12th International Conference on Transparent Optical Networks, ICTON 2010.
IEEE, 2010. 1-4

DOI: http://dx.doi.org/10.1109/ICTON.2010.5549267

Application of Internet Traffic Characterization to All-Optical Networks

Pedro M. Santiago del Río, Javier Ramos, Alfredo Salvador, Jorge E. López de Vergara, Javier Aracil, Senior Member IEEE*
Antonio Cuadra, Mar Cutanda**

**High Performance Computing and Networking, Universidad Autónoma de Madrid,*
Calle Tomás y Valiente, 28049, Madrid, Spain Tel: (+34)914972268,
e-mail: {Javier.aracil, pedro.santiago,Javier.ramos,alfredo.salvador,Jorge.lopez_vergara }@uam.es
***Telefónica I+D*
Emilio Vargas, 6
28043 Madrid

ABSTRACT

In this paper we present an analysis of several large Internet traffic traces and focus on large transfers that are suitable to send through all-optical links featuring burst switching technologies. We analyze the payload size for such optical bursts and the possible impact in network performance.

Keywords: Internet traffic analysis, optical networks, large transfers.

1. INTRODUCTION

In the recent past, all-optical technology has emerged as a cost-effective means to provide connectivity for the Internet. However, the optical network granularity may be too coarse to support the fine-grain Internet services as we know them nowadays. As it turns out, the setup time of a lightpath (or even optical burst) is relatively large in comparison to the typical transfer time for IP packets.

On the other hand, the new cloud computing paradigm is advocating for computing and storage resources to be deployed within the communications networks. The availability of server farms, possibly located close to the user, paves the way for the provision of Internet services with a better quality of experience for the user, and constitute a better match to the particular transfer capabilities of the optical network.

Figure 1 shows the reference network for the provision of Internet services in the optical cloud. On the one hand, servers are located close to the network edges, which makes it possible to reduce latency. On the other hand, the interface to the optical network provides a *higher aggregation level* in terms of users, which enables to launch end-to-end connections server-to-server at the *flow level*. For example, a YouTube video could be transferred within the YouTube content distribution network to the server which is located closer to the user. Then, the video could be streamed to the user with the minimum possible latency, resulting in a very good quality of experience. Furthermore, the server-to-server transfer would be optimized to the particular optical network features (one file per burst, for example). Consequently, increased quality of experience and optical network utilization could be achieved.

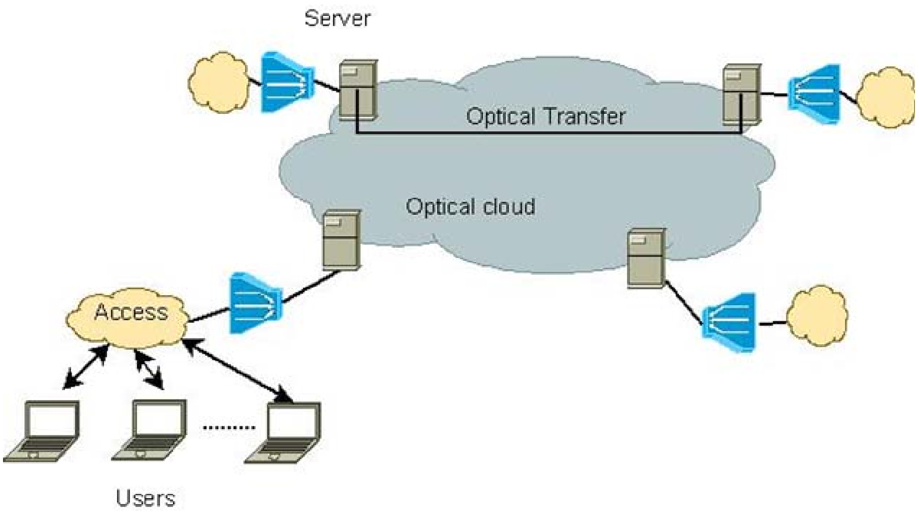

Figure 1. Network scenario

Actually, this is the network architecture provided by the FP7 STREP MAINS. In this paper we focus on the *optical transfer* segment in the figure and provide an ad-hoc statistical analysis of a large Internet traffic trace. The aim is to understand to which extent can we use the optical network with large transfers as opposed to small flows. We analyze the flow size and duration and extract the subset of flows which are amenable for optical transfer (lightpath or burst). Then, we calculate the transfer time in the new scenario, assuming that the optical transfer is penalized by a setup time (or burstification time in the optical burst switching case). Finally, we draw the conclusions and recommendations for the transport of Internet services in the optical cloud.

