

# The State of Enterprise Network Traffic in 2012

David Murray  
Murdoch University  
D.Murray@murdoch.edu.au

Terry Koziniec  
Murdoch University  
T.Koziniec@murdoch.edu.au

**Abstract**—Timely and accurate studies on the composition and nature of the Internet are crucial for continued research and innovation. The aim of this research is to aid and service emulated and simulated research methods where realism is dependent on the accuracy of Internet statistics. This study has captured anonymised, and analysed Internet traffic entering and exiting a university network. Passive measurement techniques are used to discover the mean and distribution of packet sizes. A range of Network and Transport layer statistics, such as DCSP, SACK, ECN, MSS usage, Window Scaling and Timestamps are also investigated. Active measurement techniques are used to determine Round Trip Times (RTTs) and the number of lost and misordered packets. Prior work is leveraged to explore changes in the composition and nature of the Internet over the last 10 years.

## I. INTRODUCTION

A large amount of Internet research is performed on simulators such as NS2, OPNET and NetSim. Many entrenched Internet mechanisms have been proposed, validated [1] and then iteratively improved, solely on the use of network simulators [2], [3]. It is crucial that simulations and experimental emulations of the Internet are performed using models that accurately reflect modern networks. The importance of measurement studies has been highlighted by key figures within the Internet research community [4].

Measurement studies have been run many times over the past decade. These studies, [5], [6], [7], [8], [7], [9], [10], [11], [12], [13], [14], must be run regularly to track the constantly changing environment. The findings of this research can be used to increase the accuracy of assumptions used in simulated and experimental research areas.

## II. BACKGROUND

### A. Prior Work

Measurement papers can be grouped into passive and active measurement studies. These research techniques measure different but partially overlapping attributes.

Passive measurement [5], [6], [7], [8], [7], [9], [10], [11], [12] captures, anonymises and then analyses real Internet traffic. In this type of study, end user terms of service issues as well as the legal and ethical ramifications of capturing user data must be carefully navigated [15]. In many cases, it is permissible to capture traffic so long as the IP addresses are irreversibly modified and payloads are zeroed.

Active measurement is performed with a single PC which is used to probe servers. Active [13], [14] measurement studies probe large numbers of Internet servers to determine, TCP

congestion control types, IETF TCP standards conformance, ECN, SACK, IP options and PMTUD (Path MTU Discovery) techniques.

The Internet measurement research literature has been able to highlight numerous changes in the composition of Internet traffic. This study investigates the composition of the Internet from the Data-link layer, the Network layer and the Transport layer. The structure of this paper reflects our methodology. Sections III, IV and V analyse the findings through the Data-link, Network and Transport layer of the OSI model. The individual findings in this study are compared and contrasted with prior work. Following these passive measurements, section VI uses active measurement techniques to establish the upper and lower bounds of RTTs (Round Trip Times), packet loss and misordering.

### B. Methodology

A SPAN (SwitchPort ANalyser) port was setup on a Cisco switch to mirror all traffic entering and exiting the university. As these interfaces operate at 1000 Mb/s, no special hardware was required to capture these packets. A rackmount server running Ubuntu 10.04 and TCPdump was then used to capture traffic. Packets were captured over a period of 7 days. After the packets were captured, we used pktanon [16] to anonymise all payloads and IP addresses to a separate hard drive. The anonymised hard drive was used for analysis in this study. This research underwent ethics approval with the Murdoch University Human Research Ethics Subcommittee (Approval Number: 100111).

A range of programs were used to analyse the anonymised data. High level statistics were gathered by feeding the anonymised pcap file into NTOP [17]. The majority of measured attributes and features were calculated with custom written C programs. Active measurement was also used to obtain measurements for variables such as RTTs, packet loss and mis-ordering. These measurements were taken using clients attached to the Murdoch University network, an ADSL home broadband connection and the Amazon EC2 cloud. The methods used for each test were more varied and will be further detailed in their respective section.

