

Understanding Internet Traffic Streams: Dragonflies and Tortoises*

Nevil Brownlee

CAIDA, SDSC, UC San Diego and

The University of Auckland, New Zealand, e-mail: nevil@caida.org

kc claffy

CAIDA, SDSC UC San Diego, e-mail: kc@caida.org

Abstract— We present the concept of network traffic streams, and the ways they aggregate into flows through Internet links. We describe a method of measuring the size and lifetime of Internet streams, and use this method to characterise traffic distributions at two different sites.

We find that although most streams (about 45% of them) are *dragonflies*, lasting less than 2 seconds, a significant number of streams have lifetimes of hours to days, and can carry a high proportion (50% to 60%) of the total bytes on a given link. We define *tortoises* as streams that last longer than 15 minutes. We point out that streams can be classified not only by lifetime (*dragonflies and tortoises*) but also by size (*mice and elephants*), and note that stream size and lifetime are independent dimensions.

We submit that Service Providers (ISPs) need to be aware of the distribution of Internet stream sizes, and the impact of the difference in behaviour between short and long streams. In particular any forwarding cache mechanisms in Internet routers must be able to cope with a high volume of short streams. In addition ISPs should realise that Long-Running (LR) streams can contribute a significant fraction of their packet and byte volumes – something they may not have allowed for when using traditional ‘flat rate user bandwidth consumption’ approaches to provisioning and engineering.

I. BACKGROUND

A. Measuring Internet traffic

The Internet is a global internetwork, sharing information among millions of computers throughout the world. Internet users send packets of information from one machine to another using various Internet protocols; TCP and UDP are the most common transport protocols, but other newer standard protocols are starting to appear.

Packets are carried through various links, e.g. from user host to regional Internet Service Provider (ISP), regional to backbone ISP, etc. Between such links packets are forwarded by routers using the Internet protocol (IP); IP version 4 is most common, but version 6 is now beginning to be deployed [1].

Typical users are not interested in packets on the Internet; they simply run application programs such as web browsers, which exchange packets with other computers as they carry out user requests. Groups of packets exchanged in this way are commonly referred to as *traffic flows*.

To measure traffic flows one examines packet headers as they pass by on a given link, determines which flow each packet belongs to using information extracted from its header, and counts packets and bytes for each flow. A system that gathers flow data in this way is called a *traffic meter*. Such a meter may be free-standing or may be built into a device such as a router.

*Support for this work is provided by DARPA NMS Contract N66001-01-1-8909, NSF Award ANI-9711092 ‘CAIDA: Cooperative Association for Internet Data Analysis’ and The University of Auckland.

For our investigations we use RTFM, an Internet standard Real-time Traffic Flow Measurement system [2]. The RTFM architecture defines three entities:

Meters gather data from packets so as to produce flow data.

Meter Readers collect flow data from meters.

Managers specify real-time data reduction by downloading configuration data (called *rulesets*) to meters, and also specify the intervals at which meter readers read flow data.

We make our flow measurements with NeTraMet [3], an open-source implementation of RTFM. NeTraMet includes an RTFM meter, a combined manager/meter reader, and a compiler for SRL (the Simple Ruleset Language, RTFM’s high-level language for specifying rulesets).

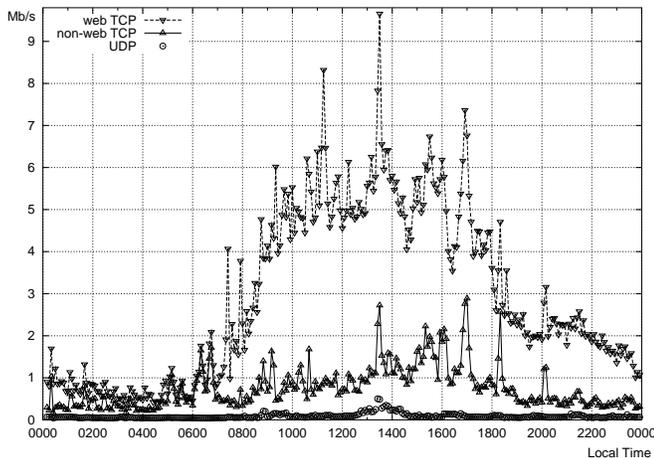
B. Traffic mix: mice and elephants

On any Internet link there is always a mix of flows from a variety of applications, carried by various transport protocols, especially TCP and UDP. UDP provides unreliable datagram delivery, i.e. an application sends UDP packets, but UDP itself provides no feedback to the sender. UDP is therefore unaware of any network congestion; streaming applications often continue to send data at constant high byte rates.

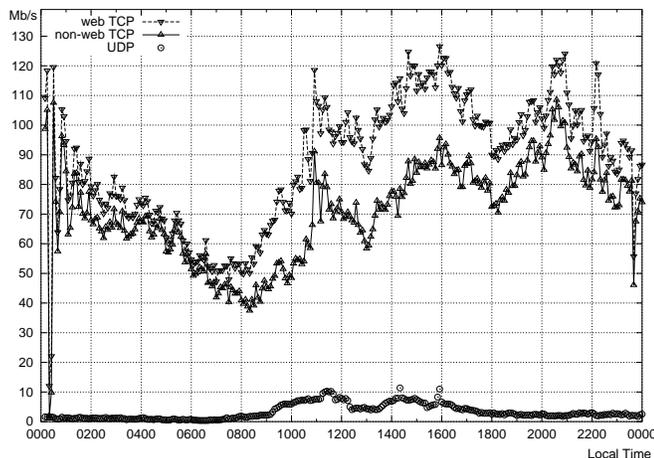
TCP, on the other hand, not only provides reliable byte stream delivery, but also uses feedback from receiving hosts to control its sending rate. TCP’s congestion management algorithms allow a TCP stream to vary its byte rate, seeking to use the highest possible rate but lowering the rate when the network becomes congested. For this reason, TCP is considered *network-friendly*.

