to the artificial timeouts per-hour of long-term UDP flows. Recall that the OC3MON artificially expires flows that last longer than one hour so they can be reported. This reporting mechanism produces artificial spikes in the traffic volume on one-hour intervals that represent the long-lived flows. We have observed that the vast majority of such long-lived flows are UDP flows. When the artificially-expired long-lived UDP flows are reported as a large spike in usage each hour, TCP's proportion of the total traffic mix drops at that instant, creating the dips seen in Figure 9c. The TCP average flow characteristics on the domestic link are much the same as on the international trunk shown in the previous figure. For TCP flows on the regional link, again we see on average 17-21 packets per flow, 5-9 kilobytes per flow, and about a 300-byte average packet size.

5.2.2 UDP

The graphs in Figure 10 show User Datagram Protocol (UDP) traffic on the international link. Aggregate flow count patterns, shown in Figures 10a and 10b, look similar to those of TCP (though TCP shows 4 or 5 times the number of flows). UDP byte and packet volumes, however, exhibit a different pattern than TCP, with a more gradual increase in volume over the course of the day, well into 19:00 hours before tailing off more abruptly into midnight. As a fraction of overall traffic, UDP is inversely proportional to the TCP load, making up less than 5% of the bytes, between 5%-15% of total packets, and 15%-25% of total flows. UDP flows average 5-15 packets per flow, with the longer flows occurring in the evening. They average 1-2 kilobytes per flow and vary from 10-18 seconds in duration.

The graphs in Figure 11 show UDP traffic on the U.S. regional link. The fractional graphs (Figures 11c and 11d) are interesting as they show UDP rise in the late evening and early morning hours as a proportion of overall traffic. The apparent burstiness in the UDP packet volume graphs is misleading. The periodic spikes are a result of the monitor's artificial expiration of long-lived UDP flows, as discussed previously. UDP traffic on the regional link averages 5-10 packets, about 15 seconds, and 1-2 kilobytes per flow. The average UDP packet size ranges from 200-500 bytes per packet.

5.2.3 ICMP

Internet Control Message Protocol (ICMP) traffic generally comprises less than 2% of the overall packets and 0.5% of the overall bytes on the two measured links. ICMP flows peak as a percentage of traffic at nearly 1.5% of total flows during late night and early morning hours. ICMP packets and flows increase in volume from the early morning hours U.S. East Coast time up to 20:00, and drop correspondingly to half their peak levels toward midnight.

5.2.4 IP-in-IP Encapsulation

We generally attribute IP-in-IP encapsulated traffic, or tunneled IP, to Mbone traffic. IP-in-IP shows up in our measurements as what appears to be bursty traffic on each time scale, but this misleading observation is again a result of the monitor's artificial expiration and reporting of long-lived flows each hour. As it is well known that Mbone sessions frequently last more than 1 hour in duration, it is a weakness in our method of reporting that skews the data into appearing bursty. On the domestic link, we see bursty traffic patterns without a discernable daily pattern. On a percentage basis, the brief bursts are as large as 20% of the link's byte traffic and 10% of the packets during several collection periods. The IP-in-IP flow percentages are negligible. On the international link, we see fewer IP-in-IP flows than on the domestic link, suggesting that there are fewer Mbone tunnels across this link.

5.2.5 Encapsulated IPv6

We see a small amount of encapsulated IP version 6 (IPv6) traffic on the measured links. The aggregate volumes show occasional spikes of IPv6 traffic, indicating the presence of a few long-lived flows. There is no discernable daily or weekly pattern in the IPv6 traffic.

5.3 TCP Application Traffic

In this section, we present data on TCP application traffic. We focus on Web client and server traffic, but also include a description of FTP (data connection) server traffic, and NNTP server traffic.

5.3.1 Web Client Traffic

The graphs in Figure 12 show Web client data for the international link. The daily bell-shaped pattern seen in the aggregate Web-client plots (Figures 12a and 12b) is nearly identical to the daily pattern seen for all TCP traffic on this trunk, shown in Figures 8a and 8b. Considering client and server traffic together, the Web dominates among applications, comprising 75% of the overall bytes, up to 70% of the overall packets, and 75% of the overall flows during daytime hours. Figures 12c and 12d show Web client traffic as a proportion of total traffic. These graphs show an increase in the fraction of Web traffic during daytime hours and a decrease overnight. While Web client traffic and server traffic are roughly equivalent in terms of packet fractions (30%-38%) and flow fractions (35%-42%), there is an asymmetry in terms of the byte fractions. Web client traffic comprises only about 6-8% of overall bytes, which, as we will see in the next section, is only about one ninth of the Web-server byte proportion. This observation is explained by the nature of the application client-server behavior, which usually consists of a short HTTP request from the client followed by a significantly longer HTTP reply from the server [Stevens96]. Web client traffic shows a steady 1 kilobyte per flow and 14-16 packets per flow for an average of about 67 bytes per packet. Web client flows last 10-15 seconds on average. Web client traffic on the domestic link looks similar to the international Web client traffic, taking into account the observed 5-6 hour shift.

5.3.2 Web Server Traffic

The graphs in Figure 13 show international link Web server traffic. Web server traffic accounts for 55%-70% of overall byte traffic, with the minimum fraction occurring at or just after midnight East Coast time. The highest percentages happen during business hours on the U. S. East Coast. Web server packet and flow volumes are generally 30-35% of the totals on the link. The aggregate traffic graphs show a 3:1 ratio between the highest and lowest traffic levels over the 24-hour period for packets and bytes. Web server traffic averages 10-15 seconds per flow, 14-18 packets per flow, and 9-12 kilobytes per flow.

5.3.3 FTP Server Data Traffic

FTP data traffic on both measured links exhibits less of a daily pattern than Web traffic. FTP byte and packet volumes rise in the early morning U.S. East Coast time hours, indicating longer FTP flows during that time. FTP accounts for a higher percentage of the bytes and packets overnight than during the day. FTP packet fractions are 3% or less, the flow fractions are less than 1%, and the byte fractions are 2 to 8%. We see a wide range in average flow duration for FTP from 20 to 500 seconds. An average FTP data connection transfers approximately 200 kilobytes, varying according to the time of day.

5.3.4 NNTP Server Traffic

Network News Transfer Protocol (NNTP) traffic accounts for 1-4% of the bytes and packets on the East Coast link depending on the time of day, with the higher percentages occurring during overnight hours. The average number of packets per flow varies from 200 to 800, with flows generally lasting 100-200 seconds and with 50-300 kilobytes transferred per flow. NNTP flows on the international link comprise 0.1-2.5% of all packets and bytes, and a negligible percentage of flows. Flow averages in packets, bytes, and seconds exhibit high variability on the international link.

5.4 UDP Application Traffic

In this section, we present data on UDP application traffic. Specifically, we focus on DNS name server traffic, and RealAudio/RealVideo traffic.

5.4.1 DNS Traffic

The graphs in Figure 14 show UDP Domain Name Service (DNS) traffic on the domestic link. In analyzing DNS traffic, we aggregate client-to-server, server-to-client, and server-to-server traffic into a single category. We note that 90% of the DNS traffic measured is server-to-server, with client-to-server and server-to-client making up approximately equal shares at 5% each. Clients are typically configured to use a local nameserver/resolver, so the majority of DNS traffic on a backbone link is expected to be server-to-server. Figures 14a and 14b show a clear daily pattern to the DNS traffic on the domestic link that follows the link's overall utilization pattern. DNS traffic comprises only 2-5% of the packets and 1-2% of the bytes on the link, but as much as 15-25% of the flows. We note that server-to-client and client-to-server flows average closer to 1 packet in length, but that server-to-server flows average slightly longer, bringing up the overall average. The average flow consists of 2-3 packets with occasional bursts over 5. The average number of bytes per flow is steady at roughly 500. We see similar characteristics for the DNS traffic on the international link.

5.4.2 RealPlayer Traffic

The graphs in Figure 15 show RealPlayer traffic on the domestic link. The RealPlayer application uses a single TCP connection (on port 7070) for control traffic and UDP streams (on ports ranging from 6970 through 7170) for its audio and video traffic. Figures 15a and 15b show aggregate TCP and UDP RealPlayer traffic with a steady volume that averages about 1000 flows. There is a daily pattern to RealPlayer traffic, with higher packet and byte levels appearing during business hours. Figures 15c and 15d show negligible flow percentages and byte and packet percentages between 0.5% and 2.5%. Figures 15e and 15f show average RealPlayer traffic per UDP flow, without counting the TCP control connection. We removed the control connection to allow us to examine only the RealPlayer UDP data streams. These UDP flows transfer about 20 kilobytes on average.

5.5 Protocol Overhead Analysis

We presented detailed packet-size information in Figures 4 and 6. We use this exact packet-size information, with one-byte granularity, to compute accurate bit-efficiency estimates for ATM and SONET trunks carrying IP traffic. In this analysis, we use protocol overhead and efficiency formulas supplied by Tony Li of Juniper Systems. For ATM, each IP packet acquires 16 bytes of ATM overhead (8 bytes for LLC/SNAP and 8 bytes for AAL5 framing), is segmented into an integral number of 48-byte cells (with padding in the last cell, if necessary), and is prepended with a 5-byte ATM cell header. SONET framing includes a 7-byte per packet PPP header and additional SONET framing overhead of 90 bytes per 2340 bytes of payload. These formulas allow us to compute the total number of bytes that would result from IP-over-SONET and IP-over-ATM encapsulation for the given packet sizes. We use packet size information for 16,794,689,797 observed packets, which represent 5,777,615,952,291 bytes of IP payload for the underlying protocol, either ATM or SONET. In the case of ATM, we assume LLC/SNAP encapsulation and the use of AAL5 framing, which are in prevalent use on the Internet today.

The results, shown in Table 1, are similar to those reported by Tony Li for FIX-West packet traces [Li97]. We see that, as expected, SONET (at 94% bit-efficiency) is more efficient than ATM (at 80%). The ATM protocol overhead, about 20% of trunk capacity, must be balanced against advantages of ATM, including the capability to multiplex among service classes, support for multiple QoS requirements, and VC routing, in assessing ATM for carrying IP traffic versus IP/SONET or other alternatives such as frame relay. We note that a more efficient encapsulation, such as "VC Based Multiplexing" [Heinanen93], that allows 40-byte packets to be carried by a single ATM cell would significantly improve IP-over-ATM efficiency. VC Based Multiplexing would improve ATM's bit-efficiency to 84.47% for the packet size data presented here.

Table 1: ATM and SONET Bit-Efficiency

	ATM	SONET
IP payload bytes	5,777,615,952,291	5,777,615,952,291
Total bytes	7,233,116,621,345	6,121,916,426,350
Bit-Efficiency	80%	94%

5.6 Optimization Analysis

In this section, we perform an analysis of the flows statistics obtained from the domestic trunk that is intended to help evaluate schemes for optimizing large-volume flows. Examples of possible flow-optimization schemes include Ipsilon's IP switching [NLM96], Cisco's Tag Switching [RDKRS97], and other proposals currently being discussed by the IETF's Multiprotocol Label Switching working group. Some schemes, including Tag Switching, base the optimization of flows on IP routing information or some administrative procedure. Others, such as IP Switching, base the creation of optimized paths on observed flow behavior. Most commonly the flow-observing strategies set a threshold packet count as a trigger for flow optimization. The assumption in such a scheme is that most flows are composed of either a small packet count or a large packet count. The short flows, which would not benefit from an optimized path, never trigger any optimization. The long-lived flows make the strategy profitable through optimization of the remaining packets on the flow after the path set-up is done.

We sought to discover for current Internet traffic, how many flows would need to be optimized in order for a significant percentage of the total packets on a trunk to have their switching optimized. Figures 16 and 17 show histograms of flow duration plotted on logarithmic and linear-scales. We see that a large number of flows are short in duration (less than 50 packets) and that the flow duration distribution has a long tail. Figure 18 addresses the question of how many packets are optimized by setting the trigger threshold at a particular value. This graph shows on the *y*-axis the percentage of flows and the corresponding percentage of packets that are switched optimally. The *x*-axis shows the threshold value for making the decision to optimize the flow. We assume that 3 packet times are needed for set-up time to optimize a flow. We choose this number somewhat arbitrarily; however, we find that using other numbers close to this value yield similar results. In our model, the number of packets whose switching is optimized for a given flow is the (total packet count for the flow) - (threshold value + 3). Figure 18 shows that for our measured data, optimizing the switching of 50% of the packets would require optimization of about 20% of the flows. Optimizing 10% of the flows results in optimization of just over 40% of the packet switching. However, optimizing just the longest 5% of the flows still yields an advantage for more than 30% of the packets.