## III. PASSIVE MEASUREMENT OF THE DATA-LINK LAYER

### A. Packet Size Distribution

Studies performed prior to 2004 described the Internet packet size distribution as trimodal [11], [12], [10], [9]. These studies found large numbers of small, <100-byte segments

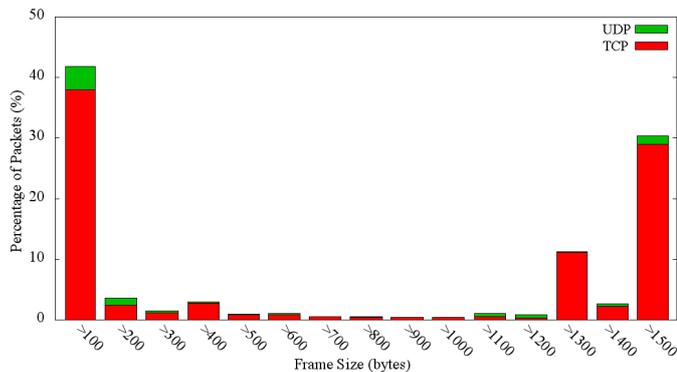


Fig. 1. Distribution of packet sizes

TABLE I  
MARKING OF IP PACKETS

	UDP	TCP
Total packets	709682609	4300376201
DSCP packets marked	1786013	1531201
DSCP % marked	0.251 %	0.036 %
ToS packets marked	33734052	557895017
ToS % marked	4.753 %	12.973 %

produced from TCP acknowledgements. A large number of 576-byte segments and a large proportion of full sized 1400-1500-byte frames. This trimodal distribution was a product of legacy RFC 879 [18] conformance.

Subsequent studies in 2004 [5] and 2006 [8] found that the distribution of packet sizes changed to bimodal. The composition of Internet packets changed with high numbers of small, <100-byte packets, and high numbers of large, 1400-1500-byte, packets. Specifically, John et al [8] found that small packets, between 40 and 100-bytes, constituted 44% of all IPv4 packets. Large packets, between 1400 and 1500-bytes constituted 37% of all IPv4 packets, with the remaining 19% falling between between 1400 and 128-bytes respectively. Pentikousis and Badr found similar results in 2004 [5].

The research in this paper confirms that the results seen in 2004 [5] and 2006 [8] are still accurate. The average TCP frame size, including the Ethernet headers, was 736-bytes which is consistent with the findings of prior research [7], [8]. Although an average exists, researchers should not use it because the actual values are at the extremes, rather than centred. A histogram highlighting the distribution is shown in Fig 1.

The large percentage of UDP packets seen under 100-bytes is predominantly DNS traffic. The results of this study concur with prior research [5], [8], suggesting that packet size distribution is still heavily bimodal.

#### IV. PASSIVE MEASUREMENT OF THE NETWORK LAYER

##### A. IP Flags and Fragmentation

Packet fragmentation on the Internet is low. In 2000 and 2001, fragmentation rates of 0.15% [11] to 0.68% [10] were found respectively. Later studies by John [8] in 2007 reported that only 0.06% of traffic was fragmented. The measurement results from this study suggest that an even lower rate, 0.046%, of traffic was fragmented. These results indicate that the amount of fragmented traffic has been on the decline for over a decade. This outcome is expected given the ubiquity of Ethernet and the use of increasingly robust PMTUD (Path MTU Discovery) [19] mechanisms.

Fragmentation rates are so low because of its detrimental affect on performance [20]. Many hosts set the DF (Don't

Fragment) bit in the IP header to prevent fragmentation from occurring. If a packet being received on a router is too big, the router will drop the packet and return a ICMP "Fragmentation Required but DF Bit Is Set" message. This ICMP message will signal that the sender should use a lower frame size. Newer PMTUD mechanisms [19] have since been standardised to increase the robustness.

John's study [8] suggests that 91.3% of IP packets have the DF (Don't Fragment) bit set. This study found that 85.1% of IP traffic has the DF bit set. The current recommendation is that fragmentation should be performed on the end nodes and therefore all IP packets should be sent with the DF bit set. A potential explanation for why our measured network has a lower DF rate is because there is a significant proportion of IPsec traffic. Many IPsec configurations do not set the DF bit [21].

