

Characterization of IP Flows Eligible for Lambda-Connections in Optical Networks

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

University of Twente

Faculty of Electrical Engineering, Mathematics and Computer Science (EEMCS)

Department of Design and Analysis of Communication Systems (DACs)

Enschede, The Netherlands

Email {t.fioreze, m.oudewolbers, r.vandemeent, a.pras}@utwente.nl

Abstract—The advance on data transmission in optical networks has allowed data forwarding decisions to be taken at multiple levels in the protocol stack (e.g., at network and optical levels). With such capability, big IP flows can be moved from the network level and switched completely at the optical level over lambda-connections, where they get better Quality of Service (QoS). Meanwhile, the regular IP routing level is offloaded and can serve smaller flows better. With the continuous growing of traffic on the Internet, the selection of big IP flows can become difficult to be done by using current management approaches (conventional management and Generalized Multiprotocol Label Switching (GMPLS) signaling). The University of Twente (UT) is researching the use of self-management as an alternative to overcome this issue. In order to properly identify IP flows eligible to be moved to the optical level, the characteristics of these flows must be known, though. In this context, this paper analyses some of the characteristics of IP flows eligible to the optical level by observing their size, duration, throughput, and recurrence. In this analysis, we observe those characteristics while using various definitions for an IP flow as well as using different time intervals. The main contribution of this paper is to show the behavior of IP flows eligible for lambda-connections. Not in the least, we also show how this knowledge can be used in our self-management of optical networks approach.

I. INTRODUCTION

Advanced optical networks are able to make data forwarding decisions at multiple levels (e.g., at network and optical levels) via multi-service optical switches [1]. Such ability enables, for instance, data packets to be fully transported at optical level (lambda switching), which bypasses the routing infrastructure (network level), resulting in a faster delivery of those packets. Moreover, in conjunction with multiplexing technologies such as Wavelength Division Multiplexing (WDM) [2], lambda-switching enables a lambda-connection to act like a virtual circuit.

Currently, a number of network operators are moving towards lambda-switching enabled networks. Amongst others: GÉANT [3], Internet2 [4], Clara [5], SURFnet [6], Canarie [7], and HEAnet [8]. Collaboration among several of these operators aims at providing optical interconnect capabilities, resulting in an interconnection of their research and education networks.

By promoting this move towards lambda-switching, network operators offer new perspectives for services in optical networks. Nowadays, IP traffic from several specialized applications, which require huge amounts of bandwidth, already profit from lambda-switched networks capabilities. Examples are: Grid applications [9], High-Definition Television (HDTV) [10] broadcasting, and the LOFAR project [11], which aims at building an interferometric array of radio telescopes mostly distributed over The Netherlands and Germany to allow astronomers to observe objects in space.

The knowledge of the heavy-hitter¹ behavior [13] [14] [15] [16] [12] of flows originated from these applications allows network managers to establish lambda-connections in advance for such flows. However, there may be also other big IP flows in current networks that could also benefit from being moved to lambda-connections, but for some reason (e.g., the network manager may not be aware of their existence) they may not be selected. As a consequence, these IP flows may not profit from less delay and jitter as well as plenty of bandwidth found at the optical level.

In order to overcome this issue, the University of Twente is investigating the use of self-management in optical networks [17] [18]. By self-management, we mean the ability of multi-service optical switches to automatically: 1) identify which IP flows should be moved to underlying lambda-connections, and 2) establish and release the required lambda-connections for these flows.

In order to automatically identify which flows should be moved to the optical level, the optical switches must observe IP flows by considering a set of parameters, such as flow volume and duration. In this context, this paper focuses on the analysis of the behavior of IP flows eligible to be offloaded to the optical level by characterizing them regarding their size, duration, throughput, and recurrence. In addition, we observe those characteristics while using various definitions for an IP flow (which is a subject to be discussed in Subsection V-C), as well as using different time intervals.

¹According to Mori *et al.* [12], heavy-hitter flows are flows with a very large number of packets, which is the same definition used in this paper.

Our analysis can be better divided in the following research questions:

- 1) What are the most suitable definitions for flows eligible to be moved to the optical level?
- 2) What are the characteristics of those flows regarding their size, duration, and throughput?
- 3) Do those flows recur? If so, how often?

