

Recent Trends in TCP Packet-Level Characteristics

Per Hurtig

Department of Computer Science
Karlstad University, Karlstad
Email: per.hurtig@kau.se

Wolfgang John

Department of Computer Science and Engineering
Chalmers University of Technology, Göteborg
Email: wolfgang.john@gmail.com

Anna Brunstrom

Department of Computer Science
Karlstad University, Karlstad
Email: anna.brunstrom@kau.se

Abstract—Up-to-date TCP traffic characteristics are essential for research and development of protocols and applications. This paper presents recent trends observed in 70 measurements on backbone links from 2006 and 2009. First, we provide general characteristics such as packet size distributions and TCP option usage. We confirm previous observations such as the dominance of TCP as transport and higher utilization of TCP options. Next, we look at out-of-sequence (OOS) TCP segments. OOS segments often have negative effects on TCP performance, and therefore require special consideration. While the total fraction of OOS segments is stable in our measurements, we observe a significant decrease in OOS due to packet reordering (from 22.5% to 5.2% of all OOS segments). We verify that this development is a general trend in our measurements and not caused by single hosts/networks or special temporal events. Our findings are surprising as many researchers previously have speculated in an increased amount of reordering.

Keywords—traffic measurement; TCP; reordering;

I. INTRODUCTION

The flexibility and versatility of the Internet architecture allows protocols and applications to be developed and deployed quickly. These properties have thus enabled a rapid evolution of the Internet. To support further development of the Internet it is important to investigate and highlight trends in its evolution.

In this paper, we follow up on our previous observations on backbone traffic [1] by comparing general packet characteristics between 35 traces from 2006 and 35 novel traces from 2009. We have complemented the basic packet-level characteristics with an analysis of out-of-sequence (OOS) TCP segments. A TCP segment is said to be OOS when it arrives at a receiver that is expecting another segment. Segments can arrive OOS for different reasons, including retransmissions, network duplication and packet reordering. Although TCP is designed to deal with OOS segments, the performance might suffer. For instance, packet reordering is a problem as it causes receivers to emit duplicate acknowledgments (dupACKs) back to the sender. A high degree of reordering can therefore cause TCP to falsely assume packet loss, leading to unnecessary retransmissions and invocation of the congestion control, which substantially lowers the throughput. Interestingly, a number of novel networking technologies that now are being deployed use mechanisms

that create reordering as a side effect. For instance, mobile ad hoc networks (MANETs) often use routing mechanisms that create reordering in their mode of operation. It is therefore important to track the development of reordering over the last years.

Studies on packet-level characteristics have been published before, especially during the early 2000's. Figures about packet size distribution and transport protocol decomposition of Internet traces from different wide-area measurement locations (OC3-OC192) have been presented repeatedly since 1997 [2]–[5]. Also usage of TCP options has been presented on data collected until 2000 [6] and 2004 [5]. However, since our summary about packet characteristics of backbone data from 2006 [1] there have been no publications with complete packet-level details of wide-area Internet traffic. Nevertheless, there have been studies specialized towards certain aspects of TCP. For instance, Maier et al. [7] present transport protocol features such as TCP option deployment and configuration on Internet traffic data from DSL connections of a large European ISP in 2008 and 2009. Qian et al. [8] compared TCP flow sizes and also tried to infer the evolution of a number of TCP implementation details in traces from 2001 and 2008. Also measurements of TCP out-of-sequence (OOS) segments have been conducted in several studies between 1999 and 2008 [9]–[13]. To our knowledge, however, there have not been any studies that have collected data from the same location, at different points in time, to detect possible trends. We believe that it is important to keep the research community updated with this type of basic information, which enables accurate and realistic simulation models to support refinement and development of network protocols and devices.

In Section II the data collection and processing is described. Section III presents and compares general IP packet characteristics, such as packet size distributions and TCP option deployments. Section IV compares different TCP OOS deliveries, and details the trends of OOS due to packet reordering. The frequency of TCP OOS segments appears to be quite stable between the measurements, affecting about 17% of all TCP connections. However, we were surprised to see that OOS segments due to packet reordering have decreased significantly in 2009. Finally, Section V provides

a number of concluding remarks.

II. DATA COLLECTION AND PROCESSING

The two data sets compared in this paper were collected in the time from April to November 2006 and January to June 2009 (see Table I) on backbone links in the Swedish University Network (SUNET). Altogether, 70 traces were collected, each trace consisting of ten minutes of bi-directional traffic. The measurement times of the 35 traces collected in 2006 were spread out over a period of eight months, and the 2009 collection times over a period of six months. We collected the traffic on OC192 (10Gbit/s) links on two different generations of SUNET. The 2006 data was collected in GigaSUNET, a ring architecture with a central Internet exchange point in Stockholm. The measurement location was on a OC192 ring which was the primary link from the region of Gothenburg to the main Internet outside Sweden. The link mainly carried traffic from major universities and large student residential networks, but also from a regional access point exchanging SUNET traffic with local ISPs.