##### B. ToS/DSCP

ToS (Type of Service) and DSCP (Differentiated Services Code Point) headers are used as part of the DiffServ QoS (Quality of Service) model. DSCP is a newer version of ToS. Certain packets can be marked with ToS/DSCP priority levels. Cooperating or trusting routers read these marks and provide differential treatment based on the marking. Some applications may incorrectly mark their packets to get preferential treatment and consequently DSCP markings are generally ignored by Internet routers.

The results of our study, in Table I, show that 13.0% of TCP packets and 5% of UDP packets received ToS, DSCP or Expedited Forwarding (EF) markings. Table II and III shows what percentage of UDP and TCP packets are marked when compared with total UDP segments or total TCP segments respectively. In our sample, the large number of CS1 marked packets are a result of an upstream AARNet (Australian Academic and Research Network) router marking, tariff or charge free traffic. These numbers are site dependent, and thus the reader should be careful when reusing these numbers.

#### V. PASSIVE MEASUREMENT OF THE TRANSPORT LAYER

##### A. Transport Layer Protocol Use

1) *TCP*: Previous studies acknowledge that the majority of the Internet operates over TCP [8], [7]. A measurement study, performed in 2006 on an Internet backbone link, found that TCP constitutes an average of 91.85% of the packets on the Internet and 97.18% of the bytes [8]. The results in this study show that TCP accounts for 84.35% of the packets and 92%

TABLE II  
PERCENTAGE OF DSCP VALUES SEEN

DSCP Value	UDP Percentage	TCP Percentage
af11	0.015292 %	0.109203 %
af12	0.013559 %	0.013166 %
af13	0.000640 %	0.001614 %
af21	0.006180 %	0.002467 %
af22	0.000383 %	0.005656 %
af23	0.000165 %	0.006165 %
af31	0.000728 %	0.000254 %
af32	0.000042 %	0.000002 %
af33	0.203546 %	0.000859 %
af41	0.006970 %	0.001706 %
af42	0.000023 %	0.000009 %
af43	0.000032 %	0.000001 %
ef	0.004104 %	0.003707 %

TABLE III  
PERCENTAGE OF PRECEDENCE/TOS VALUES SEEN

Precedence/ToS	UDP Percentage	TCP Percentage
cs1	4.686206 %	12.945971 %
cs2	0.037698 %	0.003901 %
cs3	0.001881 %	0.001808 %
cs4	0.001947 %	0.002124 %
cs5	0.003052 %	0.002503 %
cs6	0.020429 %	0.000360 %
cs7	0.002188 %	0.016501 %

of the bytes. This reduction in TCP traffic is linked with an increase in UDP traffic. The results are shown in Table IV.

2) *UDP*: Zang et al [7] produced a longitudinal study on UDP traffic stretching from 2002 to 2009. This study shows an increase in the proportion of UDP packets from 11% in 2002 to 21% in 2009. Our results also suggest an increase in UDP traffic, but not to the extent suggested by Zang et al [7]. Our study found that UDP traffic constitutes 13.92% of packets and 6.3% of the bytes. Despite the discrepancies in exact values between studies [7], [8], the results suggest that there has been an increase in the amount of UDP traffic over the past decade.

3) *Other*: Table IV shows the percentage of Transport layer protocol usage. It should be noted that some of the more recently proposed Transport layer protocols such as DCCP (Datagram Congestion Control Protocol) or SCTP (Stream Control Transmission Protocol) were not found in the captures.

### B. TCP Specific Features

TCP has numerous optional features that can significantly enhance performance. This section investigates, TCP data/ack

TABLE IV  
USAGE OF TRANSPORT LAYER PROTOCOLS

Protocol	Percent
TCP	84.35 %
UDP	13.92 %
ICMP	0.21 %
IGMP	0.01 %
IPv6 in IPv4	0.013 %
Other	1.49 %

ratios, as well as the uptake of TCP options such as: SACK (Selective Acknowledgements), MSS (Maximum Segment Size), Window Scale, Timestamps and ECN (Explicit Congestion Notification).