In order to answer these questions, the approach outlined below will be used.

- 1) Study of the literature in order to find out various approaches to characterize flows;
- 2) Collect measurements from SURFnet6 [19], the Dutch research network. In order to determine possible periodicity of eligible flows, these traces consider two weeks of collecting period;
- 3) Define a criterion that existing flows must satisfy in order to make them eligible to be moved to the optical-level. We will derive this criterion by discussing the characteristics of lambda-switched enabled networks (Subsection III);
- 4) Analysis of the collected traces in order to answer the research questions.

The remainder of this paper is structured as follows. Section II presents the study of the literature with respect to the characteristics of IP flows. Section III shows a general review of the characteristics of lambda-switched enabled networks. Following that, our research work on self-management of optical networks is introduced in Section IV. Subsequently, Section V shows in more details about the collect and analysis of network traces from SURFnet6. Finally, Section VI presents the results of our analysis, which allows us to give the answers for our research questions in the concluding Section VII.

II. RELATED WORK

The characterization of big IP flows in large scale optical networks has also been addressed in other research works [13] [14] [15] [16] [12]. From this related work, we highlight the following two investigations, due to some similarities to this paper.

Papagiannaki *et al.* [14] focus on the analysis of the volume (size) of big IP flows as well as on their historical behavior. The authors argue that the flow volume may be very volatile, which results in different load values captured at different time intervals. Similarly to our investigation presented in this paper, the aforementioned authors performed a significant variation in the time intervals (1, 5 and 30 minutes) and they concluded that the perceived load is the same independent of the time interval used. We also proved that in a previous work [20]. In spite of the significant variation in the time intervals, Papagiannaki *et al.* only used a single definition for a flow based on the network prefixes exported by Border Gateway Protocol (BGP) [21] routers.

In another related work, more similar to ours, Lan *et al.* [16] characterized big flows by observing their size, duration, throughput (rate), and burstiness, as well as by examining

their correlation. Although a considerable amount of flow characteristics were observed and correlated, the authors did not use various definitions for an IP flow, considering only the traditional 5-tuple flow definition ².

In a previous work of ours [22], we concluded that there is no substantial variation in the definition of a flow in most of the investigated related works. As a result of that, we decided to perform our analysis by considering various definitions for a flow [22]. In that analysis we characterized IP flows eligible to the optical level only with respect to their size while observing the influence of using different definitions for a flow. In addition, we collected only 1 day of network traces divided in 30 minutes intervals.

In a more recent work [20], we extended our previous investigation [22] by observing the percentage of IP traffic and the amount of IP flows moved to the optical level. At the same time, we used the same variation in the definition of a flow presented in our previous work [22], but, in contrast, we collected 2 weeks of traces divided in different time intervals (5 and 30 minutes).

In the present paper, we make use of the same collected data divided in the same time intervals (5 and 30 minutes) as well as the same variation in the definition of a flow. However, we focus on the characteristics of flows eligible for lambda-connections by observing their size, duration, throughput and recurrence.

III. OVERVIEW OF LAMBDA-SWITCHED ENABLED NETWORKS

Lambda-switched enabled networks use lambda switching³ technology to transfer huge amounts of data via lambda-connections. One example of a lambda-switched network is SURFnet6 [19], which is a hybrid optical and packet switching research network developed within the GigaPort-RoN project [23].

Although the ability to redirect specific wavelengths is a technological advance, lambda switching similarly works as traditional routing and switching. Optical switches take in a single wavelength of light from a specific fiber optic strand and recombine it into another strand that is set on a different path.

Lambda switching relies in multiplexing techniques such as Wavelength Division Multiplexing (WDM) [2] and Time Division Multiplexing (TDM) [24]. The WDM technique consists of modulating each wavelength onto a different part of the light spectrum of a single fiber. The amount of multiplexed wavelengths onto a single fiber can divide WDM into Coarse WDM (CWDM), for fibers carrying less than 8 wavelengths, and Dense WDM (DWDM), for those fibers carrying from 9 up to 160 wavelengths. On its turn, the TDM technique

²5-tuple flow definition: group packets with same source and destination IP addresses, same transport protocol, and same transport protocol source and destination port numbers.

