

# Offloading IP flows onto lambda-connections

Tiago Fioreze, Mattijs Oude Wolbers, Remco van de Meent, and Aiko Pras

University of Twente, Enschede, The Netherlands  
{t.fioreze, r.vandemeent, a.pras}@utwente.nl

**Abstract.** Optical networks are capable of switching IP traffic via lambda connections. In this way, big IP flows that overload the regular IP routing level may be moved to the optical level, where they get better Quality of Service (QoS). At the same time, the IP routing level is off-loaded and can serve smaller flows better. Within this context, this paper analyses the eligibility of IP flows to be moved to the optical level. In this analysis, we observe the percentage of IP traffic as well as the amount of IP flows moved to the optical level while using various definitions for an IP flow. The main contribution of this paper is to show how the amount of IP traffic transferred over lambda-connections considerably depends on the definition used for a flow.

## 1 Introduction

The use of lambda-switching in optical networks allows huge amounts of data to be transmitted over lambda-connections via multi-service optical switches. These optical switches are capable of making data forwarding decisions at different levels in the protocol stack [1]. Such capability enables therefore data packets to be fully switched at optical-level (lambda switching) instead of being forwarded (routed) at packet-level (packet forwarding).

Today, IP traffic from several specialized applications already profit from lambda-switched networks capabilities. Examples are grid applications [2] and High-Definition TeleVision (HDTV) broadcasting [3]. The question that may arise is whether there is also ordinary IP traffic, currently running over the existing IP network, that would be able to take advantage of these new lambda-switching capabilities. This paper investigates this question by focusing on the following research questions:

1. In current networks, how many IP flows would already be large enough to profit from being moved to the optical-level?
2. What fraction of the total IP data traffic do these flows represent? In other words, by what percentage would it be possible to offload the current IP network?

In order to answer these questions the following approach will be used: 1) Collect measurements from SURFnet6 [4], the Dutch research network; 2) Define a criterion that existing flows must satisfy to make them eligible to be moved to the optical-level; and 3) Analyze the collected traces in order to answer the research questions.

### 1.1 Related work

The analysis of big IP flows in optical networks has already been addressed in other works [5] [6], but no substantial variation in the definition for a flow is found, as concluded in a previous work of ours [7]. In [7], we showed that, in practice, IP flows

are generally characterized by using the 5-tuple flow definition (source/destination address/port and protocol). Within this context, we performed our analysis by considering various definitions for a flow while observing the volume of IP traffic moved to the optical level. In contrast to this paper, in our previous work we collected 1 day of network traces divided in 30 minutes intervals and we only analyzed the flows with respect to their volume. In this paper, we collected 2 weeks of traces divided in different time intervals (5 and 30 minutes) as well as we observe the percentage of IP traffic and the amount of IP flows moved to the optical level.

## 2 Measurement and analysis setup

For our analysis SURFnet routed a NetFlow stream from the NetFlow-enabled core routers to a NetFlow collector located at our domain, where the incoming NetFlow stream was dumped into *pcap* files, using the *tcpdump* tool. This allowed us to adjust our analysis without the need for SURFnet to retransmit the data. The total amount of collected data was about 81 GB. The total amount of NetFlow reported bytes was 4.0 TB. However, since SURFnet uses 1:100 sampling, this accounts for some 0.40 PB of actual network traffic.

Our analysis consisted of varying the definition of a flow by using different levels of granularity. The higher the level of granularity is, the more restrictive the flow definition will be when grouping IP packets. Our flow definitions take into account different end-points and have the following descending order of granularity: **App2App**: the 5-tuple flow definition; **HstHst**: set of packets with the same source and destination IP addresses; **Sub2Sub /24**: set of packets matching the 24 most significant bits of the source and destination addresses; **Sub2Sub /16**: set of packets matching the 16 most significant bits of the source and destination addresses; **Sub2Sub (NetFlow)**: set of packets matching the most significant bits of the source and destination addresses reported by NetFlow; **AS2AS**: set of packets with the same source and destination autonomous systems; and **Sub2Sub /8**: set of packets matching the 8 most significant bits of the source and destination addresses.

In order to check whether a certain IP flow is eligible to be moved to the optical level, we used the following criterion : *an IP flow is eligible if its total consumed bandwidth is equal or bigger than the minimal unit of transmission in SONET networks in a certain time interval: average throughput  $\geq 50.112$  Mbit/s*. Different time intervals (5 and 30 minutes) were used to check if they have an influence on the percentage of IP traffic as well as the amount of IP flows moved to the optical level. With such time intervals the threshold values for a flow to be considered eligible are 1.8 and 11 GBytes for 5 and 30 minutes intervals, respectively.

