

Differences between in- and outbound Internet Backbone Traffic

Wolfgang John and Sven Tafvelin

Department of Computer Science and Engineering, Chalmers University of Technology, Göteborg, Sweden
e-mail: {johnwolf,tafvelin}@chalmers.se

Abstract

Contemporary backbone-traffic is analyzed with respect to behaviour differences between inbound and outbound Internet traffic. For the analysis, 146 traffic traces of 20 minutes duration have been collected in April 2006, carrying 10.7 billion frames and 7.5 TB of data. Significant directional differences, among others found in IP fragmentation, TCP termination behaviour and TCP options usage, are pointed out and discussed on different protocol levels (IP, TCP and UDP). The analysis includes a focus on TCP connection properties, yielding P2P and malicious traffic as main reasons for the differences. The results are relevant for a better understanding of how applied network protocols are used in an operative environment. Furthermore, a quantification of malicious traffic provides related research fields, such as network security or intrusion detection, with important insights.

Keywords

Internet Measurement; Directional Traffic Differences; TCP Connection Analysis; Network Anomalies;

1. Introduction

The Internet, as emerging key component for commercial and personal communication, has in the recent years undergone a fast development and is still expanding. Unfortunately, this rapid development has left little time or resources to integrate measurement and analysis possibilities into Internet infrastructure, applications and protocols. However, the Internet community needs to understand the nature of Internet traffic in order to support research and further development [3]. One way to acquire better understanding is to measure real Internet traffic. In the MonNet project [10][24], the technical and legal complications of the measurement task were overcome resulting in packet-level traces of contemporary Internet traffic.

The MonNet traffic traces analyzed in this article have been taken from the OC192 backbone of the Swedish University Network (SUNET) during 20 days in April 2006. The links tapped provide not only a backbone for two major Universities, but also for a substantial number of student dormitories and research institutes. Additionally, the links carry exchange traffic with commercial providers due to a local exchange point inside Göteborg. Because of the high aggregation of the measured links, we believe that this recent data provides a valid footprint of Internet traffic characteristics in Sweden at the current time.

The chosen measurement point on the outermost part of a ring architecture makes the traces specifically suitable for highlighting directional differences. Put simply, the measurements were taken on links between the region of Göteborg and the rest of the Internet. This work therefore analyzes the contemporary data with respect to behaviour differences between in- and outbound backbone traffic. The presented traffic constitutes a medium level of aggregation, between campus-wide traffic and tier-1 backbone traffic. We believe that this type of network, with smaller local exchange points, represents an upcoming class of networks.

1.1. Related work

There are numerous articles about general Internet measurements [7][16][25], with only a few of them partly dealing with directional differences. Thompson [25] e.g. presented wide-area Internet traffic characteristics on nowadays rather outdated data in 1997. The data was recorded on a core-backbone and a transatlantic link, including figures about directional differences in packet sizes and byte volumes.

* This work was supported by SUNET, the Swedish University Network

In recent years, a few studies included discussions about directional differences, but typically only regarding specific properties. These articles are often based on unidirectional flow data and analyze a variety of datasets. The analyzed datasets are either collected at Tier-1 backbone level or on small campus or institute Internet gateways, so with either a low or very high level of aggregation. In his article about rapid model parameterization, Lan [14] showed differences between inbound and outbound traffic in terms of protocol mix and flow statistics, like flow size and duration. The data was recorded on the 100 Mbit/s Internet gateway of the USC Information Science Institute in 2001. Saroiu [22] analyzed different types of HTTP flows, recorded on two border routers of the University of Washington on 9 days in 2002. In this paper, WWW and P2P traffic carried over HTTP are contrasted, including a comparison of inbound and outbound flows. In his study about P2P properties in 2003, Gerber [8] was able to show that the IN/OUT traffic balance for P2P traffic on the border of a Tier-1 backbone is close to one. Kim [12] compared inbound and outbound flow statistics for different transport protocols, including flow, packet and byte ratios. The analysis was based on flow data collected in 2004 on the Internet routers of the POSTECH campus, a 2x100 Mbits/s Ethernet.

An interesting study based on packet-level traces was presented by Mellia [17]. Mellia analyzed traces collected on the Internet access link of the Politecnico di Torino campus LAN in 2000-2002. Besides presenting an automatic tool for statistical analysis of network traces, interesting results for IP and TCP characteristics are given, including a connection-level analysis of TCP.

1.2. Contribution of this work

Updated measurement results are crucial for a better understanding of how the applied technologies and protocols are used in an operative environment. In the present study, significant directional differences are pointed out and discussed on different protocol levels (IP, TCP and UDP). For TCP, the bi-direction packet level traces are reassembled to connections, in order to be able to conduct a detailed connection-level analysis. The presented results are destined for network engineers, network application developers and protocol designers in order to be able to optimize bandwidth efficiency and stability of future networks. The paper furthermore highlights network anomalies and inconsistencies, like attacking or scanning traffic. This is important knowledge, since improving the robustness of network applications and protocol implementations is gaining special importance. In fact, increasing bandwidth and growing numbers of Internet users have also lead to increased misuse and anomalous behaviour [9][13]. Knowledge of real-life traffic properties is also important for establishing more realistic simulation models [6]. Finally, some of the insights might as well bring up new research issues in related research fields, such as network security and intrusion detection. The contributions of this work are relevant, because:

- the analysis is based on updated, contemporary data
- the data was collected on links representing medium traffic aggregation, a class of networks not previously studied in the same extent
- packet-level traces allow a more detailed analysis than sampled flow-level data (e.g. TCP options)
- the presented bi-directional TCP connection analysis reflects real connections more closely than traditional flow level analysis
- the results provide a complete view of directional differences on different levels (IP, TCP, UDP)
- the special focus on network anomalies is especially important in the light of increasing amounts of network attacks

The paper is outlined as follows. Section 2 describes the methodology of collecting, pre-processing and analyzing the traces. Then some general traffic properties are presented in section 3. Next, sections 4, 5 and 6 quantify directional differences observed on different protocols levels (IP, TCP and UDP). Finally, in section 7, different traffic anomalies and inconsistencies found on the protocol levels are summarized, followed by concluding remarks about the main findings.

