TCP Use and Performance on Internet2

Stanislav Shalunov, Benjamin Teitelbaum

Abstract— **Sampled NetFlow data from a core Internet2 router are analyzed to characterize the current use and performance of Internet2, with particular emphasis on bulk TCP transfers. The distribution of throughput, transfer size, duration, and average packet size for 48,301 observed bulk TCPs is presented. The top 10% of bulk transfer TCPs achieved throughputs of 3.9 Mbps or greater, while the top 1% of bulk transfer TCPs achieved throughputs of 23 Mbps or greater. The median bulk transfer TCP throughput observed was 880 Kbps. Summaries of application and IP protocol mixes are presented. Popular applications included: NNTP (19.3% packets, 22.8% octets), active FTP (11.9% packets, 14.2% octets), HTTP (10.5% packets, 9.5% octets), and multicast (6.6% packets, 6.2% octets). Most used IP protocols were TCP (85.6% packets, 88.7% octets), UDP (12.5% packets, 10.0% octets), and ICMP (1.4% packets, 0.8% octets).**

Keywords— **Internet2, Abilene, OC-48c, backbone, TCP, throughput, NetFlow.**

I. Introduction

THE Abilene backbone network [1] forms the core of the high-performance Internet2 network infrastructure, connecting over 190 U.S. universities and research centers through 50 access circuits connecting at twelve core routers. Access circuits run at 155 Mbps (OC-3c), 622 Mbps (OC-12c), or 2.4 Gbps (OC-48c) speeds, while interior circuits connecting core routers run at 2.4 Gbps. All interior circuits and most access circuits are packet-over-SONET (POS). The remaining access circuits are IP-over-ATM. Abilene peers with over 20 networks including the vBNS, U.S. federal agency networks (ESnet, NREN, DREN), and numerous non-U.S. high-performance research and education networks. Abilene routing policy assures that Abilene only carries traffic between research and educational institutions and not to or from the commercial Internet. Abilene is truly at the core of Internet2 and represents an appropriate locus for macroscopic study of Internet2 performance.

Cisco NetFlow export was used to capture over 13 million "flows" over a 24-hour period from a single Abilene core router. Our analysis, however, focuses mostly on bulk TCP transfers. We present distributions and analyses of throughputs, transfer durations, and transfer sizes for over 48 thousand bulk TCPs.

The structure of the paper is as follows. Section II describes how sampled NetFlow data were collected, the methodology employed for extrapolating these data to characterize the use and performance of TCPs in Internet2, and an experimental validation of the methodology applied to estimating the throughput of bulk TCPs. Section III summarizes the complete data set with an emphasis on IP protocol and application mix. The heart of the paper is in Section IV, which presents an analysis of the distribution of throughput, transfer size, duration, and average packet size for 48,301 observed bulk TCPs. Section V concludes by highlighting our most significant findings and raising several questions for future work.

II. Methodology

THE data presented in this paper are based on sampled Cisco NetFlow [3] records captured from the Abilene Cleveland (CLEV) core router over a 24-hour period beginning 19 June 2001 at 23:00 UTC. These data were scaled to adjust for sampling and then used to estimate throughput and other metrics. This section presents an overview of NetFlow, our NetFlow-based measurement methodology, results from a validation exercise, and a discussion of limitations.

A. NetFlow

Originally a side-effect of route caching for performance, NetFlow monitoring has become a widely-supported feature of Cisco routers. NetFlow monitors "flows", which Cisco defines as "a unidirectional sequence of packets between given source and destination endpoints"[5]. Flow endpoints are identified by source and destination IP addresses, as well as by transport layer port number. Other defining attributes of a flow include: IP Protocol type, Type of Service (ToS), and input interface. New flows are cached in a "flow table" and accounting information for them maintained while the flow is active. A flow may be expired and evicted from the flow table for any of the following reasons:

- Flow has been idle 15 seconds[1];
- Flow has been active for longer than 30 minutes[1];
- Flow is TCP and either a FIN or an RST was captured;

S. Shalunov is with Internet2. E-mail: shalunov@internet2.edu.

B. Teitelbaum is with Internet2 and Advanced Network & Services. E-mail: ben@internet2.edu.

[1] NetFlow is highly configurable; the specific values mentioned here were those configured to collect the data for this paper.

• Flow is among the oldest in a full flow table (NetFlow applies various heuristics to age groups of flows).

Groups of evicted flow records are collected into UDP datagrams, which are exported from the router and collected externally.

