TKN Telecommunication Networks Group

Technische Universität Berlin
Telecommunication Networks Group

Search and analysis of backbone traffic measurements in the Internet

Filip Idzikowski
idzikowski@tkn.tu-berlin.de

Berlin, March 2009

TKN Technical Report TKN-09-004

TKN Technical Reports Series Editor:

Prof. Dr.-Ing. Adam Wolisz

Abstract

Results of an extensive search for measurements of traffic in backbone networks are presented in this report. Publicly available traffic statistics as well as traffic traces are considered. Moreover a temporal analysis of traffic data rate and packet size distribution is performed over selected data sets.

Contents

1 Introduction

Design of telecommunication networks is strongly dependent on the traffic load it has to carry. Different design assumptions are taken if traffic is of constant bit rate, and if it is very bursty. However, the realistic traffic is a stochastic process, so it is essential to analyse the traffic measurements in order to develop a traffic model with a finite set of parameters, which can be used in simulation experiments. Moreover new Internet applications like Voice over IP, or video services influence characteristics of the Internet traffic, which are expected to change in the long-term [3].

It is not a trivial (if feasible) task to get access to a complete set of traffic measurements (we focus on backbone networks in this work) including source and destination addresses of packets in the backbone network, their size and timestamps, accompanied by a full network topology at different layers (including link capacities and routing schemes). We present the results of an extensive search for traffic measurements available on the web. The aim of this search is to use the measurements as input to simulation experiments and to develop a traffic model which:

- can be used to generate end-to-end packet traffic in a simulated network
- does not model each microflow separately because of scalability reasons, but still allows investigation of a few microflows with the mechanisms like congestion control
- is based on the real measurements

The traffic measurements are available in various forms. We can mainly distinguish two groups: traffic statistics and traffic traces. The first group contains statistical information about the traffic over a given period of time, e.g. average data rates or packet loss rates. Traffic matrices are a special kind of traffic statistics. They contain average data rates of traffic flowing between each pair of nodes in the network, and therefore they are especially important for network modelling. The second group reports on each packet traversing the measurement point. Traces contain different information about each packet, e.g. a timestamp, source and destination IP addresses, source and destination port numbers, various information from packet headers. Both traffic statistics and traffic traces are stored in various formats, which include e.g. XML (eXtensible Markup Language) or CSV (Comma-Separated Values) format for the traffic statistics, and DAG/ERF (Extensible Record Files) and pcap (packet capture) for the traffic traces. The data can also be presented as graphs or network weathermaps (graphical presentation of network topology with marked traffic data, e.g. data rates on each network link, or link utilisation).

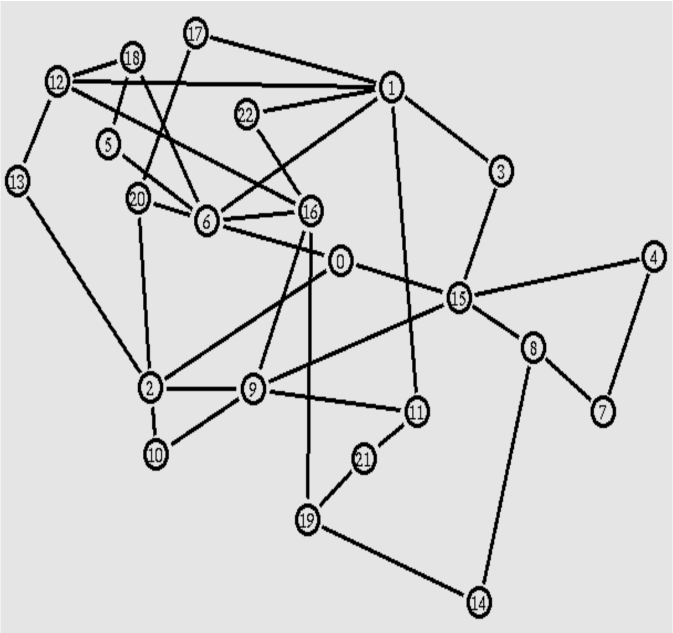

Figure 1: Anonymised topology of the GÉANT network [11]

2 Available traffic measurements

This section reports the results of an extensive search for traffic measurements we performed. We note that the availability of the data described in this report may change in time. Sometimes registration or special access permission is necessary to access data.

2.1 Traffic matrices

We report on the traffic matrices in a separate section due to their importance for network modelling (end-to-end traffic dependencies).

2.1.1 GÉANT

A set of traffic matrices with information about the data rates of end-to-end traffic in GÉANT network is available at the TOTEM Project website [10]. The values in the traffic matrices are averaged over 15 minute periods, and stored in the XML format. The units used for the matrices are kbps. All matrices cover the time period of 119 days starting on 2005-01-01 00:00 (anonymised time). A corresponding network topology is provided and consists of 23 nodes (see Figure 1).

2.1.2 Abilene

Yin Zhang collected and made available [14] the traffic matrices of the Abilene network. The matrices cover 6 months (not continuous) with a 5-minute step (24 x 7 x 24 x 12 = 48384 traffic matrices of 5 minute granularity), and are stored in an ASCII format. The units used for the matrices are (100 bytes / 5 minutes), which is 8/3 bps.

Not only real traffic measurements, but also various traffic estimates are available. This includes simple gravity model, simple tomogravity model, general gravity model and general tomogravity model [15]. A network topology together with the link capacities and OSPF weights is provided as well (12 nodes). It is identical to the one shown in Fig. 6.

2.2 Traffic statistics

2.2.1 Indiana University GNOC Weathermaps

A set of network weathermaps is provided by the Indiana University Global Network Operations Center. Graphical presentation of traffic data of the National LambdaRail network can be found under `http://weathermap.grnoc.iu.edu/`. Network weathermaps contain the following information:

- Layer 1 - topology, utilisation of the links, amount of wavelengths used on each link (Fig. 2)
- Layer 2 - topology, utilisation of the links, data rates in bits per second and packets per second, and number of errors per second (Figures 3 and 4)
- Layer 3 - topology, utilisation of the links, data rates in bits per second and packets per second, and number of errors per second (Fig. 5, note that a link between New York and Chicago used to exist in the past)

Similar information as for NLR Layer 3 is available also for the Abilene (Internet 2) Network (Figures 6 and 7 show supposedly Layer 3 topology - note that recently the nodes Indianapolis and Chicago have been merged to form one node, and that the Los Angeles node disappeared). Moreover, the traffic data rates at each link vs. time can be plotted for Layer 2 (NLR) and Layer 3 topologies (NLR and Abilene). The timescales include last 5 minutes, last hour, last day, last week and last month. Network weathermaps of I-Light (Indiana's Optical Network) and Indiana University's Core and WAN weathermaps are also available.

Raw but incomplete traffic traces in the .rrd format (Round Robin Database) can be found in [12]. It is unspecified in which network the data was caputred, however from the names of the trace files we suspect Layer 2 of the NLR network.