Early analyses and simulations of TCP behaviour focused on steady state behaviour, using ‘infinite source’ workloads, e.g. large file transfers, and assumed that high volume TCP streams (network *elephants*) would not be significantly affected by the presence of small TCP streams (network *mice*). The fundamental difference between network mice and elephants is that an elephant’s TCP session extends past TCP’s *slow start* phase, so its behaviour, including the way it interacts with other TCP sessions, is controlled by TCP’s feedback-based congestion management algorithms. However, mice cannot be controlled by feedback since they are sent and received in their entirety before TCP has an opportunity to apply feedback control.

More recent models of TCP behaviour have increasingly focused on interactions between elephants and mice. For example, Joo *et al.* [4] analyzed the expected throughput of TCP streams and how concurrent streams interact. They found that multi-



1(a): Traffic summary at Auckland (UA), kb/s vs time for 5 m intervals



1(b): Traffic summary at San Diego (UCSD), kb/s vs time for 5 m intervals

Fig. 1. Cumulative rate (kb/s) for various traffic kinds vs time of day (HHMM) at Auckland (UA) and San Diego (UCSD) for 24 hours from midnight local time on Wed, 12 Jun 02.

ple elephants can synchronise with each other, which may cause routers to drop packets. They state, “although elephants are responsible for a major proportion of the bytes on the network, the number of packets generated by mice can be sufficient to create losses from time to time.” Joo *et al.* [4] also examined the dynamics of packet drops and concluded that mice can break up synchronization effects, leading to a more efficient use of network resources. This breakup effect may explain why best-effort datagram delivery has served the Internet so well as a low-cost common denominator of network service.

As an alternative to classifying flows by size (number of bytes), i.e. as *elephants* or *mice*, one can also classify flows by their lifetime (in seconds). Shaikh, Rexford and Shin [5], using a 60-second timeout, observed flow lifetimes up to 2000 seconds and found that such ‘long-lived’ flows accounted for a high proportion of bytes on a link. They propose that ‘load-sensitive’ routers might attempt to find better routes for long-lived flows, thus improving overall link utilisation.

C. Streams, flows and torrents, traffic kinds

The term *flow* has various meanings in different contexts. For example, in routing a *flow* is a set of packets with the same source and destination IP addresses, all travelling in the same direction. Internet researchers often use 5-tuples, i.e. protocol, source and destination IP address and port number; they refer to these as *microflows*.

In this paper, we use the terminology proposed by Brownlee & Murray [6], i.e.:

- *Streams* are individual IP sessions (e.g. TCP or UDP) between ports on pairs of hosts.
- *Flows* are sets of packets travelling in either direction between a pair of end-points (which may be hosts, networks, etc.).
- A *torrent* refers to all the traffic on a given link.

For our investigations we classify all traffic within a torrent into four *kinds*, distinguished by transport protocol: *UDP*; *TCP* (the most common); and *other*. Because web is a dominant application at some sites, we subdivide *TCP* into *web (TCP)* and *non-web TCP*. We aggregate all streams for each *kind* of traffic into one of four *flows*; we present data for these flows in figures 1 and 4.

II. MEASUREMENT METHODOLOGY

The NeTraMet implementation of streams creates a data structure for each stream within a flow, and counts each stream’s packets and bytes until the stream times out, i.e. no packets are observed for a dynamically specified *timeout* interval [6]. When a stream times out, its packet and byte counts are used to add a point to its flow’s stream size and lifetime distributions. Streams that remain active for long periods of time make no contribution to the stream distributions.

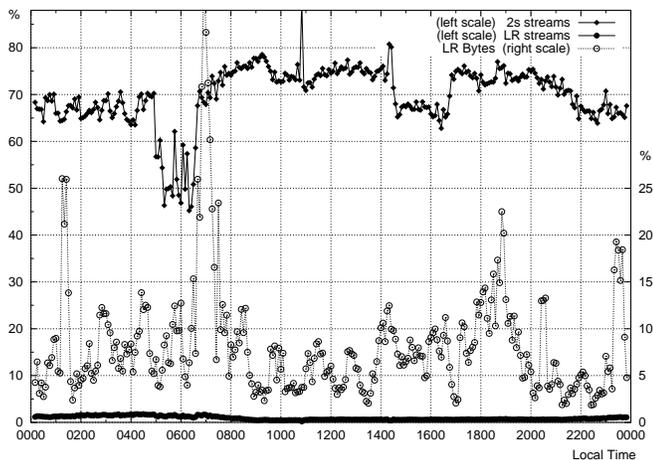
We have extended the NeTraMet meter to monitor stream lifetimes, and to automatically create flows in the meter’s flow table when a stream remains active for more than a specified time. Flows created for such long-running (LR) streams have their packet and byte counters updated for every packet. We collect data for these LR-stream flows with every meter reading, so as to obtain time series of packet and byte counts for each LR stream.