B. Data Collection

Sampled version 5 NetFlow records were captured from the Abilene Cleveland (CLEV) core router over a 24-hour period beginning 19 June 2001 at 23:00 UTC. NetFlow version 5 flow record attributes used in this paper include: the IP addresses, transport layer port numbers, and AS numbers for the source and destination; the values of the IP protocol and TOS fields; SNMP indices of the input and output interfaces on the collecting router; start and stop times; and counters for observed packets and octets. Because NetFlow's cache expiration policy can keep a long-lived flow in the flow table for up to thirty minutes, over twenty-five hours of data were captured and reduced to a data set of flows having start times within the above-mentioned 24-hour period.

The CLEV core router has access circuits for four "GigaPoPs"—PSC, NYSERNet, Merit, and OARnet—and provides 2.4 Gbps transit between Abilene's New York City and Indianapolis core routers. CLEV was selected for this study because its transit circuits are usually the busiest circuits in the network (see [6]).

A total of 13,589,626 flows were captured. Most of the analysis below, however, was performed on a subset of these flows representing 48,301 flows each of whose total transfer exceeded 10 MB. The goal in looking at this smaller set of flows was to select only those flows that attempted to transfer a significant amount of data. This selection criterion removed most interactive TCP sessions (e.g., Telnet, SSH), as well as most reverse ACK flows.[2]

NetFlow sampling was set to 1% for all CLEV input interfaces, allowing all interfaces to be monitored without unduly taxing the router's CPU. To account for sampling, all throughput, packet, and octet figures cited in this paper have been scaled up by 100. For example, given a flow record that reports *octets*, *start_time*, and *end_time*, we compute the average TCP throughput (in bits per second) as $8 \times 100 \times octets/(end\_time - start\_time)$.

C. Validation

Informally, we validated this technique for estimating TCP throughputs by comparing actual and estimated TCP throughputs for several known flows. Outside the recorded test window, 11 TCP test flows were generated using `iperf` configured for various window sizes (and consequent bandwidths). These flows were run across a transcontinental path with a round-trip time of 74.4 ms and ranged in bandwidth from 389 Kbps to 9.1 Mbps. A twelfth flow of much higher bandwidth (86 Mbps) was run across a shorter path having a round-trip time of 20 ms. The actual throughputs of these TCPs were compared against estimates based on the sampled NetFlow records captured at CLEV. The result (plotted in Fig. 1) yields a close fit to the dashed line ($y = x$).

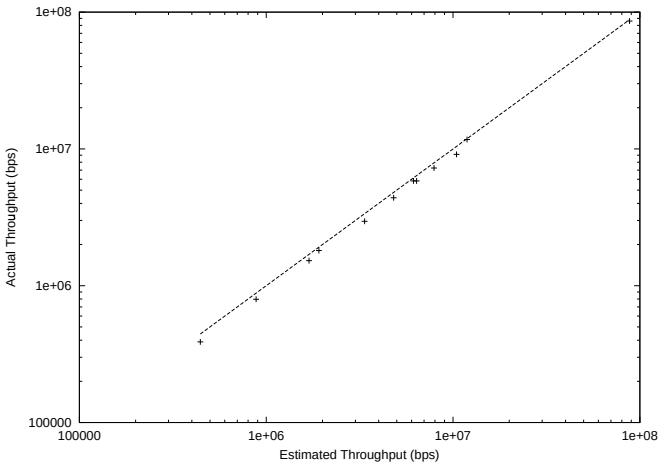

Fig. 1. Estimated *versus* Actual TCP Throughputs (log-log scale)

D. Limitations

Although we are confident that our methodology is appropriate for the study of bulk transfer TCPs, sampled NetFlow data has certain limitations. Most notably and fundamentally:

1. For flows with small numbers of packets, sampling distorts any estimates of the number of packets, octets transfered, or duration of a given flow and therefore the distributions of these parameters; a valid estimate of the number of packets in a flow (and its consequent throughput), requires that NetFlow's 1% sampling have captured at least several packets (i.e. the flow should have transfered several megabytes of data).
2. NetFlow's 1-second time resolution on start and stop times exacerbates error in the estimates of duration for short-lived flows.
3. Certain applications (such as passive FTP and Napster) use connections between dynamically allocated (or not well-known) high-numbered ports for data transfer. If one could reliably observe the thin control flows of these applications, one could hope to be able to detect them.

[2] Nevertheless, roughly 2% of the flows in the bulk transfer dataset had average packet sizes smaller than 200 bytes.

However, sampling makes this problematic.[3]

Due to the first and second limitations, it is impractical to present distributions that include short flows. We hope that ongoing work to install passive monitors in the backbone will allow this study to be extended to include short-lived flows.

Finally, the third limitation precludes useful characterization of passive FTP flows. However, since most command-line FTP clients use active FTP mode by default and there are (anecdotally) very few NATs or firewalls on Internet2, passive FTP traffic is likely limited to sessions established by Web browsers.