2. Methodology

2.1. Collection of traces

We collected our traces on the outermost part of an SDH ring running Packet over SONET (PoS). The traffic passing the ring to (outbound) and from (inbound) the main Internet is primarily routed via our tapped link, as confirmed by SNMP statistics. Simplified, we regard the measurements to be taken on links between the region of Göteborg, including exchange traffic with the regional access point, and the rest of the Internet as discussed earlier in section 1.

We use optical splitters on two OC-192 links, one for each direction. The splitters are attached to two Endace DAG6.2SE cards sitting in identical Dual-Opteron servers. The servers use a 6 disk SCSI Raid0 to keep up with the speed of the 10 Gbit/s links. The DAG cards are configured to capture the first 120 bytes of each frame to ensure that the entire network- and transport header information is preserved. The two DAG cards are chained together with help of the DAG Universal Clock Kit (DUCK), with one card serving as synchronisation input for the second card, resulting in time synchronisation typically between ± 30 ns [5].

The collection of the data was performed between the 7th of April, 10:00 and the 26th of April 2006, 10:20. During this period, we simultaneously for both directions collected four traces of 20 minutes each day at identical times. The times (02:00, 10:00, 14:00 and 20:00) were chosen to cover business, non-business as well as night time hours. Due to measurement errors in one direction at four occasions we have excluded these traces and the corresponding traces in the opposite direction.

2.2. Processing and analysis

After storing the data on disk, the payload beyond transport layer was removed and the traces were sanitized and desensitized. This was mainly done by using available tools like Endace's dagtools and CAIDA's CoralReef, accompanied by own tools for additional consistency checks, which have been applied after each pre-processing step to ensure sanity of the traces. Trace sanitization refers to the process of checking and ensuring that the collected traces are free from logical inconsistencies and are suitable for further analysis. During our capturing sessions, the DAG cards discarded a total of 20 frames within 12 different traces due to receiver errors, which includes HDLC CRC errors. Surprisingly, another 71 frames within 30 different traces had to be discarded after the sanitization process due to IP checksum errors.

By desensitization we mean the removing of all sensitive information to ensure privacy and confidentiality. The payload of the packets was removed earlier, so we finally anonymized IP addresses using the prefix preserving CryptoPAN [27]. After desensitization, the traces were moved to a central storage server. First, an analysis program was run on each trace to extract cumulated statistical data. As a second step, per-connection TCP analysis was conducted on merged, then bidirectional traces. More details on the connection analysis are described in beginning of section 5.

3. General traffic characteristics

As summarized in table 1, the 146 analyzed traces sum up to 10.68 billion PoS frames, containing a total of 7.53 TB of data. In his study on campus wide traffic, Kim [12] reported about a 1:1 ratio between outbound and inbound traffic for packets numbers, but an 1:1.38 inequality for traffic volume due to the "net provider" status of University networks. In the our data, no significant difference between neither, packet counts nor volumes, can be observed. This even distribution of traffic proves the higher level of aggregation and underlines the relevance of the presented data, representing Internet backbone traffic.

The frames contain in 99.97% of the cases IPv4 packets, which sum up to 99.99% of the carried data. The remaining traffic consists constantly of around 1200 IPv6 BGP Multicast messages, 8 CLNP routing updates (IS-IS) and 1 Cisco Discovery Protocol (CDP) message per minute. The results in the remainder of this paper are based on the IPv4 traffic only.

	Packets		Data	
	Count	%	Volume	%
Total	10.68E+9	100.00%	7.53 TB	100.00%
Outbound		48.74%		49.16%
Inbound		51.26%		50.84%

Table 1: Traffic amount of data captured

	Total		Outbound		Inbound	
	Inside	Outside	Source	Dest.	Dest.	Source
Total	634E+3	22.0E+6	275E+3	19.2E+6	490E+3	19.8E+6
TCP	408E+3	05.0E+6	176E+3	04.3E+6	310E+3	04.5E+6
UDP	484E+3	19.2E+6	175E+3	16.4E+6	384E+3	16.9E+6
Rest	155E+3	01.9E+6	024E+3	01.1E+6	146E+3	01.0E+6

Table 2: Distinct IP addresses seen

4. IP level

In this section out- and inbound traffic on the network layer level of the Internet Protocol (IP) is compared. This comparison includes the transport protocol mix, IP packet size distribution and IP fragmentation.

To start with, table 2 gives some scale to the aggregation level of the links. The numbers of distinct IP hosts seen within (inside) and outside the region of Göteborg are summarized, where outbound sources and inbound destinations are regarded as inside, and the opposite way around as outside. Note that the sum of the numbers exceeds the total numbers, since one host can obviously be both source and destination for packets of several transport protocols. As expected, the amount of hosts inside the region is outnumbered by hosts seen outside “in the Internet”. Nevertheless, there is a surprisingly high number of hosts inside, considering that the numbers of hosts at the three main customers of the links (2 major universities and the regional network for student dormitories) do not exceed 7,000 each. Indeed, these main customers sum up to about 21,000 sources of outbound TCP connections. The remaining 150,000 outbound sources belong to different providers connected to the exchange point. The amount of inbound destinations is much larger due to incoming scanning traffic. As an example, the 16 bit address ranges of the two Universities are scanned in their entirety (2x65,534). The vast amount of UDP hosts outside was found to be due to short UDP sessions caused by P2P overlay networks, which will be discussed in more detailed in section 6.

It has to be noted that even though the hosts of the three main customers represent a minor part (13%) of the observed IP addresses inside the region of Göteborg, a majority of the traffic (around 85%) consists of packets to or from these hosts.

4.1. Transport protocol breakdown

The protocol breakdown, summarized in table 3, once more confirms the dominance of TCP traffic. Compared to earlier measurements [7][16][23][25], the fractions of both TCP data volume and packet counts have even increased slightly. In the table, fractions of packets and data carried in the respective protocol are in % of total IPv4 traffic for the corresponding direction. Ratios between out- and inbound traffic are shown in parentheses, summing up to 100 for each protocol.