2. TRACE ANALYSIS

For our analysis three different packet-level traces have been used. The first, named as Trace 1 in the following sections, belongs to a Spanish provider countryside access network. This kind of networks uses GPRS technology to provide broadband access to low coverage areas. The trace contains 21 hour traffic both from residential households and small businesses. Near 70M packets were captured containing full payload. The second trace used, named as Trace 2 in the following sections, belongs to an OC192 (9953 Mbps) backbone link located between Chicago and Seattle [1]. The trace contains traffic between 5 and 6 AM (UTC) of the 15 of January of 2009. All the packets in the trace, near 1530M, are anonymized and captured without payload. The last trace used, named as Trace 3 in the following sections, belongs to DITL (A Day in the Internet Life) 2008 experiment[6]. This trace was captured on 19 March of 2008 and contains traffic between 4 and 6 PM on an academic network at USC. The trace contains 501M packets without payload. Table 1 summarizes the main features of the used traces.

Table 1. Main features of used Traces.

	Duration (hh:mm:ss)	Number of TCP flows	Number of UDP flows	Mean data per flow (bytes)	Mean duration per flow (s)
Trace 1	*21:14:18*	5515257	8663364	1966	10.63
Trace 2	*01:00:00*	30736929	33522668	10264	8.27
Trace 3	*02:00:00*	8680828	6384359	26045	11.41

Figure 1 shows the ratio of flows smaller than X bytes. This ratio is expressed in total volume data (left figure) and in number of flows (right figure). It can be observed that more than 90% of TCP/UDP flows have less than 10Kbytes of data in the three analyzed traces. However, in traces 2 and 3, more than 80% of volume traffic belongs to flows with more than 1MB, and, in trace 1, more than 50% of volume traffic belongs to flows with more than 10MB. Thus, the most resources are consumed by a small amount of flows. Transferring these large flows to the optical layer, reduces load in IP network and provide Quality of Service (QoS) of these flows [2].

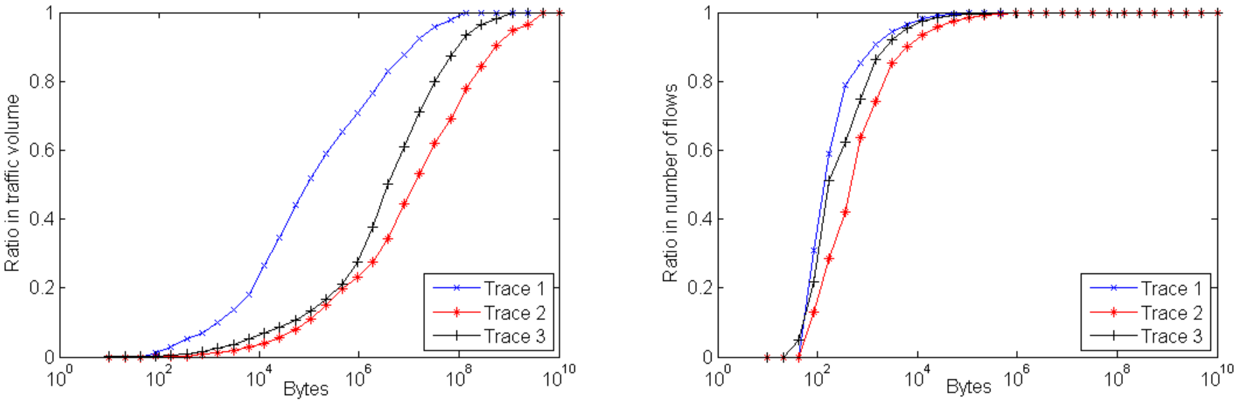

Figure 2. Ratio of flows smaller than X bytes.