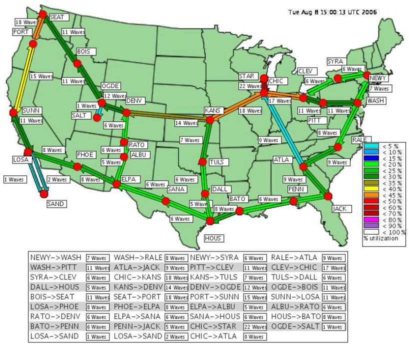

NEWY->WASH	7 Waves	WASH->RALE	8 Waves	NEWY->SYRA	6 Waves	RALE->ATLA	9 Waves
WASH->PITT	11 Waves	ATLA->JACK	9 Waves	PITT->CLEV	11 Waves	CLEV->CHIC	17 Waves
SYRA->CLEV	6 Waves	CHIC->KANS	18 Waves	KANS->TULS	7 Waves	TULS->DALL	6 Waves
DALL->HOUS	5 Waves	KANS->DENV	14 Waves	DENV->OGDE	12 Waves	OGDE->BOIS	11 Waves
BOIS->SEAT	11 Waves	SEAT->PORT	18 Waves	PORT->SUNN	15 Waves	SUNN->LOSA	11 Waves
LOSA->PHOE	8 Waves	PHOE->ELPA	8 Waves	ELPA->ALBU	5 Waves	ALBU->RATO	6 Waves
RATO->DENV	6 Waves	ELPA->SANA	6 Waves	SANA->HOUS	6 Waves	HOUS->BATO	8 Waves
BATO->PENN	6 Waves	PENN->JACK	5 Waves	CHIC->STAR	22 Waves	OGDE->SALT	1 Waves
LOSA->SAND	1 Waves	LOSA->SAND	2 Waves	CHIC->ATLA	0 Waves		

Figure 2: Layer 1 NLR Network Weathermap, Tue Aug 8 15:00:13 UTC 2006 (from [13])

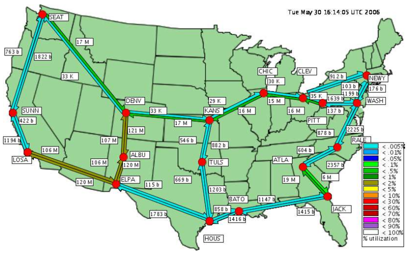

Figure 3: Layer 2 NLR Network Weathermap, Tue May 30 16:14:05 UTC 2006 (from [13])

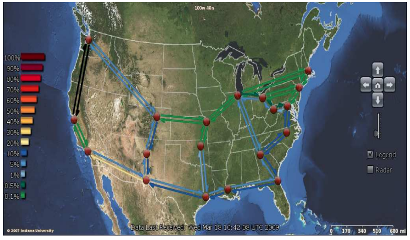

Figure 4: Layer 2 NLR Network Weathermap, Wed Mar 18 10:42:08 UTC 2009 (from [13])

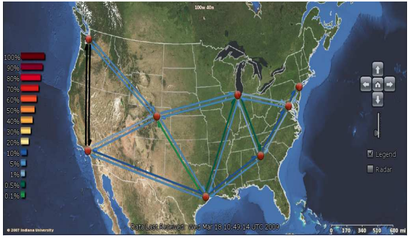

Figure 5: Layer 3 NLR Network Weathermap, Wed Mar 18 10:49:14 UTC 2009 (from [13])

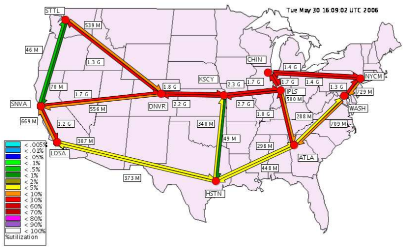

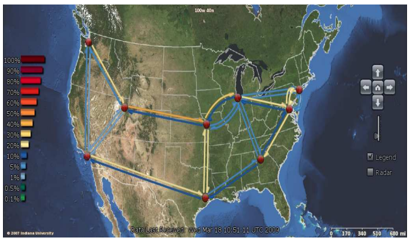

Figure 6: Abilene Network Weathermap, Tue May 30 16:09:02 UTC 2006 (from [13])

Figure 7: Abilene Network Weathermap. Wed Mar 18 10:51:11 UTC 2009 (from [13])

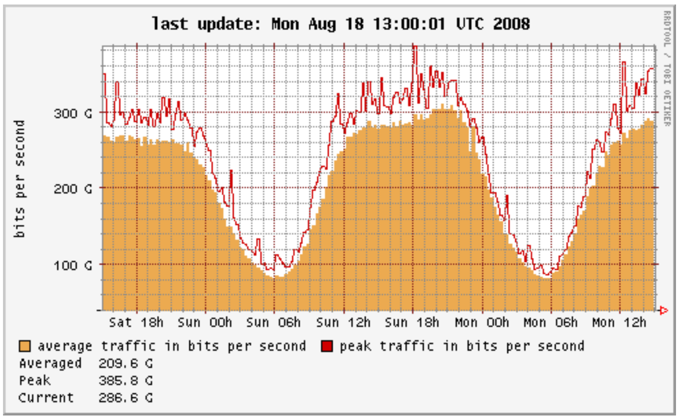

Figure 8: Average and peak data rates in bits per second, daily graph, (from [3], German Internet Exchange DE-CIX)

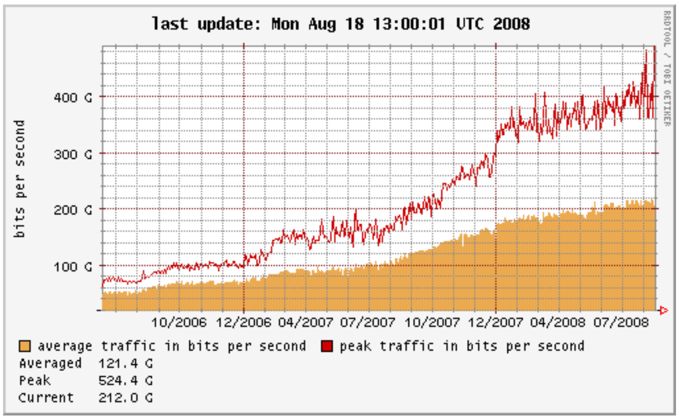

Figure 9: Average and peak data rates in bits per second, yearly graph, (from [3], German Internet Exchange DE-CIX)

2.2.2 European Internet Exchange Points

The European Internet Exchange Points publish graphs of traffic data rates over a day, week, month and year. Examples for a German Internet Exchange are presented in Figures 8 and 9. The graphs are regularly updated every 5-10 minutes.

A list of European Internet Exchange Points is presented below (following `http://www.ep.net/naps_eu2.html` and `http://www.bnix.net/other.php`):

- Austria: VIX - Vienna Internet eXchange
- Belgium: BNIX - Belgium National Internet eXchange, FREEBIX - Free Belgium Internet eXchange
- Bulgaria: SIX - Sofia Internet eXchange

