
Experiences of Internet Traffic Monitoring with Tstat

Alessandro Finamore, Marco Mellia, Michela Meo, and Maurizio M. Munafò, Politecnico di Torino
Dario Rossi, TELECOM ParisTech

Abstract

Since the early days of the Internet, network traffic monitoring has always played a strategic role in understanding and characterizing users' activities. In this article, we present our experience in engineering and deploying Tstat, an open source passive monitoring tool that has been developed in the past 10 years. Started as a scalable tool to continuously monitor packets that flow on a link, Tstat has evolved into a complex application that gives network researchers and operators the possibility to derive extended and complex measurements thanks to advanced traffic classifiers. After discussing Tstat capabilities and internal design, we present some examples of measurements collected deploying Tstat at the edge of several ISP networks in past years. While other works report a continuous decline of P2P traffic with streaming and file hosting services rapidly increasing in popularity, the results presented in this article picture a different scenario. First, P2P decline has stopped, and in the last months of 2010 there was a counter tendency to increase P2P traffic over UDP, so the common belief that UDP traffic is negligible is not true anymore. Furthermore, streaming and file hosting applications have either stabilized or are experiencing decreasing traffic shares. We then discuss the scalability issues software-based tools have to cope with when deployed in real networks, showing the importance of properly identifying bottlenecks.

Since the early days of the Internet, network traffic monitoring has always played a strategic role in understanding and characterizing users' activities. Nowadays, with the increased complexity of the Internet infrastructure, applications, and services, this role has become more crucial than ever. Over the years, a number of methodologies and tools have been engineered to assist in the daily routines of traffic monitoring and diagnosis, and to understand network performance and users' behavior [1].

To analyze a system, researchers can follow experimental science principles, and devise controlled experiments to induce and measure cause-effect relationships, or observational science principles, and study an unperturbed system. In the specific field of network traffic measurement, the above two disciplines are referred to as *active* and *passive* measurements, respectively. The active approach aims at interfering with the network to induce a measurable effect, which is the goal of the measurement itself. Active approaches generate traffic, for example, by injecting specifically crafted probe packets or altering the network state, say, by enforcing artificial packet loss. A number of Internet monitoring tools are based on active probing, ranging from simple operation management or network tomography via *ping* or *traceroute*, to more complex delay and capacity estimation via *capprobe* or *pathchar*. Finally, large and controlled

testbeds can be set up easily using tools like *netem* or *dummynet*. For the passive approach, pure observations are performed by means of dedicated tools, named sniffers by the Internet metrology community, which simply observe and analyze the traffic that flows on links. Several passive measurement tools are available. Some tools, such as *tcpdump* or *Wireshark*, are designed to let researchers interactively analyze the captured packets. Other tools are automated instead so that human interaction is minimized; examples are the flowlevel monitoring tool *NetFlow*, intrusion detection systems like *Snort* or *Bro*, and the traffic classification tool *CoralReef*. A comprehensive list of both active and passive tools can be found in [1].

Tstat is an example of an automated tool for passive monitoring. It has been developed by the networking research group at Politecnico di Torino since 2000 [2]. Tstat offers live and scalable traffic monitoring up to gigabits per second using off-the-shelf hardware. It implements traffic classification capabilities, including advanced behavioral classifiers [3], while at the same time offering performance characterization capabilities of both network usage and users' activities [4]. After more than 10 years of development, Tstat has become a versatile and scalable application, used by several researchers and network operators worldwide. In this article, we report our experience with Tstat development and use. We illustrate as a case study traffic evolution as observed during the last year at different vantage points in Europe, and discuss some issues about the feasibility of Internet traffic monitoring with common PCs that can help researchers avoid common pitfalls we have faced in the past.

This work has been supported by the European Commission through the NAPAWINE Project (Network-Aware P2P-TV Application over Wise Network), ICT Call 1 FP7-ICT-2007-1.




Figure 1. a) Tstat monitoring probe setup; b) analysis workflow.