Our NeTraMet ruleset produces one flow for each of our four traffic *kinds*. For each flow we build five distributions:

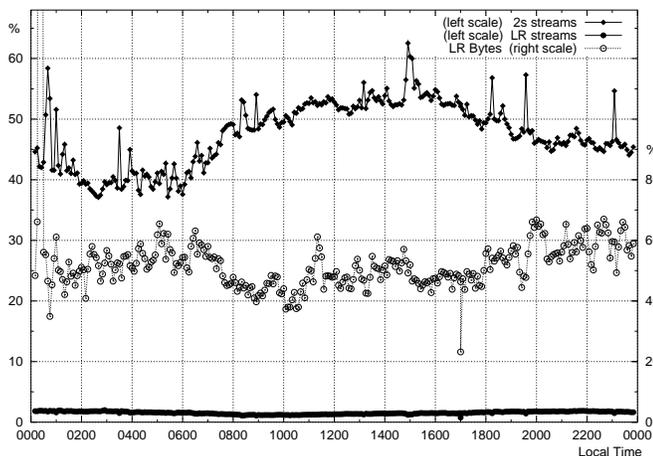
- 1-2 To/From Bytes:** Bytes in each direction of streams.
41 bins, log scale, 30 B to 600 kB
- 3-4 To/From Packets:** Packets in each direction of streams.
41 bins, log scale, 1 to 32768 packets
- 5 Flow Time:** Stream lifetime.
41 bins, log scale, 2 s to 15 minutes

We read our NeTraMet meters every five minutes, producing a set of five distributions for each five-minute interval. The counters for each distribution bin are never reset – instead we compute the distribution for each reading interval as the difference between successive readings. We can therefore observe variations in the distributions over periods of hours to days, with a time resolution of five minutes.

In order to provide higher resolutions for small and/or short-lived streams we use log scales for the distributions. Streams with size or lifetime greater than a distribution’s upper limit are counted in that distribution’s ‘overflow’ bin.



2(a): Stream and Byte % at Auckland (UA) vs time for 5 m intervals



2(b): Stream and Byte % at San Diego (UCSD) vs time for 5 m intervals

Fig. 2. Percentage of Short and Long-Running streams (black symbols using left-hand axis), and percentage of Bytes in LR streams (grey symbols using right-hand axis) vs time of day (HHMM) at Auckland (UA) and San Diego (UCSD) for 24 hours from midnight local time. Note that:

- 40% to 70% of streams are *dragonflies*, (black diamonds) lasting 2 seconds or less.
- About 1.5% of streams (black dots) are *tortoises*, i.e. LR streams lasting more than 15 minutes. They contribute 5% to 50% of a torrent's bytes (grey circles).

Although we collect byte and packet stream size distributions, in this paper we only present 'Bytes From' distributions, mainly due to space limitations, but also because the other distributions are generally similar to the 'Bytes From' ones. Similarly, the percentage of 'other' traffic we observe is negligible compared to our other three traffic *kinds*; we have not shown it on our plots.

Both our NeTraMet meters run on Linux systems located at two sites. Our OC3 meter (Auckland) observes packet headers using a commodity 100BaseT Ethernet card via *libpcap* [7]; our OC12 meter (UCSD) uses a Dag 3.2 card [8] via direct Linux drivers.

A. Short vs long-Running (LR) streams; dragonflies and tortoises

When configuring a NeTraMet meter one must specify a *stream-to-flow* time; the meter creates Long-Running (LR)

flows for all streams lasting longer than *stream-to-flow* seconds. In choosing a *stream-to-flow* value one must balance the desire to observe streams lasting for shorter periods against the costs of collecting and working with larger data sets. For this investigation we experimented with 5- and 10-minute lifetimes before choosing 15 minutes as the maximum lifetime of a *short* stream, i.e. our *tortoises* are LR streams with lifetimes greater than 15 minutes. We find that 15 minutes is a reasonable compromise, generating manageable data set sizes that still yield substantial insights into LR stream behaviour.

Measuring distributions of shorter flows is complicated by an 'edge effect.' Since a stream is only counted when it times out, the distribution counts include streams that started in an earlier interval but timed out in a given reading interval. The first bin in our Flow Time distributions counts flows with lifetimes up to two seconds, which limits the 'edge effect' error to a maximum of 0.6%. We find that a high proportion of streams fall in this first bin, so we describe them as *Very-Short* streams, i.e. streams with lifetimes of two seconds or less.

To summarise, we classify streams by lifetime as:

- *Very-Short dragonflies*, lasting up to two seconds,
- *Short*, lasting up to 15 minutes, and
- *Long-Running (LR) tortoises*, lasting more than 15 minutes.

We emphasise that stream size is independent of stream lifetime.

B. Observation sites

In this section we discuss the Internet traffic observed at two sites, the University of Auckland (UA), and UC San Diego (UCSD). We have collected data at both sites for eight complete days. Traffic patterns vary little over those days; our figures show data for Wednesday, 12 June 2002.

The University of Auckland (UA) runs a campus network serving about 35,000 users. The campus is connected to the Internet via an OC3 (155 Mb/s) ATM link, however this link is rate-limited to 9 Mb/s by the university's Internet provider. Our NeTraMet meter is connected to a 100 Mb/s Ethernet hub located between the university's access router and its firewall.