1) *TCP Data/ACK Ratios*: TCP can acknowledge every segment or alternatively, every second segment. This operation depends on the TCP congestion control mechanism and the stage of operation. When modelling TCP, many researchers assume either one ack for every data segment or one ack for every two data segments. The real ratio is somewhere between these two assumptions. The analysis of packets in our study indicates that there are 1.77 TCP data packets for every acknowledgement. An alternative perspective is that 34% of TCP segments are data-less acknowledgements.

2) *SACK*: TCP SACK are a modification of the cumulative acknowledgement behaviour of TCP. Traditionally, the acknowledgement field within the TCP header holds the sequence number of all the data that has been successfully received. Under this scheme, a burst of packet losses will take multiple round trip times to recover. SACK is a TCP option that can additionally specify which blocks have been successfully received. SACK enables TCP to more efficiently and quickly recover from multiple packet losses within the same window.

John's 2006 [8] study found that 91% of all TCP SYN segments applied the SACK opportunity. This study also analysed TCP SYN segments and found that 94% of SYNs used the SACK permitted TCP option. This result is expected because modern OSs enable SACK by default. Our testing also showed that SACK was also enabled by default in iOS and many of the Android devices tested.

3) *MSS*: The MSS is a TCP option defined in RFC 793. It is used by end nodes to state the maximum supported segment size. Unless a MSS is negotiated in the options of a TCP SYN, both ends of the connection use the 536-byte MSS value defined in RFC 1122.

Previous studies found the MSS option in 99% of TCP SYN messages [8]. This is the default behaviour of all major OSs due to the performance ramifications from using unnecessarily small packet sizes [22]. This study found the MSS option in 96.6% of TCP SYNs.

As the majority of segments carry the MSS field, the values advertised are also of interest. Prior research has shown that the majority, 93.7%, of the MSS values lie between 1400-1460-bytes [8]. Values larger than 1460-bytes account for only 0.06% of TCP SYNs and values smaller than 536-bytes are carried by another tiny fraction; 0.05%. The remaining 6.19% were values between 536 and 1400 [8].

The results of our study are shown in Table V. These results show that 26.5% of hosts advertise a MSS of 1460. The majority of hosts, 46.2%, advertise MSS between 1300 and 1460. A surprisingly large number of hosts advertised an MSS between 1000 and 1301, this is likely to be due to the prevalence of VPNs.. Only 0.0368% of packets advertised a MSS larger than 1460 which is representative of the poor uptake of Jumboframes on the Internet [22].

TABLE V  
MSS OPTION VALUES

MSS Size	Percent
Did not Advertise	3.43 %
0-1000	0.0377 %
1001-1300	23.6898 %
1301-1459	46.2704 %
1460	26.5415 %
1461-1600	0.0240 %
1601-8000	0.0045 %
8001-9000	0.0014 %
> 9000	0.0069 %

4) *Window Scale*: The *Window Size* field in the TCP header is only 16-bits long and limits the advertised size of a TCP window to 65,535-bytes. The *Window Scale* TCP option was defined in RFC 1323 and operates as a multiplier on the *window size*. The Window Scale TCP option is required for fast transfers over high latency networks. Prior research found the Window Scale option in 17.9% [8] of TCP SYN messages. This study found the Window Scale option in 63.9% of TCP SYNs. This increased uptake in the Window Scale option is expected because TCP performance is restricted in its absence.

5) *Timestamps*: Timestamps were defined alongside the Window Scale option in RFC 1323. Timestamps are used to calculate the RTT and are used for PAWS (Protection Against Wrapped Sequence numbers). PAWS is necessary because the TCP sequence number is stored in a 32-bit field, which only allows it to address up to 4-GB. Sequence number ambiguity must be resolved when sequence numbers reach the limit and wrap back around to zero. PAWS determines to which 4-GB sequence a replayed packet belongs. Prior research found that 14.5% [8] of SYNs advertised the Timestamps option. This study found that 39.2% of SYNs were advertising the Timestamps option.

6) *ECN*: The most common mechanism used to determine packet loss is through TCP duplicate acknowledgements. A sender transmitting too quickly will congest a router, causing the router to drop packets. Packet loss will be detected by the TCP receiver first, who will then inform the TCP sender through duplicate acknowledgements.