TCP packets and data show an equal ratio, as it was the case for the total traffic. In Kim’s report [12], outbound traffic carried 1.44 times more data than inbound traffic. We believe that this behaviour is not observed in our data since the traffic of diverse network types aggregating on the links measured cancel out the typical client-server imbalance (small requests, large data replies). UDP data on the other hand shows almost the same ratio (38:62) in favour of inbound data volumes in our data as previously reported by Kim. This is caused by multimedia traffic (mainly RTP) over UDP, which is more common to be served on hosts on the Internet. An interesting observation can be made for UDP packets, with an unexpected large amount of outgoing packets. A closer look reveals that three consecutive measurements carried up to 58% UDP packets, as shown in table 4. This indicates a potential single UDP burst of 14-24 hours of time. A detailed analysis shows that the packet length for the UDP packets causing the burst was just 29 bytes, leaving a single byte for UDP payload data. These packets were transmitted between a single sender and receiver address with varying port numbers. After reporting this network anomaly, the network support group of a University could identify the culprit host. This was a web server that had been exploited through a known vulnerability in a PHP script. Consequently, a UDP DoS script was installed and could run undetected, since

the network management tool was monitoring amount of per-flow data only, but not the number of packets. Although TCP data was still predominant, we believe that a dominance of UDP packets over such a time span could potentially lead to TCP starvation and raise serious concerns about Internet stability and fairness. When removing the three traces with this outstanding network event from our data, UDP packets showed the same ratio as the TCP and the overall data. Due to the small packet sizes, the ration of UDP data kept almost unchanged (36:64).

ESP traffic seems to experience a typical client-server pattern with even packet ratio, but uneven data proportions. The hosts mainly responsible for this type of traffic will be discussed again in section 4.3. An explanation for the dominance of outbound traffic for ICMP could be the large amount of incoming network attacks as shown later, triggering ICMP responses from routers and firewalls.

	Fraction of Packets in %		Fraction of Data in %	
	outbound	inbound	outbound	inbound
TCP	90.62 (48.1)	93.14 (51.9)	97.76 (49.5)	96.57 (50.5)
UDP	8.87 (56.8)	6.40 (43.2)	2.03 (37.8)	3.23 (62.2)
ESP	0.23 (52.5)	0.20 (47.5)	0.12 (66.5)	0.06 (33.5)
ICMP	0.22 (61.9)	0.13 (38.1)	0.02 (60.7)	0.02 (39.3)
GRE	0.05 (51.8)	0.05 (48.2)	0.07 (73.0)	0.02 (27.0)

Table 3: Protocol mix (ratios per protocol in parenthesis)

Packet size	total	outbound	inbound
20-39	1.50%	2.96%	0.11%
40-60	38.72%	37.26%	40.12%
576	0.96%	0.60%	1.29%
628	1.76%	2.05%	1.49%
1300	1.11%	1.20%	1.01%
1400-1500	38.01%	37.66%	38.34%

Table 5: Major modes of IPv4 packet size distribution for all data (left) and without UDP burst (right)

Date	Time	outbound	
		Packets	Data
2006-04-16	14:00	6.8%	1.7%
2006-04-16	20:00	40.6%	5.1%
2006-04-17	02:00	51.9%	6.1%
2006-04-17	10:00	58.1%	7.1%
2006-04-17	14:00	5.7%	1.8%

Table 4: UDP burst

Packet size	total	outbound	inbound
20-39	0.14%	0.18%	0.11%
40-60	39.25%	38.41%	40.02%
576	0.98%	0.63%	1.30%
628	1.79%	2.12%	1.49%
1300	1.13%	1.25%	1.01%
1400-1500	38.53%	38.62%	38.45%

4.2. Packet size distribution

While cumulative distribution of IPv4 packet sizes was reported to be trimodal in earlier measurements [7][16][23][25], more recent studies showed that it has changed to be rather bimodal [21]. The two major modes are small packet sizes just above 40 bytes (TCP acknowledgements) and large packets around 1500 (Ethernet MTU). The previous third mode of 576 bytes (default size according to RFC 879) has in our data decreased to less than 1%. Furthermore, we found that two other notable modes appeared at 628 bytes and 1300 bytes. In table 5 the major modes are summarized, with an extra table excluding the above mentioned UDP burst. As discussed in a prior study on the SUNET datasets [10], the mode at 628 bytes is an artefact of 'TCP layer fragmentation' applied by file sharing protocols like Bittorrent or DirectConnect, where 628 byte large packets typically appear after full sized packets in order to add up to 2KB blocks of data. The mode at 1300 bytes could be explained by the recommended IP MTU for IPsec VPN tunnels [4].

The studies of Thompson, Kim and Mellia [12][17][25] report about directional differences in packets sizes on two different levels of link aggregation, both caused by the classical client-server pattern. In contrast, in the SUNET data the two main modes for small and large packets show no significant directional differences. This might be due to two different reasons:

- since Thompson's report of 1997, network applications have undergone some fundamental developments
- compared to the campus-wide data of Kim and Mellia, our backbone data contains a higher aggregated traffic mix

Directional differences however can be observed for two other packet sizes. The differences between fractions of 628 byte sized packets are likely to be caused by popular P2P servers inside Göteborg's student network. It is well known that DirectConnect, but also Bittorrent are especially popular in Sweden, and

consequently also in the region of Göteborg. The cause for the difference in the default datagram size of 576 bytes is not obvious, but we think it might be caused by a better utilization of the Path MTU discovery feature in the comparable well configured hosts inside University and student networks.

4.3. IP fragmentation

Earlier studies of McCreary and Shannon [16][23] indicated an increase in the fraction of IP packets carrying fragmented traffic of up to 0.67%. We found a much smaller fraction of only 0.065% of fragmented traffic in the analyzed data, as shown in table 6. It can be noted that 72% of the fragmented traffic in our data is transmitted during office hours, at 10AM and 2PM. While Shannon, analyzing data of three different locations in 2001, found that fragmented data was equally distributed between out-and inbound data, the amount of fragmented traffic on the SUNET inbound link is about 9 times higher than on the outbound one. Where UDP and TCP is responsible for 97% and 3% respectively of all incoming fragmented segments, they just represent 19% and 18% of the outgoing. The remaining 63% outgoing fragmented traffic turned out to be IPsec ESP traffic (RFC 4303) between exactly one source and one receiver at working hours on weekdays. Each fragment series in this connection consists of one full length Ethernet MTU and one additional 72 bytes fragment. This could easily be explained by an unsuitably configured host/VPN combination transmitting 1532 byte (1572-40 byte additional IP and TCP header) instead of the Ethernet MTU due to the additional ESP header. The dominance of UDP among fragmented traffic is not surprising since Path MTU Discovery is a TCP feature only.