Figure 1(a) shows stacked bar plots of *web TCP*, *non-web TCP* and *UDP* traffic at Auckland versus local time. We read our meters every five minutes, then plot the average bit rate for each five-minute interval. The bit rate for each *kind* of traffic is indicated by the distance between its trace and the one below it. Furthermore, the topmost trace shows the whole torrent's traffic.

Figure 1(a) shows a typical day for a small enterprise site. There is little traffic in the early morning hours, except for a few brief spikes; these probably indicate periods during which mirroring servers at Auckland were updating their content. From about 7 a.m. the load grew; there was a slight dip around lunchtime, then the load remained steady through the afternoon. At about 6 p.m. the load decreased sharply, rose again slightly in the early evening, then decreased toward midnight. On this link web traffic was dominant; the ratio of web to non-web TCP traffic remaining fairly steady at about 80%. UDP traffic contributed few bytes to the load, and there was almost no traffic other than TCP and UDP.

One distinctive feature of this site is that the University of Auckland recovers the costs of Internet connectivity directly from its users, charging for each MB of data sent or received.

Departments are billed monthly for staff usage and students are charged by a real-time access control and billing system. Clearly the Internet usage patterns of Auckland users are influenced by the knowledge that they are paying for it. For example, Auckland has not (yet) seen widespread use of peer-to-peer file sharing.

UC San Diego (UCSD) is connected to the Internet via an OC12 (622 Mb/s) ATM link to the San Diego Supercomputer Centre (SDSC). SDSC has high speed links to three research networks as well as a lower speed connection to the commodity Internet. Our NeTraMet meter is attached to UCSD's OC12 link via a passive optical splitter.

Figure 1(b) shows stacked bar plots of UCSD traffic. UCSD does not attempt to recover per-byte Internet usage costs, nor does it impose rate limits on individual Internet connections. In this environment Internet usage is limited only by congestion, which will increase as the university's total Internet usage increases. Occasionally UCSD upgrades its commodity Internet capacity, historically allowing the load and congestion cycle to build up again.

UCSD's total traffic byte volume is 16 times greater than Auckland's nearly all the time, but still utilises only about 15% of their OC12 link's maximum capacity. UCSD also has more UDP traffic than Auckland; 1 Mb/s at night, rising to around 10 Mb/s during the day.

At Auckland web was the dominant TCP application, but UCSD's web to non-web TCP ratio is only about 50% (much less than Auckland's 80%), indicating that at UCSD a higher proportion of TCP traffic is generated by non-web applications.

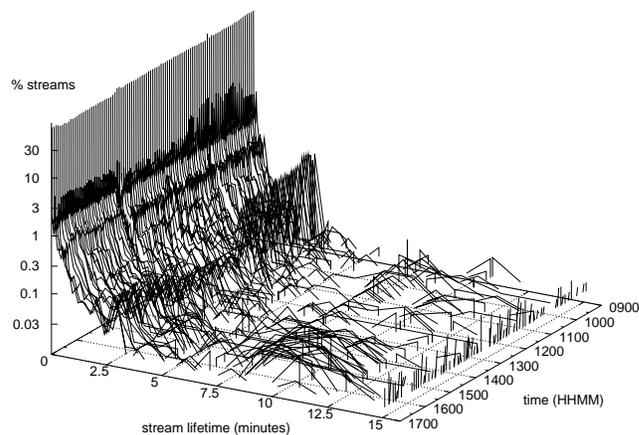
Overall, the diurnal variations for web and UDP traffic suggest that these kinds of traffic follow human activities. Non-web TCP however has a fairly high background level (about 0.5 Mb/s at Auckland and 70 Mb/s at UCSD); it varies about this level, with its UCSD minimum around 8 a.m. and its maximum around 10.30 p.m.

C. Percentage of streams and bytes in a torrent

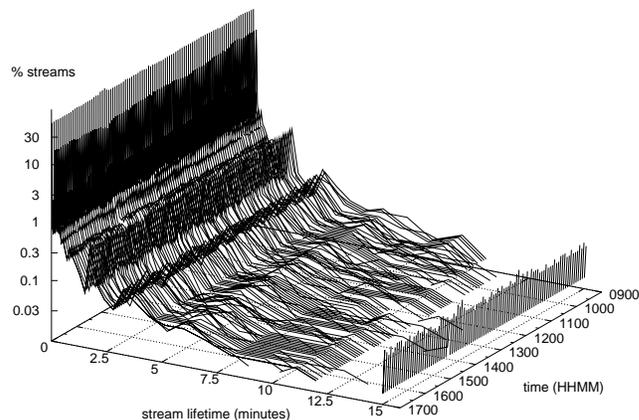
As well as producing stream packet and byte size distributions, our NeTraMet ruleset records the total number of packets and bytes observed in each traffic flow. For each meter reading we sum the LR-stream byte counts; from these sums we compute the percentage of bytes in LR streams. In this section we discuss the LR byte percentages, together with the percentages of Very-Short ($\leq 2s$) and LR ($> 15m$) streams at our two observation sites. Figure 2 plots these three measures at five-minute intervals over a day for Auckland (UA) and San Diego (UCSD).