The 2009 data was collected in the upgraded SUNET (OptoSUNET), a star structure over leased fiber. All customers are redundantly connected to a central Internet access point. Besides some local exchange traffic, the traffic routed to the main Internet outside Sweden is carried on two links (40Gb/s and 10Gb/s) between SUNET and NorduNet. The data used in this study was collected on the 10Gb/s link, which according to SNMP statistics carried 50% of all inbound but only 15% of the outbound traffic volume.

We collected data by using optical splitters attached to two Endace DAG6.2-SE cards (i.e., one measurement card for each direction). We configured the DAG cards to capture the first 120 bytes of each frame to ensure that the network and transport headers were preserved. After recording the traces, the IP-addresses were anonymized using the prefix preserving CryptoPAN software [14] and the remaining payload, beyond the transport layer, was stripped to ensure privacy and confidentiality. During data collection, the DAG cards were synchronized with each other using Endace's DUCK time synchronization [15], allowing us to merge the data into well-aligned bi-directional packet-header traces. Further details on the data collection and pre-processing can be found in [16].

To process and analyze general traffic properties we used available tools like CAIDA's CoralReef [17], as well as own specialized tools for additional sanity checking and result processing. For packets of special interest, the corresponding TCP flows were extracted and manually analyzed. For detection and classification of OOS TCP segments we used Tstat 2.0 [18]. Tstat's OOS detection and classification method is described in [13], and is a refinement of the methodology used by Jaiswal et al. [19]. We further describe this method in Section IV.

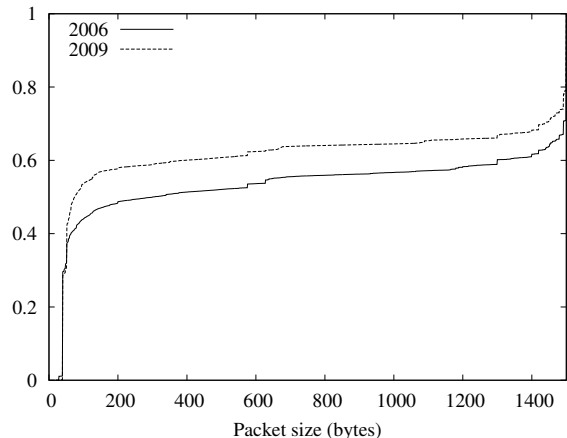


Figure 1. CDF for IP packet sizes for the 2006 and the 2009 measurements.

III. GENERAL RESULTS

The 70 packet traces consist mainly of IPv4 packets (99.98% and 99.99% of all frames observed in 2006 and 2009, respectively). The remainder of the traffic is mainly IPv6 routing traffic (BGP). In the rest of this paper, only IPv4 traffic is considered. Note that the 2006 data is a subset of the data set analyzed in [1].

A. IP Traffic Characteristics

1) *IP Packet Size Distribution*: Earlier Internet measurements, conducted between 1997 and 2002 [2]–[4], [20], reported of cumulative IPv4 packet distributions being trimodal, with major modes at about 40 bytes (TCP acknowledgments), 576 bytes (the default datagram size [21]) and 1500 bytes (the Ethernet MTU). Default datagram sizes represented about 10 – 40% of all packets. However, later measurements have reported of a much smaller fraction of default datagram sizes (e.g., 3.8% in 2004 [5]).

Figure 1 shows the cumulative distribution function (CDF) of packet sizes in our measurements. The packet size distributions are bimodal with major modes at around 40 bytes and 1500 bytes. The percentage of packets having the default datagram size is about 1% in both measurements, not even being among the three largest modes anymore. As we already reported earlier [1], this can be explained by the common use of PathMTU Discovery. We can also see that the fraction of small packets has increased significantly.

2) *Transport Protocols*: Table II shows the fractions of packets/bytes carried by each protocol compared to total IPv4 traffic. The figures confirm the domination of TCP as transport protocol, but also indicates an increasing trend in UDP traffic. The percentage of UDP packets has increased from 8.2% to 16.27%, and the amount of bytes from 3.4% to 8.53%. This is in line with other measurements that also have reported increased UDP traffic, especially in terms of flow numbers [22], [23]. The reason for this has been

Table I
SUMMARY OF THE DATA SETS.

Dataset	Collection Period	#Traces	Trace Dur.	Total Volume	Total #Packets
GigaSUNET	Apr.-Nov. 2006	35	10 min	2.3 TB	3.3×10^9
OptoSUNET	Jan.-Jun. 2009	35	10 min	4.6 TB	7.9×10^9

Table II
PROTOCOLS AT TRANSPORT LEVEL (IN %).