III. Analysis of Full Data Set

ALTHOUGH this study focuses on bulk transfer TCPs, we digress briefly to present a macroscopic summary of all flow data captured during the 24-hour observation period. A total of 13,589,626 flows were captured representing 5,708,878,100 packets and 4,124,494,613,400 octets (3.75 TB).[4]

In this section, all percentages given are relative to the total number of reported packets, octets, and flows.

A. Transit Traffic

Traffic that transits the CLEV router (i.e., not sourced or destined for a CLEV connector) comprises 52.6% of all captured bytes, 56.9% of all captured packets, and 68.0% of all captured flows.

B. Application Mix

Some popular applications are characterized in table I. Notice that while figures representing packets and octets transfered should be reliable, the number of flows in tables I and II should only be compared with data collected using a similar methodology. Since the characterization is based on NetFlow data, it is necessarily limited to observation of packet headers. NNTP, active FTP, HTTP, HTTPS, SSH, SMTP, and Telnet all have well-known TCP port numbers. NTP was identified by UDP port 123 and DNS by UDP port 53 or TCP port 53 (to account for zone transfers). NFS was identified by port 2049 (regardless of transport layer protocol). The ICMP row includes all ICMP traffic. Finally, multicast traffic was identified by selecting flows with a class D destination IP address (first four bits equal to 1110).

[3]Note that even with full information from the Internet2 core, we would not be able to detect Napster traffic since the control connection would not traverse the Abilene backbone. Leaf sites, however, have successfully used NetFlow to monitor applications such as Napster [4].

[4]Packet and octet figures have been scaled by 100 to account for the configured 1% NetFlow sampling rate.

Traffic type	Packets	Octets	Flows
NNTP	19.36%	22.79%	10.95%
Active FTP	11.90%	14.24%	3.10%
HTTP	10.50%	9.50%	24.86%
Multicast	6.61%	6.21%	1.91%
SSH	1.86%	1.41%	1.73%
ICMP	1.44%	0.85%	3.29%
DNS	1.42%	0.25%	4.19%
SMTP	1.05%	0.65%	2.51%
Telnet	0.61%	0.09%	1.61%
HTTPS	0.24%	0.14%	0.76%
NTP	0.09%	0.01%	0.38%
NFS	0.04%	0.11%	0.00%
Other	44.88%	43.75%	44.71%

TABLE I
FRACTION OF POPULAR APPLICATIONS (FULL DATA SET)

C. Number of Communicating Hosts

In the full data set, there were 381,430 unique source IP addresses, 540,986 unique destination IP addresses, and 635,731 unique IP addresses (without regard of source or destination). One might expect these three numbers to be close, because any host usually both sends and receives. The bulk of the difference is likely explained by the presence of a significant number of hosts that send or receive only a small amount of traffic, and for some of these hosts, their traffic goes undetected by sampled NetFlow. An additional reason for the different sets of senders and receivers could be some routing asymmetry.

D. IP Protocol Mix

IP protocols distribution is presented in table II. The table doesn't include 244 protocols that carried less than 1/1,000,000th of all octets.

E. Packet Sizes

Packet sizes were estimated by averaging within each flow. The most frequent packet sizes are presented in table III, while the overall distribution is presented in table IV and Fig. 2. There was also a very small number of flows with average packet size in the 4200–4399 range. There were no flows with average packet size in the 1600–4199 range.

In table III, 40-, 41-, and 42-octet average packet sizes are likely explained by the presence of TCP ACKs; 1500-octet packets are likely full-Ethernet-frame packets; 576-octet packets use the default MSS and we are glad that their number is lower than what used to be the norm in

Protocol	Packets	Octets	Flows
ICMP [1]	1.4352%	0.8475%	3.2898%
IGMP [2]	0.0037%	0.0011%	0.0139%
IP-ENCAP [4]	0.2558%	0.3373%	0.0299%
TCP [6]	85.5653%	88.7433%	83.3340%
UDP [17]	12.4660%	10.0249%	12.3365%
IPV6 [41]	0.0074%	0.0028%	0.0277%
GRE [47]	0.0343%	0.0140%	0.0318%
ESP [50]	0.0043%	0.0026%	0.0060%
AX.25 [93]	0.0015%	0.0002%	0.0054%
PIM [103]	0.0088%	0.0008%	0.0150%
Unknown [169]	0.2154%	0.0215%	0.9006%
Other	0.0023%	0.0040%	0.0094%

TABLE II
IP PROTOCOLS DISTRIBUTION (FULL DATA SET)