- Croatia: CIX - Croatian Internet eXchange
- Czech Republic: NIX - Neutral Internet eXchange
- Cyprus: CYIX - Cyprus Internet Exchange
- Denmark: DIX - Danish Internet eXchange
- England: LINX - London Internet eXchange, LIPEX - London Internet Providers eXchange, LoNAP - London Network Access Point, MaNAP - Manchester Network Access Point, Manchester Commercial Internet eXchange, RBIEX - A Peering Gateway, SOVEX - Sovereign House Exchange, Xchangepoint - Multi-National
- Estonia: Elion TIX Tallinn Internet eXchange, Linxtelecom TLLIX - Tallinn Internet Exchange
- Finland: FICIX Finnish Commercial Internet eXchange, Tampere Region Internet eXchange - TREX
- France: EuroGix - A Peering Point, FNIX6- eXchange in Paris, FreeIX - A Free French eXchange, LYONIX - Lyon Internet eXchange, MAE - Paris, PARIX - A Paris Internet eXchange, PIES - Paris Internet eXchange Service, PIX - Paris Internet eXchange, POUIX - Paris Operators for Universal Internet eXchange, SFINX - Service for French Internet eXchange, GNI - Grenoble Network Initiative
- Germany: BECIX - Berlin Internet eXchange, BCIX - Berlin Commercial Internet Exchange, DE-CIX - Deutsche Commercial Internet eXchange, ECIX - European Commercial Internet eXchange (formally BLNX) Berlin, ECIX - Dusseldorf HHCIX - Hamburg, INXS - Munich and Hamburg, Franap - Frankfurt Network Access Point, KleyRex - Kleyer Rebstcker Internet eXchange (Frankfurt), MAE - Frankfurt, MANDA - Metropolitan Area Network Darmstadt, M-CIX - Munich Commercial Internet eXchange, N-IX - Nurnberger Internet eXchange, S-IX Stuttgarter Internet Exchange Work-IX Peering Point - Hamburg, Xchangepoint - Multi-National
- Greece: AIX - Athens Internet eXchange
- Hungary: BIX - Budapest Internet eXchange
- Iceland: RIX - Reykjavik Internet eXchange
- Ireland: INEX - Internet Neutral eXchange
- Israel: IIX - Israel Internet eXchange
- Italy: MIXITA - Milan Internet eXchange, NaMex - Nautilus Mediterranean Exchange Point Rome, TOPIX - Torino Piemonte IX, TIX / Tuscany Internet eXchange

- Latvia: Latvian GIX
- Luxembourg: LIX - Luxembourg Internet eXchange
- Malta: MIX - Malta Internet eXchange
- Netherlands: AMS-IX - Amsterdam Internet eXchange, GN-IX Groningen Internet eXchange, NDIX - A Dutch German Internet eXchange, NL-IX - NL- Internet eXchange
- Norway: NIX - Norwegian Internet eXchange
- Poland: GIX - Polish Global Internet Exchange, KIX - Krakowski Internet Exchange, SIX Silesian Internet Exchange, WIX - Warsaw Internet eXchange, WRIX - Wroclaw Internet eXchange
- Portugal: GIGAPIX - Gigabit Portuguese Internet eXchange
- Romania: BUHIX - Bucharest Internet eXchange, Ronix - Romanian Network for Internet eXchange
- Russia: MPIX - Moscow Internet eXchange, NSK - IX, RIPN Home Page (MSK - IX / M9 - IX / SPB - IX), Samara IX, SIMIX - Simbirsk Internet Exchange (Ulyanovsk, Russia)
- Scotland: WorldIX - European Commercial IX (Edinburgh), ScotIX - Scottish Internet Exchange
- Slovakia: Slovak Republic, SIX - Slovak Internet eXchange
- Spain: Catnix - Catalunya Neutral Internet Exchange, ESPANIX - Spanish Internet Exchange, GALNIX - Galicia Internet eXchange, MAD-IX - Madrid Internet eXchange Punto Neutro Español de Internet
- Sweden: Linkoping Municple Exchange, NorrNod, NETNOD Internet eXchange, PolarIX (formally LIX), RIX -GH Gaveleborg Regional Internet Exchange, SOL-IX - Stockholm
- Switzerland: CIXP - CERN Exchange for Central Europe, SWISSIX - Swiss Internet Exchange, TIX - Equinix Zurich Exchange
- Ukraine: UA-IX - Ukrainian Internet Exchange

2.2.3 GRNET Network Monitor

The Network Monitor of Greek Research & Technology Network (GRNET) [9] provides among others the following tools and information:

- Nagios Monitoring Tool (restricted access)

- Database Visualization - visualises the GRNET architecture (including the structure of GRNET nodes)
- Network Weathermap - shows in a graphical way the utilisation of GRNET links, the traffic data rates on each link as well as CPU and memory usage of each node
- GRNET2 Athens MAN - shows the link utilisation and traffic data rates on the links for the GRNET2 Athens MAN (a network weathermap)
- IPv4 and IPv6 Looking Glass - allows the queries of GRNET routers
- Traffic Statistics Graphs - provides the plots of average and maximum traffic data rates (in bits per second) for the GRNET links (including the link to GÉANT network) over last 32 hours (other timescales used to be available (as for August 2008) - a day with granularity of 5 minutes, a week with granularity of 30 minutes, a month with granularity of 2 hours and a year with granularity of 1 day).
- Packet Traffic Statistics Graphs - the same as above in packets per second
- QoS diagrams - the dropped bit rate (average and maximum) for each GRNET link is provided for the same periods as above
- Multicast Weathermap - shows the traffic load for a specific group of IP addresses
- SSMping looking glass - allows pinging various destination addresses using ipv4 and ipv6
- Ping Delay and Packet Loss - provides the delay and packet loss vs time curves plotted using ping measurements from GRNET links to GÉANT, Root Name Servers, Greek Universities and international hosts. The plots show the daily, weekly, monthly and yearly time dependencies (minimum and maximum values).
- GRNET Router Status - shows the status of a router. This includes cpu load, memory used, utilisation and status of router interfaces. The first two parameters can also be plotted over a day, week, month and year.

The network topology, but no end-to-end traffic data rates are provided.

2.2.4 CAIDA traffic monitors

The Cooperative Association for Internet Data Analysis (CAIDA) supports the following realtime traffic monitors [1]:

- equinix-chicago monitor - passive network monitor on OC192 backbone link of a Tier1 ISP between Chicago, IL and Seattle, WA in both directions

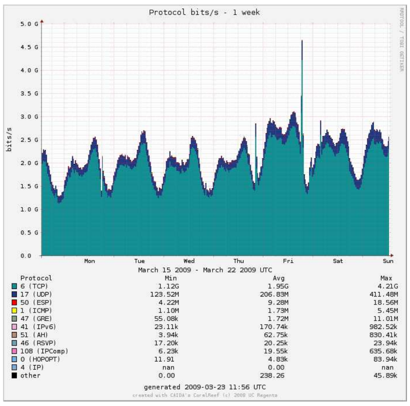

Protocol	Min	Avg	Max
🟦 6 (TCP)	1.12G	1.95G	4.21G
🟦 17 (UDP)	123.52M	206.83M	411.48M
🟥 50 (ESP)	4.22M	9.28M	18.56M
🟨 1 (ICMP)	1.10M	1.73M	5.45M
🟩 47 (GRE)	55.08k	1.72M	11.01M
🟪 41 (IPv6)	23.11k	170.74k	982.52k
🟫 51 (AH)	3.94k	62.75k	830.41k
🟦 46 (RSVP)	17.20k	20.25k	23.94k
🟪 108 (IPComp)	6.23k	19.55k	635.68k
🟦 0 (HOPOPT)	11.91	4.83k	83.94k
🟦 4 (IP)	nan	0.00	nan
⬛ other	0.00	238.26	45.89k

generated 2009-03-23 11:56 UTC

created with CAIDA's CoralReef (c) 2008 UC Regents

Figure 10: Statistics from the Chicago passive network monitor (the week 15-22.03.2009 UTC, from [1])