	GigaSUNET 2006		OptoSUNET 2009	
	Pkts	Data	Pkts	Data
TCP	91.50	96.50	82.90	90.40
UDP	8.20	3.40	16.27	8.53
ICMP	0.17	0.02	0.23	0.04
ESP	0.13	0.06	0.47	0.93

reported to be an increase in P2P signaling traffic, generating large numbers of small flows over UDP.

Also other protocols have become more prevalent. Especially the use of Encapsulating Security Payload (ESP) [24], used to enhance security and confidentiality of data transfers, has increased from 0.06% to 0.93%. Although ESP is only responsible for about 1% of the data, the increase is substantial. We speculate that this could be caused by a rising popularity of IPsec tunnels as a reaction to the new IPRED law in Sweden¹.

B. TCP Characteristics

In order to analyze TCP characteristics, we aggregated packets into bi-directional TCP flows. In this paper, a TCP flow is the same as a TCP connection. Thus, a flow is identified by a connection initiation (3-way handshake) and a termination (by FIN/RST signaling). Flows without an observed, complete handshake have been excluded from our analysis. This strict flow definition was used to allow more accurate results.

1) *Flow Lengths*: Even if our measurements do not contain flows longer than 10 minutes we investigated TCP flow size distributions. When considering TCP flows, the classical assumption is that TCP traffic is heavy-tailed, i.e., it consists of a large number of small flows (mice) and a small number of large flows (elephants) [26], [27]. The TCP flow size distributions for our measurements are plotted in Figure 2. The graph shows the CDF of TCP flow sizes in bytes.

Consistent with the classical assumption, the distribution of flow sizes appear to be heavy-tailed. About half of all flow sizes are around 1000 bytes only, but very large flows are not negligible and are responsible for a large fraction of the total traffic volume. On average, it appears as flow lengths have increased slightly from 2006 to 2009, but no significant

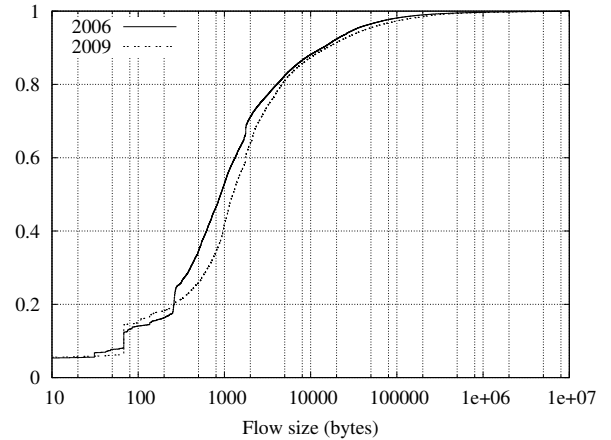


Figure 2. CDF for TCP flow lengths (in bytes).

differences are apparent. This confirms recent results by Qian et al. [8], reporting about no qualitative differences in TCP flow sizes when comparing AT&T backbone data from 2001 and 2008.

2) *TCP Options*: Earlier measurements have shown a rather significant deployment of TCP options, such as Selective Acknowledgments (SACKs), Window Scaling (WS), Timestamps (TS), and Maximum Segment Size (MSS). Allman [6], for instance, reported that about 20% of all hosts allowed the WS and TS options. SACK was shown to be more commonly deployed, about 40%. In a recent study by Maier et al. [7], WS was reported to be advertised by at least one endpoint in 32–35% of connections carrying data, and effectively used by 28–31%, TS was advertised in 11–12% and used in 8%, and SACK was advertised in 97% and used in 82% of the connections.

We have previously shown that the use of the WS, TS, SACK, and MSS options is rather widespread [1], [28]. The MSS option, for instance, was advertised in about 99% of all the TCP SYN segments in [1]. In this paper, however, we only focus on established TCP connections. Thus, we only deal with connections that have had a proper SYN-SYN/ACK exchange. Table III shows the percentage of TCP option advertisement/usage for the connections that were established during our measurements. The advertised column shows if at least one party of the connection tried to use the option, while the used column shows if the corresponding option was actually used in the connection.

As indicated by the table, TCP options are used to a

¹On April 1, 2009, an anti-piracy law based on the European directive on the enforcement of intellectual property rights (IPRED) [25] came into effect in Sweden.

Table III
TCP OPTION USAGE IN SUCCESSFUL THREE-WAY HANDSHAKES (IN %).

TCP Option	GigaSUNET 2006		OptoSUNET 2009	
	Advertised	Used	Advertised	Used
MSS	99.99	99.40	99.99	99.21
WS	21.20	18.60	44.22	37.64
SACK	96.34	80.36	98.43	86.83
TS	17.27	14.25	23.93	19.63

significant extent, and are even more commonly employed than in the DSL connections of the large European ISP studied in 2009 [7]. In our data, the MSS option is used by nearly all connections ($\approx 99\%$) in both 2006 and 2009. The use of WS have almost doubled, from about 19% to 38%. The use of SACK and TS have also increased, from 80% to 87% and from 14% to 20%, respectively.