It is important to note that the above results are dependent on how a flow is defined. Previous studies have defined a flow based solely on IP source and destination addresses and a timeout value [NLM96], in contrast to our more restrictive definition of a flow, which includes protocol and port numbers. Further aggregation of traffic into flows could be examined by defining a flow based on source and destination network, CIDR block, or Autonomous System (AS) number. The coarse-grained flow definitions reduce the number of flows that must be optimized to achieve a given percentage of switching optimization. We chose the fine-grained definition of a flow including protocol IDs and port numbers because of our interest in adding support for multiple classes of service over the Internet. If different traffic classes are to receive service with different loss and delay characteristics, the service type should be selectable by the application, which knows its loss/delay/reliability requirements. If the network sees packets of different traffic classes as belonging to a single aggregated flow, this per-application service differentiation is precluded. Efficiency gains can be achieved with fewer flow optimizations if routers see best-effort flows aggregated to source/destination address, CIDR block, or AS number while seeing other QoS flows identified by our more restrictive definition. These results indicate the usefulness of multi-level flow definition for best-effort and other traffic types.

6 Summary

We made the following observations on two OC-3 links in MCI's commercial Internet backbone over 2 time ranges (24-hour and 7-day) in the presence of up to 240,000 flows. All data is based on the monitors reporting on all expired flows, which are uni-directional, on back-to-back 5-minute intervals.

- Wide-area traffic levels follow 24-hour patterns. Traffic on the measured domestic East Coast link increases between the hours of 5:00 and 10:00 by as much as 300-500% in packets, bytes, and flows. Traffic between the

U.S. and the U.K. is dominated by the U.S.-to-U.K. direction in packets and bytes, except for bytes transferred in evening hours U.S. EDT. We observe a 15-25% reduction in total traffic over weekends.

- Average packet sizes also vary over time following a clear 24-hour pattern on our measured backbone links. Directional packet-size asymmetry is evident in the international traffic measurements. Average packet sizes vary from 175 bytes to more than 400 bytes. The statistical importance of average packet size is questioned by evidence of strong modality in packet sizes. About 40% of all packets are 40 bytes (indicative of TCP acknowledgements, FINs, and RSTs); other modes occur at 44 bytes (5%), 552 bytes (5%), 576 bytes (6%), and 1500 bytes (10%). These modes represent TCP implementations using the common maximum segment sizes of 512, 536 and 1460 bytes (for Ethernet-attached hosts). Few packets are observed beyond 1500 bytes, and ninety percent of the packets are 576 bytes or smaller.
- Of IP traffic, TCP accounts for 95% or more of the bytes, 85-95% of the packets, and 75-85% of the flows. TCP flows average fewer than 20 packets, about 7 kilobytes, and under 20 seconds in duration. UDP makes up most of the remaining IP traffic and ICMP packets account for less than 1% of all packets.
- Web traffic dominates as the single largest Internet application, with client/server traffic accounting for more than half the bytes (65-80%), packets (55-75%), and flows (65-75%) seen on the measured links. Web server traffic averages 10 kilobytes/flow, 15 packets/flow and 13 seconds/flow. Web client traffic averages 1 kilobyte/flow, 15 packets/flow and 13 seconds/flow.
- Other major TCP applications, such as FTP-data and NNTP, make up a higher percentage of the overall traffic mix in the evening and overnight hours than during the daytime; however, they rarely exceed 10% of the total traffic.
- Of UDP applications, DNS traffic, the majority of which on the backbone is server-to-server communication, accounts for 1/3 to more than 1/2 of all UDP packets and bytes, depending on time-of-day. DNS flows consist of only a few packets under our restrictive definition of flow, and result in as much as 25% of all flows at their 5:00 peak as seen on the domestic link. An Internet audio/video service called RealPlayer was measured to exhibit steady flow counts over time with packet and byte burstiness, some of which is attributed to the flow expiration method of the monitor. RealPlayer traffic averages less than 1% of all bytes and packets.
- The OC3MON monitor reports on expired flows. Streams lasting more than 60 minutes are artificially expired, their traffic for the hour aggregated, and the flows reported at the next collection query. Results for IP-in-IP traffic, normally attributed to Mbone tunnels, suffer from this reporting method by indicating large packet spikes on one-hour intervals.
- Protocol overhead analysis of precise packet size data shows approximately 80% bit-efficiency for IP-over-ATM traffic using LLC/SNAP and AAL5 encapsulation for the mix of packet sizes observed. The same IP traffic over a hypothetical IP/SONET link would be about 94% bit-efficient.
- Flow optimization analysis, aimed at evaluating schemes for optimizing switching paths for long-lived flows shows that under our flow definition and assumption of a 3-packet set-up time, optimizing the switching of 50% of the packets would require optimization of about 20% of the flows. Optimizing 10% of the flows affects roughly 40% of the packets. This type of data is relevant to analysis of various IP-level switching schemes currently in design and implementation by router vendors and the IETF.

7 Future Work

The data we present in this paper raises a number of questions in need of further study. First, our data gathering method of reporting expired flows from the monitor suffers from a flaw. Long-lived flows are artificially-expired after one hour so they can be reported. This method creates artificial traffic spikes in the data that represent the flows that last longer than an hour. A data reporting mechanism that reports active flow information instead of data on expired flows would not suffer from this weakness. We plan to investigate alternative methods for reporting flow data to avoid misrepresenting long-lived flows.

We recognize that further study is required on the class of audio and video stream applications represented in this paper only by RealPlayer. Related to these products are the IP telephony applications, which together may change the future landscape of Internet traffic.

We mention long-term traffic trends in this paper only with respect to comparative numbers from previous studies. We can conduct longer-term analyses on the order of months and years in the future with our raw data from the OC3MON collectors. In addition, we plan to deploy additional monitors on other links in the backbone, as well as the OC12 version of the monitor once it is completed.

Internet statistics based on averages have been called into question here and in other studies. We are currently limited to reporting averages for flow characteristics because of the way the monitor aggregates the flow statistics per protocol. Reporting higher moments or distribution information would require the monitor to maintain per-flow statistics for quantities such as flow duration, and flow volume in packets and bytes. Because we are anxious to report more illuminating statistics, we intend to modify the monitor to report statistics on individual flows; however, we anticipate a challenge in maintaining state on hundreds of thousands of simultaneous flows. We expect that a statistical sampling technique may facilitate this endeavor.

Finally, we are interested in using the OC3MON's capabilities to conduct more detailed analyses into areas such as TCP implementation behavior, routing dynamics, and web-caching.

8 Monitor Availability

The MCI-NLANR collaboration, dubbed Coral, is continuing in two directions as an ongoing effort to extend OC3MON's usefulness: MCI is pursuing higher speed capability in the form of OC12MON, and NLANR is supporting broader user requirements. An online, web-based query engine for real-time vBNS OC3MON data can be found at http://www.vbns.net. The latest release of the OC3MON software source code is available via FTP from http://www.nlanr.net/Coral.

Acknowledgements

The authors would like to thank Kim Claffy, Greg Minshall, and Saib Jarrar for their time and effort in reviewing early drafts of this paper and for providing invaluable feedback that improved all sections of the paper. We also are grateful to the anonymous reviewers for their helpful suggestions and insightful comments. Finally, we would like to thank our colleagues at MCI who contributed to this work, including Roy Alcala, Joel Apisdorf, Jim Boyle, Scott Huddle, Angela Roote, and Jack Waters.

Biographies

Kevin Thompson is a Senior Engineer in the vBNS Engineering Group at MCI. He supports statistics collection architecture and implementation for the vBNS. He was employed as an engineer at the MITRE Corporation in the Networking Center until 1995. He received a B.S. in Computer Science from the University of Virginia in 1987 and an M.S. in Computer Science from the George Washington University in 1992.

Gregory J. Miller has been a Senior Engineer in the vBNS Engineering Group at MCI since September 1996. His focus is on network performance measurement, traffic analysis, and IP and ATM Quality of Service mechanisms. Before joining MCI, he was a Senior Member of the Technical Staff at the MITRE Corporation. He received a B.S. degree from Loyola College in 1988, and the M.S. and Ph.D. degrees from the University of Delaware in 1990 and 1993, all in computer science.

Rick Wilder is the lead engineer for the vBNS project. He was also a member of the original design team for the internetMCI. Prior to MCI, Rick worked at Advanced Network and Services where he did IP/ATM performance studies and development of an IP security product. Rick was also a lead engineer at the MITRE Corporation where he prototyped several communications systems and performed measurement studies on congestion avoidance in connectionless packet networks. He is currently active in the implementation of differentiated IP services for advanced applications. He received an MS degree in Computer Science from the American University.

References

[ACTW97] J. Apisdorf, K. Claffy, K. Thompson, and R. Wilder, "OC3MON: Flexible, Affordable, High-Performance Statistics Collection," *Proceedings of INET '97,* Kuala Lumpur, Malaysia, June 1997.

[BC94] H. W. Braun and K. Claffy, “Web Traffic Characterization: an assessment of the impact of caching documents from NCSA’s web server,” *Second International World Wide Web (WWW) Conference ’94,* Chicagao, IL, October 1994 (http://www.nlanr.net/Papers/wtc.html).

[Caceres89] R. Caceres, “Measurements of Wide-Area Internet Traffic,” UCB/CSD 89/550, University of California, Berkeley, CA, December 1989.

[CDJM91] R. Caceres, P. Danzig, S. Jamin, D. Mitzel, “Characteristics of Wide-Area TCP/IP Conversations,” *Proceedings of ACM SIGCOMM ’91,* September 1991.

[cache] http://ircache.nlanr.net/Cache.

[CPB93] K. Claffy, G. Polyzos, and H-W. Braun, “Traffic Characteristics of the T1 NSFNET Backbone,” *Proceedings of INFOCOM ’93*, San Francisco, CA, March 1993.

[Claffy96] K. Claffy, “Internet Workload Characterization,” Ph.D. Dissertation, University of California, San Diego, CA, June 1994.

[FIXWEST] http://www.nlanr.net/NA/FIX/Stats/West/index.html.

[Frazer95] K. D. Frazer, “NSFNET: A Partnership for High-Speed Networking, Final Report 1987-1995,” Merit Network, Inc., 1995.

[Heimlich89] H. Heimlich, “Traffic Characterization of the NSFNET Backbone,” *USENIX Conference Proceedings,* Winter 1989.

[Heinanen93] J. Heinanen, “Multiprotocol Encapsulation over ATM Adaptation Layer 5,” Request for Comments 1483, IETF, July 1997.

[Kleinrock76] L. Keinrock, *Queueing Systems, Volume II: Computer Appplications,* John Wiley and Sons, 1976.

[Li97] T. Li, “IP/ATM Efficiency,” message to the ION mailing list, May 1997.

[MH97] D. McRobb and J, Hawkinson, “cflowd: a Cisco flow-export collector,” http://figaro.ans.net/cflowd, 1997.

[NLM96] P. Newman, T. Lyon, G, Minshall, “Flow Labeled IP: A Connectionless Approach to ATM,” *Proceedings of IEEE INFOCOM ’96,* March 1996.

[Paxson94] V. Paxson, “Growth Trends in Wide-Area TCP Connections,” *IEEE Network,* 8(4), pp. 8-17, July/August 1994.

[Paxson97] V. Paxson, “Measurements and Analysis of End-to-End Internet Dynamics,” Ph.D. Thesis, University of California, Berkeley, CA, April 1997.

[RDKRS97] Y. Rekhter, B. Davie, D. Katz, E. Rosen, G. Swallow, “Cisco Systems’ Tag Switching Architecture Overview,” Request for Comments 2105, IETF, February 1997.

[stats] http://www.nlanr.net/INFO

[Stevens94] W. R. Stevens, *TCP/IP Illustrated*, *Volume 1: The Protocols*, Addison-Wesley, Reading, MA, 1994.

[Stevens96] W. R. Stevens, *TCP/IP Illustrated*, *Volume 3*, *TCP for Transactions, HTTP, NNTP, and the UNIX Domain Protocols,* Addison-Wesley, Reading, MA, 1996.