ECN (Explicit Congestion Notification) [23] is an alternative congestion notification method. ECN is initiated by an endpoint setting the ECN-Echo and CWR bits and accepted when an ECN-Echo is returned. The IP header is also modified to provide routers with knowledge of ECN support. A “01” or “10” in the 6th and 7th bits of IP’s DSCP field indicate that the flow is supporting ECN. When a router on the Internet is experiencing congestion, instead of dropping the packet, it can indicate that congestion is being experienced by filling the 6th and 7th bits of the DSCP header with a “11”.

The use of ECN is optional. A 2004 study on TCP options by Pentikousis and Badr [5] found that ECN deployment was marginal, 0.15%. Furthermore, the proportion of traffic that was marked with CE (Congestion Experienced) was also very low [5]. A latter active measurement study by Medina et al in 2004 found that only 2.1% of web servers were ECN capable

TABLE VI  
USAGE OF TRANSPORT LAYER PROTOCOLS

TCP Option	Percent
MSS	96.597 %
SACK Permitted	94.009 %
SACK	0.003 %
TSOPT	39.290 %
WSOPT	63.971 %
Echo	0.007 %
Echo Rep	0.001 %
Other	0.091 %

TABLE VII  
USAGE OF THE TOP 6 WELL KNOWN PORTS

Application	Port	Total	Sent	Received
proxy	3128	113.4 GBytes	99.8 GBytes	13.6 GBytes
www	80	108.1 GBytes	101.6 GBytes	6.5 GBytes
smtp	25	14.8 GBytes	739.2 MBytes	14.1 GBytes
https	443	13.2 GBytes	10.3 GBytes	2.9 GBytes
Cisco IPsec	10000	6.4 GBytes	4.0 GBytes	2.4 GBytes
ssh	22	5.8 GBytes	5.6 GBytes	232.2 MBytes

[13].

Our study found zero TCP flows requesting ECN in the TCP SYN. By analysing all TCP messages, we found that 0.0007% were marked with ECN CE in the IP header. We also found that 0.07% of TCP packets had markings in the IP header indicating they were ECN capable. Due to the absence of ECN markings in the TCP header, the markings in the IP header may be a product of legacy RFC 791 conformance, which specified the use of these bits for QoS.

### C. Top port usage

Few prior studies have discussed port usage. One study using measurements from the Internet backbone [8] stated that large numbers of ports known to be used by P2P and file sharing networks were found. These results are obviously highly based on context of the measurement point.

This study found that WWW, SMTP, HTTPS, Cisco IPsec and SSH were the most used ports. The top 6 are shown in VII. Murdoch University has a web proxy operating on port 3128 to cache frequently requested web objects. Thus the total web traffic could be interpreted as the sum of WWW and proxy traffic (113.4+108.1 = 221.5GB).

## VI. ACTIVE MEASUREMENTS

In addition to the results produced from passive measurement techniques, active measurement was also used to ascertain RTTs and path length, loss rates, misordering.

### A. Experimental Design

This test downloaded 700-MB files from HTTP servers in Australia, Sweden, Taiwan and the United States. The location in the world where the files were downloaded to was also varied. The first two tests were downloads initiated from Perth WA over an ADSL connection and a university network link. To provide more generic and repeatable results, Amazon EC2 cloud machines were also instantiated in Europe and in the

TABLE VIII  
RESULTS OF ACTIVE MEASUREMENT TESTS FROM A ADSL2+ CONNECTION IN PERTH, WESTERN AUSTRALIA

HTTP Server Loc	Aus	Sweden	Taiwan	US
Avg BW (MB/s)	1.35	1.42	1.07	0.376
Avg BW (Mb/s)	10.8	11.36	8.56	3.008
Number of hops	7	18	13	15
RTT	44ms	440ms	150ms	342ms
Lost packets	0.029%	0.02775%	0.0445%	0.508
Misordered	0%	0%	0%	0