IV. TCP OOS DELIVERIES

This section presents and compares OOS segments found in the measurements. As mentioned in the introduction, an OOS segment is a segment that arrives unexpectedly at the TCP receiver. In our definition of OOS we also include segments that already have been received, i.e., duplicated segment arrival. We start by describing the methodology used for identifying and classifying OOS segments. We then give an overview of OOS segments observed in our measurements. Finally, we provide an extended analysis of reordered TCP segments.

A. Methodology

The methodology used was originally developed in [19] and later extended in [13]. The methodology both identifies OOS segments and classifies them. For example, if a segment is lost and is retransmitted using a retransmission timeout the methodology will identify retransmission timeout as the underlying cause.

The basic detection of an OOS segment is straightforward; given a bi-directional traffic trace the sequence and acknowledgment numbers can be used to infer if segments are arriving in-order or OOS. To further classify OOS segments is a more complicated task. We chose to use the Tstat tool [18] which implements the methodology in [13]. A short description of the classification follows below, while the exact details of the classification process can be found in [13].

If the observed segment has both the same IP identifier and the same sequence number as an already observed segment it is due to network duplication (NetDup). If the observed segment is unacknowledged and the loss recovery of the sender is likely to have triggered the transmission of the corresponding segment it might be due to a retransmission. Triggering of the loss recovery mechanisms are assumed if the estimated loss recovery time is greater than the expected RTO, or if three duplicate acknowledgments have been

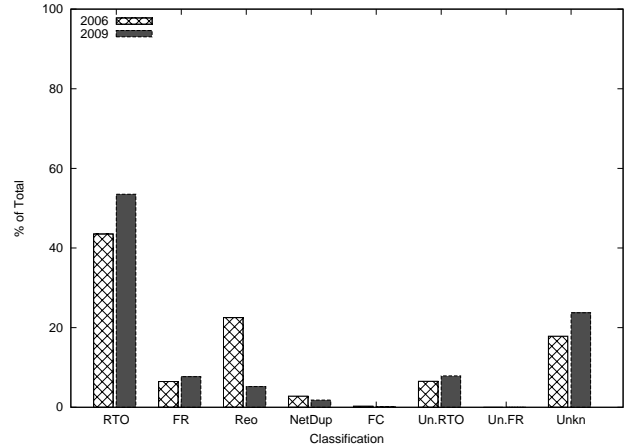


Figure 3. OOS classification of TCP segments in 2006 and 2009. The significant decrease in reordered segments, from 2006 to 2009, is the main difference between the data sets.

observed (Fast Retransmit (FR)). If the loss recovery of the sender is likely to have triggered a retransmission of the segment although both the segment and its acknowledgment have been observed already the retransmission is unnecessary (Un.RTO, Un.FR). If the segment has been observed and acknowledged previously, and TCP window probing is conducted, it is classified as flow control (FC). If the segment has not been observed yet, and it is unlikely that the sender has invoked any of its retransmission mechanisms, it is reordered (Reo).

In addition to these categories, Tstat may also classify segments as “unknown” (Unkn). This classification is used whenever Tstat is unable to infer the exact cause of an OOS segment. As Tstat’s OOS segment classifier uses IETF’s TCP standards when calculating e.g., the RTO of a connection, segments can be OOS for no obvious reason. This happens as there are variations between different TCP implementations, both in logic and in configuration.

B. OOS Overview

Figure 3 shows the classification of all OOS segments in our measurements. In total, 1.6% of all segments were OOS, for both the 2006 and the 2009 data. This is about the same as, or slightly lower than, figures reported in related measurements: in [12] 0.9% – 7.1% of all packets in seven different traces were OOS, and in [13] an average of 5% – 8% were reported. The differences between our and related measurements can be of many reasons. For instance, [12] only considered connections with at least 10 segments and [13] reported large variations between different measurement points.

Although the amount of OOS seems to be stable between our measurements, the OOS distributions vary. Three distinctive differences can be observed. First, the amount of OOS segments due to RTOs are more common in 2009 (53.5%)

Table IV
TRAFFIC VOLUME AND OOS BREAKDOWN FOR 2006 TRAFFIC, ALL FIELDS IN % OF TOTAL.