Tstat Overview

Tstat started as an evolution of *tcptrace* [5], which was developed to track and analyze individual TCP flows, offering detailed statistics. Tstat's initial design objective was to automate the collection of TCP statistics of aggregated traffic, adding real-time traffic monitoring features. Over the years, Tstat evolved into a more complex tool offering rich statistics and functionalities. Developed in ANSI C for efficiency, Tstat is today an open source tool that allows sophisticated multi-gigabit-per-second traffic analysis to be run live using common hardware. Tstat design is highly flexible, with several plugin modules offering different capabilities that are briefly described in the following. In addition, plugins can be activated and deactivated on the fly, without interrupting monitoring. Being a passive tool, live monitoring of Internet links, in which all flowing packets are observed, is the typical usage scenario. Figure 1a sketches the common setup for a probe running Tstat: on the left there is the network to monitor (e.g., a campus network). It is connected to the Internet through an access link that carries all packets originated from and destined to terminals inside the monitored network. The Tstat probe observes the packets and extracts the desired information. Note that this scenario is common to a wide set of passive monitoring tools; therefore, the problems faced when designing Tstat are common to other tools as well.

Monitored Objects

The basic objects passive monitoring tools consider are *IP packets* transmitted on the monitored link. Flows are then typically defined according to some rules to group all packets identified by the same *flow_{ID}* that have been observed in a given time interval. A common choice is to consider *flow_{ID}* = (*ipProtoType*, *ipSrcAddr*, *srcPort*, *ipDstAddr*, *dstPort*) so that TCP and UDP flows are considered. For example, in the case of TCP, the start of a new flow is commonly identified when the TCP three-way handshake is observed; similarly, its end is triggered when either a proper TCP connection teardown is seen, or no packets have been observed for some time. Simi-

larly, in the case of UDP, a new flow is identified when the first packet is observed, and it is ended after an idle time.

As Internet conversations are generally bidirectional, two opposite unidirectional flows (i.e., having symmetric source and destination addresses and ports) are then typically grouped and tracked as *connections*. This allows separate statistics to be gathered for *client-to-server* and *server-to-client* flows (e.g., the size of HTTP client requests and server replies).

Furthermore, the origin of information can be distinguished, so it is possible to separate local hosts from remote hosts in the big Internet. As depicted in Fig. 1a, traffic is then organized in four classes:

- **Incoming traffic:** The source is remote and the destination is local.
- **Outgoing traffic:** The source is local and the destination is remote.
- **Local traffic:** Both source and destination are local.
- **External traffic:** Both source and destination are remote.

This classification allows statistics to be separately collected about incoming and outgoing traffic; for example, one could be interested in knowing how much incoming traffic is due to YouTube, and how many users access Facebook from the monitored network. The local and external cases should normally not be considered, but can be present in some scenarios.

At packet, flow, and application layers, a large set of statistics can be defined and possibly customized at the user's will. For Tstat, several statistics are already available, and they can easily be customized and improved because Tstat is open source. A detailed description of all available measurement indices can be found in [2, 7].

Workflow Analysis

As far as the analysis process is concerned, each observed packet is handed over to each active plugin, as illustrated in Fig. 1b. Following the Internet naming standard and going up in the protocol stack, layer 2 (L2) frame decapsulation is first done. Then the network layer (layer 3, L3) header is processed. Given the datagram service offered by IP networks, at L3 only per-packet statistics (e.g., bit rate, packet length) are possible.

Going up to the transport layer (layer 4, L4) analysis, a set of common statistics for both TCP and UDP flows are maintained, such as packet and byte counters, round-trip time (RTT), and data download throughput.