TABLE IX  
RESULTS OF ACTIVE MEASUREMENT TESTS FROM A UNIVERSITY NETWORK IN PERTH, WESTERN AUSTRALIA

HTTP Server Loc	Aus	Sweden	Taiwan	US
Avg BW (MB/s)	6.79	1.11	1.83	5.28
Avg BW (Mb/s)	54.32	8.88	14.64	42.24
Number of hops	9	22	16	18
RTT	30ms	398ms	374ms	248ms
Lost packets	0.123%	0.0178%	0.122%	0.007%
Misordered	0.0704%	0.0207%	0.068%	0.068%

United states. The combination of 4 HTTP servers and 4 different download sites resulted in 16 different permutations. Each of these independent results is based on an average of twelve 700-MB downloads performed over a 24 hour period. The speed, RTT and number of hops required for every download was recorded. These downloads were captured using TCPdump and analysed offline. The results from this active measurement test is shown in Tables VIII, IX, X, XI.

#### B. RTT and path length

Our study found RTTs between 30-ms and 440-ms and hop lengths between 7 and 26 hops. Obviously RTTs lower than 30-ms are possible but we feel that these are reasonable assumptions for Internet style latencies. The latencies measured in this paper are unlikely to change in the future. Based on the speed of light through fibre optics over a distance of half the circumference of the world, the lowest possible RTT between nodes on opposite sides of the world is approximately 180-ms. Thus, even with the most optimistic assumptions, latencies between devices on the Internet are likely to remain relatively constant in the foreseeable future.

#### C. Loss Rates

Clear statistics on packet loss over the Internet are also missing from the literature. Packet loss can be caused by congestion on a router, or interference on a link. When a router buffer fills, packets will be dropped from the queue to alleviate congestion and notify the TCP senders that packets are being sent too fast. In the case of either congestion based or interference based losses, prior research has suggested that they typically occur in bursts and are unevenly distributed [4], [24].

To determine the number of lost and misordered packets, the packet capture of the downloads were analysed. A C program was used to count the number of segments that went missing and were later replayed. These missing segments were

TABLE X  
RESULTS OF ACTIVE MEASUREMENT TESTS FROM A AMAZON EC2 INSTANCE IN DUBLIN, IRELAND

HTTP Server Loc	Aus	Sweden	Taiwan	US
Avg BW (MB/s)	0.671	3.86	1.38	2.137
Avg BW (Mb/s)	5.368	30.88	11.04	17.096
Number of hops	19	16	26	19
RTT	329.4ms	58.6ms	320.9ms	109.7ms
Lost packets	0.173%	0.0253%	0.0012%	0.0325%
Misordered	0%	0%	0.0027%	0%

TABLE XI  
RESULTS OF ACTIVE MEASUREMENT TESTS FROM A AMAZON EC2 INSTANCE IN CALIFORNIA US

HTTP Server Loc	Aus	Sweden	Taiwan	US
Avg BW (MB/s)	1.47	1.98	4.86	2.975
Avg BW (Mb/s)	11.76	15.84	38.88	23.8
Number of hops	11	17	13	14
RTT	181ms	187ms	209ms	67ms
Lost packets	0.19%	0.03075%	0.0315%	0.04375
Misordered	0%	0%	0.0025%	0

categorised as lost or misordered based on the duration of time between a packet going missing and its replay time. Any packet being replayed in under half of the RTT was categorised as misordered.

Our study found packet loss rates between 0.4% and 0.0012%. An average of all the scenarios is deliberately absent because it is important to be mindful that loss rates are heavily context dependent. These results on loss are being presented as approximate bounds for experimental and emulated research.

Many factors make these experiments unrepeatable. The amount of congestion experienced across the entire world will vary with the time of day. The path that the packets take will also vary over time as new links are added. Perhaps the unreliability of these methods has deterred previous studies from such measurements. Despite the unavoidable problem of experimental repeatability, these results are valuable estimates in the absence of other indicators of loss rates.

#### D. Packet Misordering

Packet reordering may occur as a result of queuing, scheduling, load balancing routing protocols, or parallelisms in routers and across links. Packet reordering can significantly degrade performance because it may result in unnecessary retransmissions and TCP window reductions. Numerous studies have reported on the incidence of packet reordering on the Internet [25], [26], [27], [28]. Studies have used numerous different active and passive methods to measure misordering.