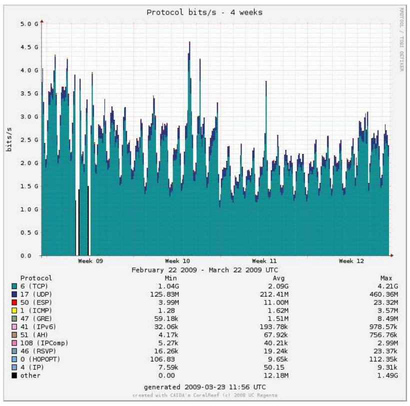

Protocol	Min	Avg	Max
🟦 6 (TCP)	1.04G	2.09G	4.21G
🟦 17 (UDP)	125.83M	212.41M	460.36M
🟥 50 (ESP)	3.99M	11.00M	23.32M
🟨 1 (ICMP)	1.28	1.62M	3.57M
🟩 47 (GRE)	59.18k	1.51M	8.49M
🟪 41 (IPv6)	32.06k	193.78k	978.57k
🟫 51 (AH)	4.17k	67.92k	756.76k
🟪 108 (IPComp)	5.27k	40.21k	2.99M
🟦 46 (RSVP)	16.26k	19.24k	23.37k
🟦 0 (HOPOPT)	106.83	9.65k	112.35k
🟦 4 (IP)	7.59k	50.15	9.31k
⬛ other	0.00	12.18M	1.49G

generated 2009-03-23 11:56 UTC

created with CAIDA's CoralReef (c) 2008 UC Regents

Figure 11: Statistics from the Chicago passive network monitor (one weeks 9-12 of 2009, from [1])

- The San Diego Network Access Point (SDNAP) - passive network monitor on SDANP - a peering point of various Organisations and Internet Service Providers in the San Diego area

- ampath-oc12 monitor (Miami,FL) - passive network monitor on OC12 link between AMPATH International Exchange Point (Florida International University) and Internet2 (Miami, Florida). It was active until the 18th of March 2008.

- equinix-sanjose - a network monitor on OC192 backbone link of a Tier1 ISP between San Jose, CA and Los Angeles, CA in both directions. The statistics are not directly linked from [1]. They can be found under `http://www.caida.org/data/passive/monitors/equinix-sanjose.xml`.

The data rate versus time traffic statistics can be obtained from these monitors (see e.g. Figures 10 and 11). The statistics can be broken down with respect to: the Layer 3 Protocol (port), Application, Source Country, Destination Country, Source AS, Destination AS (no AS breakdown for the equinix-chicago and equinix-sanjose monitors). The classification with respect to countries and ASes often returns the answer ’unknown’ though. Units of bits/sec, packets/sec and tuples/sec (flows/sec) are available. Statistics can be presented in a broad variety of views and timescales (1 day, 1 week, 4 weeks and 2 years). Since at OC192 line rates the traffic report generators used in equinix-chicago and equinix-sanjose can’t keep up with generating flow files that are needed for the realtime traffic reports, the reports are generated using flow estimation described in [4].

2.2.5 Sprint

The Sprint Academic Research Group [7] presents results of analysis of the traces collected by the IPMON systems on more than 30 bidirectional OC3/12/48 links between August 9th, 2000 and January 11th, 2005. The results show dependency of the following metrics on time: link utilisation (in Mbps and packets per second), number of active flows (a flow is defined as a set of packets with the same protocol number, destination and source IP addresses, and port numbers), traffic breakdown by protocol (in packets per second and Mbps), traffic breakdown by application (as flow, packet and byte percentage of the traffic). Moreover packet size distribution, delay statistics and routing info (BGP - Number of Prefixes, Number of ASes, Prefixes vs. AS hops, Address Space vs. AS hops) are available. Unfortunately, the network topology, at which the measurements were performed is unavailable. The traces themselves are publicly unavailable.

2.2.6 WITS: Waikato Internet Traffic Storage

Traffic data rate vs. time plots based on traffic traces can be found in the Waikito Internet Traffic Storage [8]. The data rate can be presented in packets per second, and in Mbits per second. The following traces (unavailable at WITS webpages) were analysed [8]:

- Auckland - seven measurement campaigns at the University of Auckland (in most cases an OC3 link) performed between July 1999 and December 2003. All non-IP traffic has probably been discarded, in which case there will only be TCP, UDP and ICMP traffic present in the trace. In most of the traces any user payload within the 64 byte capture record has been zeroed. Two traces contain ATM cell headers, and not the IP packets. Trace formats are ERF, DAG formats, legacy ATM.

- Local ISP A - this is a collection of traces taken at a New Zealand ISP using tcpdump on a Linux box located inside the ISP's internal network in the following periods: from Tue Nov 2 14:04:18 1999 to Wed Nov 10 06:56:16 1999 (local time), from Fri Dec 10 16:09:05 1999 to Fri Dec 17 14:51:48 1999 (local time), and from Mon Dec 20 16:02:30 1999 to Wed Jan 5 18:35:40 2000 (local time). The traces are stored in the pcap format.

- Local ISP B - this collection of traces was taken at a New Zealand ISP. The measurement periods were: from Thu Feb 24 10:45:00 2005 to Thu Feb 24 21:15:00 2005 (local time), from Thu Jun 9 16:53:50 2005 to Fri Jun 10 19:45:00 2005 (local time), and from Thu Feb 8 11:53:28 2007 to Mon Feb 12 21:17:37 2007 (local time). Packet records are truncated four or eight bytes after the end of the transport header except in the case of DNS traffic, which is snapped twelve bytes after the end of the transport header. The traces are stored in the ERF format.

- NZIX-I - this is a collection of ten-minute traces captured using a proprietary software solution at the New Zealand Internet Exchange between Thu Nov 12 10:56:44 1998 and Sun Apr 11 14:11:12 1999 (local time). Information about an ISP connection is unavailable. A customised trace format (legacy NZIX) was used that recorded a timestamp, the packet wire length and a CRC. The first 54 bytes of the packet starting from the Ethernet header was captured.

- NZIX-II - this is a collection of GPS-synchronised IP header traces captured using a DAG32E at the New Zealand Internet Exchange between Tue Jun 27 18:33:03 2000 and Mon Jul 10 15:29:05 2000 (local time). All non-IP traffic has been discarded and only TCP, UDP and ICMP traffic is present in the trace. Any user payload within the 64 byte capture record has been zeroed. The trace format is Legacy Ethernet (DAG).

- Waikato - six sets of traces captured at the border of the University of Waikato network between Sun Dec 7 00:00:01 2003 and Sat Mar 15 03:21:28 2008 (local time). The packets have been truncated at the end of the transport header or four bytes after the end of the transport header, except for DNS. No user payload is included in any of the packets. However, ICMP packets are truncated after 8 bytes of ICMP header - the IP header for the original datagram is not included. Trace format is ERF.

No information about the network topology is given. Plots showing the traffic data rates in bits per second and packets per second are available in the Waikito Internet Traffic Storage, but no traces themselves. Some traces (parts of NXIX-II and Auckland) can be found at the site of NLANR Measurement & Network Analysis (see section 2.3.2) though.

2.3 Traffic traces

Not all traces found in the Internet are presented in this section. We considered only traces recorded in the backbone networks, or at the access of the big institutions (like universities).

2.3.1 CAIDA traces

CAIDA provides three sets of traces. The first one has been captured on an OC48 link, the second one on an OC12 link, and the third one on an OC192 link.

The OC48 traces (The CAIDA OC48 Traces Dataset - Apr 24 2003, Colleen Shannon, Emile Aben, kc claffy, Dan Andersen, Nevil Brownlee, `http://www.caida.org/data/passive/passive_oc48_dataset.xml`) were collected during three time periods: from 2002-08-14 16:00 UTC (+0000) to 2002-08-14 18:59:59.999 UTC (+0000), from 2003-01-15 17:59:34.091 UTC (+0000) to 2003-01-15 19:01:36.908 UTC (+0000), and from 2003-04-24 07:00 UTC (+0000) to 2003-04-24 07:59:59.999 UTC (+0000). The traces contain anonymised packet headers collected in both directions of an OC48 west coast peering link for a large ISP.