At the application-layer (L7), the main goal of a monitoring tool is to perform traffic classification task, that is to identify the application that generated the traffic. As traffic classification is known to be prone to fallacies, several approaches have been studied in the literature [6, 8]. Each tool has its peculiarities. In the case of Tstat, three different engines are available, each relying on different technologies. They are designed to work even when the complete packet payload is not available, which is a common situation in live network monitoring, since usually only a limited portion of each packet is exposed to the sniffer due to privacy reasons. Tstat implements a deep packet inspection (DPI) technology similar to the one adopted in tools like *l7filter* and *OpenDPI*. In addition to what is done in those tools, Tstat offers a richer set of statistics that complement pure classification, making it more flexible than pure classification tools such as the ones mentioned earlier. Other commercial tools like *ntop* and *peak-flow* offer not only classification capabilities but complete solutions for traffic monitoring, anomaly detection, and security. Designed to work on operative networks, they rely on NetFlow/sFlow technology as input; that is, they do not analyze packet traces, but data records with flow-level statistics. Typically, their deployment is highly invasive, as several analyzers have to be deployed at strategic points in the network. Tstat, instead, is open source software designed to be installed on common hardware at the edge of the network where packet-level traces are analyzed.

In particular, the simplest classification engine offered by Tstat is Pure DPI (PDPI). It uniquely identifies applications by matching a *signature* in the application payload. All the application signatures are collected in a dictionary, defining a set of classification rules, and then checked against the current packet payload until either a match is found or all the signatures have been tested. In the first case, the packet/flow is associated with the matching application, while in the second case it is labeled “unknown.” Signatures cover a large set of applications, ranging from standard email protocols to peer-to-peer applications like BitTorrent, eMule, Gnutella, PPLive, and Sopcast. Extending and updating the signatures is a key issue with PDPI, as we discuss later.

The second engine, Finite State Machine DPI (FSMDPI), inspects more than one packet of a flow. Finite state machines are used to verify that message exchanges conform to the protocol standard; a specific sequence of matching rules have to be triggered to have a positive match. For example, if the first packet contains GET http:// and the response carries HTTP/1.0 OK, the flow can be considered as HTTP. Using this approach, more complex signatures can be defined, allowing web-based applications like YouTube, Vimeo, Facebook, and Flickr, or chat services like MSN, XMPP/Jabber, and Yahoo Messenger, to be identified. Finally, voice over IP (VoIP) calls based on RTP/RTCP are identified using FSMDPI as well.

To cope with applications that leverage on encryption mechanisms that make any DPI classifier useless, Tstat implements a Behavioral Classifier (BC) engine which exploits statistical properties of traffic to distinguish among applications. For example, packet size or interarrival time in flows carry information about the application generating the content, so VoIP flows exhibit different characteristics with respect to data download flows. Using this approach, Tstat identifies encrypted traffic like that generated by Skype and the obfuscated P2P file sharing of BitTorrent and eMule [3].

We present some results that exploit the traffic classification capabilities of Tstat. While the performance and accuracy of the classifier are out of the scope of this article, overall, they have been found to “outperform [other] signature based tools used in the literature” when compared by independent researchers [9].

Input Data

Software-based monitoring tools like Tstat are designed to work in real time when installed in operational networks. The software tool runs on a “probe,” a dedicated PC that “sniffs” traffic flowing on an operative link, as shown in Fig. 1a. The libpcap library is the de facto standard application programming interface (API) to capture packets from standard Ethernet linecards under several operating systems. Dedicated high-end capture devices such as Endace DAG or AITIA S1GED cards are also available on the market.¹ They offer hardware packet monitoring solutions that offload the CPU guaranteeing higher performance than software based solutions. Tstat supports both standard sniffing based on libpcap, and hardware solutions as the ones mentioned earlier.

Furthermore, Tstat can also be compiled as a “library” to allow easy integration with already existing tools such as those typically deployed by an ISP that already has a monitoring solution. In the latter case, the ISP is free to decide which packets should be processed by Tstat to cope with privacy and anonymization issues. In our experience, this approach has been very successful in facilitating the integration of Tstat with the monitoring tools of several ISPs around Europe and with other traffic analysis tools developed by the research community.

Besides live traffic analysis, monitoring tools are also commonly adopted to process packet-level traces that have been previously collected. In this case, the tool can be used to inspect specific traffic for post-mortem analysis, or to develop more complex statistical analysis for advanced performance evaluation, or to double check the accuracy of any new index that is being developed. Since several trace file formats are available on the market, a variety of file formats should be supported, such as *pcap*, *erf*, *etherpeek*, *snoop* to name a few. Besides already supporting a large set of trace input file format, Tstat allows new formats to easily be integrated thanks to its open and flexible design.