³Lambda switching gets its name from λ , the 11th letter of the Greek alphabet, which has been adopted in networking to refer to an individual optical wavelength.

consists of dividing a wavelength in time slots in order to send data frames. This approach is a basis for the today's standards used in digital communication, Synchronous Optical Networking (SONET) and Synchronous Digital Hierarchy (SDH).

SONET, which is standardized by the American National Standards Institute (ANSI), is a set of standards for synchronous data transmission over fiber optic networks that are often used for framing and synchronization at the physical layer. SONET is based on transmission at speeds of multiples of 51.840 Mbps. SDH is the international version of the standard published by the International Telecommunications Union (ITU). Transmission rates of up to 10 Gbit/s can be achieved in today's SONET systems and the 40 Gbit/s systems are possible (Table I).

Optical carrier level	Payload (rate) (Mbps)	Overhead (rate) (Mbps)	Total (rate) (Mbps)
OC-1	50.112	1.728	51.840
OC-3	148.608	6.912	155.520
OC-12	601.344	20.736	622.080
OC-24	1.202.208	41.472	1.243.680
OC-48	2.405.376	82.944	2.488.320
OC-192	9.621.504	331.776	9.953.280
OC-768	38.486.016	1.327.104	39.813.120

TABLE I
SDH & SONET CLASSIFICATIONS.

The highest rate that is commonly deployed is the OC-192 or STM-64 circuit, which operates at rate of just under 10 Gbit/s. Speeds beyond 10 Gbit/s are technically viable and are under evaluation. Where fiber exhaust is a concern, multiple SONET signals can be transported over multiple wavelengths over a single fiber by means of Dense Wave Division Multiplexing (DWDM). Such circuits are the basis for all modern transatlantic cable systems and other long-haul circuits.

IV. SELF-MANAGEMENT OF LAMBDA-SWITCHED ENABLED NETWORKS

Two approaches are currently used for the management of lambda-connections in lambda-switched enabled networks [25]: conventional management and GMPLS signaling. The former is characterized by a centralized management entity (e.g., human manager or an automated management process) that is in charge of establishing lambda-connections and deciding which IP flows should be moved to the optical level. In contrast, the latter is characterized by the fact that optical switches coordinate the creation of lambda-connections among themselves after being requested to do that. The decision which IP flows should be moved to the optical level however is taken by a centralized entity or by the entities exchanging data flows.

Both approaches, however, have an important shortcoming: they require human interaction to detect flows and manage lambda-connections. This interaction may be slow and error-prone. Currently, when a lambda-connection is requested

within one single domain (intra-domain), several steps are taken (e.g., phone calls and emails exchanges) between requesters and network domain administrators in order to establish the lambda-connections. Hence, it may take hours before a desired lambda-connection can be used. When requests for a connection span multiple domains (inter-domain), the lambda-connection provisioning may take even much longer. In addition to that, several big IP flows eligible for lambda-connections may somehow be transiting undetected by the manager's eyes, contributing therefore to the congestion of current IP networks.

Our solution to overcome these shortcomings consists of providing self-management capabilities to multi-service optical switches. Our self-management solution allows optical switches to be in charge of automatically selecting IP flows to the optical-level as well as creating/releasing lambda-connections for them. Network managers would only be required to configure the optical switches in order to define parameters for: 1) selecting IP flows to be offloaded onto lambda-connection, and 2) establishing and releasing lambda-connections for those flows. Once configured, the optical switches cooperatively work by themselves. It is worth to mention that the knowledge about the behavior of IP flows eligible for the optical level is important in this configuration process. Figure 1 depicts how our self-management solution looks like.

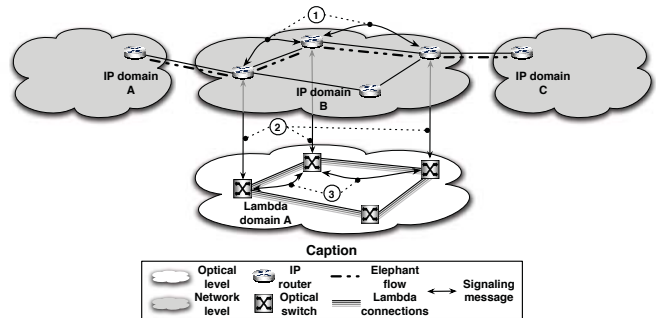


Fig. 1. Self-management of lambda-connection in optical networks.