	Total	outbound	inbound
Total	0.065% (100.0%)	0.014% (100.0%)	0.113% (100.0%)
TCP	(4.5%)	(18.0%)	(2.9%)
UDP	(88.6%)	(18.8%)	(97.1%)
ESP	(6.8%)	(63.1%)	(0.0%)

Table 6: Fractions of IPv4 fragments

The first approach to explain the differences is based on the fact that the probability for a packet to be fragmented is increasing with each hop. According to a TTL analysis of the fragmented traffic, the average hop count for outbound traffic was 6.77, whereas the average hop count for inbound traffic was 9.43. This alone does not seem to be significant enough to explain the imbalance between inbound and outbound fragments. We believe that another possible explanation could again be the fact that SUNET and its connection networks are very well configured and administered compared to Internet standards.

5. TCP level

In order to conduct a detailed connection level analysis on TCP, we merged the tightly synchronized unidirectional traces. From the resulting bidirectional traces an analysis program collected per-connection data, including packet and data counts for both directions, start- and end times, TCP flags and counters for erroneous packet headers and multiple occurrences of special flags like RST or FIN. We define a connection by the classical tuple of IP addresses and ports for source and destination. A TCP connection starts with the observation of the first SYN segment and is closed by either one FIN segment for each direction or one RST segment. Additional SYN segments for one tuple can sometimes be seen in the same direction, most commonly within scanning campaigns. In this case, further “connections” are opened within the analysis program in order to record the pure SYN packets separately. The following non-pure-SYN packets are always recorded within the most recently opened connection. We decided not to use a timeout threshold for unclosed connections, since our traces are limited to 20 min duration anyhow.

A significant part of the traffic is routed asymmetrically, due to hot-potato routing. 8% of the TCP data was sent via the outgoing link, without any corresponding TCP packets seen on the incoming. Asymmetrical

traffic on the incoming link was even more common, accounting for 20% of the observed TCP data. Knowing the prefixes of the SUNET network segments in the area of Göteborg, it was possible to show that around 14% of the TCP data is actual transit traffic with neither source nor destination being SUNET customers inside Göteborg, entering the links via the local exchange point. Of the transit traffic, 67% was asymmetrical traffic, which means that 1/3 of all asymmetrically routed traffic is transit traffic as well. In the following subsections, first, TCP connections are classified according to their connection setup and termination behaviour. Then, connection properties like packet count, byte size and lifetime are analyzed with respect to connection direction. Finally, TCP options are discussed in the rather novel approach of per-connection information for SYN requests and replies.

5.1. Connection breakdown

The following tables summarize the connection breakdown for TCP in all 146 traces. The analysis database recorded a total of 72.6 Million connections according to our definition. Additional 8.9 million bidirectional flows do not include an initial SYN segment, which means that they either start before the measurement times or have asymmetrical properties. One million of these flows include no SYN, FIN or RST segments but show packets in both directions, which means that about 3.4% of the established connections last longer than 20 minutes. However, this small number of long lasting connections carries about 34% of the total TCP data. This is not unexpected, given the observations of Brownlee [2], saying that flows longer than 15 minutes carry more than 50% of the traffic on a link. According to their destination port numbers, the long lasting connections typically carry traffic of different P2P protocols and popular messenger services. The following analysis is performed on TCP connections with initial SYN segments.

	total		outbound		inbound	
	Count	%	Count	%	Count	%
TCP connections	72.6E+6	100.00%	28.0E+6	38.56% (100.00%)	44.6E+6	61.44% (100.00%)
rejected	44.3E+6	60.99%	12.3E+6	(44.04%)	32.0E+6	(71.63%)
established	28.3E+6	39.01%	15.7E+6	(55.96%)	12.7E+6	(28.37%)

Table 7: TCP connection attempt breakdown

rejected connections	44.3E+6	100.00%	12.3E+6	27.84% (100.00%)	32.0E+6	72.16% (100.00%)
scanning - no reply	34.8E+6	78.66%	08.2E+6	(66.74%)	26.6E+6	(83.26%)
asymmetric traffic	04.8E+6	10.84%	02.2E+6	(17.94%)	02.6E+6	(8.10%)
scanning - RST reply	04.3E+6	9.81%	01.7E+6	(13.83%)	02.6E+6	(8.25%)

Table 8: Rejected connection breakdown (no 3-way handshake)

Table 7 presents the total of all TCP connections with initial SYN segments. We define established and rejected connections as connections experiencing a proper 3-way handshake or not, respectively. Outbound in this context means that the initial SYN packet was sent on the outbound link. Inbound consequently means that the connection establishment was initiated outside the region of Göteborg. The tables 8 and 9 summarize the termination properties for rejected and established connections. In the tables, the first line represents the vertically summed values for each respective column of absolute packet counts or relative fractions. The fractions of out- and inbound connections in relation to the total amount of connections are additionally given in the first line, summing up to 100% horizontally.

Arlitt [1] quantified different TCP connection states based on the campus wide traffic recorded at the University of Calgary between 2003 and 2004. He quantified rejected connections with about 25-30% of all TCP connections. Our contemporary data includes much more unsuccessful connection attempts, as shown in table 7. A major difference between the numbers of rejected outbound and inbound initiated connections is evident in table 8. The large amount of unreplied SYN packets on the incoming link was already indicated earlier, when discussing the numbers of distinct IP addresses appearing on the incoming link. These are mainly attacks trying to exploit well known vulnerabilities on ports commonly used by Trojans. The scans

often cover the entire IP ranges of the connected networks inside Göteborg and are likely to be destined for non existing endpoints. Entrance routers to the specific network typically drop this kind of packets, which explains the absence of response packets. In some cases an ICMP response might be triggered, which would explain the larger number of outgoing ICMP packets according to table 3. Regardless of the much higher number of incoming scans, there is also a substantial number of outgoing unreplied connection attempts. More than 70% of the 8.2 Million attempts are sent by hosts within the student network. Note that not all of these attempts are necessary network scans. There is a large fraction of non-malicious outbound connection attempts to non existing hosts, resulting in unsuccessful connection attempts. This is often observed for P2P traffic, where unreliable file-sharing peers are common.

In cases where scanning attempts reach existing hosts on arbitrary port numbers, host-based firewalls should preferably drop the packets, but might in some cases reply immediately with an RST packet. This behaviour is more than twice as common for hosts in the student network as compared to hosts in University networks, which indicate that private Internet hosts are less carefully configured.