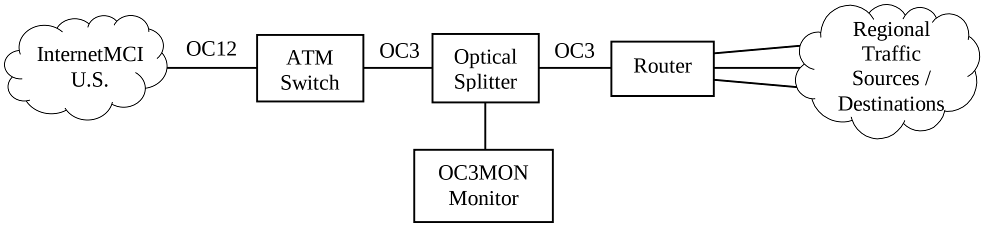

Figure 1 - Domestic Link Measurement Point

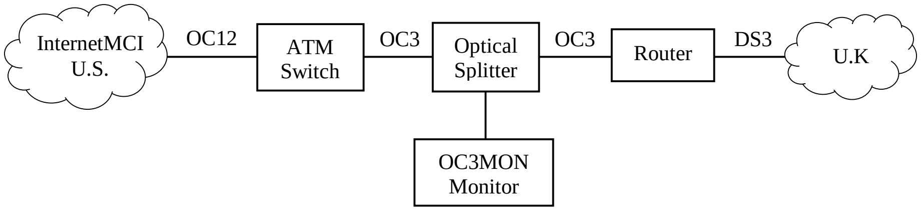

Figure 2 - U.S. - U.K. -International Link Measurement Point

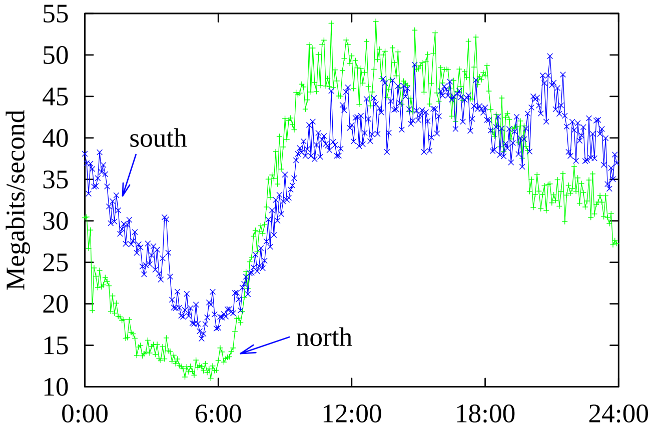

Figure 3a: Byte volume for 1 day on domestic link.

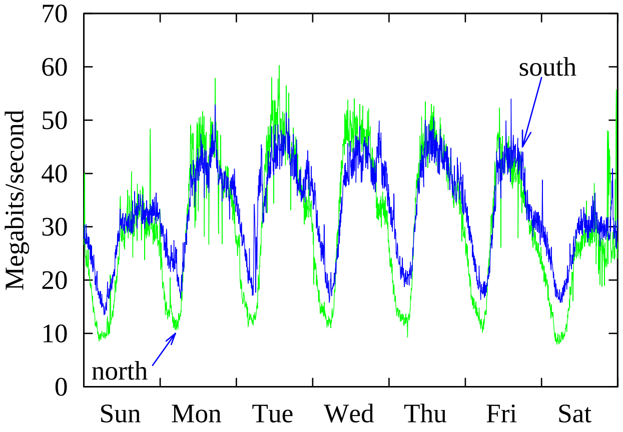

Figure 3b: Byte volume for 7 days on domestic link.

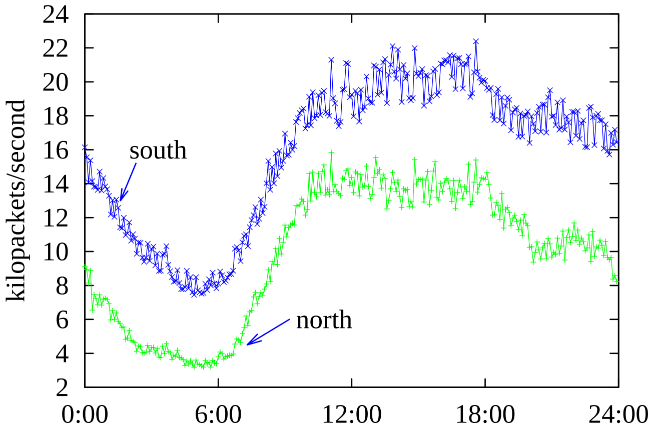

Figure 3c: Packet volume for 1 day on domestic link.

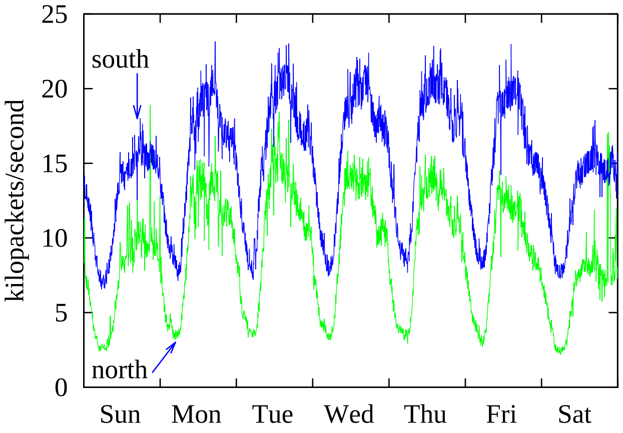

Figure 3d: Packet volume for 7 days on domestic link.

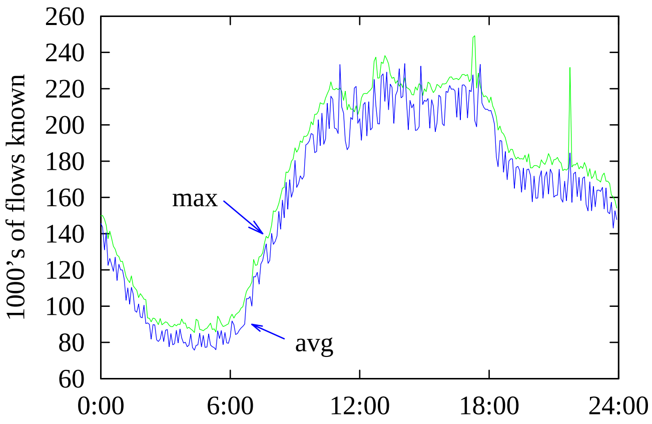

Figure 3e: Flow volume for 1 day on domestic link.

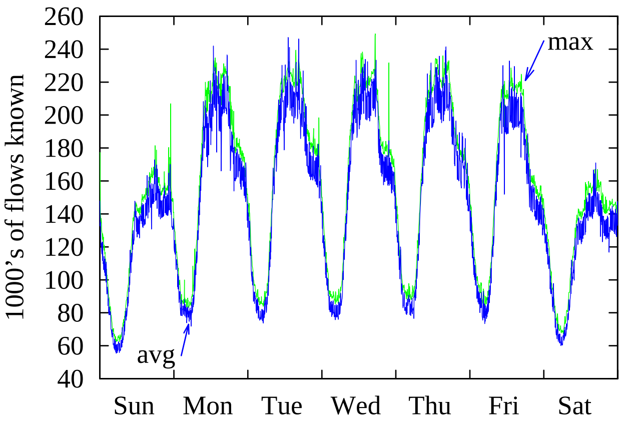

Figure 3f: Flow volume for 7 days on domestic link.

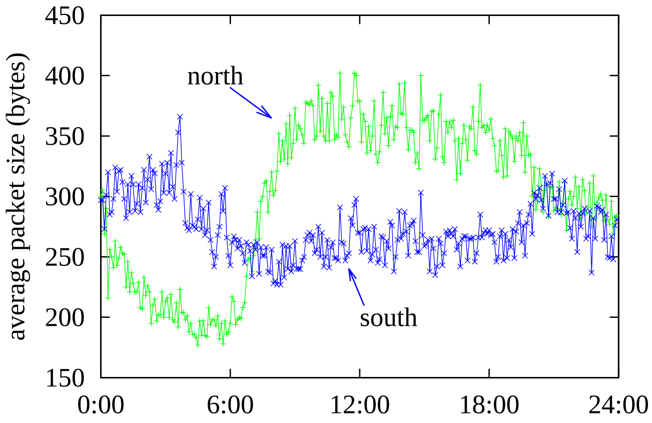

Figure 4a: Average packet size for 1 day on domestic link.

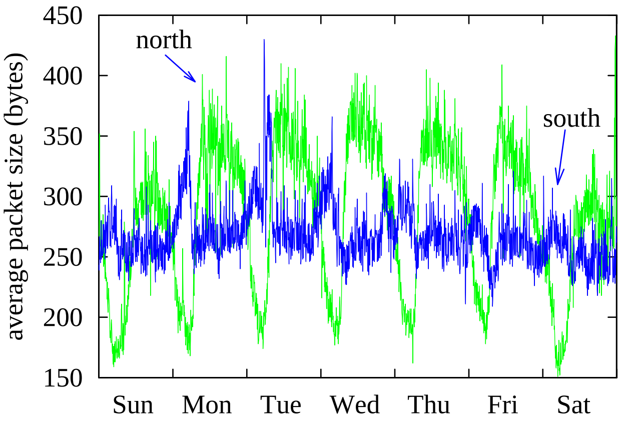

Figure 4b: Average packet size for 7 days on domestic link.

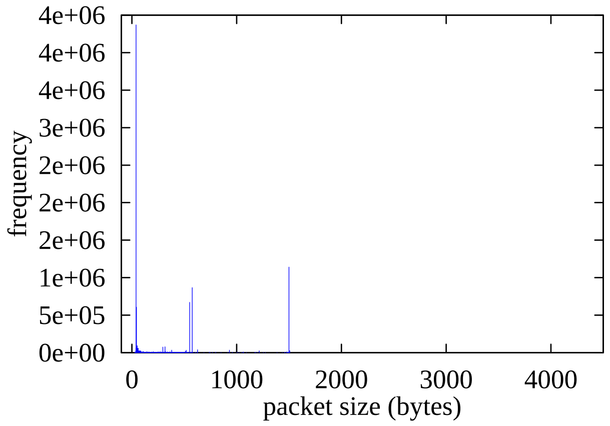

Figure 4c: Packet size histogram from a sample on domestic link.

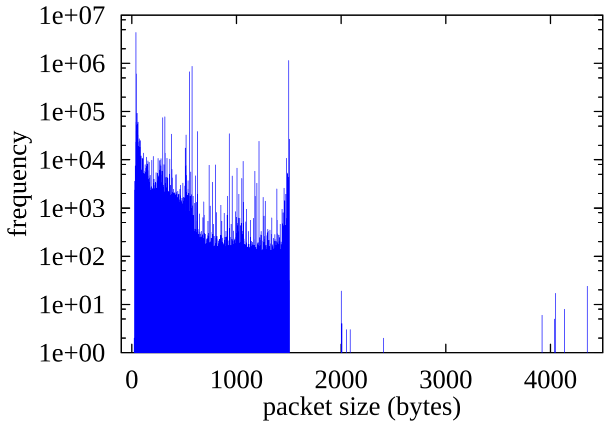

Figure 4d: Logarithmic-scale packet size histogram from a sample on domestic link.

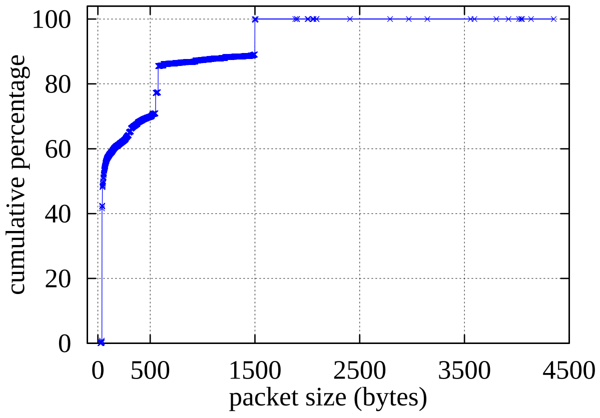

Figure 4e: Cumulative packet size distribution from a sample on domestic link.

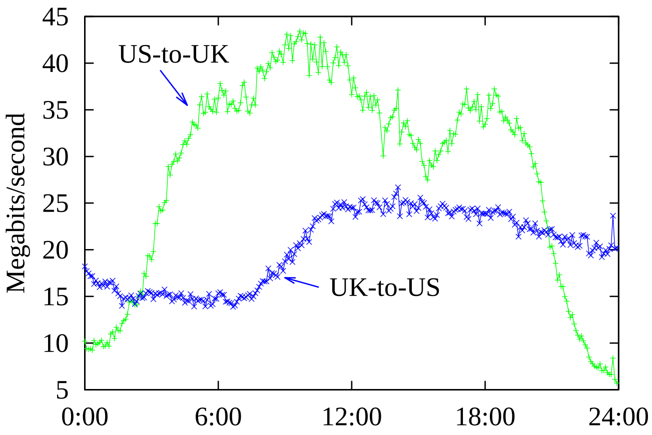

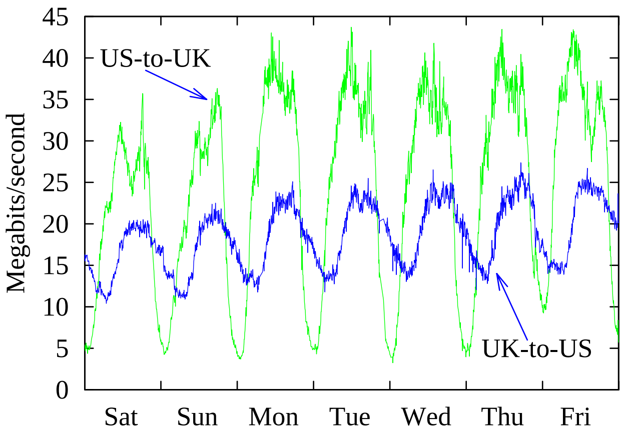

Figure 5a: Byte volume for 1 day on international link.