Packet size	40	1500	52
Frequency	18.88%	12.00%	4.05%
Packet size	552	41	42
Frequency	2.09%	1.83%	1.09%
Packet size	44	4470	576
Frequency	0.90%	0.85%	0.78%

TABLE III
MOST FREQUENT AVERAGE PACKET SIZES (FULL DATA SET)

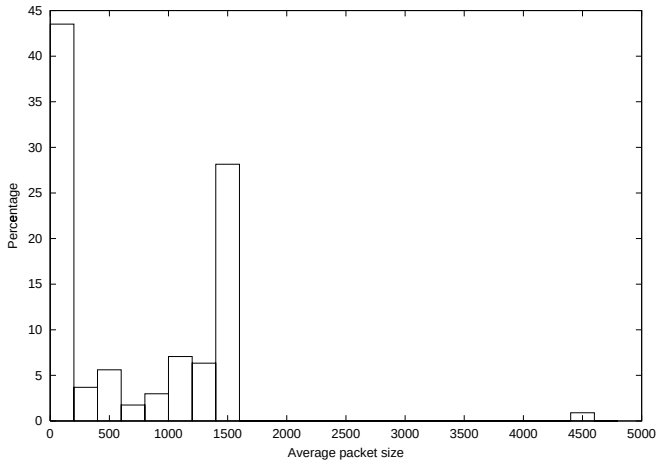

Fig. 2. Average Packet Size Histogram (Full Data Set)

Size range	0–199	200–399	400–599
Fraction	43.51%	3.69%	5.61%
Size range	600–799	800–999	1000–1199
Fraction	1.75%	2.98%	7.07%
Size range	1200–1399	1400–1599	4400–4599
Fraction	6.34%	28.15%	0.90%

TABLE IV
AVERAGE PACKET SIZE DISTRIBUTION (FULL DATA SET)

the past; 44-octet average packet size may be explained by the presence of TCP ACKs with space-occupying options; 4470 octets is the Abilene MTU; finally, 552-octet packets are likely sent by TCPs that generate 512-octet payload. We do not know whether 52-octet average packet size corresponds to any real phenomenon in the network.

IV. ANALYSIS OF BULK TRANSFER TCPS

WE have selected 48,301 TCP flows that have transfered more than 10,000,000 octets[5] for further study and analysis. It should be pointed out that while 1% sampling significantly distorts the distribution of numbers of packets and octets, as well as of durations of all flows, these "bulk transfer" flows' data are relatively safe to simply scale to obtain its more detailed characteristics.

Bulk TCPs transfered 1,508,141,000 packets (26.4% of total) and 2,128,871,432,700 octets (1.93 TB, 51.6% of total), so they represent a significant fraction of network traffic.

Subsequently in this section, all reported percentages are relative to number of flows, octets, and packets of bulk transfer flows.

A. Transit Traffic

In the bulk TCPs data set, there were 26,089 (54%) transit flows that accounted for 785,009,800 (52%) packets and 1,007,297,291,300 (0.9 TB or 47%) of octets.

B. Application Mix

The set of identified popular applications for bulk TCPs was essentially reduced to NNTP and active FTP (see table V).

C. Number of Communicating Hosts

The number of users of bulk TCPs is presented in table VI. The number of NNTP sources approximates the number of net-news servers on Internet2. The number of

[5]Taking sampling into account, we've actually selected flows with more than 100,000 octets.

Traffic type	Packets	Octets	Flows
NNTP	27.48%	27.15%	18.94%
Active FTP	21.93%	19.40%	20.14%
HTTP	2.86%	2.59%	3.52%
HTTPS	0.08%	0.06%	0.05%
Other	47.64%	50.79%	57.34%

TABLE V
FRACTION OF POPULAR APPLICATIONS (BULK TCPS)

	Sources	Destinations	Speakers
All flows	2983	4450	6987
Active FTP	278	501	772
NNTP	42	60	74

TABLE VI
NUMBER OF UNIQUE IPS (BULK TCPS)

active mode FTP sources approximates the number of FTP servers that frequently serve large files on Internet2.

D. Packet Sizes

Packet sizes were estimated in the same manner as for the full data set. Most frequent packet sizes are presented in table VII. Overall packet size distribution is shown in table VIII and Fig. 3. There was also a very small number of flows with average packet size in the 4200–4399 range. There were no flows with average packet size in the 1600–4199 range.