Asymmetric traffic was included in the summary for rejected connections (table 8) for reasons of completeness. Naturally, asymmetric traffic can not experience a bidirectional 3-way handshake, which means that we cannot consider this traffic as being established.

	total		outbound		inbound	
	Count	%	Count	%	Count	%
established connections	28.3E+6	100.00%	15.7E+6	55.21% (100.00%)	12.7E+6	44.68% (100.00%)
proper closing (2xFIN)	19.0E+6	66.99%	11.4E+6	(72.87%)	07.6E+6	(59.71%)
FIN and RST outbound	03.2E+6	11.21%	542E+3	(3.46%)	02.6E+6	(20.81%)
FIN and RST inbound	01.7E+6	6.06%	711E+3	(4.54%)	01.0E+6	(7.93%)
single RST	02.2E+6	7.71%	01.6E+6	(9.98%)	620E+3	(4.89%)
FIN, RST in counter dir.	01.2E+6	4.11%	889E+3	(5.67%)	276E+3	(2.18%)
unclosed	01.0E+6	3.63%	487E+3	(3.11%)	540E+3	(4.27%)

Table 9: Established connection termination breakdown

In table 9 finally the 28.3 Million connections with proper 3-way handshake observed are split up into different termination behaviours per direction. Considering the quite even distribution of TCP traffic volumes (table 1) it is somewhat surprising to see around 10% more outbound than inbound established connections. These differences in connection counts are cancelled out in the high level summary by differences in connection properties, as presented in the next subsection.

A major fraction (67%) of the established connections is closed properly by FIN segments in each direction, which seems to be quite low, considering that TCP resets should be a rare event according to the TCP standard (RFC 793). On the other hand, a prior study by Arlitt [1] highlighted that TCP connections are becoming more likely to be closed by RST segments (15%), mainly due to irregular web server and browser implementations. Comparing the behaviour of in- and outbound connection in our data we find that connections opened from inside Göteborg are more likely to be closed by proper FIN handshakes. This is compensated by a higher number of connections involving RST segments on the incoming link. While single RSTs in either direction can still be regarded as proper connection termination, the number of connections closed by FIN, followed by additional RST segments is surprisingly high (more than 30% on the inbound connections), even when considering Arlitts results. In fact, the fractions of connections closed by both FIN and RST segments sent by the client (the originator) are close to Arlitts numbers. (3.5% and 7.9% resp.) and are indeed mainly caused by web traffic. The main surprise is the large numbers of connections terminated by FINs and RSTs sent by the server (the responder), which are unproportionally large for inbound connections, meaning that they are closed by servers inside Göteborg. As main source of this behaviour a handful of hosts inside the student network could be identified, according to their port

numbers serving different kinds of popular P2P protocols. This reset behaviour is probably used to reduce the CPU and memory overhead introduced by connections entering the TIME_WAIT state on peers [1].

The 3.6% of unclosed connections lies close to the fraction of long-lasting connections, quantified in section 5.1. These unclosed connections are indeed mainly long lasting flows, and consequently carry almost 50% of all data carried by established connections. While 50% of these unclosed, long lasting incoming connections show destination port numbers of popular P2P protocols, the same port numbers account only for 10% of the outbound connections.

In addition to the high number of connections consisting of one SYN segment only, we also observed as many as 57 Million connections consisting of RST segments only (not shown in the tables). Of these single RST segments, 96% are seen on the outbound link, almost entirely in asymmetrical fashion, without any incoming segment triggering the resets. Only a handful source/destination pairs are responsible for these segments during short periods of time, so the first suspicion was that this could be reset attacks [26]. However, closer investigation showed no variations in sequence numbers or no other typical symptoms, so TCP reset attacks can be ruled out. We believe that the outbound link could be the return path for an asymmetrical routed denial of service (DoS) attack, generating the observed RST segments. Still, it is surprising that no similar behaviour could be observed to the same extent on the symmetrical routed data.

5.2. Quantification of P2P traffic

Since we expect P2P to have a huge impact on traffic characteristics, we tried to quantify P2P traffic for each direction with a simple port number analysis. Even though it is well known that P2P traffic is trying to hide itself and that port number methods strongly underestimate actual numbers [11][18], we believe that this analysis is still valid for comparing amounts of P2P connections between directions.

A list of common port-numbers for popular file-sharing protocols was identified, specifically for different DirectConnect, Bittorrent, Edonkey and Gnutella implementations. According to these port-numbers, outbound P2P connections carry around 13% of P2P packets and data, while for inbound connections this fraction is about twice as large with around 25%. Note, that these large volumes of data are carried by a small number of connections (less than 1%). Beside the probably quite large underestimation of these numbers, they indicate that P2P traffic is in fact at least about 2 times more common among inbound connections. The high amount of inbound established P2P connections, as already indicated in section 5.1, could be the result of a number of popular P2P peers inside Göteborg. Another possible explanation could be an increasing use of modern P2P clients (like RevConnect) inside Göteborg, triggering reverse connections from peers outside, on the Internet.

5.3. Connection properties

This section provides detailed information about different connection properties such as lifetime, size and packet count. The analysis deals only with bidirectional connections which have been established by a 3-way handshake. Ordering the TCP connections by data volume and number of packets carried shows that a small number of top connections accounts for most of the data and packets. This indicates the characteristically 'elephant and mice phenomenon', saying that the majority of Internet data is carried by a small percentage of large flows, so called elephants [15][20]. More specifically, outgoing connections appear to have less pronounced elephants, since it needs 0.08% and 0.17% to carry 50% of the total amount of data and packets respectively for outgoing connections, while only 0.07% and 0.14% are sufficient for 50% for inbound connections. This directional difference can be described even more clearly, considering that 3.9% of the outbound and as few as 0.9% of the inbound connections carry 90% of the data, and 26.3% and 12.2% respectively carry 90% of the packets seen in the particular direction.

Generally, artefacts of the client-server pattern (small requests, large data replies) can be observed for connections established in both directions. While outbound connections yield an average ratio of 1:1.6 in

favour for incoming data, inbound connections show a higher ratio of 1:1.86 in favour of outgoing data. This means that the smaller number of inbound connections (around 45% of all connections) carries more data and more packets primarily in outgoing direction, according to the client-server pattern. This imbalance is cancelled out to an almost even ratio in the high-level view of sections 3 and 4. The imbalance in connections properties is mainly caused by the larger fraction of heavy incoming P2P connections.