Length	Pkts	Flows	OOS	RTO	FR	Reo	Dups	FC	Un.RTO	Un.FR	Unkn
SHORT	5.78	64.55	21.90	25.12	1.42	18.02	87.52	13.73	38.37	0.17	10.22
MEDIUM	10.48	27.46	17.94	24.75	2.25	11.19	7.18	1.88	37.68	0.67	10.28
LONG	83.74	8.00	60.16	50.13	96.33	70.80	5.30	84.39	23.94	99.17	79.51
BREAKDOWN				43.50	6.48	22.52	2.79	0.30	6.54	0.01	17.85

Table V
TRAFFIC VOLUME AND OOS BREAKDOWN FOR 2009 TRAFFIC, ALL FIELDS IN % OF TOTAL.

Length	Pkts	Flows	OOS	RTO	FR	Reo	Dups	FC	Un.RTO	Un.FR	Unkn
SHORT	5.29	60.10	17.86	20.43	0.98	17.21	64.83	17.57	35.15	0.00	8.37
MEDIUM	10.64	30.64	26.57	33.25	3.08	12.28	21.15	1.95	44.18	5.38	17.00
LONG	84.07	9.25	55.56	46.32	95.94	70.50	14.02	80.47	20.67	94.62	74.63
BREAKDOWN				53.51	7.68	5.19	1.81	0.18	7.89	0.00	23.74

than in 2006 (43.5%). Second, segments classified as OOS due to unknown reasons have increased from 2006 (17.9%) to 2009 (23.7%). The unknown classification is for instance used when a packet seems to be retransmitted but the number of dupACKs are less than three or the estimated RTO has not yet expired. The increase of OOS segments in this category might be related to the increase in TCP implementations that use more aggressive loss recovery mechanisms than the ones standardized in [29], [30]. Linux, for example, uses a minimum allowed RTO of 200 ms, while the standard is 1 s. Furthermore, Windows XP allows segments to be retransmitted after only 2 dupACKs. Finally, the amount of OOS due to packet reordering dropped significantly between the 2006 (22.5%) and 2009 (5.2%) measurements. We will analyze this more thoroughly in Section IV-D.

Although the measured backbone traffic is highly aggregated, it might be misleading to draw conclusions based solely on packet-level observations. A few flows could skew the statistics, such as long-lived high-volume elephant flows where every other packet is OOS. Therefore, OOS segments were also classified on a per-flow basis (see Figure 4). For the measurements in 2006, 17.3% of all flows had at least one OOS segment, and in 2009 16.4% of the flows had at least one OOS segment. The OOS distribution on flow-level is also about the same as on packet-level. Thus, it is not likely that the observed trends are due to a small number of misbehaving flows.

C. OOS Details

Tables IV and V show the 2006 and 2009 OOS distributions for different flow size classes. We present the results in a similar way as Mellia et al. [13]. The figures in the tables correspond to the ratio between the number of specific OOS events occurring in a given flow class over the total number of such OOS events. For instance, 17.21% of all reordered segments in the 2009 measurements (Table V) occurred in short flows.

Three different flow size classes are used in the tables.

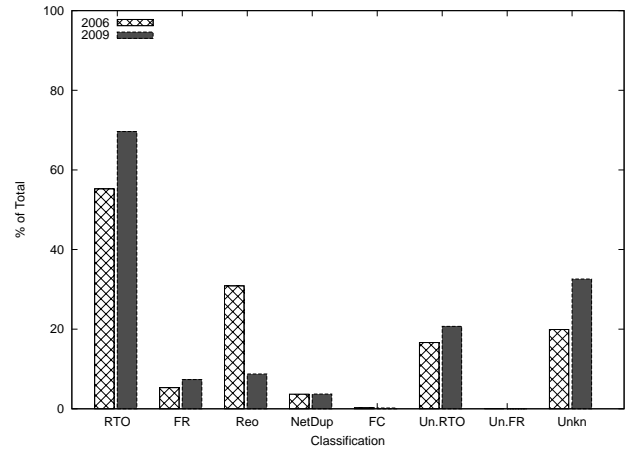


Figure 4. OOS classification of TCP flows in 2006 and 2009. The significant decrease in reordered segments, from 2006 to 2009, is the main difference between the data sets.

Short flows are flows including not more than five data segments (1-5). Medium sized flows have a minimum of six data segments and a maximum of 20 data segments (6-20). Long flows have a payload that is larger than 20 data segments (>20). The bottom lines of the two tables show the total occurrence of the different OOS classes. Thus, the bottom lines convey the same information as Figure 3.