The OC12 traces were collected from CAIDA's AMPATH monitor at the AMPATH Internet Exchange (see section 2.2.4) from 2007-01-08 23:00:30 UTC to 2007-01-11 01:00:30 UTC. They contain bidirectional packet header traces with no payload from an OC12 ATM link.

The OC192 traces were collected by equinix-chicago and equinix-sanjose monitors (see section 2.2.4) from approximately 2008-03-19 19:00 to 20:00 UTC (single direction only) and from 2008-03-19 00:00 to 06:00 UTC (the same direction) in the commercial backbone. The traces are anonymised and contain traffic from an Internet backbone link. Moreover anonymised OC192 traces (year 2009) have been made available recently.

All the traces are stored in a pcap format. Topology of the network in which the traces were collected is unavailable.

2.3.2 NLANR Measurement & Network Analysis

A set of traffic traces collected at different networks can be found under [5]. The traces are stored either in DAG or ERF formats, however the trace format is not explicitly mentioned in the description of a few traces.

- Abilene - five traffic traces, two of them collected on an OC48c Packet-over-SONET links (eastbound and westbound, towards Cleveland and Kansas City from Indianapolis). It consists of a pair of two hour contiguous bidirectional packet header traces collected at the Indianapolis router node (IPLS), which is (at the time of collection) a CISCO GSR 12015 with four OC48c uplinks, four OC12c links and one OC3c link. The traces are stored in the DAG format. One traffic trace contains information from the OC192c Packet-over-SONET link from Internet2's Indianapolis (IPLS) Abilene router node towards Kansas City (KSCY) (four-hour data collected on June 1st, 2004). The other traffic trace contains stratified random sampling header data from all three backbone links at IPLS between June 18th, 2004 and August 19th, 2004. The fifth traffic trace contains information about router delay data between T640 links. The trace is stored in the DAG PoS format. Additionally, a network topology is provided.

- AMPATH - a ten-day collection of stratified randomly sampled ten-minute IP header traces collected at AMPATH, Miami, FL, in March 2005 (OC3MON).

- University of Auckland - six traffic traces, among others from the university's Internet access link (46 days altogether), a three point measurement (three different measurement points taken simultaneously, trace records are of 64 bytes fixed length, containing full TCP/IP and UDP/IP headers in most cases, 4.5 days), and an ATM cell header trace collected at the University of

Auckland OC3c ATM link (13.5 hours). The traces are stored in the DAG format.

- Bell Labs - a one week contiguous Internet access IP header trace collected at Bell Labs research, Murray Hill, NJ, at the end of May 2002.

- CENIC-I - 48-hour contiguous data set collected on the 10 Gigabit CENIC HPR backbone link between Sunnyvale and Los Angeles between Thursday March 17th and Saturday March 19th, 2005.

- CESCA-I - a three-hour (10 am - 1 pm) GPS-synchronised IP header trace captured with an Endace DAG4.2GE dual Gigabit Ethernet network measurement card in February 2004 at the Anella Cientfica (Scientific Ring), the Catalan R&D network. A network graph as well as five minute and hourly graphs (data rate vs. time including protocol breakdown) are available.

- Leipzig - two traces - one continuous five-day GPS-synchronised IP header trace (from Thu Nov 21 20:00:00 2002 to Tue Nov 26 14:00:00 2002) and one discontinuous one-day GPS-synchronised IP header trace (between Fri Feb 21 12:13:59 2003 and Sat Feb 22 21:00:00 2003). Both traces were taken with a pair of DAG3 cards at the University of Leipzig Internet access link (OC3 Packet-over-Sonet link running at 155.52 Mbits/s, connection to the German research network (G-WiN)). Trace records are of 64 bytes fixed length, containing full TCP/IP and UDP/IP headers in most cases. Delay graphs are provided too.

- NCAR-I - a one-hour IP header trace captured by NLANR PMA with an Endace DAG4.2GE dual Gigabit Ethernet network measurement card at the end of January 2004.

- NZIX-II - a five-day IP header trace collected at the New Zealand Internet Exchange. The collection is dominated by a contiguous five-day trace starting on Wednesday 5th of July 2000, containing approximately 843 million IP headers (see also section 2.2.6). The trace is stored in a DAG format (fixed 64 bytes record format with 40 bytes of IP header (usually covering most, if not all, of the TCP/IP and UDP/IP headers)).

- SC2004 Bandwidth Challenge Collection - OC192MON was operated from Monday November 8th through to Thursday November 11th, 2004. Most of the time the OC192MON was collecting and analysing data in real time with one major gap between Tuesday night and Wednesday morning, during which the system was collecting IP packet header trace data - this data set. The OC192MON was initially tuned into the Abilene link towards New York. This configuration was changed on Tuesday night, after which the system was observing the Abilene link to Chicago, until it was turned off on Thursday afternoon. All times are Eastern Standard (Pittsburgh, PA, local time). Graphs showing statistics (data rates in bits per second and packets per second, number of active connections, number of new connections, average connection time, amount of packets per connection, amount of bits per connection, Dag loss counter, and one minute load averages) are available as well.

- San Diego-I - twelve (originally thirty) hour IP header trace captured by NLANR PMA with an Endace DAG4.2GE dual Gigabit Ethernet network measurement card at the end of January 2004.
- Tera - 10GigE traces, collected with an NLANR PMA OC192MON located on SDSC's TeraGrid Cluster during the week starting Sunday 8th of February 2004. The trace file format is Endace's ERF, fixed record sizes at 88 Bytes each. Graphs containing application breakdown in bits/sec, application breakdown in packets/sec, IP protocol breakdown in bits/sec, IP protocol breakdown in packets/sec are available as well. The traces are stored in the ERF format. The SPSC TeraGrid topology is available.

2.4 Traffic databases

The following traffic databases have been found during the search:

- MOME - Cluster of European Projects aimed at Monitoring and Measurement (`http://www.ist-mome.org/database/` - not updated any more)
- DatCat - Internet Measurement Data Catalog (`http://imdc.datcat.org/Home`)
- NLANR - Measurement & Network Analysis (`http://pma.nlanr.net/PMA/StatQuery.html`)
- WITS - Waikato Internet Traffic Storage (`http://wand.cs.waikato.ac.nz/wits`)
- SNAPP - collections of SNMP Network Analysis and Presentation Package (`http://dc-snmp.wcc.grnoc.iu.edu/i2net/raw-data.cgi`)

3 Traffic data analysis

3.1 Temporal characteristic of the Internet traffic

Fluctuations of the traffic data rates are crucial for the performance of the telecommunication networks. The bursty nature of traffic causes congestions in the networks as well as jitter at the traffic destinations. However, the temporal traffic characteristics are also a motivation for the traffic grooming algorithms and dynamic virtual topology design. We analyse some data rate vs. time plots in this section.

3.1.1 GÉANT

The GÉANT traffic matrices [11] (see section 2.1.1) contain the all end-to-end traffic data rates, what is of great advantage over single traffic traces. We visualised and analysed the temporal data rates of each end-to-end flow (some values are missing in the measurement sets).

We plotted the data rates over three time intervals: a month, a week and a day.