Output Statistics

Each monitoring tool offers a set of output statistics that are strictly bound to the goal of the tool itself. For example, intrusion detection systems like *Snort* or *Bro* output the list of triggered alarms and violations, while traffic classification tools like *Tie* or *CoralReef* report statistics about application traffic shares. Considering Tstat, statistics are available at different granularities: per-packet, per-flow, and aggregated. At the finest level of granularity, *packet traces* can be dumped into trace files for further offline processing. This output format is extremely valuable when coupled with Tstat classification capabilities; indeed, packets generated by different applications can be dumped to different files. For example, it is possible to instruct Tstat to only dump packets generated by Skype and BitTorrent applications.

At an intermediate level of granularity, *flow-level logs* are text files providing detailed information for each monitored flow. A log file is arranged as a simple table where each column is associated with specific information, and each row reports the two unidirectional flows of a connection. Several

¹ <http://www.endace.com>, <http://www.aitia.ai>

Location	Users	Technology	Type
Polish ISP	10,000	ADSL	Home
Hungarian ISP	4000	ADSL	Home
Italian ISP	15,000	ADSL	Home
Italian ISP	5000	FTTH	Home
Italian campus	10,000	LAN and WLAN	Campus

n Table 1. *Probe characteristics.*

flow-level logs are available, such as the log of all UDP flows or of all VoIP calls. The log information is a summary of the connection properties. For example, the starting time of the VoIP call, its duration, the number of suffered packet losses, and jitter are all valuable metrics that allow monitoring VoIP quality of service. Flow-level logs use much less space than the original packet-level traces, and can be collected for much longer periods of time.

At an even higher level of granularity, Tstat gathers statistics about flows aggregates. Two formats are available in this case. *Histograms* are empirical frequency distributions of collected statistics over a set of flows. For example, the distribution of the VoIP call duration is automatically computed by considering all VoIP flows that were observed during each 5-min time interval. To overcome the problem of storage space explosion of packet traces, flow-level logs, and histograms over time, the second available format is represented by Round Robin Database (RRD) [10]. It allows a database to be built that spans several years by limiting the amount of disk space. RRD handles historical data at different granularities: newer samples are stored at higher frequencies, while older data are averaged in coarser timescales. This dramatically reduces the requirements in terms of disk space (a priori configurable); and thanks to the tools provided by the RRD technology, it is possible to visually inspect the results. For example, RRD data collected by a Tstat probe can be queried in real time using a simple web interface [2]. Results presented earlier are obtained from the corresponding RRD data.

Traffic Trends from Different Vantage Points

After having presented the main Tstat features and characteristics, we now show Tstat capabilities through a few results and discuss some conclusions drawn from our long experience in using it.

We have been collecting measurement data since 2005 in collaboration with several Internet service providers (ISPs). A Linux-based Tstat probe has been installed and properly configured at different points of presence (PoPs). The results presented in this article refer to five European PoPs and characterize 20 months of traffic collected from May 2009 to December 2010.

Probe Description

The main characteristics of the five probes are summarized in Table 1, which reports the PoP location, the approximate number of aggregated users, the access technology, and the type of customers distinguishing between home or campus users. As can be observed, the set of probes is very heterogeneous: it includes home users in three different countries, with asymmetric digital subscriber line (ADSL) or LAN and wireless LAN (WLAN) access technologies. Depending on the type of contract with the ISP and the quality of the physical medium, ADSL technology offers users different bit rates,

ranging from 2 to 20 Mb/s downstream and up to 1024 kb/s upstream. Fiber to the home (FTTH) customers are offered 10 Mb/s full duplex Ethernet connectivity, while campus users are connected to a 10Gb/s-based campus network using either 100 Mb/s Ethernet, or IEEE 802.11a/b/g WiFi access points. The campus network is connected to the Internet via a single 1 Gb/s link, and a firewall is present to enforce strict policies, block peer-to-peer (P2P) traffic (unless obfuscated), and grant access only to official servers inside the campus.