The flow size distributions are quite similar in 2006 and 2009, moving slightly towards longer flows in 2009. About 60% of all flows are short (containing 5–6% of all packets), 27–30% of the flows are medium sized (containing 10% of all packets) and 8–9% of the flows are long (containing 83–84% of all packets). The amount of OOS segments in the different flow size classes have also shifted slightly, making medium sized flows subjected to more OOS in 2009 (from 17.94% to 26.57%), while the OOS in short and long flows has decreased somewhat. This is also visible when inspecting the different OOS categories where the medium sized flows now have a larger portion of almost every OOS

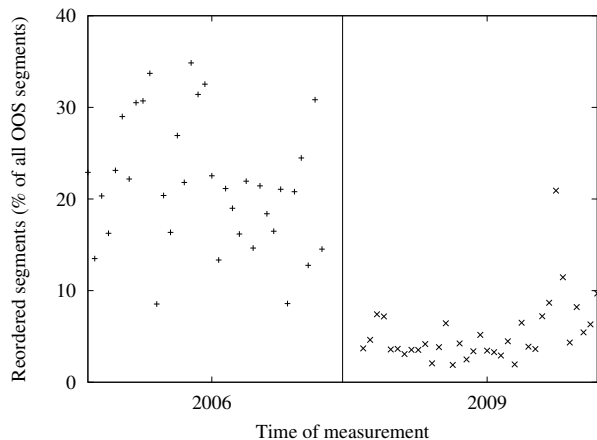


Figure 5. Reordered segments as a fraction of all OOS segments, per measurement.

category. In general, it is evident that the relatively small amount of packets belonging to short flows account for a large portion of all OOS segments (17 – 22%). The large flows, which represent a vast majority of all the traffic do only account for 55–60% of the OOS segments. Thus, short flows seem to be more subjected to OOS than medium size and long flows.

When comparing 2006 to 2009, the most interesting difference is the total distribution of OOS segments (the bottom line in the tables). For instance, the portion of RTOs and unknown events has increased while reordering has decreased significantly (from 22.52% to 5.19%). While the portion of reordering in 2006 are comparable with previous studies, e.g., Mellia et al. [13] found reordering to be responsible for 28.12% of all OOS segments in traces from 2004, the 2009 results displays a surprisingly low amount of reordering. It is very hard to find specific reasons to why these OOS categories have changed so significantly between 2006 and 2009. The increase in RTOs and unknown events might be a consequence of more recent, and aggressive, retransmissions schemes; RTO timers that expire before fast retransmit can be invoked; fast retransmit algorithms that requires less than three dupACKs. For the decrease in reordering, it is even harder to speculate about causes.

D. Packet Reordering

Packet reordering can occur for a number of reasons including e.g., multi-path routing and parallelism within routers [9], [10]. As mentioned in the introduction, novel networking technologies that now are being deployed use some of these techniques and are thus believed to create packet reordering in their mode of operation [31]. To mitigate the negative effects of reordering, a number of reordering robust TCP versions have been developed during the last years (e.g., [31]–[33]). Most of these proposals do, however, not inhibit the actual reordering but merely the neg-

ative effects reordering poses on TCP performance. It would therefore be intuitive that reordering might have increased the last years. Our measurements, on the other hand, indicate a significant *decrease* in packet reordering. It is, however, important to remember that novel networking technologies that might lead to increased reordering often are employed in specialized networks, and that such deployments do not affect backbone traffic that much.

Figure 5 shows the fraction of reordered segments (in % of all OOS segments) for our 70 collected traces in 2006 and 2009. As shown in the graph, the amount of reordering, and also the variance between the traces, is much smaller in the 2009 traces. For the 2006 traces the fraction of reordered segments goes from approximately 10% to 35%. The 2009 traces display a rather stable amount of reordering around 1 – 8%, except for a few outliers.

To rule out that reordering events are induced by a few specific networks (or a few specific routers), we looked at the amount of reordering events per /16 and /24 network prefix² (class B and class C networks). In 2006, the average number of reordering events per /16 network was 177, with an 95% confidence interval of [139, 215]. In 2009, the corresponding figure was only 45 with an 95% confidence interval of [32, 57]. In 2006, the average number of reordering events in /24 networks was 45, with a 95% confidence interval of [33, 58]. In 2009, the average number of reordering events per /24 network was only 15, with a corresponding confidence interval of [12, 19]. The mean values together with the confidence intervals lead us to the conclusion that the significant decrease in reordering events, between 2006 and 2009, was not caused by a few misbehaving routers/networks.

To further rule out temporal events as the cause for the decrease in packet reordering, we investigated the temporal distributions. Figure 6 shows OOS events during one measurement in 2009. The y -axis shows the number of OOS segments during each millisecond of the measurement (x -axis). The white bottom part of the graph shows the fraction of reordered segments. The frequency of OOS (and reordering) events does not follow any particular mode, but rather appear as noise. Although the graph only shows one trace, this type of distribution of OOS events is representative for all traces. Since we are not able to attribute the decrease in packet reordering to any specific network event, we speculate that modern networking equipment has been more carefully designed not to introduce packet reordering, e.g., by taking routing decisions on flow or IP-pair level rather than on individual packet level.

V. CONCLUSIONS

To reveal recent trends in TCP packet-level characteristics we have measured and compared highly aggregated back-

²Note that the applied prefix preserving IP address anonymization allows this type of analysis.