We classified the monthly traffic (march 2005 of anonymised time) between each node pair into the following categories:

(a) characteristic with one base level and occurring positive peaks (Fig. 12)

(b) characteristic with one base level and occurring positive and negative peaks (Fig. 13)

(c) a mix of the above characteristics with a rapid transition from one to the other (Fig. 14)

(d) characteristic with two base levels and occurring positive and negative peaks and hops between these two levels (Fig. 15)

(e) characteristic with a decreasing base level (Fig. 16)

The first two characteristics can be found very frequently in the GÈANT network (the topology is shown in Fig. 1). The temporal distances between the peaks (a) differ, however the peaks are usually 4-5 times higher than the base level. The peaks at the characteristic (b) are 2-3 times higher (taking absolute values of both the positive and negative peaks) than the base level and occur very often. The mix (c) has the properties of (a) and (b) and the hops between both take place rarely (or never - we observed just a single characteristic with one hop in the March 2005 traffic). The characteristics (c), (d) and (e) are almost singular cases.

The ratio between the maximum average traffic data rate between a pair of nodes (71 MBps from node 10 to node 16) and the minimum positive average traffic data rate (1.01 MBps from node 17 to node 17 - internal traffic; 1.31 MBps from node 14 to node 21 - external traffic) in March 2005 (anonymised time) equals over 70:1. Looking at the dynamics of the traffic between nodes 0 and 16 (Fig. 20) it can be observed that the data rate often varies by 50-100 MBps. Three parameters describing the traffic were found out after analysis of the 1 month of the data measurements: base level of the traffic data rate, the gap between two peaks and the average value of a peak in relation to the base level.

The analysis of the weekly traffic resulted in the following classification:

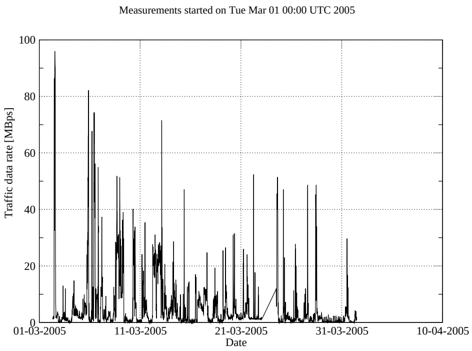

Figure 12: Monthly traffic on the 0-1 GÉANT Path [11]

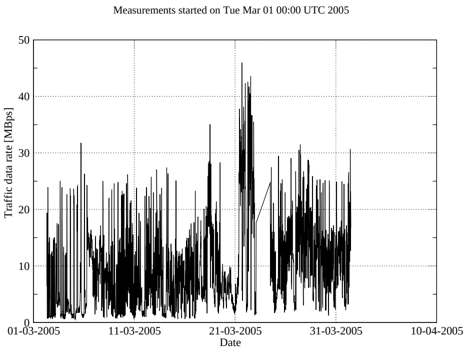

Figure 13: Monthly traffic on the 3-0 GÉANT Path [11]

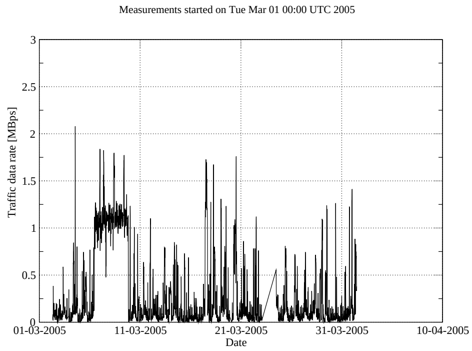

Figure 14: Monthly traffic on the 9-22 GÉANT Path [11]

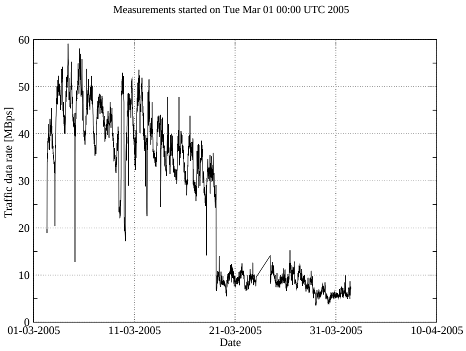

Figure 15: Monthly traffic on the 8-16 GÉANT Path [11]

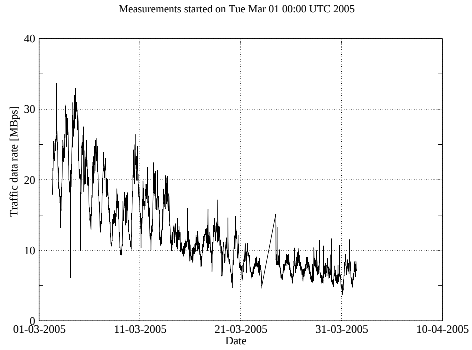

Figure 16: Monthly traffic on the 7-8 GÉANT Path [11]

(a) characteristic with periodic deviation that reflects the 7 days of a week (i.e. 5 working days, and a weekend - Fig. 17)

(b) characteristic with a base level and occurring positive peaks (Fig. 18)

(c) characteristic with a base level and occurring positive and negative peaks (Fig. 19)

(d) characteristic with a hop between two base levels of traffic data rate and diverse deviations (Fig. 20)

The above mentioned characteristics can be observed in all investigated weeks of the traffic data. The set of parameters describing the traffic is almost identical to the parameters from the monthly analysis. However, the values of the parameters are different (e.g. the gaps between two peaks are stochastically distributed in the monthly analysis, while quite regular periods can be observed in the weekly analysis).

The following characteristics resulted from the daily analysis of traffic:

(a) characteristic with two base levels and two hops between these two levels (from the higher level to the lower one and backwards, see Fig. 21)

(b) characteristic with a base level and two peaks (Fig. 22)

(c) characteristic with continuously changing base level (highest increase of the traffic data rate at the evening time) and interfering peaks (Fig. 23)

(d) characteristic with a constant base level, and positive and negative peaks (Fig. 24)

The maximum values of traffic can be usually observed in the same time of a day. The traffic on Monday is different from the traffic on Sunday (different amount of peaks and their distribution). Fig. 25 shows a sample traffic characteristic over 24 hours on a Sunday between two edge nodes. The increase of traffic data rate on evening and night time (anonymised!) is quite noticeable. Such a characteristic is not a common one though. The daily changes of the traffic data rate has very often completely different shapes (see e.g. Fig. 26). Therefore general modelling of the traffic seems to be quite complex.

Summing up, the end-to-end traffic can be defined using the following parameters:

- amount of base levels of the traffic data rate
- value of the base level of the traffic data rate
- number of hops between the base levels
- number of peaks in a time period

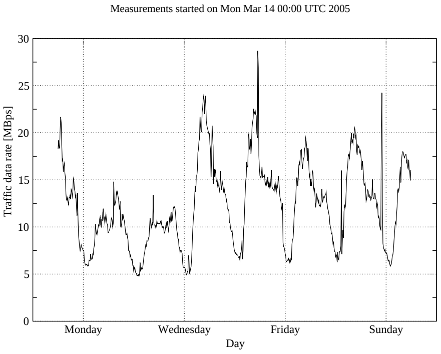

Figure 17: Weekly traffic on the 15-7 GÉANT Path [11]

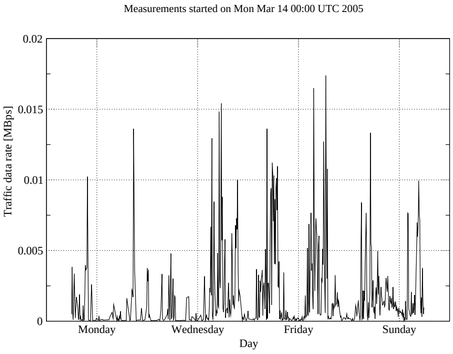

Figure 18: Weekly traffic on the 1-22 GÉANT Path [11]