We felt that the fact that some hosts have taken advantage of packets that had exactly the size of Abilene MTU deserved special attention. Notice that since no actual packet can be larger than 4470 octets, the fact that the average packet size for a given flow was 4470 octets (when rounded off to an integer) indicates that an overwhelming majority of packets within the flow actually had sizes of 4470 octets. We have selected all flows from the set of bulk TCPs with average packet size 4470; there were 370 such flows. It turned out that all of them went from the same host to the same host. The source AS was NSFNETTEST14-AS [237] (Merit); the destination as was AMES-NAS-GW [24] (NASA AMES). The destination port number was 5900; source port number was different for each flow. Applications that use the unregistered port 5900 include Virtual Network Computer (VNC), Bergen Sound Server, Jabber, and SharedFXP.

Packet size	1500	552	4470
Frequency	17.42%	4.32%	3.23%
Packet size	1499	1496	1497
Frequency	2.37%	2.34%	2.24%
Packet size	1498	1495	1494
Frequency	2.21%	2.05%	1.86%

TABLE VII
MOST FREQUENT AVERAGE PACKET SIZES (BULK TCPS)

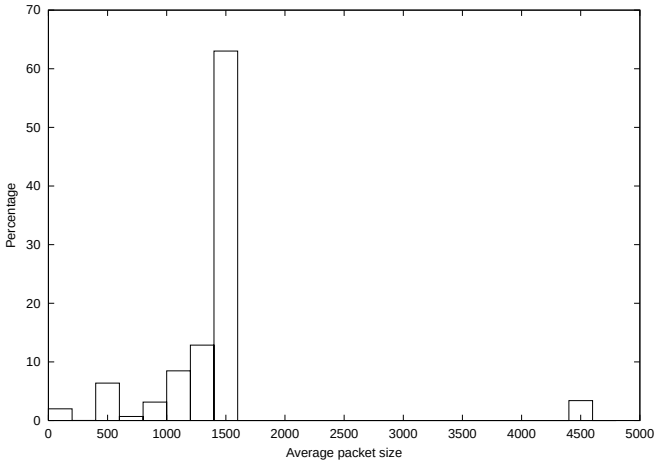

Fig. 3. Average Packet Size Histogram (Bulk TCPs)

Size range	0–199	200–399	400–599
Fraction	2.00%	0.01%	6.39%
Size range	600–799	800–999	1000–1199
Fraction	0.69%	3.14%	8.48%
Size range	1200–1399	1400–1599	4400–4599
Fraction	12.86%	63.02%	3.40%

TABLE VIII
AVERAGE PACKET SIZE DISTRIBUTION (BULK TCPS)

In table VII, 1500-, 1499-, 1496-, 1497-, 1498-, and 1495-octet average packet size might be all explained by the presence of a large number of 1500-octet packets along with a much smaller number of shorter packets.

E. Flow Concurrency

Sampling limitations prevent one from reliably computing flow concurrency for the full data set. At the same time, bulk TCPs are longer and have larger numbers of packets, so the error in their reported start and stop times (relative to their duration) is smaller. Therefore, for bulk TCPs it is possible to find out their average concurrency. The mean concurrency (computed as the sum of durations of all bulk TCP flows divided by the number of seconds in

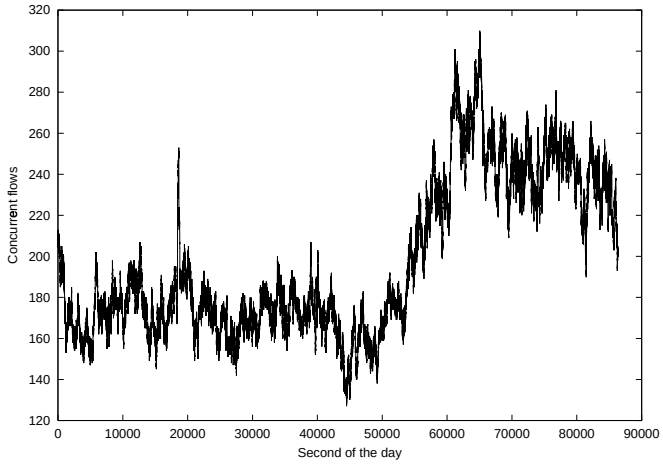

Fig. 4. Number of Concurrent Bulk TCPs

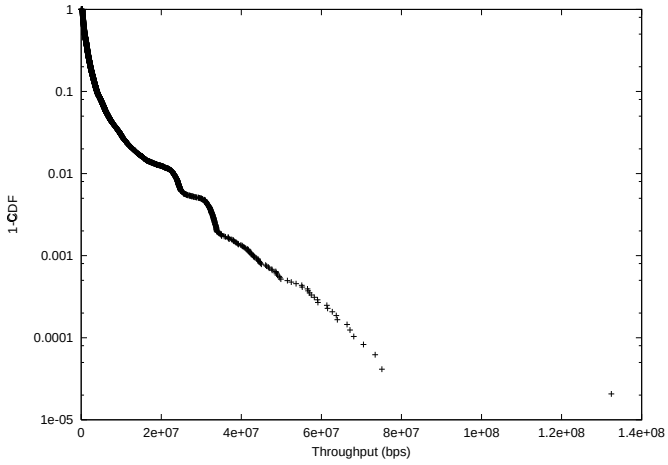

Fig. 5. Throughput Distribution of Bulk TCPs (log scale)

a day) was 197.4.