Probes were upgraded several times to update the Tstat version and include advanced features to enhance traffic classification accuracy and augment the number of protocol signatures. All probes are configured to continuously collect RRD information.

Traffic Share and Trends

We first present results from May 1, 2009 to December 22, 2010. Figure 2 shows the traffic breakdown for incoming traffic (i.e., traffic received by customers). The applications generating the largest amount of traffic are highlighted using different colors. Gaps in the figures correspond to outage periods of the probes. Over time, we enhanced the classification portfolio of Tstat by adding both PDPI/FSM DPI and behavioral rules. For example, since June 2009 we have been collecting statistics about both streaming applications, such as YouTube, Vimeo, Google video, and other flash-based streaming services, and file hosting web-based services like RapidShare and MegaUpload that allow users to share large files. Light and dark pink colors highlight them in the plots. Double checked in the campus network first, we then deployed these capabilities into other probes. Similarly, since December 2009 the BitTorrent obfuscated traffic (plotted in light green) is correctly identified by Tstat, and the more recent BitTorrent UDP-based data transport protocol uTP [11] is correctly classified since July 2010 (dark red). This latter classifier was developed while investigating the cause of the sudden increase of UDP traffic share clearly visible in the Hungarian vantage points during February 2010. This is an example of the usage of Tstat to effectively support traffic monitoring.

Several considerations can be derived from the presented results.

- Before the uTP protocol was adopted by BitTorrent, the volume of UDP traffic was marginal in all vantage points except in Italian ISP. This is due to this ISP offering video on demand (VoD) services over UDP, which makes the volume of VoD UDP traffic in this network about 10 percent of the total. Customers of the same operator are offered native VoIP service using standard RTP/RTCP protocols over UDP. Still, the volume of traffic due to VoIP is almost negligible, accounting for less than 2 percent of total traffic (dark purple in the figure). Nowadays, UDP traffic can top 30 percent of total volume, depending on the popularity of BitTorrent-uTP or VoD applications. Therefore, the widely popular statement that UDP traffic is negligible does not hold anymore.

- Applications' usage is very different at different places. For example, in Poland the fraction of HTTP traffic is predominant, with more than 60 percent of traffic due to several applications adopting HTTP. In both the Italian ISP PoPs, instead, P2P applications amount to more than 50 percent of traffic, with eMule clearly being preferred over BitTorrent. In Hungary, on the contrary, BitTorrent is more popular (with a traffic share above 20 percent), while an almost negligible amount of traffic is due to eMule. Finally, note that in the Italian campus network the fraction of P2P traffic is marginal in traffic share.

- Some long-term trends are clearly visible. For example, until July 2010, P2P traffic share was generally decreasing,




Figure 2. Comparison of traffic as observed on five different traffic probes. Gaps in the data were due to temporary outages of the probes.

while streaming and file hosting applications were gaining popularity. More precisely, Table 2 reports the average per month variation of the traffic share. Specifically, we compute the linear increase/decrease of traffic during each month, and then we average the results over the first and second six months of 2010 separately. We do not report results for the campus network since P2P traffic is blocked, and this biases results. Interestingly, in the first semester data confirm the trends reported in [12] where a decrease of P2P traffic was noticed, while both file hosting and video streaming usage were increasing. During the second semester, we observe an unexpected change in this trend: from June 2010, P2P traffic starts increasing while file hosting traffic is either stable or decreasing. Further investigations revealed that this might be related to changes in the policies of RapidShare [14], the most popular file hosting service in these countries. In particular, RapidShare started enforcing more limitations to non-paying customers to incentive them to subscribe to their service. This caused users to switch back to P2P content download. The Italian probes are not affected by this change since RapidShare is not popular in Italy. Overall, these indices are a clear indication of very fertile and dynamic scenarios that call for continuous and persistent traffic monitoring and classification.

• While the above mentioned changes in traffic shares are typically slow, sudden changes are possible due to changes in the application. For example, as already mentioned, the popular μ torrent, the official BitTorrent client application, was updated during February 2010 to use by default uTP instead of TCP. Correspondingly, there is an increase of UDP traffic clearly visible in some probes.