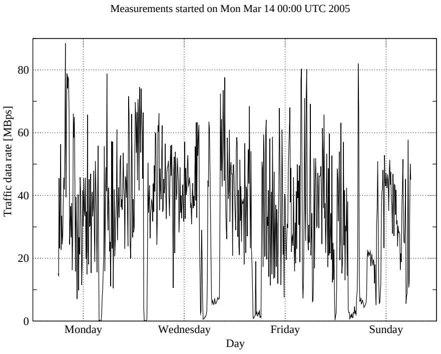

Figure 19: Weekly traffic on the 0-6 GÉANT Path [11]

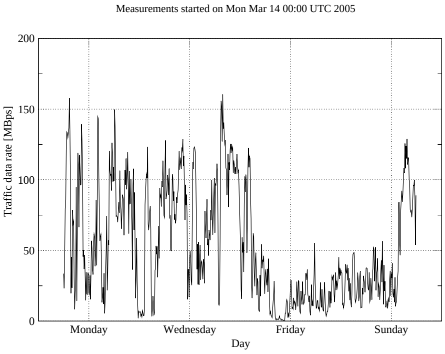

Figure 20: Weekly traffic on the 0-16 GÉANT Path [11]

- values of the peaks in relation to the base level
- gap between the peaks
- periodicity of the traffic

The continuous changes of the traffic in the GÉANT network occur quite rarely. Therefore the list above does not include any corresponding parameter.

3.1.2 Abilene

We analysed also the Abilene traffic matrices [14] (see section 2.1.2). We considered the measured traffic matrices and not their estimates. The patterns of days can be observed in most of the traffic demands (see e.g. Fig. 27 showing the data rate between New York and Atlanta on the week 01-07.03.2004) unless the data rate is very low (few Mbps). In that case some positive peaks can be observed (similar to

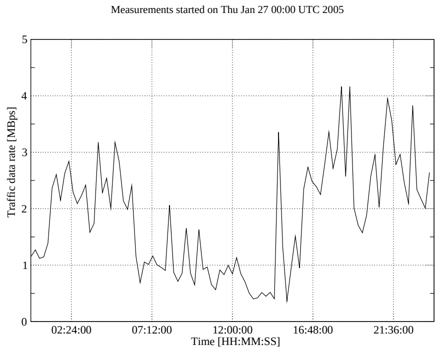

Figure 21: Daily traffic on the 11-21 GÉANT Path [11]

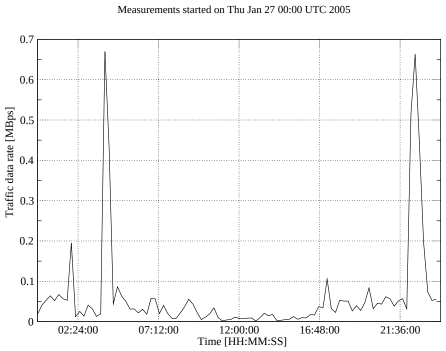

Figure 22: Daily traffic on the 15-10 GÉANT Path [11]

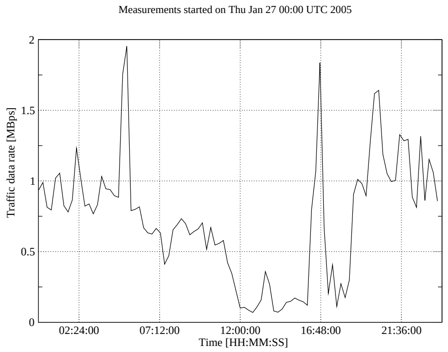

Figure 23: Daily traffic on the 11-10 GÉANT Path [11]

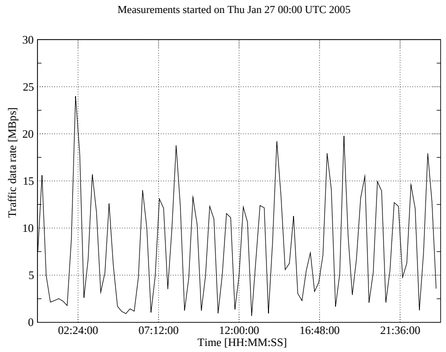

Figure 24: Daily traffic on the 3-2 GÉANT Path [11]

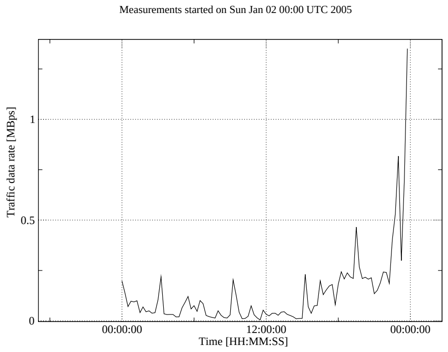

Figure 25: Daily traffic on the 0-10 GÉANT Path [11]

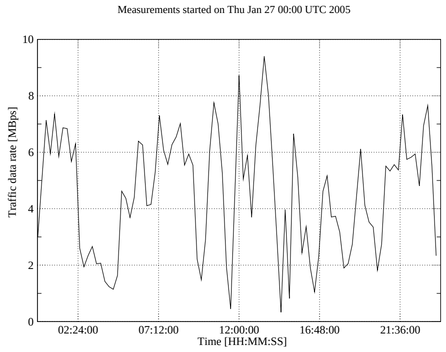

Figure 26: Daily traffic on the 3-6 GÉANT Path [11]

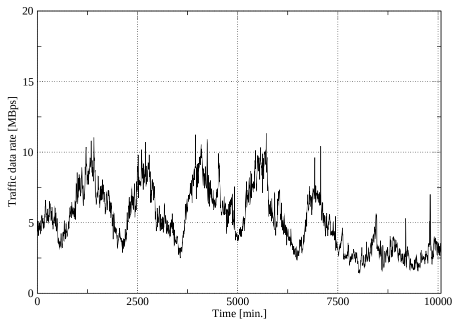

Figure 27: Traffic between New York and Atlanta in the Abilene Network [14] on the week 01-07.03.2004

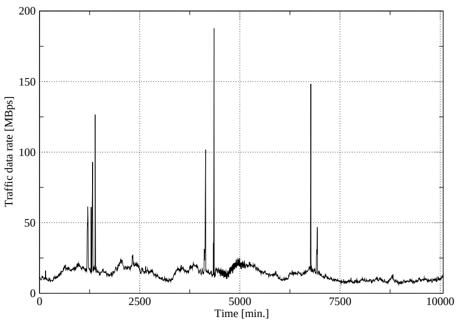

Figure 28: Traffic between Los Angeles and Chicago in the Abilene Network [14] on the week 01-07.03.2004

the GÉANT Network). Fig. 28 shows extreme peaks that reach even 1.5 Gbps (traffic between Los Angeles and Chicago on the week 01-07.03.2004). Rapid changes of the data rate of traffic demand is not an odd phenomenon though. Since Abilene is a research network, these rapid changes may be coupled to starting a new experiment and transmitting big amount of data from one research lab to another one. Silence periods can be observed as well, where almost no data is transfered between nodes.