First, the percentage of Long-Running streams is plotted with black dots (bottom trace) in figure 2, using the left-hand y-axis scale. At UCSD (figure 2(b)) about 1.5% of the streams were LR, and this level varied little during the day. In Auckland (figure 2(a)) there were only about 0.5% LR streams during the day, but nearly 1.5% from midnight until dawn.

Second, the percentage of Very-Short streams is plotted with black diamonds (top trace) in figure 2, also using the left-hand y-axis scale. At UCSD (figure 2(b)) there was a clear diurnal variation of the Very-Short stream percentage from 38% around 4 a.m. to about 55% around 3 p.m., corresponding well with



3(a): Stream lifetimes at Auckland (UA) at 5-min intervals



3(b): Stream lifetimes at UC San Diego (UCSD) at 5-min intervals

Fig. 3. Stream lifetime distributions at Auckland (UA) and UC San Diego (UCSD). Percent streams vs lifetime (minutes) for eight hours, 9 a.m. to 5 p.m. local time on Wed, 12 Jun 02. Note that although the link data rates differ, the shape of the distribution is similar for both sites.

UCSD's diurnal variation in web traffic shown in figure 1(b). At Auckland however (figure 2(a)) about 70% of the streams were Very-Short nearly all day. Between 5 a.m. and 7 a.m. the Very-Short stream percentage dropped to 50%; during that interval there were considerably fewer TCP web bytes than non-web bytes.

Third, the percentage bytes in LR streams is plotted with open circles (middle trace) in figure 2, using the right-hand y-axis scale. At UCSD (figure 2(b)) about 50% of all bytes were in Long-Running streams, and this percentage varied little during the day. At Auckland (figure 2(a)) only about 5% of all bytes were in Long-Running streams, most likely because Auckland had much less non-web TCP traffic than UCSD.

To summarise:

- About 1.5% of UA and UCSD streams were Long-Running (LR).
- 40% (UCSD) to 70% (UA) of the streams were Very-Short, and most of them appear to be web traffic.
- 5% (UA) to 50% (UCSD) of all bytes were in LR streams, and most of them appear to be non-web traffic.

A. Lifetime distributions

Figure 3 shows lifetime distributions for short streams, i.e. those with lifetimes up to 15 minutes, at Auckland (figure 3(a)), and UCSD (3(b)). In both these plots at least 45% of the counted streams lie in the first bin, i.e. they had lifetimes of two seconds or less. We use a logarithmic scale for the distribution's y-axis, to reveal the whole range of percentages. The rest of the distribution falls away quickly at Auckland, where there are few streams with lifetimes above 2.5 minutes. At UCSD the lifetime distribution falls away more slowly, but there are few streams with lifetimes above 5 minutes. However, above these lifetimes, the lifetime distributions slope down gently toward 15 minutes, our maximum short stream lifetime.

As defined in section II-A, our LR streams have lifetimes greater than 15 minutes. When an LR stream times out, the meter increments the overflow bin for its distributions (i.e. the distributions for that stream's flow). The lifetime distribution overflow counts appear as the high values plotted for y values above 15 minutes, i.e. the spikes at the right-hand edge of figures 3(a) and 3(b). As we saw on the LR byte percentage plots (figure 2), although there are few LR streams, they can account for a high percentage of a flow's total bytes throughout the day.

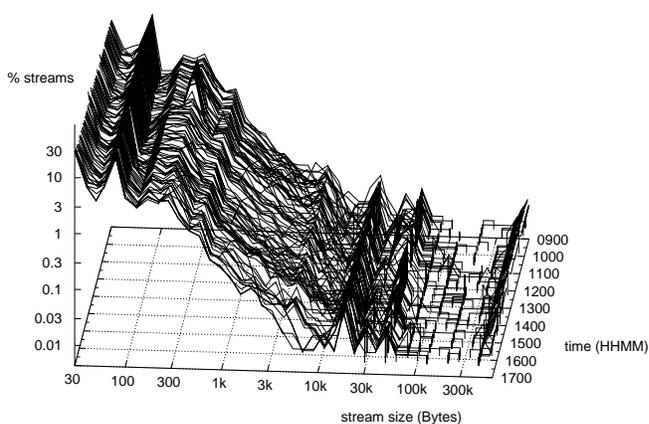
The most striking feature of the distributions in figure 3 is that their shapes are similar. At both sites we observe that the distributions do not change rapidly with time. At UCSD (3(b)) there was little change in the shape of the distributions over the eight hours shown. At Auckland (3(b)), where the link capacity is lower, the proportion of streams with lifetimes between 7.5 and 12.5 minutes increased during the afternoon. This increase could indicate that when the Auckland's link's byte load is high, not only are there more users, but those users are working with larger files.