Further, we were able to compute with an acceptable degree of reliability the number of concurrent bulk TCPs going through the CLEV router at any given point in the day. To avoid any bias at the beginning and the end of the observation period we have considered the time circular for the purposes of concurrency computations (e.g., second 86537 was considered the same as second 137). This wrap-around only affects initial 30 minutes of the observation interval. The number of concurrent flows is relatively stable and doesn't depend much on the time of the day (see Fig.4; note that the scale on the ordinate was chosen to show the maximum amount of information).

F. Throughput Distribution

Observing throughput distribution of bulk TCP flows was the primary motivation of this paper. Its graphical representation is shown on Fig. 5 (logarithmic scale on the ordinate) and Fig. 7 (logarithmic scale on both axes). On

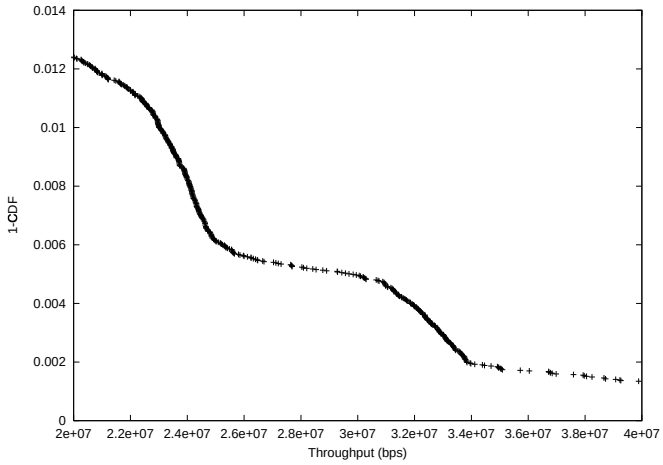

Fig. 6. Throughput Distribution of Bulk TCPs (20–40 Mbps Interval Enlarged)

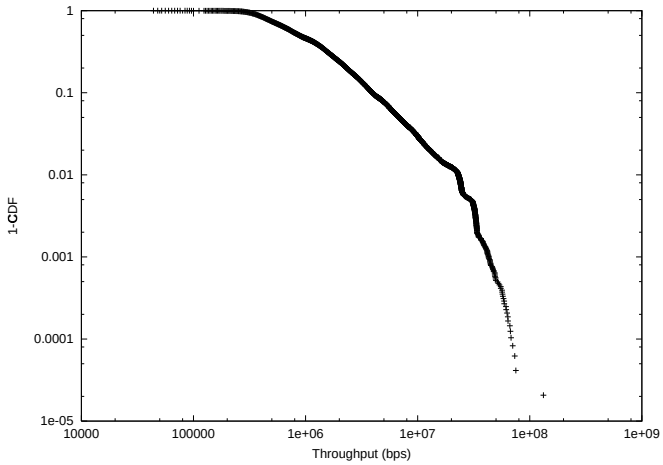

Fig. 7. Throughput Distribution of Bulk TCPs (log-log scale)

	Throughput (Kbps)	Size (KB)	Duration (s)
Min	44	9,766	3
50%	880	16,259	190
10%	3,942	98,111	950
5%	6,780	166,098	1,796
1%	23,000	442,892	1,801
Max	132,416	3,879,384	1,803

TABLE IX
SELECTED POINTS FROM DISTRIBUTION GRAPHS (BULK TCPS)

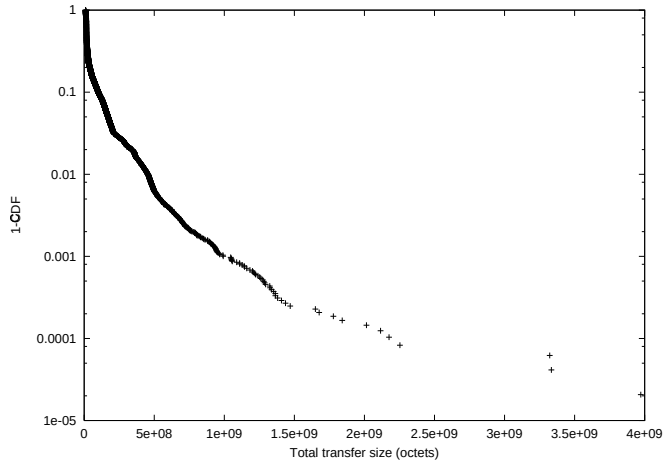

Fig. 8. Transfer Size Distribution of Bulk TCPs (log scale)