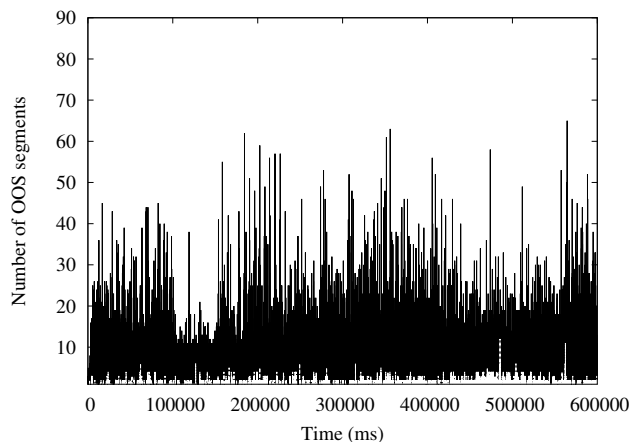


Figure 6. OOS segments during one of the 10 minute measurements in 2009.

bone traffic from 35 traffic traces collected in 2006 with 35 corresponding traffic traces collected in 2009. The analysis shows that although TCP is still the dominating transport protocol, the use of UDP has increased significantly from 2006 to 2009. Furthermore, the use of TCP options like WS, SACK and TS have also continued to increase over these years. The most frequently used TCP option is the MSS option, which is used in over 99% of all TCP connections.

We have also looked at TCP OOS deliveries and found that although the relative amount of OOS deliveries is stable, OOS caused by packet reordering has decreased significantly from 2006 to 2009. The change does not seem to be due to a few misbehaving hosts/routers or due to any major temporal event, but rather a general development. This is an interesting result, as many researchers have speculated in the increase of packet reordering due to novel networking technologies that create packet reordering as a side effect. These novel technologies are, however, mostly deployed in specialized networks and maybe therefore not prone to affect Internet backbones significantly. It is, however, important to continue considering the levels of packet reordering in backbones, as specialized networks and their mechanisms will be incorporated into the regular Internet.

ACKNOWLEDGEMENTS

This work was supported by SUNET, the Swedish University Computer Network.

REFERENCES

- [1] W. John and S. Tafvelin, "Analysis of internet backbone traffic and header anomalies observed," in *IMC '07: Proceedings of the 7th ACM SIGCOMM Conference on Internet Measurements*, San Diego, CA, USA, October 2007, pp. 111–116.
- [2] K. Thompson, G. Miller, and R. Wilder, "Wide-area internet traffic patterns and characteristics," *IEEE Network*, vol. 11, no. 6, pp. 10–23, 1997.
- [3] S. McCreary and K. Claffy, "Trends in wide area IP traffic patterns - a view from ames internet exchange," CAIDA, SDSC, UCSD, Technical Report, 2000.
- [4] C. Fraleigh, S. Moon, B. Lyles, C. Cotton, M. Khan, D. Moll, R. Rockell, T. Seely, and C. Diot, "Packet-level traffic measurements from the sprint IP backbone," *IEEE Network*, vol. 17, no. 6, pp. 6–16, 2003.
- [5] K. Pentikousis and H. Badr, "Quantifying the deployment of TCP options - a comparative study," *IEEE Communications Letters*, vol. 8, no. 10, pp. 647–649, October 2004.
- [6] M. Allman, "A web server's view of the transport layer," *SIGCOMM Comput. Commun. Rev.*, vol. 30, no. 5, pp. 10–20, 2000.
- [7] G. Maier, A. Feldmann, V. Paxson, and M. Allman, "On dominant characteristics of residential broadband internet traffic," in *IMC'09: Proceedings of the 9th ACM SIGCOMM Conference on Internet Measurements*, 2009.
- [8] F. Qian, A. Gerber, Z. M. Mao, S. Sen, O. Spatscheck, and W. Willinger, "TCP revisited: a fresh look at TCP in the wild," in *IMC'09: Proceedings of the 9th ACM SIGCOMM Conference on Internet Measurements*, 2009.
- [9] V. Paxson, "End-to-end internet packet dynamics," *IEEE/ACM Trans. Netw.*, vol. 7, no. 3, pp. 277–292, 1999.
- [10] J. C. R. Bennett, C. Partridge, and N. Shectman, "Packet reordering is not pathological network behavior," *IEEE/ACM Trans. Netw.*, vol. 7, no. 6, pp. 789–798, 1999.
- [11] L. Gharai, C. Perkins, and T. Lehman, "Packet reordering, high speed networks and transport protocol performance," in *13th International Conference on Computer Communications and Networks (ICCCN'04)*, Chicago, IL, USA, October 2004.
- [12] S. Rewaskar, J. Kaur, and D. F. Smith, "A passive state-machine approach for accurate analysis of TCP out-of-sequence segments," *SIGCOMM Comput. Commun. Rev.*, vol. 36, no. 3, pp. 51–64, 2006.
- [13] M. Mellia, M. Meo, L. Muscariello, and D. Rossi, "Passive analysis of TCP anomalies," *Computer Networks*, vol. 52, no. 14, pp. 2663–2676, 2008.
- [14] J. Xu, J. Fan, M. Ammar, and S. B. Moon, "On the design and performance of prefix-preserving IP traffic trace anonymization," in *IMW '01: 1st ACM SIGCOMM Workshop on Internet Measurement*, San Francisco, CA, USA, 2001, pp. 263–266.
- [15] S. Donnelly, "Endace DAG Timestamping Whitepaper," 2007, Endace Technology Limited.
- [16] W. John, S. Tafvelin, and T. Olovsson, "Passive internet measurement: Overview and guidelines based on experiences," *Computer Communications*, vol. 33, no. 5, pp. 533–550, 2010.
- [17] K. Keys, D. Moore, R. Koga, E. Lagache, M. Tesch, and K. Claffy, "The Architecture of CoralReef: an Internet Traffic Monitoring Software Suite," in *A workshop on Passive and Active Measurements*, 2001.