Figure 5b: Byte volume for 7 days on international link.

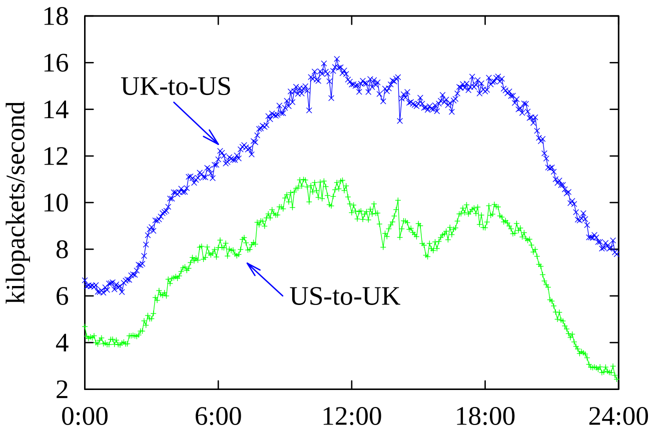

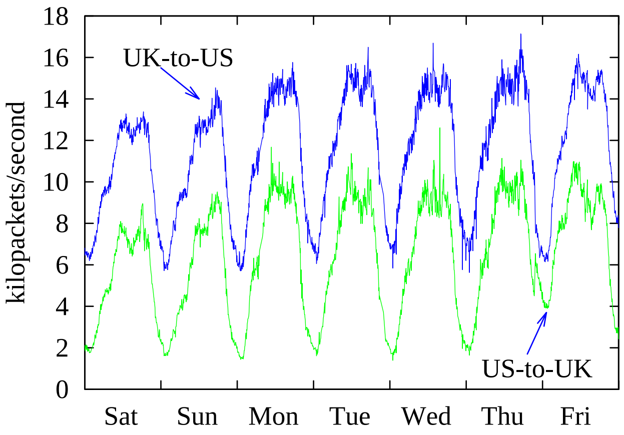

Figure 5c: Packet volume for 1 day on international link.

Figure 5d: Packet volume for 7 days on international link.

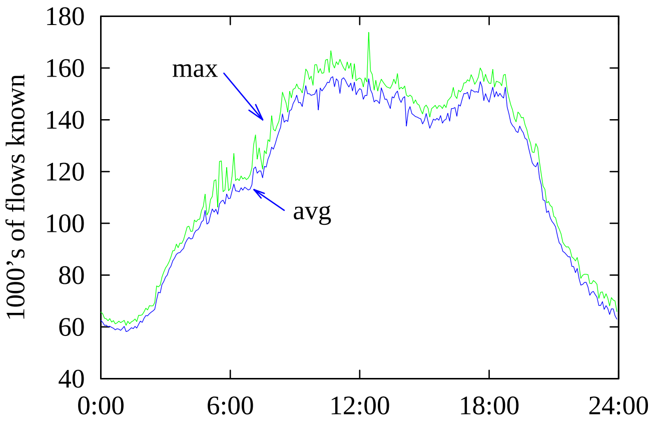

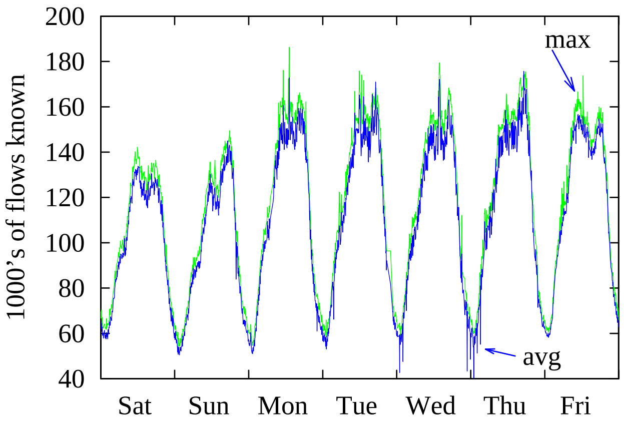

Figure 5e: Known flows for 1 day on international link.

Figure 5f: Known flows for 7 days on international link.

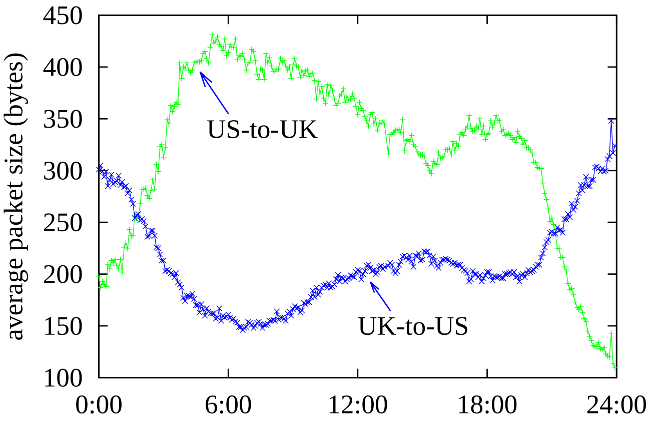

Figure 6a: Average packet size for 1 day on international link.

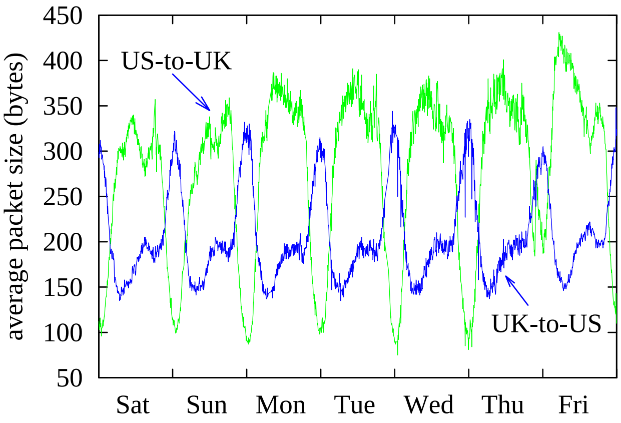

Figure 6b: Average packet size for 7 days on international link.

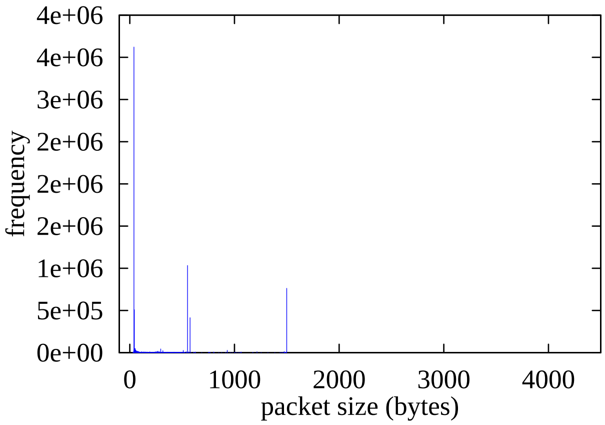

Figure 6c: Packet size histogram from sample on international link.

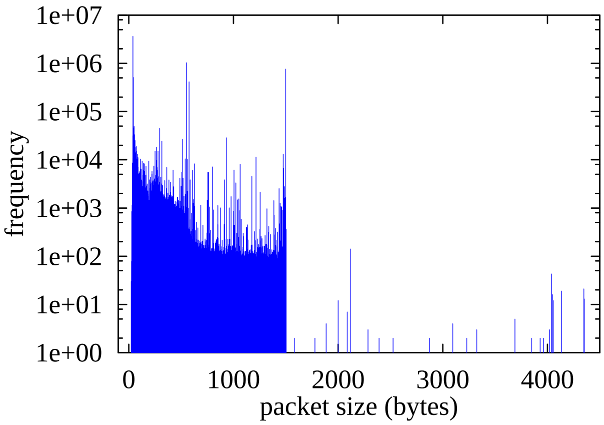

Figure 6d: Logarithmic-scale packet size histogram from sample on international link.

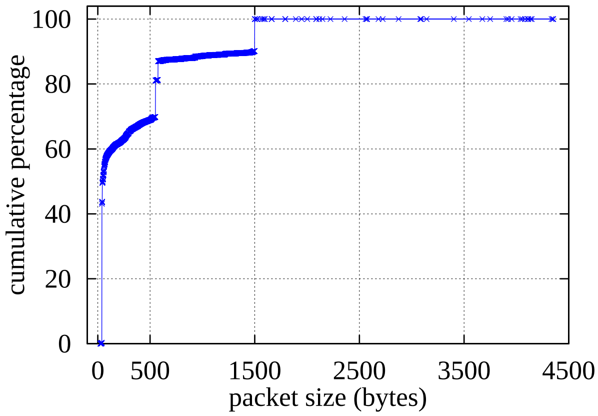

Figure 6e: Cumulative packet size distribution from a sample on international link.

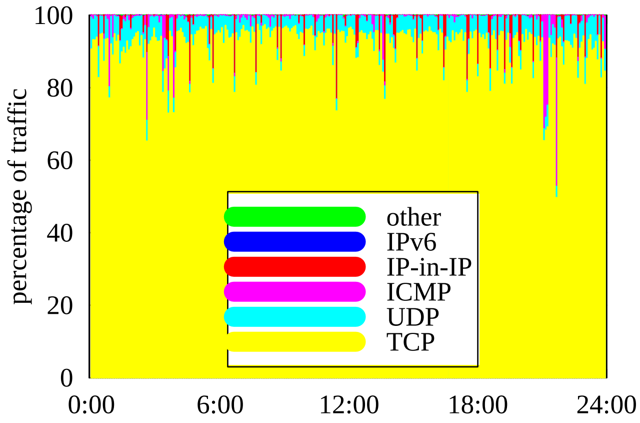

Figure 7a: Traffic composition in bytes by IP protocol for 1 day on domestic link.

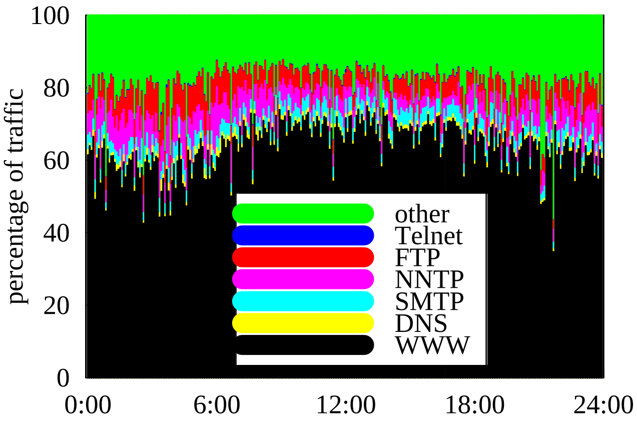

Figure 7b: Traffic composition in bytes by application for 1 day on domestic link.

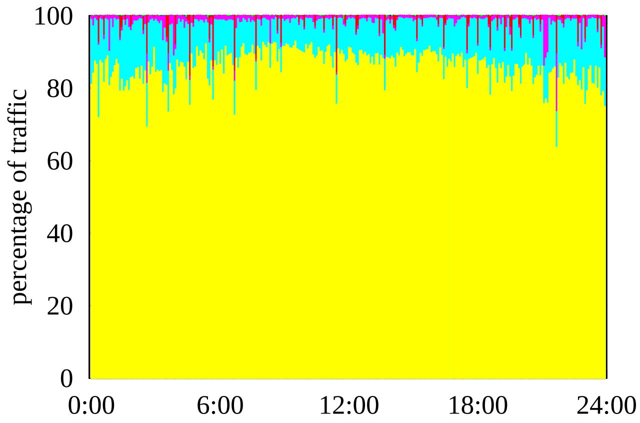

Figure 7c: Traffic composition in packets by IP protocol for 1 day on domestic link.