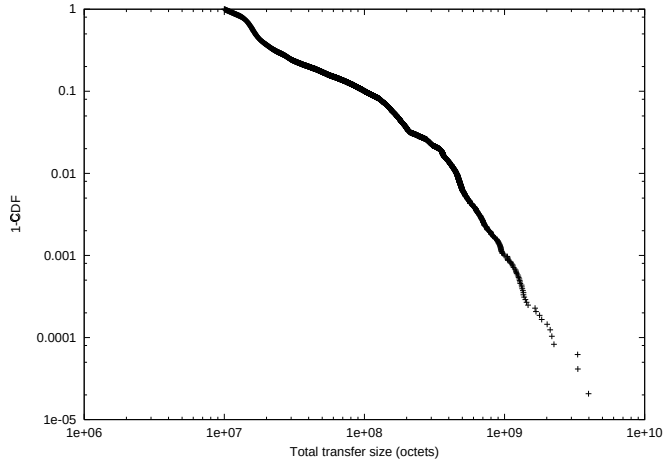

Fig. 9. Transfer Size Distribution of Bulk TCPs (log-log scale)

all distribution plots, we use $1 - \mathrm{CDF}$ on the ordinate; the reason for that (rather than using just the cumulative distribution function) is that we're mostly interested in the tail of the distribution, therefore, having density function (not readily available) or $1 - \mathrm{CDF}$ reveals more information when the ordinate has logarithmic scale.

Notice that the 20–40 Mbps interval has some interesting structure. It is shown in normal coordinates on Fig. 6.

Actual numbers for characteristic values (median, 10%, 5%, and 1% more than) can be found in the first column of table IX.

G. Transfer Size Distribution

Transfer sizes (Fig. 8, Fig. 9) are generally affected by the 30-minute cut-off interval in NetFlow. The main effect is for those 5.3% of flows that were evicted from the NetFlow cache at 30-minute cut-off.

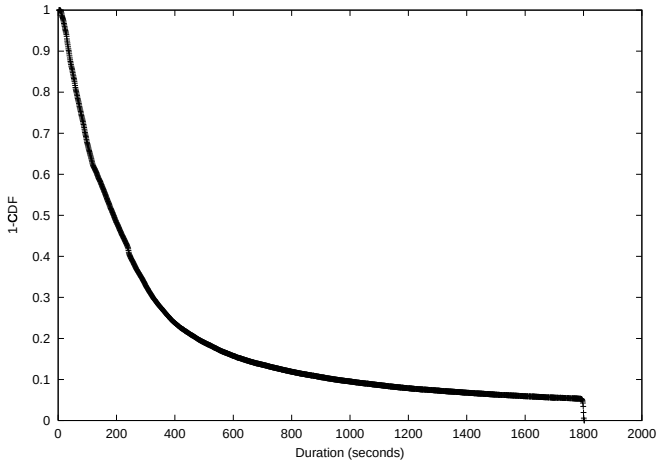

Fig. 10. Duration Distribution of Bulk TCPs

H. Duration Distribution

NetFlow cuts off flows that are longer than 30 minutes. This is the default timer, and we didn't change it. Fig.10 shows the observed distribution of duration of bulk TCPs.

I. Top 10 List

A NetFlow study would be incomplete without a Top 10 list of users. Table X excludes flows that lasted 120 seconds or less (because the throughput gauge is less reliable for them) and lists at most a single (fastest) flow from any given pair of source and destination AS numbers.

V. Discussion

INTERNET2 has a relatively high percentage of ICMP traffic. This is likely explained by a well-developed measurement infrastructure. Higher fraction of UDP than in some earlier studies of various networks probably is explained by deployment of multicast and streaming of multimedia content. We expect that commercial backbones would have lower multicast numbers, even if multicast is enabled (within Internet2 community, a major factor that holds back multicast growth is support of multicast in campus networks).