In Figure 1, IP and optical domains coordinate with one another in order to detect IP flows and manage lambda-connections. Both domains are assumed as already been configured by network managers. IP routers located at IP domain B are exchanging information (e.g., flow throughput) regarding the existence of a big IP flow transiting between IP domains A and C (step 1). Based on this exchanged information and on the configuration performed by network managers, the IP and optical domains make decisions on whether a flow is eligible for a dedicated lambda-connection. If the decision is in favor of creating a lambda-connection (i.e., the IP flow is eligible to be moved to the optical-level), the IP routers signal the optical switches in lambda domain A (step 2). Then, the optical switches coordinate among themselves to create a dedicated lambda-connection to the detected IP flow (step 3). From that point on, the IP flow is completely switched via a lambda-

connection in lambda domain A. Further information about our self-management solution can be found in [17] [18].

V. MEASUREMENT AND ANALYSIS SETUP

This section presents how network data was collected from SURFnet6 as well as how our analysis was performed.

A. Collecting network information

For our analysis we measured the SURFnet6 network for two weeks. The core routers — NetFlow-enabled — of SURFnet6 are located in Amsterdam. These core routers export information about all network traffic within SURFnet6.

The NetFlow-enabled core routers export NetFlow records to various collectors, including ours. Since SURFnet6 network is a high-speed network (10Gbps links and up), technological issues such as limited processing resources prohibit exporting NetFlow records for all packets. To reduce the amount of packets that are represented in NetFlow records, packet sampling is used. This methodology is also known as Sampled NetFlow; it entails selecting 1 packet out of n packets. The sampling currently used in SURFnet6 is 1 per 100 packets, hence providing NetFlow data representing 1% of the total traffic in SURFnet6.

To facilitate our measurements, SURFnet routed a NetFlow stream from two of the core routers to a NetFlow collector located at the University of Twente (UT) domain. This stream was sent via an encrypted tunnel using Zebedee [26]. The use of Zebedee was needed since NetFlow does not provide any kind of encryption. At the UT domain, a machine was setup to dump the incoming NetFlow stream into so-called *pcap* files, using the *tcpdump* tool [27]. We decided to use *tcpdump* instead of real NetFlow collector software: this way, we could store the NetFlow stream in its “raw state”. This allowed us to adjust our analysis without the need for SURFnet to retransmit the data. The collecting procedure is depicted in Figure 2.




Fig. 2. The collecting NetFlow data scenario.

The total amount of data that was received from SURFnet was about 81 GB. This includes all UDP, IP, and pcap overhead (e.g., packet timestamps). The total amount of NetFlow reported bytes was 4.0 TB. However, since SURFnet uses 1:100 sampling with NetFlow, this accounts for some 0.40 PB of the actual network traffic.

B. Filtering and storing network information

Some of the NetFlow fields exported by SURFnet (e.g., IP ToS) were not needed in our analysis. The decision was

therefore to filter those fields also to decrease the size of the analyzed data. Table II shows the NetFlow fields that were considered in our analysis. For a complete description of the fields existing in NetFlow version 9, see [28].

Field name	Description
Start time	The system uptime at which the first packet of a flow was switched
End time	The system uptime at which the last packet of a flow was switched
Bytes	The number of bytes associated with an IP flow
Packets	The number of packets associated with an IP flow
IPv4 source address	The IPv4 source address of a flow
IPv4 destination address	The IPv4 destination address of a flow
L4 source port	TCP/UDP source port number
L4 destination port	TCP/UDP destination port number
Source AS	The source BGP autonomous system number
Destination AS	The destination BGP autonomous system number
Source Mask	The number of contiguous bits in the source address subnet mask
Destination Mask	The number of contiguous bits in the destination address subnet mask

TABLE II
THE CONSIDERED NETFLOW FIELDS.

As a last step before starting our analysis, all the considered data was stored into a MySQL database. The reason to use MySQL was the familiarity of the authors with the tool and also because MySQL has several aggregate functions (e.g., AVG(), MAX(), MIN(), and so on) that allowed us to perform fast queries.

C. Analyzing the collected data

The research analysis consisted of defining IP flows by using different levels of granularity (Figure 3).