In 1997 Paxon found that 2% of packets were misordered [25]. In 2002 Jaiswal et al found that packet reordering was approximately 5% [29]. In 2004 Wang et al found that 3.2% of all packets were misordered [28]. Wang et al also recorded the distribution or extent of reordering and stated that, as re-ordering was usually minor, it would not trigger a fast retransmit and therefore would not heavily effect TCP performance [28].

The results of this measurement study show that reordering is heavily site dependent. It also shows that packet reordering has reduced from the levels previously recorded. The results, shown in Tables VIII, IX, X, XI, suggest that Taiwan, or the network attached to our download site in Taiwan, may have two links that load balance. The Murdoch network is also responsible for causing packet reordering. The highest level of reordering was occurring over the university network, at a rate of 0.074%.

A closer analysis of reordering indicates that many of these packets were only arriving approximately 0.250- $\mu$ s microseconds after the detection of a misordered packet. Thus the results indicate that the extent of reordering, in the Murdoch network and in Taiwan, is so minor that it does not trigger a TCP fast retransmit. The conclusions drawn from these results are that; although packet reordering is common, it is site dependent. The study also suggests that the extent of misordering is so slight that it is not detrimental/pathological to performance.

## VII. CONCLUSION

This study has used a combination of passive and active measurement techniques to measure the state of Internet communications. We found that packet sizes are still heavily bimodal and fragmentation rates are decreasing. TCP is still the major transport protocol on the Internet, however, the use of UDP has risen in recent years. The use of TCP SACK, Window Scaling and Timestamps are have all increased whereas ECN adoption remains non-existent. Internet latencies were between 30-ms and 440-ms. Loss rates were measured between 0.4% and 0.0012%. Most sites did not experience packet reordering. The sites that did experience reordering were minor and were not detrimental to performance. Internet measurement studies are hugely important [4] for simulated and emulated research methods that depend on accurate statistics to build realistic Internet models. It is hoped that further studies are performed in the future to track potential changes.

## ACKNOWLEDGMENT

The authors would like to acknowledge Amazon for their research grant. This grant was used to provide a consistent and repeatable platform for the active measurements produced in this paper. The authors would also like to acknowledge Jarren Beveridge from Murdoch University Communications Services for his assistance collecting, handling and anonymising the data.