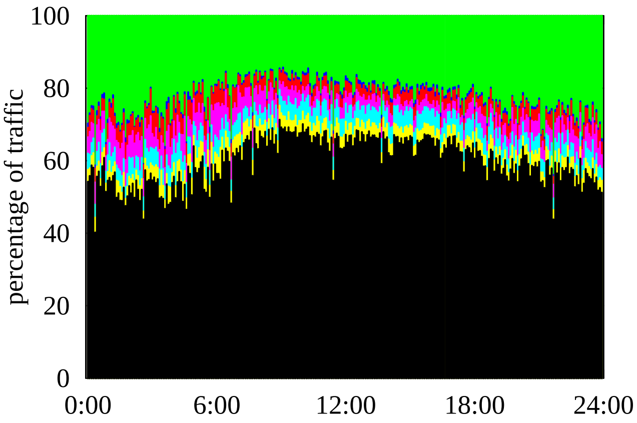

Figure 7d: Traffic composition in packets by application for 1 day on domestic link.

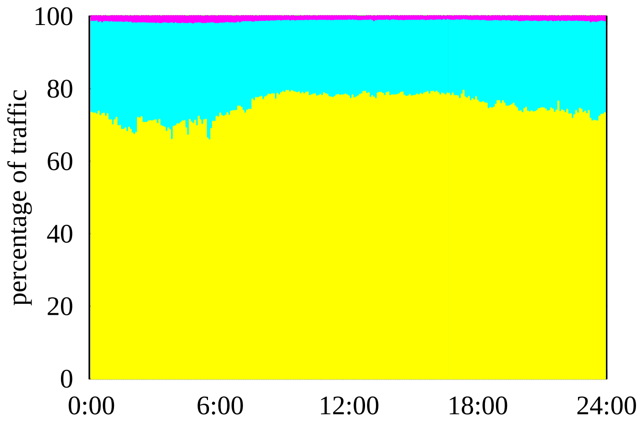

Figure 7e: Traffic composition in flows by IP protocol for 1 day on domestic link.

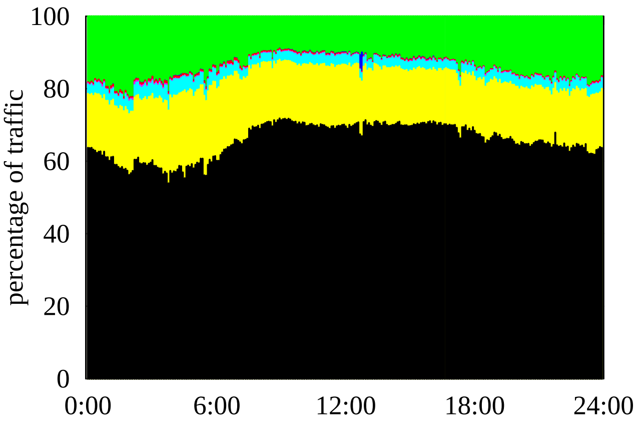

Figure 7f: Traffic composition in flows by application for 1 day on domestic link.

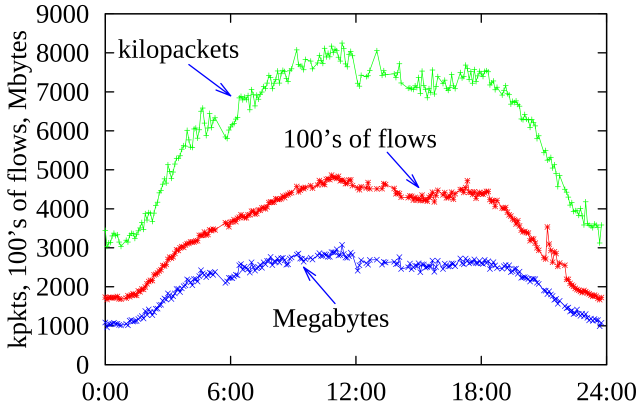

Figure 8a: Aggregate TCP traffic for 1 day on international link.

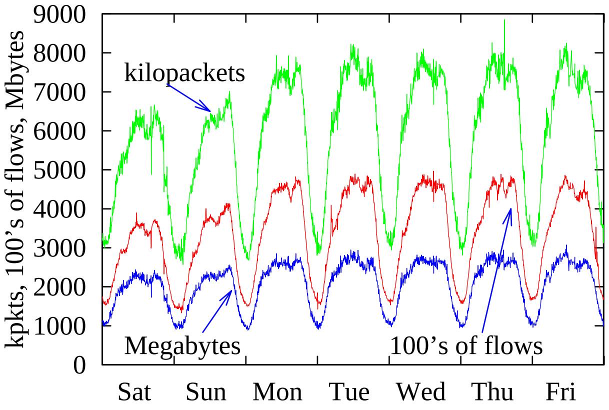

Figure 8b: Aggregate TCP traffic for 7 days on international link.

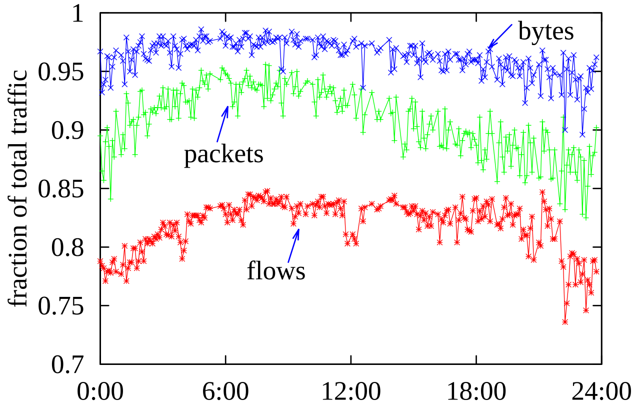

Figure 8c: Fraction of traffic offered by TCP for 1 day on international link.

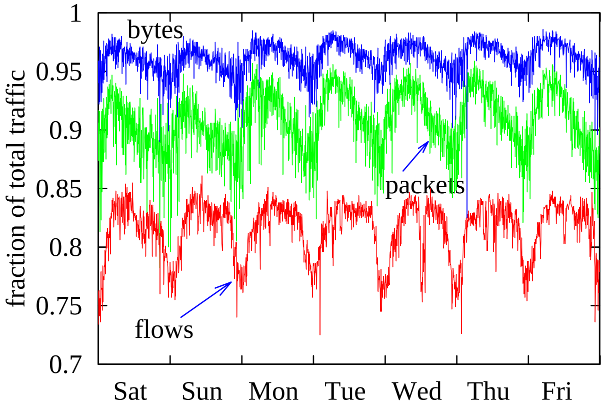

Figure 8d: Fraction of traffic offered by TCP for 7 days on international link.

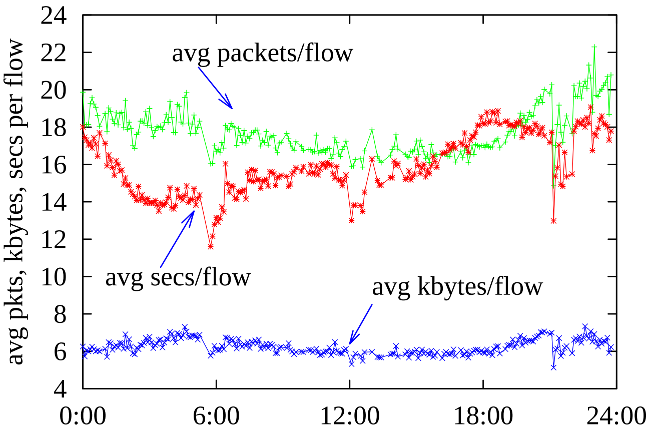

Figure 8e: Average traffic per TCP flow for 1 day on international link.

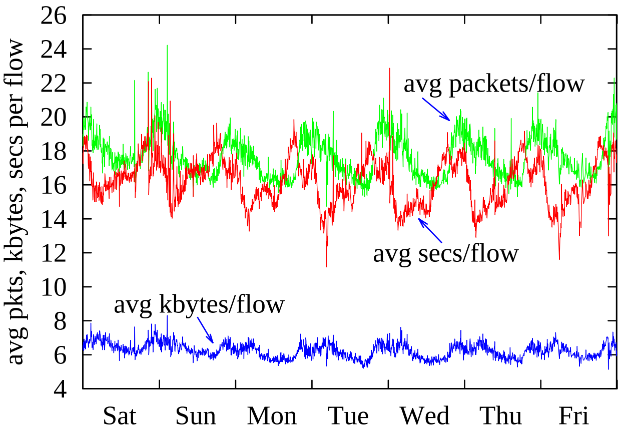

Figure 8f: Average traffic per TCP flow for 7 days on international link.

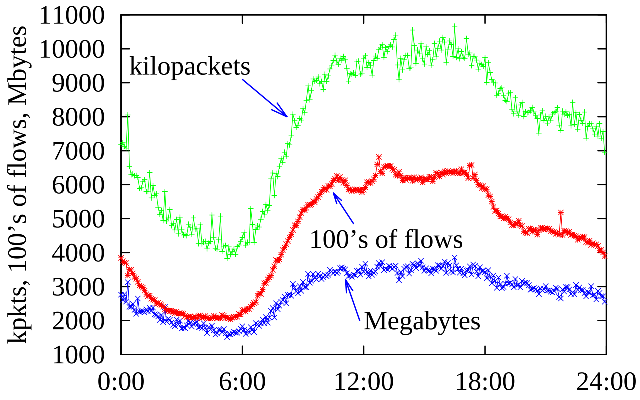

Figure 9a: Aggregate TCP traffic for 1 day on domestic link.

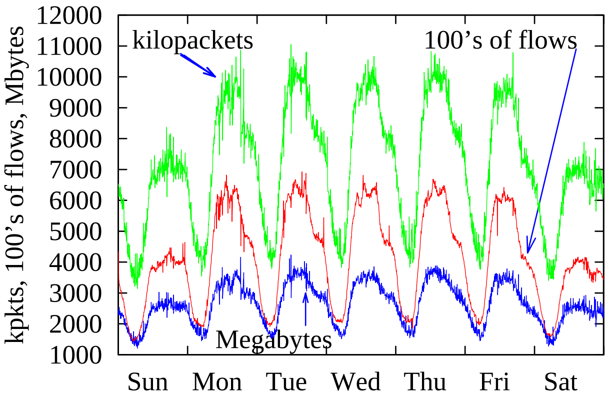

Figure 9b: Aggregate TCP traffic for 7 days on domestic link.

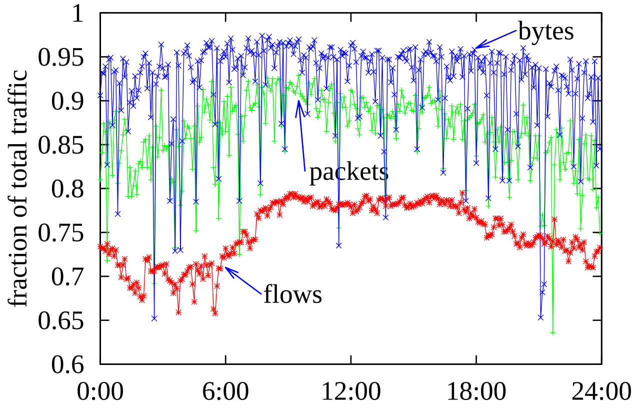

Figure 9c: Fraction of traffic offered by TCP for 1 day on domestic link.

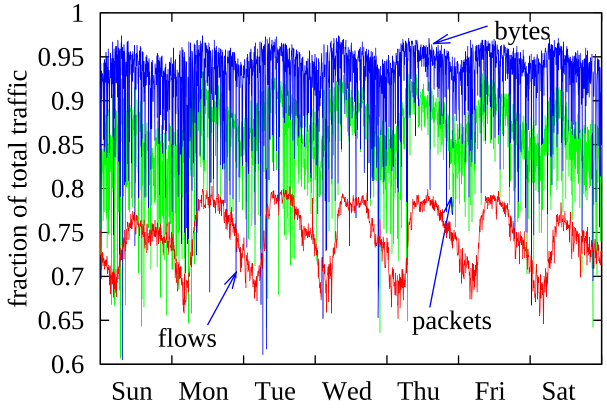

Figure 9d: Fraction of traffic offered by TCP for 7 days on domestic link.

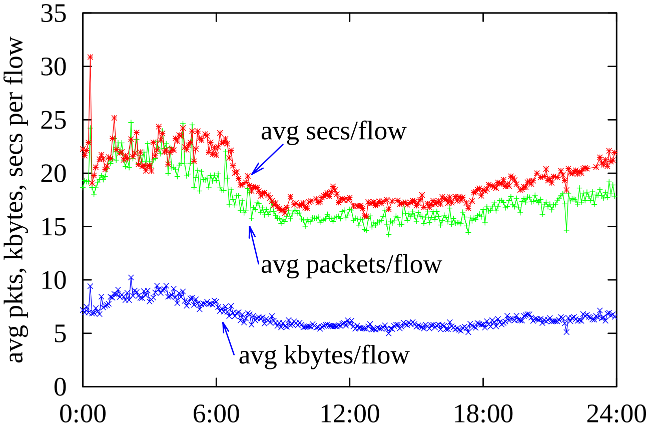

Figure 9e: Average traffic per TCP flow for 1 day on domestic link.