3.1.3 Statistics from the European Internet Exchange Points

The numerous Exchange Internet Points presented in section 2.2.2 usually publish plots similar to the ones presented in the previous two sections. The analysis of these plots (dated 22.08.2008) has shown differences in the total traffic data rates that can be observed at various exchange points. On one hand, the AMSIX (Amsterdam Internet eXchange), DE-CIX (Deutsche Commercial Internet eXchange - see Figures 8 and 9) or LINX (London Internet eXchange) handle traffic in range of hundreds of Gbps (the observed peak of input traffic on 22.08.2008 in AMSIX was 416,637 Gbps), and on the other hand the majority of the Exchange Points handle traffic around 10 Gbps. The minimum was observed at the CIX (Croatian Internet eXchange), where the traffic was on average 111.486 Mbps on 22.08.2008.

The day of a week pattern could be observed at most of the Exchange Points, and so could the yearly increase of traffic data rate. An interesting observation could be made at the plots of CATNIX (Cataluya Neutral Internet eXchange) traffic data rates. Peaks reaching 72 Gbps (in daily and weekly graphs) or even 10 Tbps (yearly graph) can be observed, while the yearly average data rate was 623.4 Mbps.

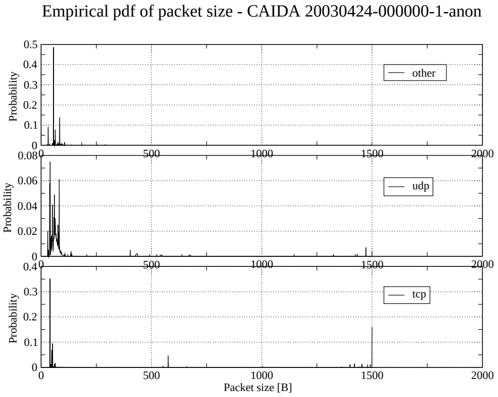

Figure 29: Packet size distribution of one of CAIDA OC48 traces (about 90% of the traffic is tcp) [2] separately for layer 4 protocols

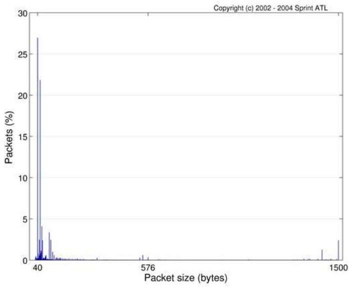

Figure 30: Packet size distribution of one of Sprint traces (from [6])

3.2 Packet size distribution

We investigate the distribution of packet sizes recorded in traffic traces.

3.2.1 CAIDA OC48

In particular, we analyse the CAIDA OC48 traces [2] (see section 2.3.1). We take the five-minute trace starting on 2003-04-24 07:00 UTC (+0000) [`http://imdc.datcat.org/collection/1-0018-N=CAIDA+OC48+Traces+2003-04-24`] (2003-04-24 09:00:00 stored in the trace itself) as a representative example of the CAIDA OC48 traces. Its empirical distribution function of packet sizes in is shown in Fig. 29. Note that about 90% of the packets are TCP packets. The trimodal distribution commonly assumed in the literature [6] is confirmed (peaks at 40 Bytes, 576 Bytes and 1500 Bytes). However, the weights of the peaks have changed. The frequency of occurrence of packets of size 1500 Bytes got higher, while the frequency of occurrence of packets of size 576 Bytes got lower. The most frequently seen packets are the ones of size 40 Bytes, which corresponds to the TCP SYN and ACK packets.

3.2.2 Sprint

We present the results of the Sprint Academic Research Group (see section 2.2.5) as a reference point to the results presented in the previous section (CAIDA OC48 traces). We show the packet size distribution of one of the Sprint packet traces (sj-00.0-0511-0 taken on the 10th of January 2005) in Fig. 30, where the tcp packets constitute almost 80% of all the packets in the traces. The majority of the packets is of size 40 Bytes or slightly bigger. However peaks at 1500 Bytes and 576 Bytes can be observed as well.

4 Conclusions

The amount of traffic measurements publicly available in the Internet is limited. A single traffic trace provides information about the packet size distribution, packet interarrival times etc., however it is just a snapshot of the network. On the other hand, traffic matrices contain the end-to-end traffic dependencies, but their granularity is usually much lower than that of packet traces. Moreover, the behaviour of single packet microflows, is dependent on the network topology, link capacities, routing algorithms and congestion in the network. It can be concluded that the most complete set of publicly available information including the network topology, link capacity, OSPF link weights, set of traffic matrices, as well as single traffic traces can be found for the Abilene network.

After analysing different kinds of traffic measurements, coming from different networks, we found out that data rates plotted out of the traffic matrices are significantly smaller than the ones published by the European Internet Exchange Points. The data rates in the considered measurements are relatively small.

As for the packet size distribution, there are fewer and fewer 576-byte packets in the Internet. This decrease is compensated by the increasing number of small packets (around 40 Bytes in size).

More efficient usage of network resources is highly dependent on the traffic. More efficient grooming and routing policies depend on the knowledge of the real traffic. Therefore more complete data sets than these available today are essential.

References

[1] CAIDA Internet Data. Realtime Monitors. `http://www.caida.org/data/realtime/index.xml` (accessed on 23.03.2009).

[2] CAIDA OC48 Trace Project. CAIDA OC48 Traces 2003-04-24 (collection). `http://imdc.datcat.org/collection/1-0018-N=CAIDA+OC48+Traces+2003-04-24` (accessed on 19.03.2009).

[3] German Internet Exchange DE-CIX. DE-CIX Traffic Statistics. `http://www.de-cix.net/content/network/Traffic-Statistics.html` (accessed on 19.03.2009).

[4] C. Estan, K. Keys, D. Moore, and G. Varghese. Building a Better NetFlow. In *In ACM SIGCOMM*, pages 245–256, 2004.

[5] National Laboratory for Applied Network Research (NLANR). NLANR PMA: Special Traces Archive. `http://pma.nlanr.net/Special/` (accessed on 22.08.2008).

[6] C. Fraleigh, S. Moon, B. Lyles, C. Cotton, M. Khan, D. Moll, R. Rockell, T. Seely, and C. Diot. Packet-Level Traffic Measurements from the Sprint IP Backbone. *IEEE/Network*, November/December 2003.

[7] Sprint Academic Research Group. Packet Trace Analysis. `https://research.sprintlabs.com/packstat/packetoverview.php` (accessed on 22.08.2008).

[8] WAND Network Research Group. WITS: Waikato Internet Traffic Storage. `http://wand.cs.waikato.ac.nz/wits/` (accessed on 19.03.2009).

[9] Greek Ministry of Development. Greek Research & Technology Network. `http://netmon.grnet.gr/` (accessed on 19.03.2009).

[10] The Totem Project. GÉANT Traffic Matrices. `http://totem.info.ucl.ac.be/dataset.html` (accessed on 19.03.2009).

[11] S. Uhlig, B. Quoitin, J. Lepropre, and S. Balon. Providing Public Intradomain Traffic Matrices to the Research Community. *ACM SIGCOMM Computer communication Review*, 36(01):83–86, 2006.

[12] Indiana University. Collection of SNMP Network Analysis and Presentation Package. `http://dc-snmp.wcc.grnoc.iu.edu/i2net/raw-data.cgi` (accessed on 19.03.2009).

[13] Indiana University. Global Network Operations Center Weathermaps. `http://weathermap.grnoc.iu.edu/` (accessed on 19.03.2009).

[14] Yin Zhang. Abilene traffic matrices. `http://www.cs.utexas.edu/~yzhang/research/AbileneTM/` (accessed on 19.03.2009).

[15] Yin Zhang, Matthew Roughan, Nick Duffield, and Albert Greenberg. Fast Accurate Computation of Large-Scale IP Traffic Matrices from Link Loads. In *In ACM SIGMETRICS*, pages 206–217, 2003.