Table 2 and 3 show the main features of large flows (which are greater than 10 and 100 Mbytes, respectively). It is worth noting that the mean duration and mean volume are much greater than when all flows are considered, as well as the ratio of TCP flows with respect to UDP flows is also much greater than in general case.

Table 2. Main features of flows greater than 10 Mbytes.

	Number of TCP flows	Number of UDP flows	Mean data per flow (MBytes)	Mean duration per flow (s)	Ratio in traffic volume	Ratio in number of flows
Trace 1	118	1	23	7057	10%	0.0008%
Trace 2	7174	586	30	1289	36%	0.01%
Trace 3	4370	199	45	1041	53%	0.03%

Table 3. Main features of flows greater than 100 Mbytes.

	Number of TCP flows	Number of UDP flows	Mean data per flow (MBytes)	Mean duration per flow (s)	Ratio in traffic volume	Ratio in number of flows
Trace 1	4	0	104	1726	1%	0.00005%
Trace 2	262	37	203	2559	9%	0.0005%
Trace 3	364	12	276	1635	26%	0.003%

Finally in this section, we have analyzed the types of traffic in trace 1 using L7-filter-based detection. Note that it is not possible to use these classifiers in traces 2 and 3 because there is no payload in them. Fig. 3 shows the traffic volume separated in different services considering only flows larger than 10 Mbytes (left) and considering all flows (right). It can be observed that more than 55% of large flows (all sessions except P2P) could be sent by the optical layer. Because of these contents are from well known service providers (such as rapidshare.com, youtube.com, gmail.com, etc), specific agreements could be made to establish direct lightpaths.

We have removed from the pie chart sessions corresponding to irrelevant http services (e.g. local online newspapers, message boards, etc) for this study. This traffic involves 19% of traffic volume when only large sessions are considered and 42% when all sessions are considered.

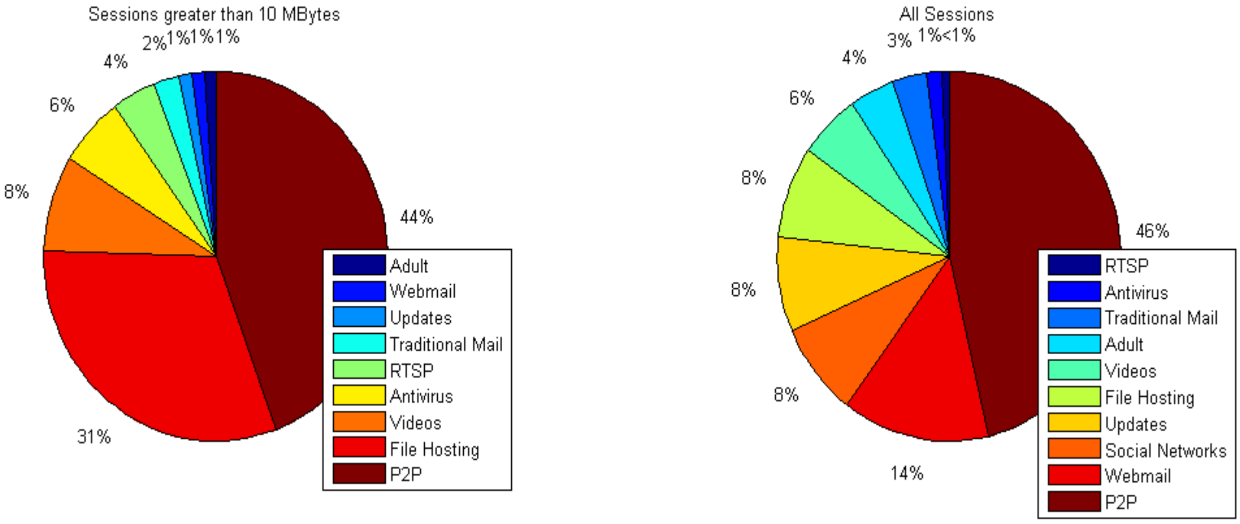

Figure 3. Traffic volume separated in different services.