B. Byte size distributions

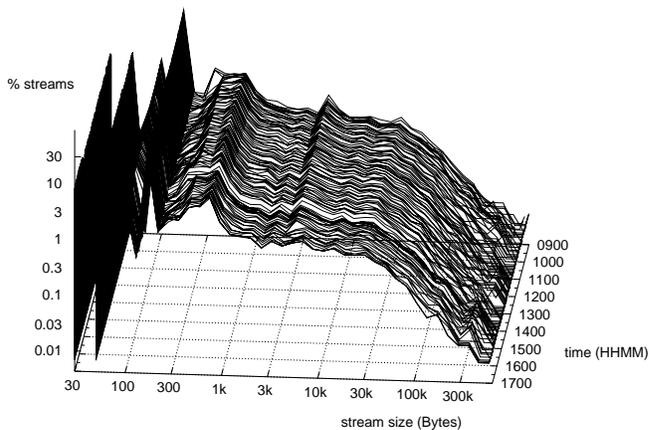
Figure 4 shows byte size distributions for short streams at UCSD; byte size distributions at Auckland (not shown) are similar. The distributions are collected using 41 bins in a log scale from 30 bytes to 600 kB. The jagged appearance for small stream sizes, i.e. below about 300 bytes, is an artifact of the limited bin resolution for streams with only a few packets.

For short UDP streams (figure 4(a)), the byte size distribution has peaks at about 30 and 80 bytes, indicating that a high proportion of UDP streams have only one or two packets in their *from* (i.e. destination to source) direction. Short UDP stream sizes fall steadily (on a log scale) from about 300 bytes to 20 kB; they also have a noticeable percentage of streams in the overflow bin, i.e. with sizes above 600 kB. Successive five-minute distributions vary somewhat over periods of about 15 minutes, producing the corrugated effect on the plot.

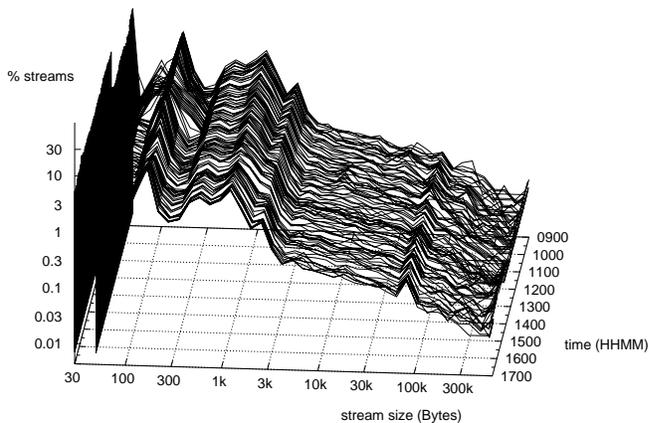
Short Web streams (figure 4(b)) have high peaks at 30, 50, 100 and 200 bytes, a local maximum from about 300 to 800 bytes, a plateau from 1 kB bytes to about 40 kB, and a fairly steep fall from there. Since web streams use TCP, they require at least two packets in each direction, hence the size distribution's peaks below 300 bytes indicate streams that most likely failed to establish a TCP session. The local maximum from 300 to 800 bytes is



4(a): UDP stream sizes at UC San Diego (UCSD) at 5-min intervals



4(b): Web stream sizes at UC San Diego (UCSD) at 5-min intervals



4(c): Non-web TCP stream sizes at UC San Diego (UCSD) at 5-min intervals

Fig. 4. Size distributions for short streams at UCSD. Percent of streams vs stream size (kB) for eight hours, 9 a.m. to 5 p.m. local time on Wed, 12 Jun 02. Note the clear differences between the three flow kinds.

probably for streams carrying small web objects (buttons, 'file not found' messages, etc.), and the plateau suggests that web objects have a fairly flat file size distribution below about 40 kB, with a power law fall (linear on the log plot, figure 4(b)) for large files. The five-minute web stream size distributions are remarkably steady, suggesting that web usage patterns at UCSD are stable over long periods.