- [18] Politecnico di Torino, <http://tstat.tlc.polito.it/index.shtml>, October 2008.
- [19] S. Jaiswal, G. Iannaccone, C. Diot, J. Kurose, and D. Towsley, "Measurement and classification of out-of-sequence packets in a tier-1 IP backbone," *IEEE/ACM Trans. Netw.*, vol. 15, no. 1, pp. 54–66, 2007.
- [20] C. Shannon, D. Moore, and K. Claffy, "Beyond folklore: Observations on fragmented traffic," *IEEE/ACM Trans. Netw.*, vol. 10, no. 6, pp. 709–720, 2002.
- [21] J. Postel, "The TCP maximum segment size and related topics," *Internet RFCs, ISSN 2070-1721*, vol. RFC 879, November 1983.
- [22] M. Zhang, M. Dusi, W. John, and C. Chen, "Analysis of UDP Traffic Usage on Internet Backbone Links," in *Ninth Annual International Symposium on Applications and the Internet (SAINT)*, Seattle, USA, 2009, pp. 280–281.
- [23] W. John, "Characterization and Classification of Internet Backbone Traffic," Department of Computer Science and Engineering, Chalmers University of Technology, Göteborg, SE, Tech. Rep., 2010, Doctoral Thesis, ISBN 978-91-7385-363-7.
- [24] S. Kent, "IP encapsulating security payload (ESP)," *Internet RFCs, ISSN 2070-1721*, vol. RFC 4303, December 2005.
- [25] European Parliament, "Directive 2004/48/EC of the European Parliament and of the Council," 2004, <http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=OJ:L:2004:157:0045:0086:EN:PDF> (accessed 2010.01.18).
- [26] Y. Zhang, L. Breslau, V. Paxson, and S. Shenker, "On the characteristics and origins of internet flow rates," *SIGCOMM Comput. Commun. Rev.*, vol. 32, no. 4, pp. 309–322, 2002.
- [27] K. Lan and J. Heidemann, "A measurement study of correlations of internet flow characteristics," *Computer Networks*, vol. 50, no. 1, pp. 46–62, 2006.
- [28] W. John, S. Tafvelin, and T. Olovsson, "Trends and Differences in Connection-behavior within Classes of Internet Backbone Traffic," *Passive and Active Network Measurement Conference (PAM)*, pp. 192–201, 2008.
- [29] V. Paxson and M. Allman, "Computing TCP's retransmission timer," *Internet RFCs, ISSN 2070-1721*, vol. RFC 2988, November 2000.
- [30] M. Allman, V. Paxson, and E. Blanton, "TCP congestion control," *Internet RFCs, ISSN 2070-1721*, vol. RFC 5681, September 2009. [Online]. Available: <http://www.rfc-editor.org/rfc/rfc5681.txt>
- [31] S. Bohacek, J. P. Hespanha, J. Lee, C. Lim, and K. Obraczka, "A new TCP for persistent packet reordering," *IEEE/ACM Trans. Netw.*, vol. 14, no. 2, pp. 369–382, 2006.
- [32] S. Bhandarkar, A. Reddy, M. Allman, and E. Blanton, "Improving the robustness of TCP to non-congestion events," *Internet RFCs, ISSN 2070-1721*, vol. RFC 4653, August 2006.
- [33] A. Sathiaselalan and T. Radzik, "Reorder notifying TCP (RN-TCP) with explicit packet drop notification (EPDN)," *International Journal of Communication Systems*, vol. 19, no. 6, pp. 659–678, 2005.