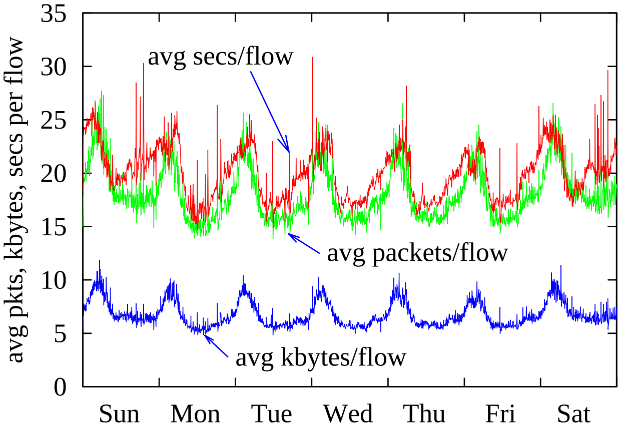

Figure 9f: Average traffic per TCP flow for 7 days on domestic link.

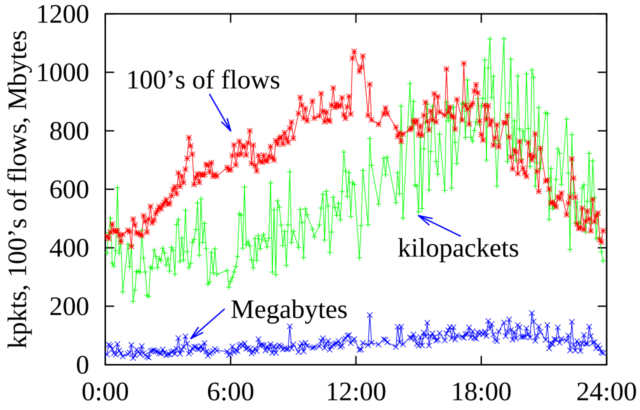

Figure 10a: Aggregate UDP traffic for 1 day on international link.

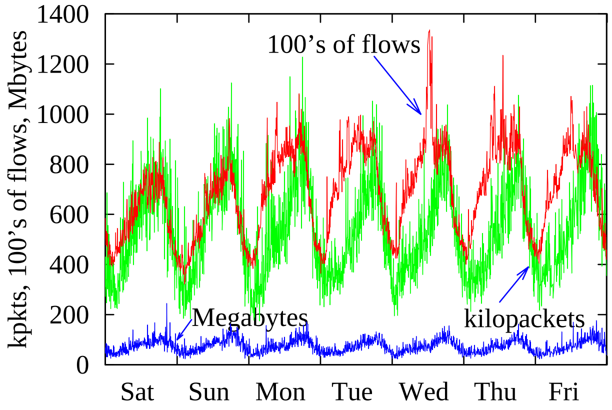

Figure 10b: Aggregate UDP traffic for 7 days on international link.

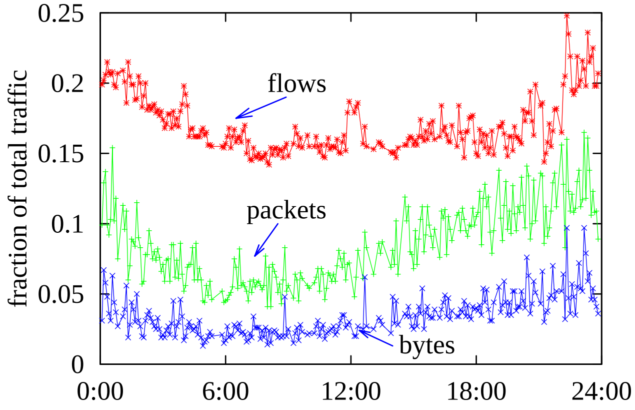

Figure 10c: Fraction of traffic offered by UDP for 1 day on international link.

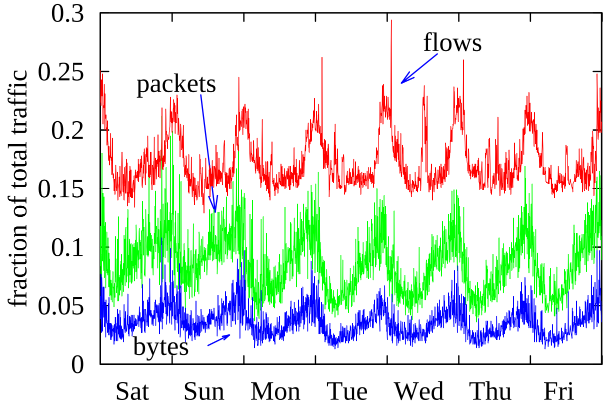

Figure 10d: Fraction of traffic offered by UDP for 7 days on international link.

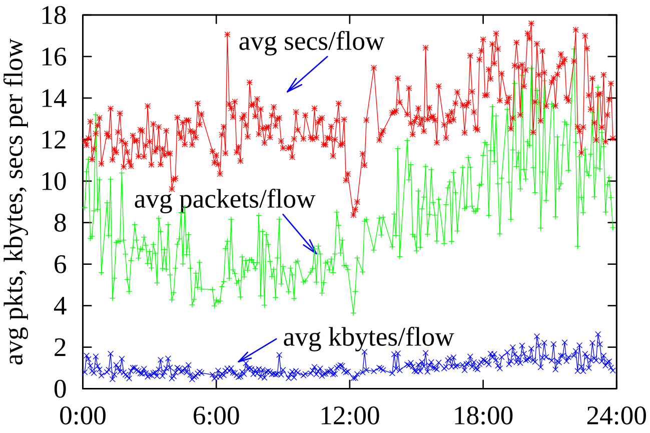

Figure 10e: Average traffic per UDP flow for 1 day on international link.

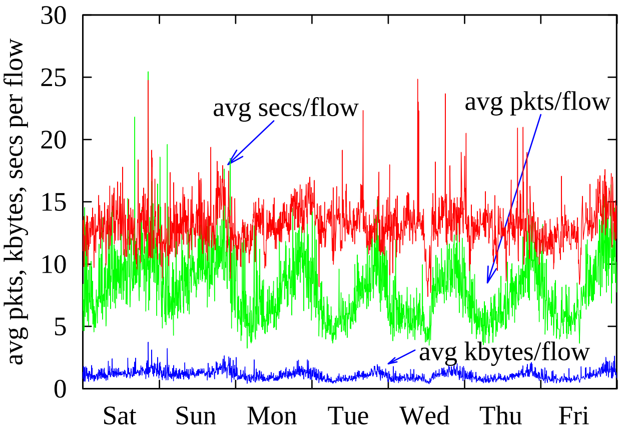

Figure 10f: Average traffic per UDP flow for 7 days on international link.

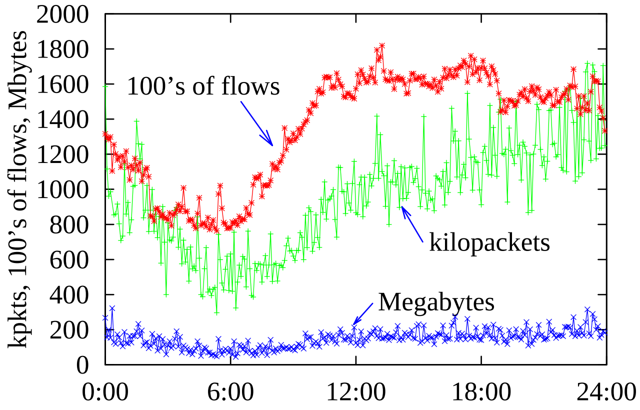

Figure 11a: Aggregate UDP traffic for 1 day on domestic link.

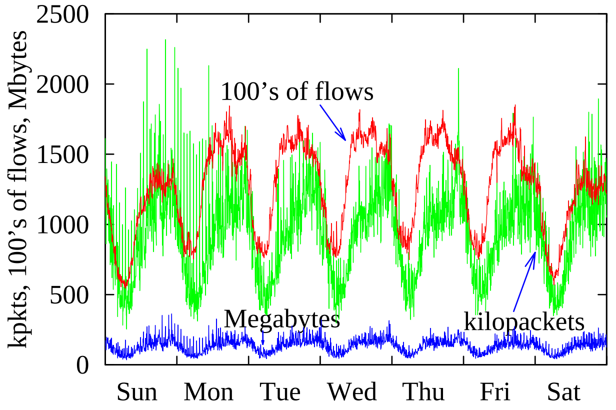

Figure 11b: Aggregate UDP traffic for 7 days on domestic link.

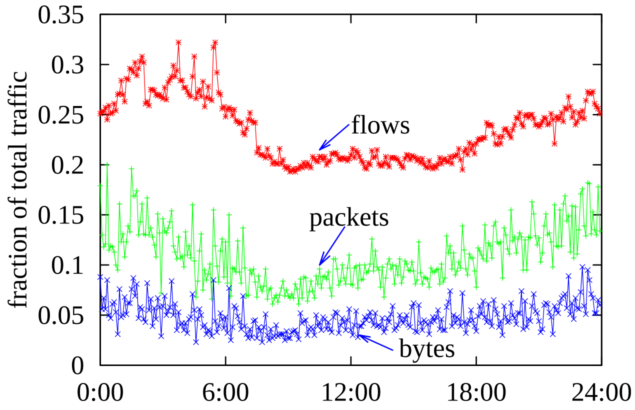

Figure 11c: Fraction of traffic offered by UDP for 1 day on domestic link.

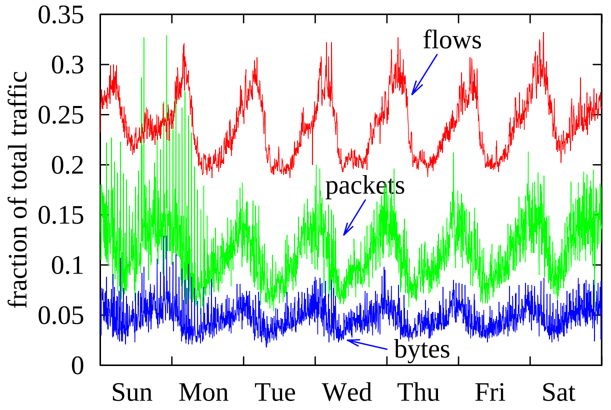

Figure 11d: Fraction of traffic offered by UDP for 7 days on domestic link.

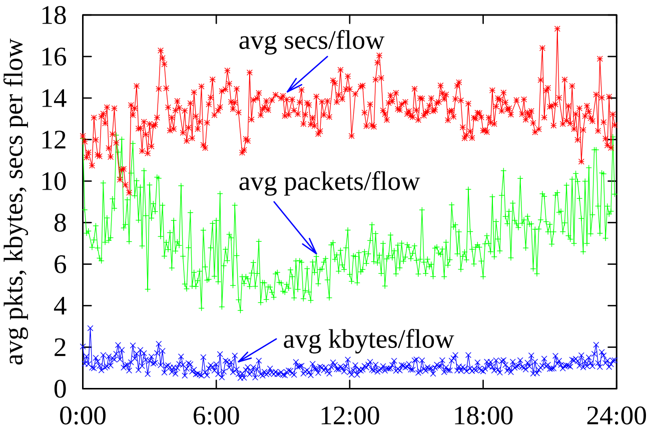

Figure 11e: Average traffic per UDP flow for 1 day on domestic link.

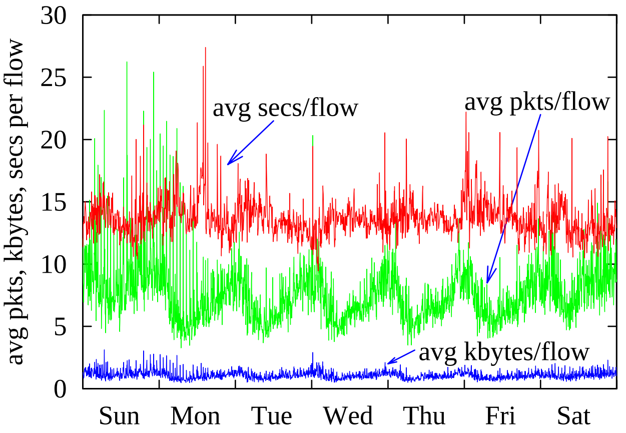

Figure 11f: Average traffic per UDP flow for 7 days on domestic link.