We expect that TCP throughput on Internet2 should be higher than on commercial networks because of better connectivity of end-users and lack of backbone congestion. It should be said, however, that the observed bulk TCP throughput is surprisingly low given the infrastructure in place (10BaseT half-duplex minimal end-host connectivity and 100BaseT full-duplex typical for anything that requires high network performance; an OC-48c backbone that has virtually no loss and very low jitter, OC-12c or OC-3c access circuits). We hypothesize that the observed bulk TCP throughput is to a large extent explained by in-

App	Mbps	pkt size	secs	Src AS	Src port	Dest AS	Dest port
VNC?	132	4470	240	NSFNETTEST14-AS [237]	2049	AMES-NAS-GW [24]	5900
Unknown	55	1500	239	PSCNET-HS-TEST-AS [1207]	2224	AMES-NAS-GW [24]	5789
Unknown	29	1490	375	CORNELL [26]	4550	NCSA-AS [1224]	10196
Unknown	26	1493	223	PENN-STATE [3999]	17695	ARGONNE-AS [683]	51008
FTP	24	1499	1122	PSCNET-HS-TEST-AS [1207]	2492	NSFNETTEST14-AS [237]	20
FTP	16	1406	260	NLM-GW [70]	20	INDIANA-AS [87]	39130
FTP	12	1400	837	NLM-GW [70]	20	UIUC [38]	9949
Unknown	11	1500	144	NCSA-AS [1224]	45513	BUFFALO-ASN [3685]	1528
FTP	10	1481	293	NASA-HPCC-ESS [7847]	14167	NIST-BOULDER [2648]	20
FTP	9	1423	374	SWCHSC [19925]	3448	UPENN-CIS [55]	20

TABLE X
FASTEST FLOWS WITH UNIQUE AS SOURCE AND DESTINATION (BULK TCPS)

adequate TCP window sizes and failed duplex negotiations with Ethernet switches.

It should be pointed out that we could only identify just over half of the observed traffic. We expect passive mode FTP and Napster to be major contributers to the unidentified fraction.

Mark Fullmer of the Ohio Internet2 Technology Evaluation Center (ITEC) is producing nightly reports [2] based on Abilene NetFlow data.

The questions that we could not answer include:

1. What is the distribution of throughput, duration, and transfer size of all TCP connections?
2. What fraction of the difference between the number of IP sources and destinations (subsection III-C) is explained by routing asymmetry?
3. How are TCP options used on Internet2?
4. What applications are using the network? It would be desirable to identify more than 55% of traffic (subsection III-B); we could not identify passive mode FTP and Napster.
5. For what fraction of bulk TCP connections is throughput limited by the window size and what fraction of packets and octets do these connections carry? What are the other major limiting factors; can one identify duplex mismatch from the backbone?

Having full (and not sampled) NetFlow data from the network would answer questions 1 and 2.

In order to answer question 3, one would need access to more than just NetFlow data. Sampled passive monitoring that captures enough of each packet (e.g., 50 octets) would be sufficient to answer this question.

To answer question 4, full passive monitoring data is required (e.g., one would be able then to observe passive mode FTP control connections and be able to identify data connections). To identify traffic from Napster, more than data from Internet2 core is required, since one would need to observe control connections that do not traverse Internet2. Having a Napster usage study performed at a few leaf sites would contribute to the answer to this question. Observation of packet payloads can help identify more applications.

Finally, to answer the first part of question 5, more than just passive monitoring is required, but rather a combination of passive (to learn about advertized window sizes and achieved throughput) and active (to learn the round-trip time) monitoring; since minimal round-trip time doesn't change very often (only due to routing changes) active and passive measurements can be separated in time. The second part of this question would likely require non-trivial active measurements combined with passive observation; active measurements would have to happen *while the connection is active.* We do not expect that a study to answer question 5 would consider all TCP connections that go through Internet2 core; sampling of connections would be appropriate.

Passive monitoring devices are currently being deployed in Internet2 core to help answer these and other questions.

Internet2 welcomes additional research based on these and related measurement data. Interested researchers should contact the authors.

ACKNOWLEDGMENTS

The authors would like to thank Matt Mathis of Pittsburgh Supercomputing Center for valuable discussion and hands-on help with the calibration process, Claudia DeLuna of NASA Jet Propulsion Laboratory for al-

lowing us to use her `iperf` server across the continent, Mark Fullmer of Ohio State University for operational support of Abilene NetFlow data collection, and Guy Almes and Matt Zekauskas of Internet2 and Advanced Network & Services for numerous suggestions.

REFERENCES

[1] *Abilene Home Page*, http://www.internet2.edu/abilene/
[2] *Abilene NetFlow Nightly Reports*, http://www.itec.oar.net/abilene-netflow/
[3] *Cisco IOS NetFlow*, http://www.cisco.com/warp/public/732/Tech/netflow/
[4] *FlowScan—Network Traffic Flow Vizualization and Reporting Tool*, http://www.caida.org/tools/utilities/flowscan/
[5] *NetFlow Services and Applications*, http://www.cisco.com/warp/public/cc/pd/iosw/ioft/neflct/tech/napps_wp.htm
[6] *Indiana University Abilene NOC Weathermap*, http://hydra.uits.iu.edu/ abilene/traffic/