## REFERENCES

- [1] Sally Floyd and Van Jacobson, "Random early detection gateways for congestion avoidance", *IEEE/ACM Trans. Netw.*, vol. 1, pp. 397–413, August 1993.
- [2] Wu chang Feng, Dilip D. Kandlur, Debanjan Saha, and Kang G. Shin, "A self-configuring red gateway", in *INFOCOM*, 1999, pp. 1320–1328.
- [3] Nabhan Hamadneh, David Murray, Michael Dixon, and Peter Cole, "Weighted RED (WTRED) Strategy for TCP Congestion Control", in *Informatics Engineering and Information Science*, vol. 252, pp. 421–434. Springer Berlin Heidelberg, 2011.
- [4] Sally Floyd and Eddie Kohler, "Internet research needs better models", *SIGCOMM Comput. Commun. Rev.*, vol. 33, pp. 29–34, January 2003.
- [5] Kostas Pentikousis, Hussein Badr, and Trace Analysis, "Quantifying the Deployment of TCP Options - A Comparative Study", *IEEE Communications Letters*.
- [6] Theophilus Benson, Aditya Akella, and David A. Maltz, "Network traffic characteristics of data centers in the wild", in *Proceedings of the 10th annual conference on Internet measurement*, New York, NY, USA, 2010, IMC '10, pp. 267–280, ACM.
- [7] Min Zhang, Maurizio Dusi, Wolfgang John, and Changjia Chen, "Analysis of udp traffic usage on internet backbone links", in *Proceedings of the 2009 Ninth Annual International Symposium on Applications and the Internet*, Washington, DC, USA, 2009, pp. 280–281, IEEE Computer Society.
- [8] Wolfgang John and Sven Tafvelin, "Analysis of internet backbone traffic and header anomalies observed", in *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*, New York, NY, USA, 2007, IMC '07, pp. 111–116, ACM.
- [9] 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", *IEEE Network*, vol. 17, no. 6, pp. 6 – 16, nov.-dec. 2003.
- [10] C. Shannon, D. Moore, and KC. Claffy, "Beyond Folklore: Observations on Fragmented Traffic", *IEEE/ACM Transactions on Networking*, pp. 709–720, Dec 2002.
- [11] Sean McCreary and KC Claffy, "Trends in Wide Area IP Traffic Patterns - A View from Ames Internet Exchange", ITC Specialist Seminar, 2000.
- [12] Kevin Thompson, Gregory J. Miller, and Rick Wilder, "Wide-area internet traffic patterns and characteristics", *IEEE Network*, vol. 11, pp. 10–23, 1997.
- [13] Alberto Medina, Mark Allman, and Sally Floyd, "Measuring the evolution of transport protocols in the internet", *ACM Computer Communication Review*, 2005.
- [14] Steven Bauer, Robert Beverly, and Arthur Berger, "Measuring the state of ecn readiness in servers, clients, and routers", in *Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference*, New York, NY, USA, 2011, IMC '11, pp. 171–180, ACM.
- [15] Wolfgang John, Sven Tafvelin, and Tomas Olovsson, "Review: Passive internet measurement: Overview and guidelines based on experiences", *Comput. Commun.*, vol. 33, pp. 533–550, March 2010.
- [16] Marcus Schoeller Christoph P. Mayer, Thomas Gamer, "Pktanon packet trace anonymization", Online: <http://www.tm.uka.de/software/pktonon/>, 2012.
- [17] Luca Deri, "Ntop", Online: <http://www.ntop.org/>, 2012.
- [18] J. Postel, "Tcp maximum segment size and related topics", IETF RFC 879, 1983.
- [19] M. Mathis and J. Heffner, "Packetization Layer Path MTU Discovery", RFC 4821, 2007.
- [20] Christopher A. Kent and Jeffrey C. Mogul, "Fragmentation considered harmful", in *ACM SIGCOMM*, 1987, pp. 390–401.
- [21] Cisco Systems, "DF Bit Override Functionality with IPsec Tunnels", Online: [http://www.cisco.com/en/US/docs/ios/12\\_2t/12\\_2t2/feature/guide/ftdfipsc.html](http://www.cisco.com/en/US/docs/ios/12_2t/12_2t2/feature/guide/ftdfipsc.html), 2012.
- [22] David Murray, Terry Koziniec, Kevin Lee, and Michael Dixon, "Large MTUs and Internet Performance", in *IEEE 13th Conference on High Performance Switching and Routing*, 2012, pp. 390–401.
- [23] K. Ramakrishnan, S. Floyd, and D. Black, "The Addition of Explicit Congestion Notification (ECN) to IP", RFC 3168, 1990.
- [24] David Murray, Terry Koziniec, and Michael Dixon, "Solving ack inefficiencies in 802.11 networks", in *IMSAA-09: IEEE International Conference on Internet Multimedia Systems Architecture and Applications*, 2009.
- [25] Vern Paxson, "End-to-end internet packet dynamics", *SIGCOMM Comput. Commun. Rev.*, vol. 27, no. 4, pp. 139–152, Oct. 1997.
- [26] Jon C. R. Bennett, Craig Partridge, and Nicholas Shtetman, "Packet reordering is not pathological network behavior", *IEEE/ACM Trans. Netw.*, vol. 7, no. 6, pp. 789–798, Dec. 1999.
- [27] John Bellardo and Stefan Savage, "Measuring packet reordering", in *Proceedings of the 2nd ACM SIGCOMM Workshop on Internet Measurement*, New York, NY, USA, 2002, IMW '02, pp. 97–105, ACM.
- [28] Yi Wang, Guohan Lu, and Xing Li, *A Study of Internet Packet Reordering*, 2004.
- [29] Sharad Jaiswal, Gianluca Iannaccone, Christophe Diot, Jim Kurose, and Don Towsley, "Measurement and classification of out-of-sequence packets in a tier-1 ip backbone", *IEEE/ACM Transaction on Networking*, vol. 15, no. 1, pp. 54–66, Feb. 2007.