• Residential probes have stable shares over time, even if trends are present. Instead, in the campus network the traffic share changes over time, so a weekly pattern is clearly visible (also see the figure and related comments in the next section). Indeed, during the weekend, few users are present on the campus, and little traffic flows on the link.

• Obfuscated traffic for P2P applications is not very common in the monitored probes. We double checked that this is not a Tstat classification problem (i.e., Tstat not correctly identifying obfuscated P2P traffic). Consider a host a with IP address IP_a running BitTorrent on $Port_a$. Since the application uses the same $Port_a$ to receive both plain and obfuscated connections, we count the number of connections going to $(IP_a, Port_a)$ and see how many of those are not labeled as BitTorrent (i.e., neither plain nor obfuscated BitTorrent). Those are possibly BitTorrent connections not correctly classified by Tstat. We consider all hosts that are running BitTorrent in a 2 h trace on January 21, 2010 in Poland. Results show that less than 0.5 percent of flows/bytes are not classified as BitTorrent (those are called “false negatives” using classification terminology). Similarly, we computed the percentage of flows that are classified as BitTorrent obfuscated but going to some host which is not likely to run BitTorrent on that port. They account for less than 0.01 percent of flows/bytes (those are called “false positives” using classification terminology). Similar results are obtained considering eMule. This shows that Tstat’s obfuscated classification engine is very reliable.

In conclusion, the presented results highlight the importance of constantly monitoring the network with a flexible tool that has to be constantly upgraded and enhanced to follow its changes.

	P2P		HTTP stream		File hosting	
Location	Jan/June	July/Dec	Jan/June	July/Dec	Jan/June	July/Dec
Polish ISP	-0.81	0.78	0.47	0.57	0.73	-1.06
Hungarian ISP	-0.62	1.98	0.72	0.04	0.45	-0.53
Italian ISP ADSL	-1.1	0.28	0.78	0.06	0.4	0.01
Italian ISP FTTP	-1.3	0.59	0.67	0.03	0.08	0.01

Table 2. Average trends of the two semesters in 2010.

Scalability Issue of Software Based Monitoring Tools

When implementing a live monitoring tool, the knowledge of the maximum sustainable load the probe can handle is one of the most critical issues that must be faced. Indeed, as seen in the previous section, Internet traffic changes widely over both time and space. On a finer timescale, traffic is known to exhibit even larger variability considering both the packet and flow levels. For example, packet-level burstiness can stress the sniffing hardware so that packet bursts can arrive at very high speed. Packet capturing, filtering, and timestamping are then critical, especially if implemented in software. Similarly, bursts of new flows can stress the per-flow operations, so memory management becomes typically a bottleneck.

While Tstat is as an example of advanced traffic monitoring tool, most of the operations it handles are common to any flow-level sniffer and monitoring tool. Indeed, similar data structures must be used to store basic per-flow information such as flow identifier, packets and bytes counters, timestamp, and classification status. Notice that flow structures must be accessed and updated for each packet: hence, efficient data structures like hash tables must be considered, where collisions are minimized and eventually handled using chaining. Further optimizations of memory management are also needed; freed structures should be handled manually as reuse lists by a garbage collector so as to avoid generic and expensive garbage collection routines to kick in and slow down the main analysis operations.

In [13] we extensively analyzed the computational complexity of the Tstat analysis workflow, showing that even with off-the-shelf hardware it is possible to run advanced analysis techniques on several gigabits per second worth of traffic in real time.

To provide some examples of the typical workload Tstat has to support and highlight some critical points in the design of a flow sniffer, Fig. 3 shows the evolution over one week of the total link bit rate (gray line), number of tracked flows (black line) and maximum CPU utilization (dotted line); that is, the total time spent by the CPU in running Tstat, including both kernel and user space CPU time. Measurements refer to a time window of 5 min. Results for the Italian ISP FTTH and Italian campus probes are reported in the top and bottom plots, respectively; results from other probes are not reported for the sake of brevity.