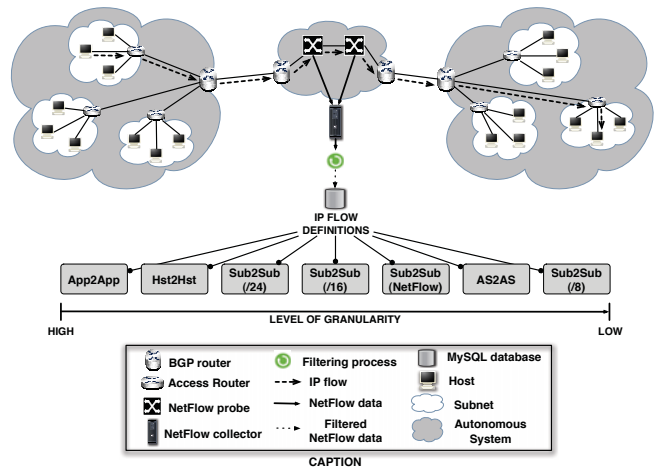


Fig. 3. The different definitions for a flow.

The higher granular a flow definition is, the more details a flow definition has regarding the IP packets grouped into flows. In addition, the higher the level of granularity is, the more restrictive the flow definition will be when grouping IP packets. Our flow definitions go from a high level of granularity to a low level of granularity. Moreover, flows were defined in this work by using the NetFlow fields listed in Table II.

The flow definitions consider different end-points and have the following descending order of granularity:

- 1) *App2App*: the 5-tuple flow definition;
- 2) *HstHst*: group of packets with the same source and destination IP addresses;
- 3) *Sub2Sub /24*: group of packets matching the 24 most significant bits of the source and destination IP addresses;
- 4) *Sub2Sub /16*: group of packets matching the 16 most significant bits of the source and destination IP addresses;
- 5) *Sub2Sub (NetFlow)*: group of packets matching the most significant bits of the source and destination IP addresses reported by BGP-enabled routers;
- 6) *AS2AS*: group of packets with the same source and destination autonomous systems;
- 7) *Sub2Sub /8*: group of packets matching the 8 most significant bits of the source and destination IP addresses.

Given the abovementioned definitions for a flow, we define an evaluation criterion to check whether a certain IP flow is eligible to be moved to the optical level, viz.:

Criterion: an IP flow is eligible to be moved to the optical level if its bandwidth is equal or bigger than the minimal unit of transmission in SONET networks in a certain time interval: average throughput ≥ 50.112 Mbit/s.

Variants were also defined for this criterion by varying the time interval: 5 and 30 minutes. The variation in the time interval was used to check if the accuracy of the flow characteristics changes when different time intervals are used. With such time intervals the threshold values for a flow to be considered eligible for a lambda-connection are 1.8 and 11 GBytes for 5 and 30 minutes intervals, respectively. The following flow characteristics were observed:

- *size*: the amount of bytes of a flow during the collecting period;
- *duration*: the lifetime of a flow during the collecting period;
- *throughput*: the average bandwidth consumed by a flow during its duration; and
- *recurrence*: The periodicity of a flow during the collecting period. Considered flow recurrence periods: every 6hs, 12hs, 24hs, 48hs, and 168hs.

VI. RESULTS

This section presents the results of our analysis. The results shown in this section only consider flows that satisfied our

evaluation criterion stated in subsection V-C. In addition, we also used a confidence interval of 95% to calculate the confidence limits of our mean values.

A. Average flow size

Figure 4 shows the average size of flows using the various flow definitions at different time intervals.




Fig. 4. Average flow size with 5 and 30 minutes time interval.

The average flow size with 30 minutes interval is higher than the average flow size with 5 minutes interval, as, clearly, our evaluation criterion with 30 minutes interval is more restrictive and selects only flows that are big in size. Note that there is a rather large variability in the flows size when using 30 minutes interval. For example, we saw huge flows such as 24 TB, but also smaller flows such as 12 GB in *Sub2Sub /8* flows. Five of our definitions for a flow present a big average flow size. Three of them, *Sub2Sub /24*, *Sub2Sub /16*, and *Sub2Sub /8*, can be considered relevant in theory; but the other two, *Sub2Sub (NetFlow)* and *AS2AS*, can be used in practice.