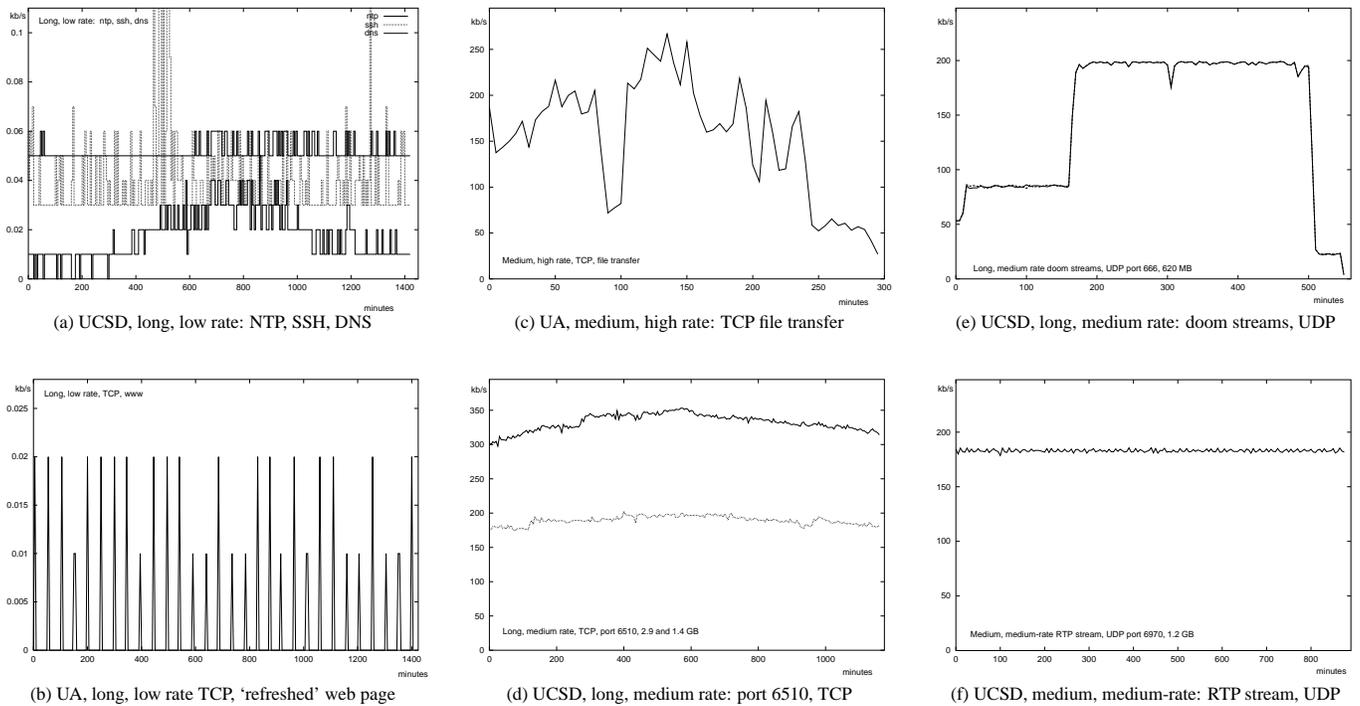


Fig. 5. Long-running (LR) stream histories, Auckland (UA) on Thu 4 Apr 02 (b,c) and San Diego (UCSD) on Thu 28 Mar 02 (a,d,e,f). Bit rate (kb/s) vs elapsed time (minutes)

Short non-web TCP stream distributions (figure 4(c)) have a stream size peak at 30 and 50 bytes, smaller peaks at 180 and 1500 bytes, and a steady fall from there to 600 kB. The steady fall suggests that non-web stream sizes also have a power law distribution, at least for stream sizes from about 1500 bytes to 300 kB.

There are also noticeable local maximum at 5 kB for web streams and 90 kB for non-web streams, indicating transmission of many objects this size; we have not yet determined what they were.

To summarise:

- Short stream size distributions for UDP, non-web and web TCP traffic are distinctly different and stable over periods of hours.
- Our study only reflects data collected at Auckland and UCSD, but at this stage we believe these sites are representative examples of medium to large Internet-edge networks.

IV. LONG-RUNNING STREAM (*tortoise*) BEHAVIOUR

We find that long-running streams, i.e. streams lasting more than 15 minutes, occur frequently at both Auckland and UCSD. To gain a better understanding of their behaviour we selected a representative set of LR streams and produced ‘thumbnail’ plots of their bit rates (kb/s) vs time (minutes), as shown in figure 5. In this section we comment on the behaviour of these streams.

- Figure 5(a) shows three streams with bit rates below 60 b/s. At such low bit rates our bit rate resolution is poor, producing the stepped effect in the traces. The top trace was an NTP stream.

These are always present, serving to keep end-system clocks synchronised. The middle trace was an SSH stream that was mostly quiescent but had occasional intervals reflecting user activity. The lower trace was a DNS stream with a diurnal bit rate pattern, highest during the afternoon.

- Figure 5(b) shows a 24-hour web stream that consisted entirely of 20 b/s bursts at half-hour intervals, suggesting that it was carrying a web page that was ‘refreshed,’ i.e. re-displayed, every half hour. The apparent size of the bit rate bursts depends on whether they happen to fall in one or two of our five-minute reading intervals.
- Figure 5(c) shows an FTP file transfer, taking five hours at a high but variable rate, 50 to 200 kb/s. Its high bit rate variance suggests that this stream was traversing severely congested Internet paths.
- Figure 5(d) shows two streams that ran at 180 and 320 kb/s nearly all day, moving about 2.9 and 1.4 Gigabytes. Although they were TCP streams their data rates did not vary much, indicating that there was little congestion on their Internet paths. These streams effectively reduced the link’s available capacity by several hundred kb/s.
- Figure 5(e) shows two near identical Doom streams that ran for nine hours at 80 to 200 kb/s. They were UDP streams, hence their steady rate for long periods. Their sudden rate changes presumably correspond to changes in the state of the game.
- Figure 5(f) shows a RealAudio stream that ran at 180 kb/s for 15 hours, transferring 1.2 Gigabytes. This was a UDP stream; its rate varied only about ± 5 kb/s, much less than the TCP streams in figure 5(d).

To summarise:

- Long, continuous LR streams may be low-rate (service support or user interaction), or high-rate (audio/video data streams).
- Brief and medium-duration LR streams tend to be high-rate, running until some user-initiated activity is completed.
- TCP LR streams show rate variation as Internet congestion changes over time, with rate variations similar for streams sharing congested links in their Internet paths.
- UDP streams tend to run at fairly constant bit rates, but these rates change in response to application dynamics.