Considering the total link bit rate, the two probes handle approximately the same amount of traffic, which climbs to nearly 500 Mb/s at the peaks. Notice that the peak hour occurs at different times, reflecting the different user habits of home and campus users. The number of active flows is also very different, with the campus probe having to handle a per-flow load about two times higher. This is due to the different traffic mix generated by campus users, as previously shown in

Fig. 2. Therefore, hash table sizes must be correctly tuned to support the various values of the load.

Consider now CPU load curves. We observe very different behavior: the Italian ISP probe shows very low CPU utilization, which is not correlated with the traffic load pattern. On the contrary, the campus maximum CPU utilization is always above 30 percent, and tops 100 percent during sustained traffic load. Investigating further, we pinpointed this to be due to the packet capturing input module, which is based on a common Gigabit Ethernet linecard in the campus probe, while the Italian ISP probe relies on a dedicated Endace linecard. Based on our experience, the major bottleneck is due to the linecard-to-memory communications, which can overload CPU by generating a large number of interrupt requests (IRQs) per second (i.e., one for each received packet). Dedicated traffic capturing devices solve this problem by imple-




Figure 3. Total link bit rate, number of flows and maximum CPU utilization during a typical week: a) Italian ISP FTTH probe; b) Italian campus probe.

menting timestamping functionalities and direct memory access (DMA)-based transfers of packet batches. The CPU utilization figures of the other probes, not shown in the article due to lack of space, confirm this. All ISP probes are indeed equipped with dedicated hardware capturing linecards, so the maximum CPU utilization remains very limited even if they have to handle a large volume of traffic, topping at about 1.5 Gb/s.

To sum up, with common hardware it is possible to monitor several gigabits per second of traffic volume in real time, provided the packet capturing is performed with efficient hardware that offloads from the CPU the per-packet memory copy and timestamping operations. Similarly, efficient memory management algorithms must be adopted to perform both the per flow operations, which optimize the flow lookup performed for every packet, and the garbage collection mechanisms required to avoid memory starvation.

Conclusions

In this article we describe our experience in engineering and using Tstat, a software-based Internet traffic monitoring tool we have been developing for the past 10 years. Presenting measurements collected from several ISP networks, we have shown that Internet traffic widely changes over both time and space: application shares are different at different networks even if common trends are visible due to slow changes in applications popularity; however, sudden changes are observed after the deployment of disruptive technologies made by applications themselves. This calls for the development of automatic mechanisms that continuously update the classification capabilities of a tool, a challenging goal the research community is currently facing.

References

- [1] L. Cottrell, "Network Monitoring Tools Collection," <http://www.slac.stanford.edu/xorg/nmtf/nmtf-tools.html>.
- [2] Tstat Homepage, <http://tstat.tlc.polito.it>.
- [3] A. Finamore *et al.*, "KISS: Stochastic Packet Inspection Classifier for UDP Traffic" *IEEE/ACM Trans. Net.*, vol. 18, no. 5, pp. 1505-1515, Oct. 2010.
- [4] M. Mellia, R. Lo Cigno, and F. Neri, "Measuring IP and TCP Behavior on Edge Nodes with Tstat," *Computer Networks*, vol. 47, no. 1, Jan. 2005, pp. 1-21.
- [5] TCPTrace Homepage, <http://www.tcptrace.org>.
- [6] T. Nguyen and G. Armitage, "A Survey of Techniques for Internet Traffic Classification Using Machine Learning," *IEEE Commun. Surveys & Tutorials*, vol. 10, no. 4, 2008, pp. 56-76.
- [7] A. Finamore *et al.*, "Live Traffic Monitoring with Tstat: Capabilities and Experiences," *8th Int'l. Conf. Wired/Wireless Internet Commun.*, Lulea, Sweden, 1-3 June, 2010.
- [8] H. Kim *et al.*, "Internet Traffic Classification Demystified: Myths, Caveats, and the Best Practices," *ACM CoNEXT 2008*, Madrid, Spain, Dec. 2008, p. 12.
- [9] M. Pietrzyk *et al.*, "Challenging Statistical Classification for Operational Usage: the ADSL Case," *ACM Internet Measurement Conf.*, Chicago, IL, Nov. 2009.
- [10] RRDtool Homepage <http://oss.oetiker.ch/rrdtool/>.
- [11] S. Shalunov, G. Hazel, and J. Iyengar, "Low Extra Delay Background Transport (LEDBAT)," IETF draft, Oct. 2010.
- [12] C. Labovitz *et al.*, "Internet Inter-Domain Traffic," *ACM SIGCOMM*, New Delhi, India, Aug. 2010.
- [13] D. Rossi and M. Mellia, "Real-Time TCP/IP Analysis with Common Hardware," *IEEE Int'l. Conf. Commun. (ICC'06)*, Istanbul, Turkey, June 2006.
- [14] "Reward programme: RapidPoints and RapidDonations to be Discontinued," <http://www.rapidshare.com/#!rapidshare-ag/rapidshare-ag-news>, June 2010.