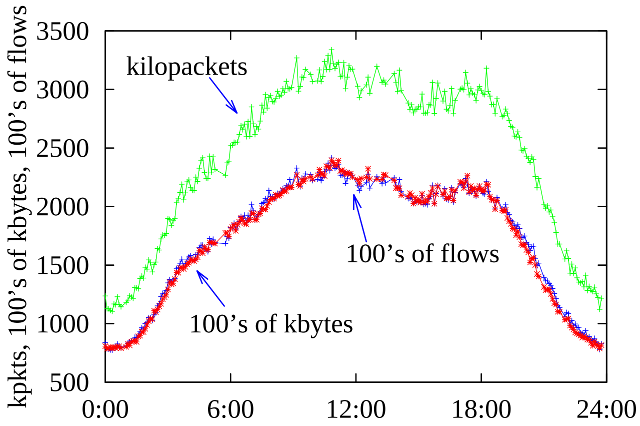

Figure 12a: Aggregate Web client traffic for 1 day on international link.

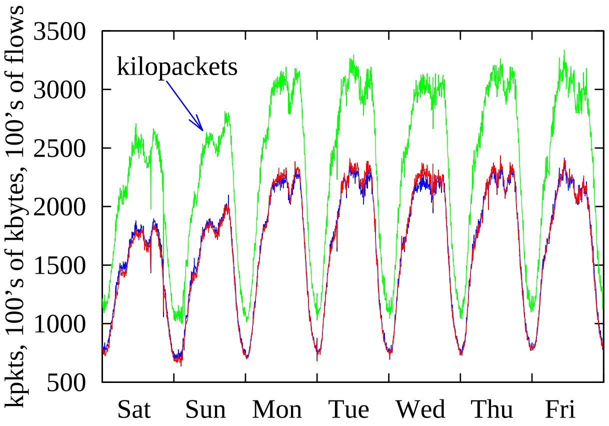

Figure 12b: Aggregate Web client traffic for 7 days on international link.

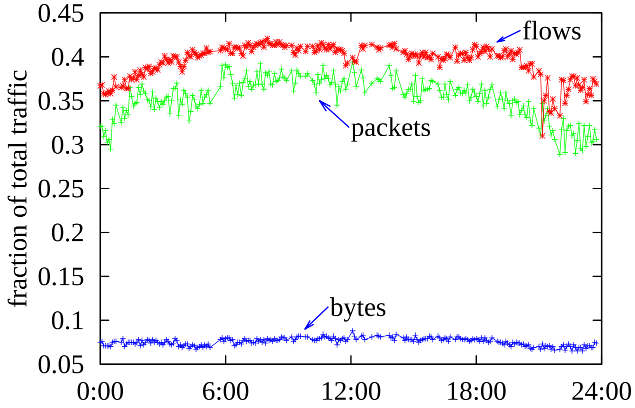

Figure 12c: Fraction of traffic offered by Web clients for 1 day on international link.

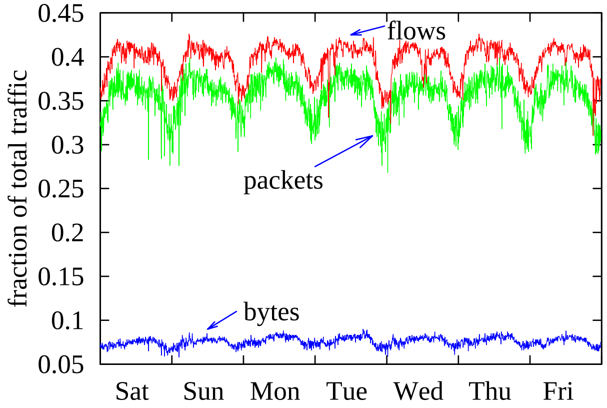

Figure 12d: Fraction of traffic offered by Web clients for 7 days on international link.

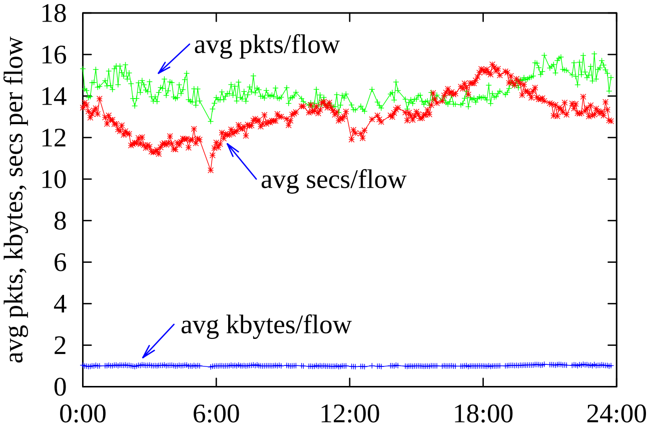

Figure 12e: Average traffic per Web client flow for for 1 day on international link.

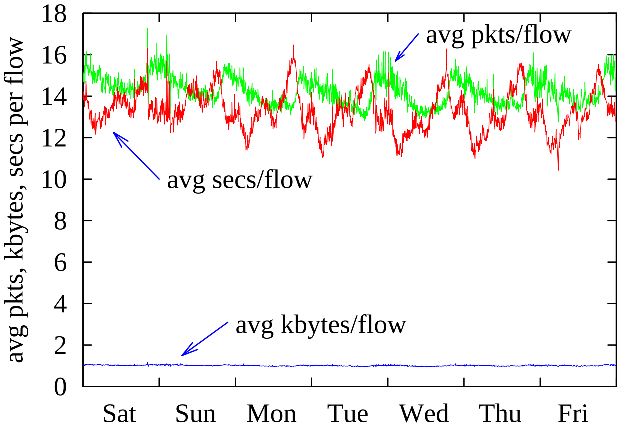

Figure 12f: Average traffic per Web client flow 7 days on international link.

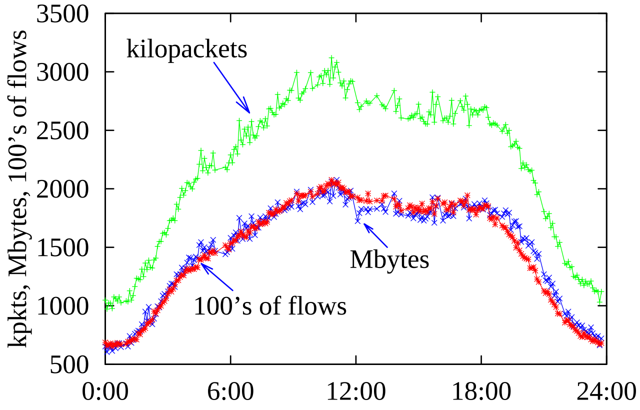

Figure 13a: Aggregate Web server traffic for 1 day on international link.

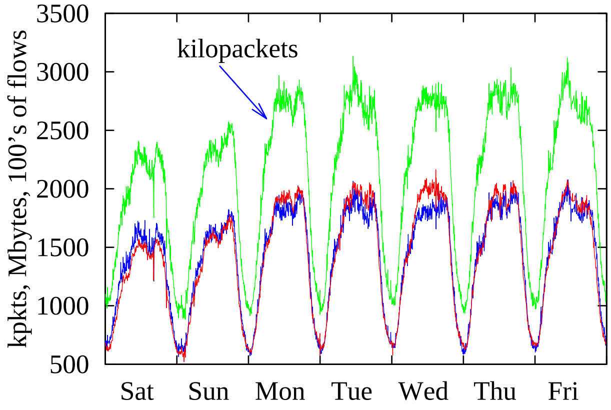

Figure 13b: Aggregate Web server traffic for 7 days on international link.

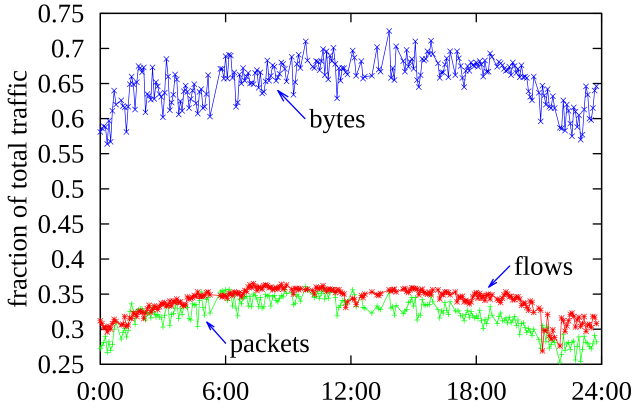

Figure 13c: Fraction of traffic offered by Web servers for 1 day on international link.

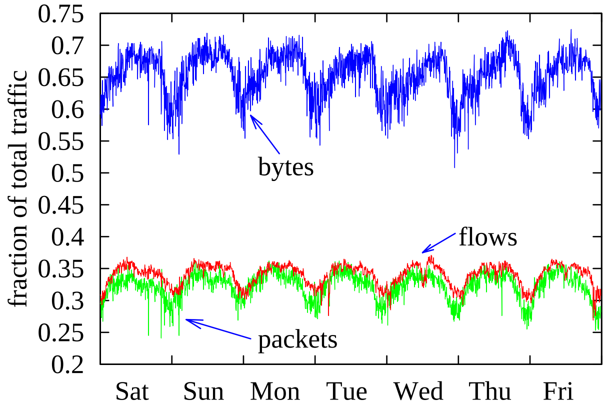

Figure 13d: Fraction of traffic offered by Web servers for 7 days on international link.

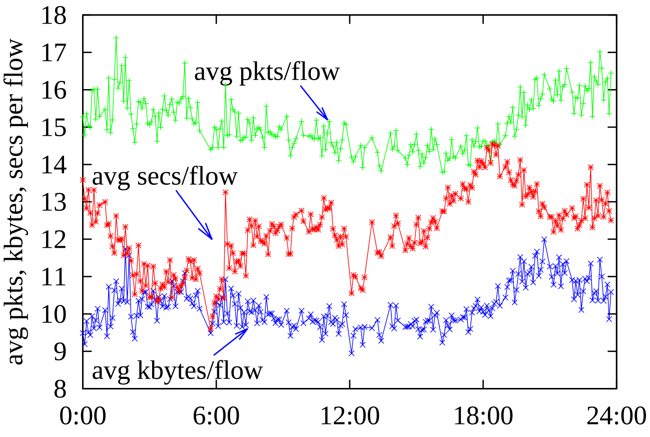

Figure 13e: Average traffic per Web server flow for 1 day on international link.

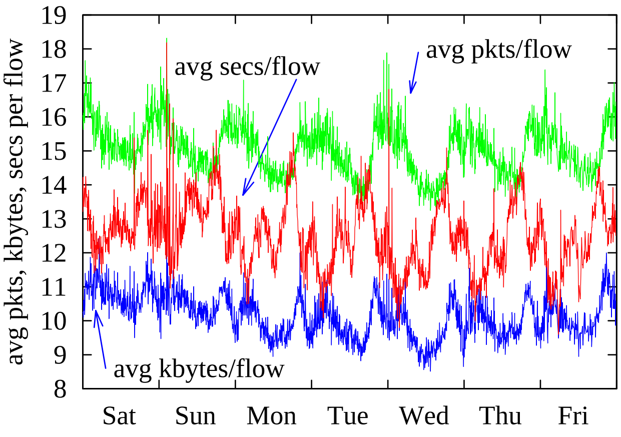

Figure 13f: Average traffic per Web server flow for 7 days on international link.

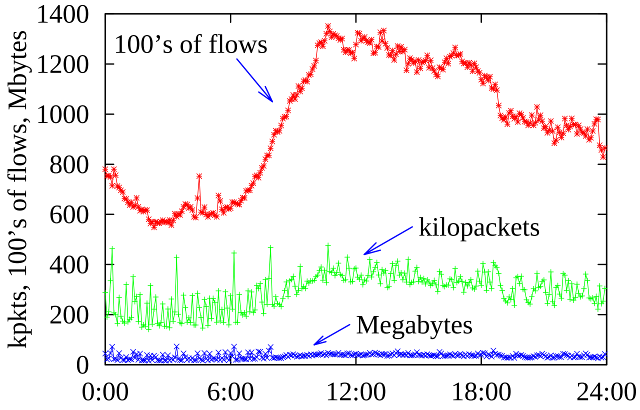

Figure 14a: Aggregate DNS traffic for 1 day on domestic link.

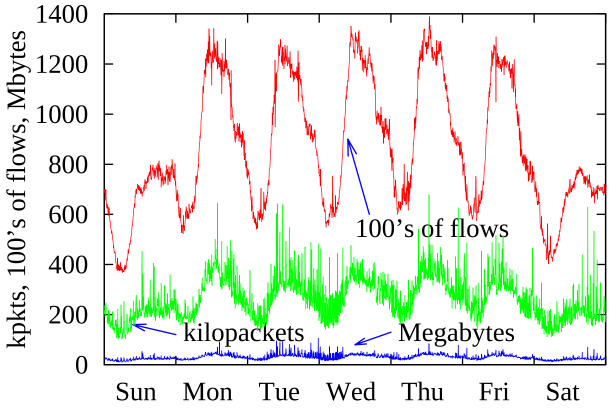

Figure 14b: Aggregate DNS traffic for 7 days on domestic link.