The differences between in- and outbound connections are summarized in table 10 by means of statistical properties. In the table, mean, standard deviation (σ), median and 80th percentile (P80) are given for different connections properties per direction of the initial connection establishment. While mean and σ of connection lifetimes appear to be quite similar for both directions, the values for sizes and packet counts are significantly larger for inbound connections. It needs to be noted that some of the values and figures in this subsection are somewhat biased since the traces are limited to 20 min of duration. Long-lasting connections, which are likely to be elephants, are therefore not taken into account to the full extent. Especially the values for mean and σ can therefore to be considered as an underestimate, while median and P80 are less biased. In order to be able to better interpret median and 80th percentile, we included figures for the distributions of connection lifetimes, sizes and packet counts. Figure 1 illustrates distribution of lifetimes in bins of 1sec. The magnified figure presents the first 25 seconds, with higher resolution of 15.6 ms bins. Figure 2 shows connection size distribution, summarized in bins of 1Kbyte. The insert magnifies the first 9 Kbytes with 20 Byte bin-size. Figure 3 finally illustrates packets counts, including magnification for the first 100 packets.

Property		mean	σ	median	P80
Lifetime in sec	out	18.2	60.7	1.8	16.6
	in	17.3	65.8	0.6	24.8
Size in Kbytes	out	61.0	2362	1.1	2.9
	in	81.5	3298	1.9	8.9
Packet Count	out	81.5	2289	11.5	22.0
	in	113.0	3538	11.5	21.0

Table 10: Statistical properties of TCP Conn.




Figure 1: Conn. lifetimes with 1sec bins (15.6ms bins)

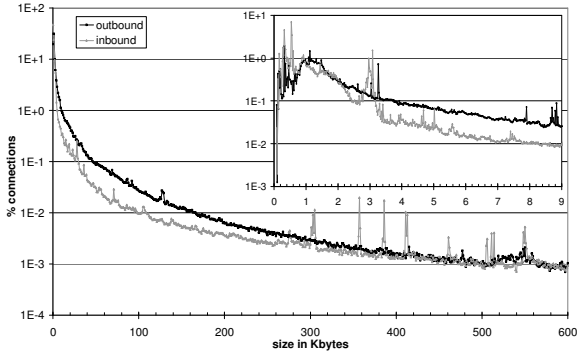


Figure 2: Conn. sizes with 1Kbyte bins (20 Byte bins)




Figure 3: Packets per connection

Mori [19] presented mean values for flow durations on web and P2P flows extracted from inbound campus traffic in 2002. Web traffic yielded 9.5 sec mean, while P2P flow result in a mean of 307 sec. Projecting the values to our data, it can be concluded that the mean values of around 18 sec are a hybrid between web and P2P traffic, which is in fact the case due to University traffic on one hand and private student traffic on the other hand. Considering the underestimated nature of our values, it again indicates a quite substantial amount of long lasting P2P traffic on the measured links. Other studies, including Kim, Lan and Zhang

[12][15][28], presented cumulative distribution figures, reporting of median values of about 1sec and P80 values of around 10 sec. In the SUNET data especially the 80th percentile is significantly larger, again proofing that connections in contemporary traces tend to be significantly longer due to an increased amount of P2P traffic. This property is more pronounced for inbound connections when comparing the P80 values for connection lifetime. Surprisingly, inbound connections do not only tend to be longer, but are also more likely to be shorter than 5 seconds compared to outbound connections. This is indicated by the median values, but can be seen nicely in the magnification of figure 1. The large number of incoming connections in this region can be explained by rejected login attempts on application level, like SSH or SMTP. In general, figure 1 shows a number of protocol timeouts, typically close to half minute or minute borders. For most of the times, the fractions of inbound connections lie below the outbound ones, which is compensated by a higher number of long lasting flows, as discussed earlier.

Regarding connection sizes, Mori [19] also presented mean values with 20.6 Kbytes for web flows, and as large as 5.8 Mbytes for P2P flows. As for lifetimes, the mean values for the presented data lie in between these extreme values. Earlier studies reported about median values of around 1 Kbytes and P80 values of between 1 and 10 Kbytes [15][28], which is similar to our findings. Even though in contrast to connection lifetimes, both median and P80 value are larger for inbound connections, there are peaks in the magnification of figure 2 for incoming connection sizes below 1Kbyte and around 3 Kbytes. According to a port analysis, the former stems from connections trying to exploit a known security hole in some MS SQL server versions on a handful of hosts inside Göteborg, while the latter can be explained by unsuccessful SSH login attempts, probably mainly intrusion attempts as well. Generally, figure 2 shows that inbound connections tend to be less likely to carry small amounts of data, which again indicates that there is a higher number of “elephants” carrying a lot of data on the incoming link. This seems to be connected to the similar behaviour found for connection lifetimes, even though there is not necessarily a strong correlation between duration and size, as reported by Lan and Zhang [15][28]. The spikes for inbound traffic seen in figure 2 between 300 and 550 Kbytes are results of connections from a single host to one host on destination port 2135. This is rather a special application than another security exploit, since except these small connections there are also a larger number of connections carrying a large amount of data between these hosts.

As illustrated in the magnification of figure 3, connections with less than 20 packets show very similar patterns for both directions, consequently resulting in similar median and P80 values. Nevertheless, the differences in the mean values as well as the lower values for the inbound graph in figure 3 shows that packet counts are to some degree correlated with connection sizes. As for connection sizes, the spikes between 20 and 30 packets are artefacts from unsuccessful SSH logins, and the spikes between 300 and 360 stem from the unidentified connections to port 2135.

5.4. TCP option usage

In earlier work, TCP options analysis was typically done by counting occurrences of different TCP options in all SYN and SYN/ACK segments seen in packet-level traces [10][21]. In our current work, the thorough connection analysis allows us to give better insight into options advertisements between clients and servers within single TCP connections. Since this analysis is focused on proper established connections only, attacking and scanning traffic, which might bias simple counts of SYN segments, are filtered out.

Table 11 summarizes TCP option employment for the four major TCP options types typically advertised during connection establishment. Fractions of connections carrying the particular option in SYN or SYN/ACK segments are given, split up for outbound and inbound established connections. The third column presents the fractions of connections advertising the option in both initial segments, hence actually establishing the connection with the specific optional feature.