A. LR stream lifetimes

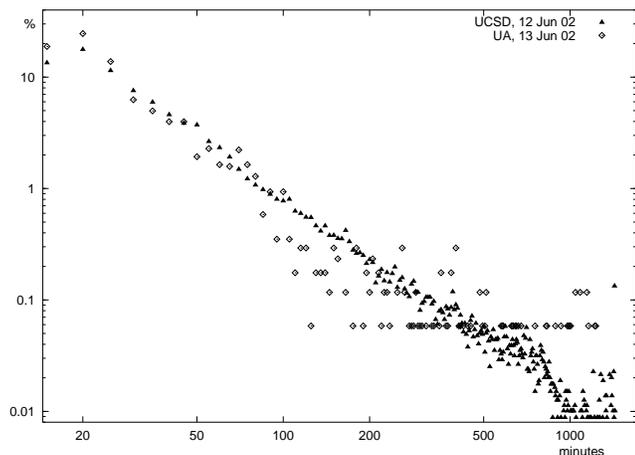


Fig. 6. Stream Lifetime Distributions, UCSD and UA. Percentage of LR streams in torrent vs stream lifetime (minutes).

Figure 6 is a log-log scatter plot showing percentage of LR streams vs LR stream lifetime (minutes) for Auckland and UCSD. The two plots cluster around a line, suggesting that stream sizes follow a power law distribution. The spread of points around this line is narrow for lifetimes from 20 to about 100 minutes; the spread increases for longer lifetimes.

More points are plotted for UCSD than for Auckland, providing more detail for lifetimes above 500 minutes. Similarly, because there are fewer points for Auckland (reflecting the lower traffic rate at Auckland), the y-axis has lower resolution for the Auckland plot, producing the line at stream percentage 0.06%. The two plots are nonetheless similar, indicating that users at the two sites are running similar application cross-sections.

V. CONCLUSION

As recently as July 2000, Zhang *et al.* [9] observed that “Internet traffic is now dominated by mice, i.e. small objects 10-20 kB in size; the average web document is only around 30 kB,” but in contrast reported that “the majority of the packets and bytes belong to elephants.” Similarly, in April 2001 Brownlee *et al.* [10] measured stream byte size distributions and found that TCP data streams had a 95th percentile of approximately 15 kB.

Since 2000 Internet link speeds increased as users migrated to cable modem and DSL connections, backbone links were upgraded from OC3 (155 Mb/s) toward OC48 (2.4 Gb/s), and ISPs installed newer, faster routers to handle increasing packet loads.

At the same time computer hardware improved; systems with 1 GHz processors, 512 MB memory, 20 GB disk drives and ever-increasing I/O bus speeds became common. This dramatic increase in network and computer capability has allowed users to work with ever larger files. As a result we now observe that the average size of web objects has increased considerably, with web objects up to 50 kB becoming common.

Along with increasing file size, the last few years have seen the rapid growth in usage of an ever increasing set of peer-to-peer file sharing systems. e.g. Napster, Gnutella, E-Donkey, etc. These peer-to-peer applications have significantly changed the traffic mix, so that a higher overall proportion of their streams have large numbers of bytes. In addition to streaming protocols carrying audio and video programs, VoIP or multimedia conferencing are increasingly common. Clearly these trends will continue.

Our current observations confirm that most streams are very short. At least 45% of streams have lifetimes of 2 seconds or less, (*dragonflies*), and about 98% of them last less than 15 minutes. However, the remaining one or two percent, which we call Long-Running (LR) streams (*tortoises*), have lifetimes of hours to days and can carry a high proportion (50% to 60%) of the total bytes on a link.

We submit that Internet Service Providers (ISPs) need to be aware of the behaviours of short streams. In particular any forwarding cache mechanisms in Internet routers must be able to cope with the high volume, both absolute and as a percentage, of short streams. In addition, ISPs should realise that LR streams can contribute a significant fraction of their packet and byte volumes, reducing the available bandwidth of their Internet links.

Lastly, we emphasise that streams can be classified not only by their size (*mice and elephants*), but also by their lifetime (*dragonflies and tortoises*). Furthermore, stream size and lifetime are independent dimensions; both are of interest in understanding the overall behaviour of streams in a torrent.

REFERENCES

- [1] J. Bound, *IPv6 Implementation*, ISOC Member Briefing #4, September 2001, available from <http://www.isoc.org/briefings/004/>
- [2] N. Brownlee, C. Mills and G. Ruth, *Traffic Flow Measurement: Architecture*, RFC 2722, October 1999
- [3] N. Brownlee, *Using NeTraMet for Production Traffic Measurement*, Intelligent Management Conference (IM2001), May 2001
- [4] Y. Joo, V. Ribeiro, A. Feldmann, A. C. Gilbert and W. Willinger, *On the impact of variability on the buffer dynamics in IP networks*, Proc. 37th Annual Allerton Conference on Communication, Control and Computing, September 1999, available from <http://www.dsp.rice.edu/publications>
- [5] A. Shaikh, J. Rexford and K. G. Shin, *Load-Sensitive Routing of Long-Lived IP Flows*, Proc. ACM SIGCOMM, September 1999
- [6] N. Brownlee and M. Murray, *Streams, Flows and Torrents*, PAM2001 workshop, April 2001
- [7] tcpdump/libpcap website, <http://www.tcpdump.org/>
- [8] Dag website, <http://dag.cs.waikato.ac.nz/>
- [9] Y. Zhang and L. Qiu, *Understanding the End-to-End Performance Impact of RED in a Heterogeneous Environment*, Cornell CS Technical Report 2000-1802, July 2000, available from <http://www.aciri.org/floyd/red.html>
- [10] N. Brownlee, kc. claffy, M. Murray and E. Nemeth, *Methodology for Passive Analysis of a University Internet Link*, PAM2001 Workshop, April 2001