3. PERFORMANCE EVALUATION

Given the low bit error rate in optical links, simpler transport protocols for large optical bursts like stop and wait as proposed in [3] maximize the throughput compared with complex ones such as Transmission Control Protocol by means of reducing the overhead of acknowledgements and window management. Also TCP has shown retransmission rates over 10% due to network congestion and faulty window management [4], which degrades the data transfer performance.

In our experiments we compare the performance of both protocols assuming TCP maximum segment size may be arbitrarily large. This is clearly a best case for TCP and it allows us to make a comparison in terms of burst blocking probability. The model for OBS throughput is the same as in [3] and for TCP we apply the model from

[5], both are summarized in Formula 1 and Formula 2. Note that L_{Burst} is the length of the optical burst, which is made equal to cw_{max}, the maximum window size in TCP.

OBS: Optical Burst Switching

$$Thr_{OPT} = \frac{L_{Burst}}{RTT}\left(1 - P_{error}\right)$$

Formula 1

TCP: Transmission Control Protocol

$$Thr_{TCP} = \min\left\{\frac{cw_{\max}}{RTT}, \frac{MSS}{RTT\sqrt{\frac{2}{3}P_{error}}}\right\}$$

Formula 2

The stop and wait and TCP equations have been plotted for different error rates burst error rates and the results are shown in Figure 3. It can be observed that OBS file transfer provides better throughput than TCP using burst length greater or equal than 10 Mbytes and error rates above 1e-7.