Biographies

ALESSANDRO FINAMORE [S'09] (finamore@tlc.polito.it) received an M.Sc. in computer engineering in 2008 from Politecnico di Torino. Between 2008 and 2009 he was a research assistant at Politecnico di Torino, and in 2009 he joined the Telecommunication Network Group at Politecnico di Torino as a Ph.D. student. His research interests are traffic classification and computer programming.

MARCO MELLIA [SM'08] (mellia@tlc.polito.it) received his Ph.D. degree in telecommunications engineering in 2001 from Politecnico di Torino. In 1999 he was with the CS Department at Carnegie Mellon University, Pittsburgh, Pennsylvania, and since April 2001 he has been with the Electrical Engineering Department of Politecnico di Torino. He has co-authored over 140 papers published in international journals and conferences, and has participated in the program committees of several conferences including IEEE INFOCOM and ACM SIGCOMM. His research interests are in the fields of traffic measurement, P2P applications, and energy-aware network design.

MICHELA MEO [M'02] (meo@tlc.polito.it) received her Laurea degree in electronic engineering in 1993 and her Ph.D. degree in electronic and telecommunication engineering in 1997, both from Politecnico di Torino, Italy. Since November 1999 she has been an assistant professor at Politecnico di Torino. She has co-authored more than 100 papers, about 40 of them in international journals. She has guest edited six special issues of international journals, including *ACM MONET*, *Performance Evaluation Journal*, and *Computer Networks*. Her research interests are in the field of performance evaluation and modeling, traffic classification and characterization, and peer-to-peer and green networking. She was program co-chair of two editions of ACM MSWiM, general chair of another edition of ACM MSWiM, program co-chair of IEEE QoS-IP, IEEE MoVeNet 2007, and IEEE ISCC 2009, and was on the program committees of about 50 international conferences, including Sigmetrics, INFOCOM, ICC, and GLOBECOM.

MAURIZIO M. MUNAFÒ [M'98] (munafò@tlc.polito.it) received his Dr. Ing. degree in electronic engineering and Ph.D. degree in telecommunications engineering from Politecnico di Torino, Italy, in 1991 and 1994, respectively. He is an assistant professor in the Electronics Department of Politecnico di Torino. He has co-authored about 60 journal and conference papers in the area of communication networks and systems. His current research interests are in simulation and performance analysis of communication systems, and traffic modeling, measurement, and classification.

DARIO ROSSI [M'02] (dario.rossi@enst.fr) is an associate professor with the Computer Science and Networking Department of Telecom ParisTech, Paris, France. He received his M.Sc. and Ph.D. degrees from Politecnico di Torino in 2001 and 2005, respectively, and during 2003-2004 held a visiting researcher position in the Computer Science division of the University of California, Berkeley. At Telecom ParisTech, he is responsible for several European research projects, such as FP7 NAPA-WINE, Celtic TIGER, TIGER2, and TRANS, and ANR Connect. He has co-authored over 70 papers in leading conferences and journals, holds four patents, and has participated in the program committees of several conferences including IEEE INFOCOM, ICC, IPCCC, and GLOBECOM. His research interests include green networking, P2P networks, Internet traffic measurement, and traffic engineering.