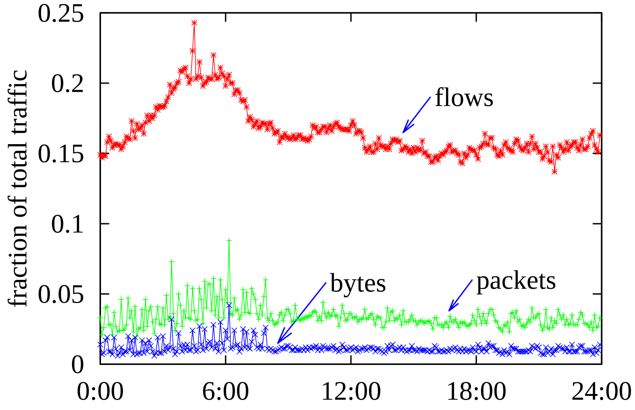

Figure 14c: Fraction of traffic offered by DNS for 1 day on domestic link.

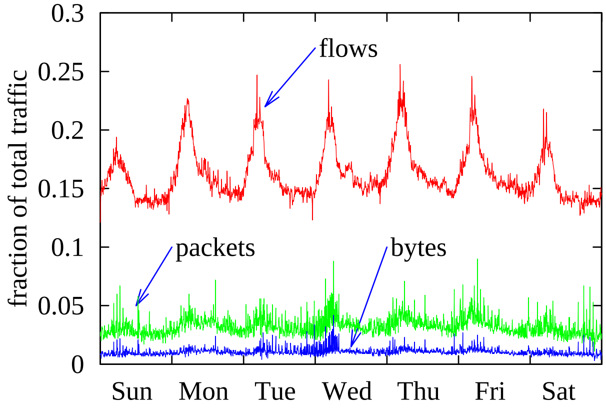

Figure 14d: Fraction of traffic offered by DNS for 7 days on domestic link.

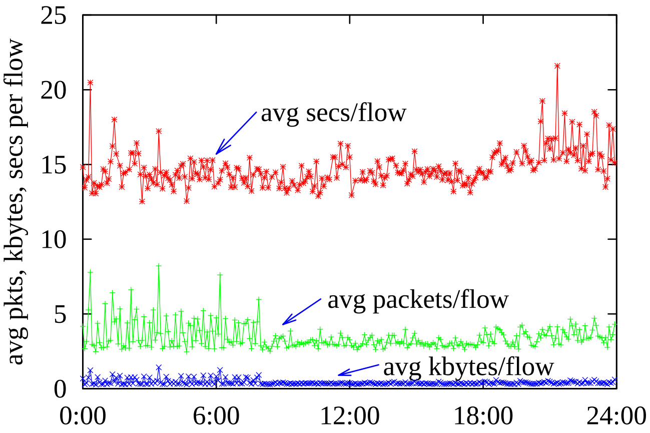

Figure 14e: Average traffic per DNS flow for 1 day on domestic link.

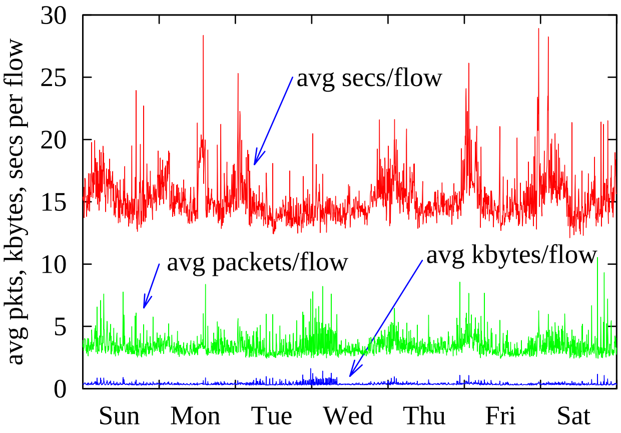

Figure 14f: Average traffic per DNS flow for 7 days on domestic link.

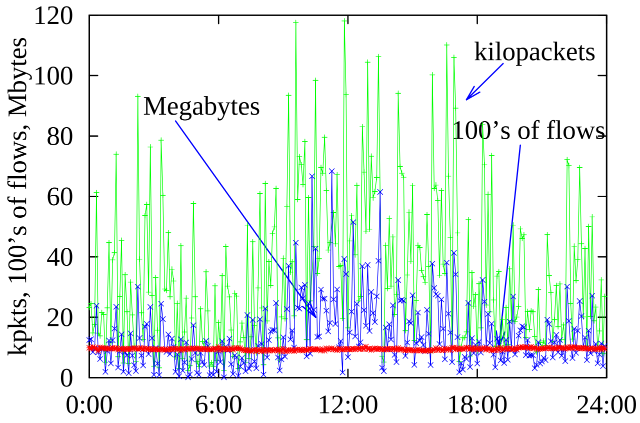

Figure 15a: Aggregate RealPlayer traffic for 1 day on domestic link.

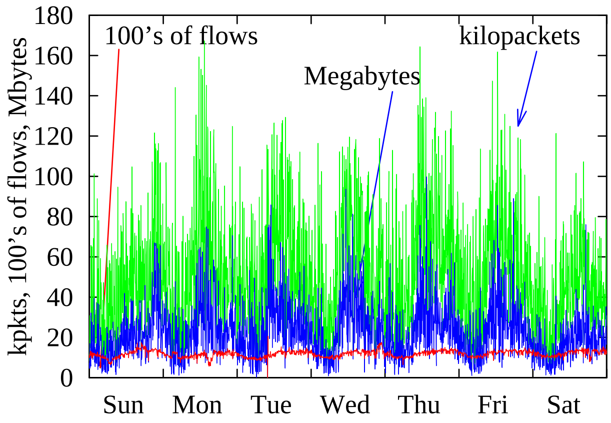

Figure 15b: Aggregate RealPlayer traffic for 7 days on domestic link.

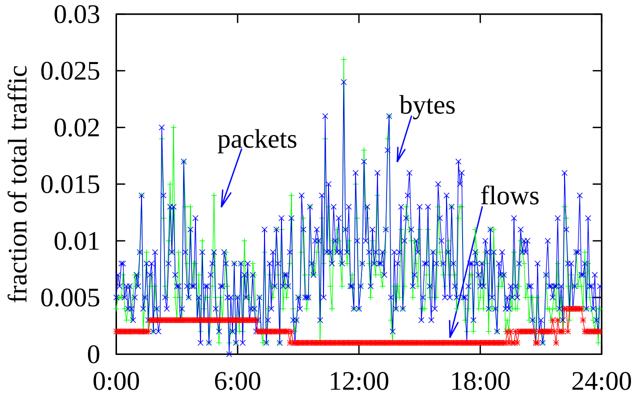

Figure 15c: Fraction of traffic offered by RealPlayer for 1 day on domestic line.

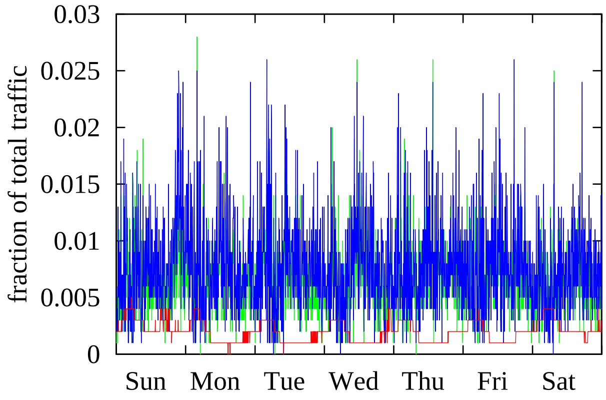

Figure 15d: Fraction of traffic offered by RealPlayer for 7 days on domestic link.

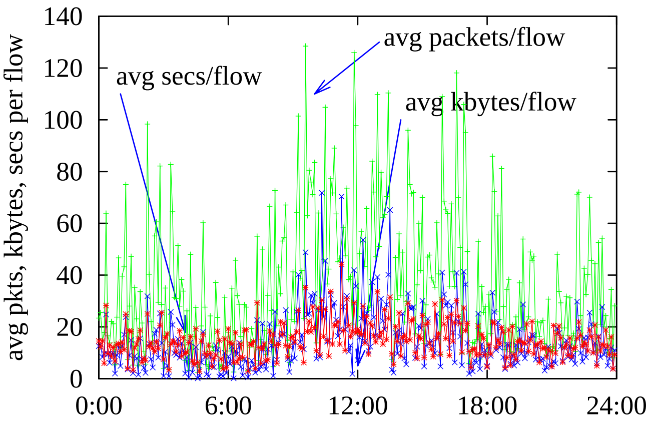

Figure 15e: Average traffic per RealPlayer UDP flow for 1 day on domestic link.

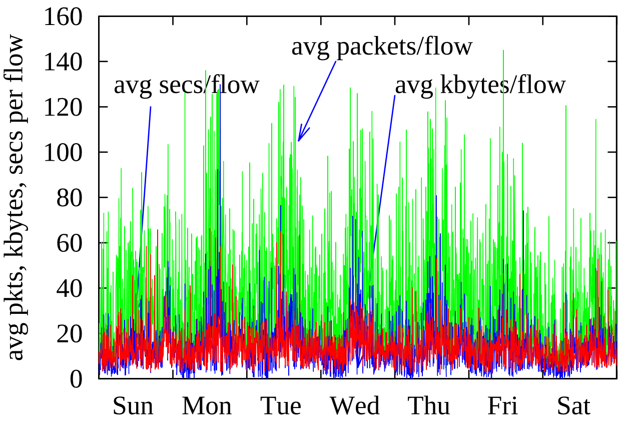

Figure 15f: Average traffic per RealPlayer UDP flow for 7 days on domestic link.

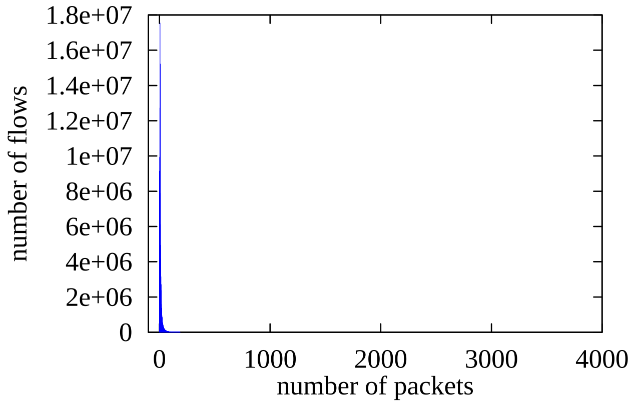

Figure 16: Linear-scale flow-length histogram from domestic link sample.

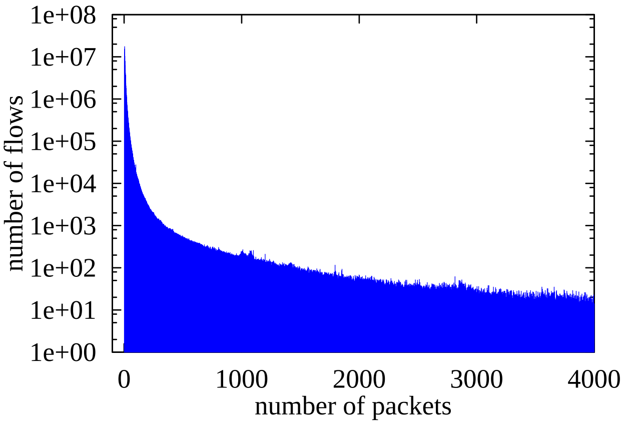

Figure 17: Logarithmic-scale flow-length histogram from domestic link sample.

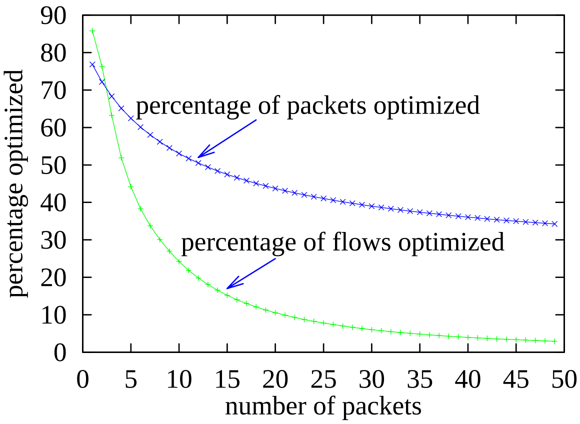

Figure 18: Percentage of flows and packets optimized when optimization is triggered after *x* packets.