	MSS			WS			SACK			TS		
	SYN	SYN/ACK	both	SYN	SYN/ACK	both	SYN	SYN/ACK	both	SYN	SYN/ACK	both
outbound	100.00%	99.59%	99.59%	19.36%	15.46%	15.46%	93.67%	69.70%	69.70%	16.50%	12.32%	12.32%
inbound	99.94%	99.92%	99.85%	24.33%	23.85%	23.83%	97.22%	90.40%	90.38%	19.72%	18.51%	18.50%

Table 11: TCP options for inbound and outbound connections

In general, the numbers are in range of the reported values of the previous studies. The maximum segment size option (MSS) is used extensively by clients and servers for both directions. To our surprise, the window scale (WS), timestamp (TS) and selective acknowledgement permitted (SACK) options on the other hand are about 1.5 times more common among inbound connections. Looking at the destination port numbers for these connections, the difference can be explained by a much more diverse mix of applications among inbound connections in favour of primarily web traffic on port 80 in outgoing connections. The incoming connections include large fractions of recognized P2P protocols, but also substantial amounts of SMTP, SSH and MS SQL sessions, which are mainly break in attempts as discussed in section 5.3. These protocols are often used to carry more data than conventional web traffic, so it seems natural that clients and servers are interested in optimizing throughput by use of these TCP options.

6. UDP level

Since UDP offers no connection establishment or termination, we defined UDP flows as the sum of bidirectional packets observed between a specific tuple of source and destination IP and port numbers, taking advantage of the timeout value of 20 min given by the trace duration. In the 2x73 network traces, 68 million such UDP flows have been observed, carrying around 7% of the packets and only 2-3% of the data. Interestingly, 51 out of the 68 Million UDP flows (76%) carry less than 3 packets in either direction. Our first guess, that classical UDP services like DNS and NTP would be primarily responsible for these flows, proved to be wrong. In fact, only 5% and 1.7% of the small UDP flows serve DNS or NTP requests, respectively. On the other hand P2P overlay networks, such as Kademia or other distributed hash table (DHT) protocols, are responsible for at least 18% of these small flows, where we expect this naïve port analysis to be a huge underestimate again. The purpose of these overlay networks is to keep the peers routing tables updated in a completely decentralized fashion. This is done periodically by sending DHT “pings” in small UDP packets, replied by the recipient. No significant difference between inbound and outbound DHT queries could be observed, which makes sense when considering the type and the nature of these overlay networks.

Based on the simple port classification, different common network attacks on UDP port numbers for MS SQL, MS messenger “spam” or Netbios were found to be responsible for at least in 8% of the 51 Million short flows. These flows consisted in more than 90% of the cases of one inbound packet only, sometimes performing scans on entire IP ranges.

The two main sources for UDP flows, P2P overlay networks and attacking traffic, finally also explain the extreme amount of distinct IP addresses seen on the outside of the links measured (presented in table 2) since P2P network span the entire globe and experience a very high fluctuation in peering partners.

7. Summary and Conclusions

We presented directional differences found on recent packet level traces taken on links with medium aggregation level, carrying traffic from two major Universities, about a dozen of large student dormitories and a local exchange point. Since access to contemporary traffic on highly aggregated links is still uncommon, we believe that this study can contribute to a better understanding of the changing behaviour of the Internet. After short discussions about the two main factors responsible for the observed directional differences in our traces, malicious traffic and P2P traffic, this paper will be closed with summarizing conclusions.

7.1. Malicious traffic

Already the protocol breakdown revealed one outstanding long-duration UDP DoS attack originated within a major University in Göteborg, due to an DoS script injected from outside by exploitation of a known vulnerability. The fact that this attack was undetected by the network management tools in operation indicates the need for continuous refinement of network monitoring policies.

Despite this UDP burst, it can be said that basically every kind of malicious traffic is much more common in traffic coming from the main Internet. Already on a very high level analysis, incoming network scans were evident when analysing distinct IP addresses seen. There are about three times more rejected connections observed among inbound connections, with a majority of them being unreplied scanning attempts, but also a substantial number of immediate reset terminations. Around 90% of the counted header anomalies appeared on the incoming link, which goes hand in hand with the above mentioned scans. These packet header anomalies include inconsistencies in the IP flags, TCP header length and TCP connection flags field, which was discussed in more detail in an earlier study on the MonNet data [10]. Even though these header anomalies are very rare compared to the total number of packets, they indicated again skewed distribution of malicious traffic towards incoming traffic. The inconsistencies were shown to stem from network attacks trying to exploit protocol vulnerabilities as well as active OS fingerprinting tools.

Also the analysis of statistical connection properties within established connections revealed a large number of inbound login attempts to SSH, SMTP or MS SQL servers. Finally, on UDP level scanning traffic and security exploits were shown to happen in more than 90% of the cases within incoming traffic, which are as well in the order of millions in absolute numbers.

This summary of malicious behaviour confirms the suspicion that the main number of anomalies indeed originates on the outside, on the "unfriendly" Internet. It was shown that anomalies are between 3 and 9 times more common among inbound data. Typical University campus networks, but even student networks, are comparable well behaving, probably due to higher configuration and administration efforts.

7.2. P2P traffic

Except the directional differences due to malicious traffic, P2P is a second source heavily influencing traffic properties. Even with a simple, underestimating port analysis, we could show that P2P traffic is a major part of the traffic, responsible for at least twice as much packets and volume among inbound traffic as compared to outbound traffic. Artefacts of P2P traffic were found in packet size distribution, TCP connection termination behaviour, TCP options and statistical connection properties. P2P traffic was also shown to be a major source for long-duration traces, especially among inbound connections. Additionally, P2P overlay traffic is responsible for the major amount of UDP flows, carrying typically less than 3 small sized packets, but being responsible for several millions of distinct IP addresses observed in the traffic. These short flows are furthermore hard to distinguish from malicious scanning or attacking traffic, which needs to be taken into consideration by network engineers and security experts working on sampled flow level analysis.