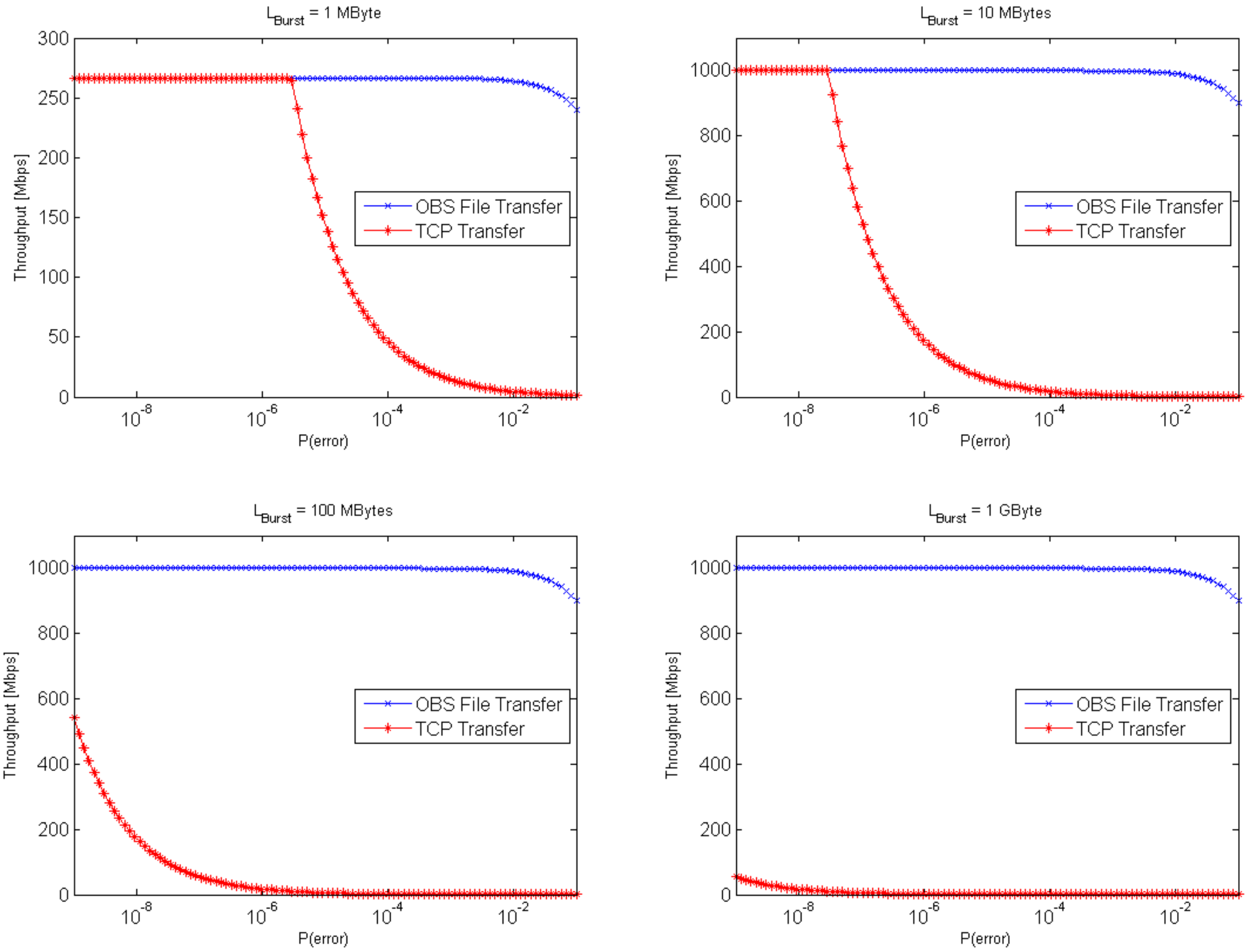

Figure 4.Throughput comparison between All-optical transfer and TCP transfer.

According to the previous sections near the 55% of the traffic volume of sessions greater than 10 MB could be sent using lighpaths as they are big enough to fit into one or more optical bursts using $L_{Burst} = 10$ MB. On the previous analysis, we observed in trace 1 that sessions bigger than 10 MB represent the 10% of the traffic volume. On the other traces the sessions bigger than 10 MB represent near the 36% and 56% of the total traffic volume. These “high load” sessions could be sent using dedicated lightpaths providing high throughput to them while assuring quality of service to the traffic in the electrical layer due to the reduction of the router queues load. Also the quality of service of the “high load” sessions could be improved establishing dedicated lightpaths to well known destinations. For example the 32% of traffic volume was directed to file hosting services as Rapidshare, Megaupload or DespositFiles. If the operator makes a commercial arrangement with these service providers this type of traffic could be sent using files over lightpaths reducing the latency and increasing throughput. Other traffic like P2P is excluded from this technique since the destinations are not fixed and change quickly.

4. CONCLUSIONS

In this paper we perform a preliminary traffic analysis to find which share of the traffic can be swiftly transferred using the optical layer solely. The results of this paper show that 55% of the total traffic volumen could benefit from this optical layer transfer. On the other hand such offloading to the optical layer is also benefitial for the electronic counterpart, because of the utilization decrease at the routers.

ACKNOWLEDGEMENTS

We would like to acknowledge the USC/LANDER project for providing the traces of the Day In the Life of The Internet (DITL) March 18-19, 2008 dataset.

REFERENCES

[1] The Cooperative Association for Internet Data Analysis (CAIDA). http://www.caida.org

[2] T. Fioreze, M. Oude Wolbers, R. Meent van de, A. Pras, Finding Elephant Flows for Optical Networks. *In: 10th IFIP/IEEE International Symposium on Integrated Network Management*, IM, 21-25 May 2007, Munich, Germany

[3] M. Izal, J. Aracil: IP-over-WDM dynamic link layer: challenges, open issues, and comparison of files-over-lightpaths versus photonic packet switching, *in Proc. SPIE Opti-Comm 2001*, Denver, CO, Aug. 2001

[4] F. Qian, A. Gerber, Z.M. Mao, S. Sen, O. Spatscheck, W. Willinger. TCP revisited: a fresh look at TCP in the wild. In *Proceeding of the 9th ACM SIGCOMM Conference on internet Measurement Conference.* Chicago, IL, 04-06 Nov 2009.

[5] M. Mathis, J. Semke, J. Mahdavi, T. Ott, The macroscopic behavior of the TCP congestion avoidance algorithm.*SIGCOMM Comput. Commun. Rev*. 27, 3 (Jul. 1997), 67-82.

[6] Day In the Life of The Internet (DITL) March 18-19, 2008 dataset, PREDICT ID
USC-LANDER/DITL_I2-20080317. Traces generated 2008-03-17.
Provided by the USC/LANDER project.
<http://www.isi.edu/ant/lander>.