## 3 Results

This section shows the amount of flows that satisfied our criterion and the percentage of IP traffic eligible to be moved to the optical level. Figure 1 shows the amount of flows eligible for a lambda-connection while using various definitions for a flow as well as different time intervals (5 and 30 minutes interval). By using a small time interval (i.e., 5 minutes) the number of flows eligible to move to the optical level is much higher than using a large time interval (i.e., 30 minutes). The reason is that in case a small time

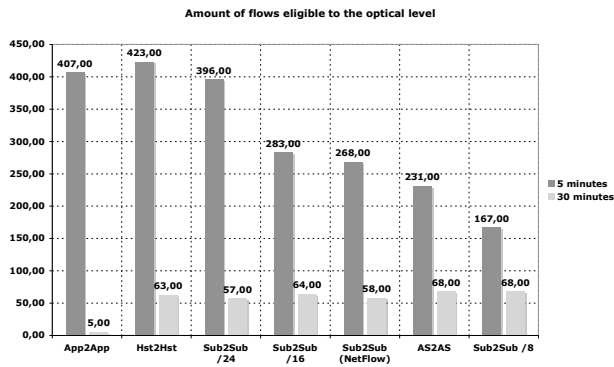


Fig. 1. The total amount of flows with 5 and 30 minutes time interval

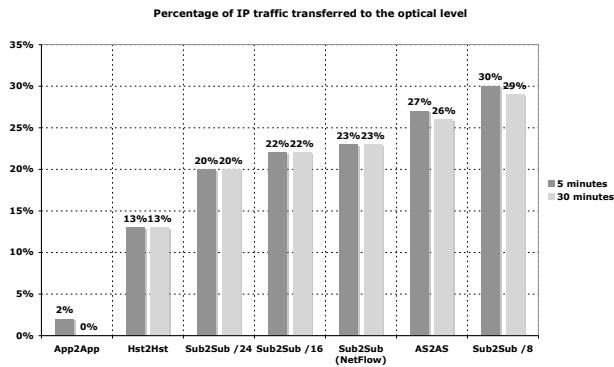


Fig. 2. Percentage of IP traffic with 5 and 30 minutes time interval

interval is used, the threshold value for a lambda-connection is also small (1.8 GBytes). As a consequence more flows satisfied our criterion. On the other hand, by using a big time interval, only a few very big flows of over 11GB in size, which last for a long period, are selected. It can also be seen in the graph — as could be expected — that lower granular flow definitions group more packets into fewer number of flows than higher granular flows.

A high number of flows does not directly imply that a large fraction of the IP traffic is eligible to be moved to the optical level. This can be seen in Figure 2, in which the percentage of IP traffic transferred to the optical level is shown, based on the total amount of collected data during 2 weeks (0.4 PB). It can also be seen that using different time intervals do not considerably change the percentage, but on the other hand the flow definition impacts the fraction considerably: the lower granular a flow definition is, the more packets will be grouped into flows, and, thus, the more traffic is selected to be moved over lambda-connections.

## 4 Conclusion

This paper presented various definitions for an IP flow and used them in conjunction with different time intervals to analyze their influence on moving flows to the optical level. The influence was observed in terms of percentage of IP traffic and amount of IP flows. Based on the analysis performed and in the obtained results, we can now answer our research questions.

1. *In current networks, how many IP flows would already be large enough to profit from being moved to the optical-level?*

Based on our criterion, the number of flows that could be moved depends on the flow definition used as well as the time interval. During our 2 weeks of data analysis, between 390 and 450 flows would already profit from being moved to the optical-level if 5 minutes interval would be used in conjunction with high granular flow definitions. In case of the 30 minutes interval, between 55 and 70 flows.

2. *What fraction of the total IP data traffic do these flows represent? In other words, by what percentage would it be possible to offload the current IP network?*

The fraction of IP traffic moved to the optical level strongly depends on the flow definition used. Higher granular flow definitions such as *App2App* and *Hst2Hst* allow between 0% and 13% of the total traffic. On the other hand, lower granular flow definitions allow between 20% and 30%.

The main contribution of this work is to show that the probability of finding IP flows eligible for lambda-connections increases when using less restrictive flow definitions. As future work, we aim at analyzing the characteristics of the eligible IP flows to the optical level by observing their size, duration, throughput, and recurrence. We would like to thank SURFnet for allowing us to perform measurements on their network. This paper was supported in part by the EC IST-EMANICS Network of Excellence (#26854).

## References

1. Leon-Garcia, A., Widjaja, I.: Communication Networks: fundamental concepts and key architectures (2nd edition). McGraw-Hill Companies, (2003)
2. Boutaba, R., Golab, W., Iraqi, Y., Li, T., Arnaud, B.St.: Grid-Controlled Lightpaths for High Performance Grid Applications. In: Journal of Grid Computing. Vol.1, No. 4, pp. 387–394, (2003)
3. Rouskas, G.N.: Optical layer multicast: rationale, building blocks, and challenges. IEEE Network Magazine. Vol.17, No. 1, pp. 60–65, (2003)
4. SURFnet: SURFnet6 lighpaths mark start of the new Internet area (press release). Available in: [http://www.surfnet.nl/info/en/artikel\\_content.jsp?objectnumber=107197](http://www.surfnet.nl/info/en/artikel_content.jsp?objectnumber=107197)
5. Mori, T., Kawahara, R., Naito, S., Goto, S.: On the characteristics of Internet traffic variability: spikes and elephants. In: International Symposium on Applications and the Internet. pp. 99–106, (2004)
6. Wallerich, J., Dreger, H., Feldmann, A., Krishnamurthy, B., Willinger, W.: A methodology for studying persistency aspects of internet flows. In: ACM SIGCOMM Computer Communication Review. Vol.35, No.2, pp. 23–36, (2005)
7. Fioreze, T., Wolbers, M.O., Meent, R., Pras, A.: Finding elephant flows for optical networks. Application session proceedings of the 10th IFIP/IEEE International Symposium on Integrated Network Management (IM 2007). pp. 627 – 640, (2007)