7.3. Conclusion

While some high-level analysis, like cumulated traffic volumes or protocol breakdown, could suggest an even distribution between inbound and outbound traffic, this study reveals that there are a number of significant directional differences found on different protocol levels. Especially the detailed TCP connection analysis, contrasting incoming and outgoing established connections by statistical means, revealed significant differences. Even though connections established in both directions show a typical client-server pattern, this behaviour is more pronounced among inbound connections. Generally, inbound connections, established from the outside, are shown to be more likely to carry larger volumes of data (elephants), larger number of packets and experience longer connection lifetimes. However, these differences, caused by the imbalance in P2P traffic, cancel out on high-level summaries because established inbound connections are on the other hand about 10% fewer than outbound connections.

First of all, the comprehensive analysis yielded required insights for network developers and traffic engineers. Furthermore, the results can be important input in order to improve quality and authenticity of future simulation models. Finally, the highlighted traffic anomalies are relevant for better understanding of security related issues like intrusion detection or detection of large scale attacks.

Acknowledgement

The authors want to thank Pierre Kleberger for his kind technical support and Tomas Olovsson for his valuable and constructive comments throughout the MonNet project.

References

- [1] M. Arlitt and C. Williamson, "An analysis of TCP reset behaviour on the Internet," *Computer Comm. Review*, vol. 35, 2005.
- [2] N. Brownlee and K. C. Claffy, "Understanding Internet traffic streams: dragonflies and tortoises," *IEEE Communications Magazine*, vol. 40, pp. 110-17, 2002.
- [3] N. Brownlee and K. C. Claffy, "Internet Measurement," *IEEE Internet Computing*, vol. 8, pp. 30-33, 2004.
- [4] CiscoSystems, "IPsec VPN WAN Design Overview," Cisco Documentation, 2006.
- [5] S. Donnelly, "Endace DAG Timestamping Whitepaper," Endace Withepapers, 2006.
- [6] S. Floyd and E. Kohler, "Internet research needs better models," *Computer Communication Review*, vol. 33, pp. 29-34, 2003.
- [7] C. Fraleigh, S. Moon, B. Lyles, C. Cotton, M. Khan, D. Moll, R. Rockell, T. Seely, and S. C. Diot, "Packet-level traffic measurements from the Sprint IP backbone," *Network, IEEE*, vol. 17, pp. 6-16, 2003.
- [8] A. Gerber, J. Houle, H. Nguyen, M. Roughan, and S. Sen, "P2P The Gorilla in the Cable," in *National Cable & Telecommunications Association(NCTA) National Show*. Chicago, IL, 2003.
- [9] A. Householder, K. Houle, and C. Dougherty, "Computer attack trends challenge Internet security," *Computer*, vol. 35, 2002.
- [10] W. John and S. Tafvelin, "Analysis of Internet Backbone Traffic with focus on Header Anomalies," submitted for publication, Chalmers, Göteborg, Sweden, 2007.
- [11] T. Karagiannis, A. Broido, M. Faloutsos, and K. Claffy, "Transport layer identification of P2P traffic," *Proceedings of the 4th ACM SIGCOMM conference on Internet measurement*, Taormina, Sicily, Italy, 2004.
- [12] M.-S. Kim, Y. J. Won, and J. W. Hong, "Characteristic analysis of internet traffic from the perspective of flows," *Computer Communications*, vol. 29, pp. 1639-1652, 2006.
- [13] K. Lan and A. Hussain, "The Effect of Malicious Traffic on the Network," *Proceedings of the Workshop on Passive and Active Measurements (PAM)*, 2003.
- [14] K.-C. Lan and J. Heidemann, "Rapid model parameterization from traffic measurements," *ACM Transactions on Modeling and Computer Simulation*, vol. 12, pp. 201-29, 2002.
- [15] K.-C. Lan and J. Heidemann, "A measurement study of correlations of Internet flow characteristics," *Computer Networks*, vol. 50, pp. 46-62, 2006.
- [16] S. McCreary and K. C. Claffy, "Trends in wide area IP traffic patterns - A view from Ames Internet Exchange," *Cooperative Association for Internet Data Analysis - CAIDA*, San Diego Supercomputer Center, San Diego 2000.
- [17] M. Mellia, R. Lo Cigno, and F. Neri, "Measuring IP and TCP behavior on edge nodes with Tstat," *Computer Networks*, vol. 47, pp. 1-21, 2005.
- [18] A. W. Moore and K. Papagiannaki, "Toward the Accurate Identification of Network Applications," *Lecture Notes in Computer Science*, pp. 41-54, 2005.
- [19] T. Mori, M. Uchida, and S. Goto, "Flow analysis of internet traffic: World wide web versus peer-to-peer," *Systems and Computers in Japan*, vol. 36, pp. 70-81, 2005.
- [20] T. Mori, M. Uchida, R. Kawahara, J. Pan, and S. Goto, "Identifying elephant flows through periodically sampled packets," *Proceedings of the 4th ACM SIGCOMM conference on Internet measurement*, Taormina, Sicily, Italy, 2004.
- [21] K. Pentikousis and H. Badr, "Quantifying the deployment of TCP options - a comparative study," *IEEE Communications Letters*, vol. 8, pp. 647-9, 2004.
- [22] S. Saroiu, K. P. Gummadi, R. J. Dunn, S. D. Gribble, and H. M. Levy, "An analysis of internet content delivery systems," *Proceedings of the 5th symposium on Operating systems design and implementation*, Boston, Massachusetts, 2002.
- [23] C. Shannon, D. Moore, and K. C. Claffy, "Beyond folklore: observations on fragmented traffic," *IEEE/ACM Transactions on Networking*, vol. 10, pp. 709-20, 2002.
- [24] S. Tafvelin, "Presentation: QoS measurements," *TERENA Networking Conference*, Poznan, Poland, 2005.
- [25] K. Thompson, G. J. Miller, and R. Wilder, "Wide-area Internet traffic patterns and characteristics," *IEEE Network*, vol. 11, 1997.
- [26] P. A. Watson, "Slipping in the Window: TCP Reset Attacks," *Technical Whitepaper*, 2003.
- [27] J. Xu, J. Fan, M. Ammar, and S. B. Moon, "On the design and performance of prefix-preserving IP traffic trace anonymization," *Proceedings of the 1st ACM SIGCOMM Workshop on Internet Measurement*, San Francisco, California, USA, 2001.
- [28] Y. Zhang, L. Breslau, V. Paxson, and S. Shenker, "On the characteristics and origins of Internet flow rates," *Computer Communication Review*, vol. 32, pp. 309-322, 2002.