B. Average flow duration

When the total period of collected data (2 weeks) is divided in 30 minutes intervals, the estimated duration of flows tends to be longer than their real duration (Figure 5).




Fig. 5. Average flow duration with 5 and 30 minutes time interval.

In this case, the use of smaller time intervals such as 5 minutes reports more precise estimation about the real duration of a flow. This can be seen in Figure 5, where the estimation regarding the duration of a flow is much bigger when 30 minutes intervals are used. There is also large variability in the flows duration when using 30 minutes interval. Regarding the comparison between the duration *versus* different definitions for a flow, flow definitions with lower level of granularity (e.g., *Sub2Sub (NetFlow)* and *AS2AS*) tend to aggregate more IP packets into longer continuous periods of time.

C. Average flow throughput

The average flow throughput appears reasonable constant independently of the flow definition used (Figure 6). However, when longer time intervals are used, the average throughput tends to be bigger as well. That happens because only big flows are selected. In contrast, when smaller time intervals are used, not only the big flows are selected, but also many small ones, which decreases the average throughput.

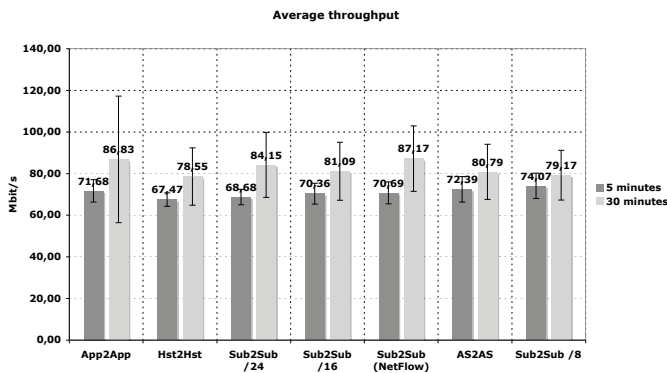


Fig. 6. Average flow throughput with 5 and 30 minutes time interval.

D. Flow recurrence percentage

The greatest recurrence percentage found in our analyses was a daily recurrence (Figure 7).

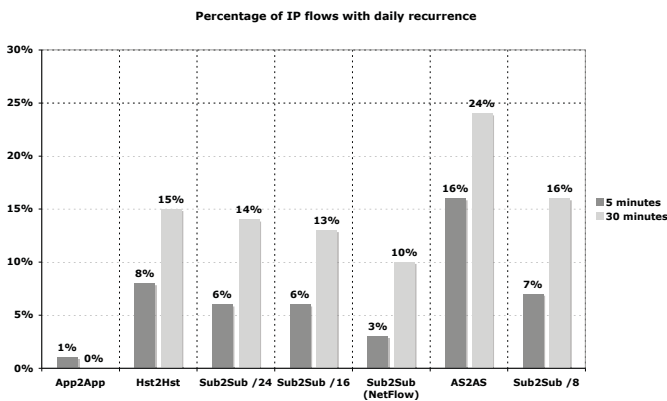


Fig. 7. Daily recurrence with 5 and 30 minutes time interval.

The other recurrence periods considered in our analyses showed results smaller than 2% and were purposely omitted

in this section due to space constraints. The use of 5 minutes interval selects several flows that have a burst behavior and are sporadic in time (i.e., they occur at irregular period of time). On the other hand, the 30 minutes interval selects flows that are longer in time and occur regularly (e.g., backup traffic). As a consequence of that, the percentage of flows with a daily recurrence is bigger when long time intervals are used (e.g., 30 minutes interval). In addition to that, the use of different definitions for a flow has certain influence over the recurrence due to the time aggregation. Flows that occur in different time intervals may overlap when less granular flow definitions are used. As a result, flows that may be sparse in time (e.g., *App2App*) can be grouped into regular occurrence intervals and therefore present a higher percentage of recurrence (e.g., *AS2AS*).

VII. CONCLUSION

This paper presented the characteristics of IP flows eligible to the optical level while using various different definitions for a flow. The characteristics observed were size, duration, throughput, and recurrence. In addition, different time intervals were used in order to study whether they have any impact on the eligibility (i.e., flow characteristics) of moving flows to the optical level.

Based on the analysis performed and in the results obtained, we can now answer our research questions.

- 1) *What are the most suitable definitions for flows eligible to be moved to the optical level?*

In order to offload IP flows that consume considerable amount of resources at the network level onto lambda-connections, and, at the same time, maximize the use of the high capacity transmission offered by the optical level, we look for definitions of a flow that are big in size (Subsection VI-A) and long in duration (Subsection VI-B). Based on that, *Sub2Sub (NetFlow)* and *AS2AS* are the best candidates for flow definition in practice. On the other hand, in theory, the *Sub2Sub /8* is the best option. However, in practice, this flow definition is not feasible as it groups flows coming from non-adjacent networks.

- 2) *What are the characteristics of those flows regarding their size, duration, and throughput?*

The characteristics of flows eligible to be moved to the optical level mostly depend on the flow definition used, as well as on the restrictiveness of our criterion (Subsection V-C). The average size and duration of a flow increases when longer time intervals and lower granular flow definitions are used. In the case of the flow throughput (Subsection VI-C), its average is constant independent of the flow definition used, but it varies if different time intervals are used due to the size of the selected flows. In average, the size of flows eligible to be moved to the optical level varies from small sizes (18 GB) to big sizes (1.7 TB). The flow duration can be a couple of hours, but also several

days. Regarding throughput, it stays between 60 and 100 Mbit/s.

3) Do those flows recur? If so, how often?

Flows eligible to the optical level present, predominantly, a daily basis recurrence, as showed in our analyses (Subsection VI-D). We also found some other flow recurrences (e.g., weekly backups within SURFNet6), but they reoccurred less than 2% of the considered analysis period (2 weeks). Within the daily basis recurrence, the use of long time intervals influences in the selection of flows that last for long periods and have a regular behavior pattern. On the other hand, small time intervals select several flows with a burst behavior whose occur in non-regular periods of time. This allows concluding that restrictive flow definitions (e.g., *App2App*) show lesser recurrence percentage than non-restrictive ones (e.g., *AS2AS*).

The main contribution of this work was to show the characteristics of the flows eligible for lambda-connections. The knowledge of such characteristics is particularly important to our research work regarding to the self-management of lambda-connections in optical networks. By providing quantification of the observed characteristics to the multi-service optical switches, they can automatically take decisions about moving flows to the optical level. For instance, the optical switches can be configured to use a definition for a flow such as *Sub2Sub (NetFlow)* and *AS2AS* that groups more IP packets into long and big flows. The use of such definitions can result in a better performance of both network and optical levels: heavy-hitter flows could be offloaded from the network level to the optical level, where they would get no jitter and plenty of bandwidth. Meanwhile, routing resources at the network level could be better used by the remaining smaller IP flows.

ACKNOWLEDGMENTS

The research on self-management has been supported by SURFnet in the context of the GigaPort-RoN project, and by the EC IST-EMANICS Network of Excellence (#26854). The work presented in this paper has also benefited from collaboration with INRIA Nancy. Special thanks to Lisandro Zambenedetti Granville for his help in this paper.

REFERENCES

- [1] A. Leon-Garcia and I. Widjaja, *Communication Networks: fundamental concepts and key architectures*, 2nd ed. New York: McGraw-Hill Companies, 2003.
- [2] C. S. R. Murthy and M. Gurusamy, *WDM Optical Networks: Concepts, Design and Algorithms*. Englewood Cliffs: Prentice Hall Publishers, 2001.
- [3] GÉANT, "GÉANT homepage," <http://www.geant.net/>, accessed on 07/2007.
- [4] Internet2, "Internet2 homepage," <http://www.internet2.edu>, accessed on 07/2007.
- [5] CLARA, "Latin American Cooperation of Advanced Networks," <http://www.redclara.net/en/index.htm>, accessed on 07/2007.
- [6] SURFnet, "SURFnet Home," <http://www.surfnet.nl/info/en/home.jsp>, accessed on 07/2007.
- [7] CANARIE, "CANARIE Home," <http://www.canarie.ca/about/index.html>, accessed on 07/2007.

- [8] HEAnet, "HEAnet Home," <http://www.heanet.ie>, accessed on 07/2007.
- [9] R. Boutaba, W. Golab, Y. Iraqi, T. Li, and B. S. Arnaud, "Grid-controlled lightpaths for high performance grid applications," *The Journal of Grid Computing*, vol. 1, no. 4, pp. 387–394, 2003.
- [10] G. Rouskas, "Optical layer multicast: rationale, building blocks, and challenges," *IEEE Network Magazine*, vol. 17, no. 1, pp. 60–65, 2003.
- [11] LOFAR, "LOFAR homepage," <http://www.lofar.org/>, accessed on 07/2007.
- [12] T. Mori, T. Takine, J. Pan, R. Kawahara, M. Uchida, and S. Goto, "Identifying heavy-hitter flows from sampled flow statistics," *IEICE Transactions on Communications*, vol. E90-B, no. 11, pp. 3061 – 3072, 2007.
- [13] A. Soule, K. Salamata, N. Taft, R. Emilion, and K. Papagiannaki, "Flow classification by histograms: or how to go on safari in the internet," in *SIGMETRICS '04/Performance '04: Proceedings of the joint international conference on Measurement and modeling of computer systems*. New York, NY, USA: ACM Press, 2004, pp. 49–60.
- [14] K. Papagiannaki, N. Taft, and C. Diot., "Impact of flow dynamics on traffic engineering design principles," in *Proceedings of the Twenty-third Annual Joint Conference of the IEEE Computer and Communications Societies*, 2004, pp. 2295 – 2306.
- [15] J. Wallerich, H. Dreger, A. Feldmann, B. Krishnamurthy, and W. Willinger, "A methodology for studying persistency aspects of internet flows," *ACM SIGCOMM Computer Communication Review*, vol. 35, no. 2, pp. 23–36, 2005.
- [16] K.-C. Lan and J. Heidemann, "A measurement study of correlations of internet flow characteristics," *Comput. Networks*, vol. 50, no. 1, pp. 46–62, 2006.
- [17] T. Fioreze and A. Pras, "Using self-management for establishing light paths in optical networks: an overview," in *Poster session proceedings of the 12th EUNICE Open European Summer School 2006 (EUNICE 2006)*, 2006, pp. 17 – 20.
- [18] T. Fioreze, R. van de Meent, and A. Pras, "An architecture for the self-management of lambda-connections in hybrid networks," in *Proceedings of the 13th EUNICE Open European Summer School 2007 (EUNICE 2007)*, 2007, pp. 141 – 148.
- [19] SURFnet, "SURFnet6 lightpaths mark start of the new Internet area (press release)," http://www.surfnet.nl/info/en/artikel_content.jsp?objectnummer=107197, accessed on 07/2007.
- [20] T. Fioreze, M. Wolbers, R. van de Meent, and A. Pras, "Offloading IP flows onto lambda-connections," in *Proceedings of the 18th IFIP/IEEE International Workshop on Distributed Systems: Operations and Management (DSOM 2007)*, 2007, pp. 183 – 186.
- [21] Y. Rekhter, T. Li, and S. Hares, "A Border Gateway Protocol 4 (BGP-4)," Request for Comments: 4271, Jan. 2006, IETF.
- [22] T. Fioreze, M. Wolbers, R. van de Meent, and A. Pras, "Finding elephant flows for optical networks," in *Application session proceedings of the 10th IFIP/IEEE International Symposium on Integrated Network Management (IM 2007)*, 2007, pp. 627 – 640.
- [23] GigaPort, "GigaPort homepage," <http://www.surfnet.nl/info/en>, accessed on 07/2007.
- [24] E. Hernandez-Valencia, "Hybrid Transport Solutions for TDM/Data Networking Services," *IEEE Communications Magazine*, vol. 40, no. 5, pp. 104–112, May 2002.
- [25] G. Bernstein, B. Rajagopalan, and D. Saha, *Optical Network Control: Architecture, Protocols, and Standards*. Reading: Addison Wesley Publishers, 2003.
- [26] Zebedee, "Zebedee: Secure IP tunnel," <http://www.winton.org.uk/zebedee/>, accessed on 07/2007.
- [27] Tcpdump, "TCPDUMP public repository," <http://www.tcpdump.org>, accessed on 07/2007.
- [28] B. Claise, "Cisco Systems NetFlow Services Export Version 9," Request for Comments: 3954, Oct. 2004, IETF.