



On the Performance of Grooming Strategies for Offloading IP Flows onto Lightpaths in Hybrid Networks

Rudolf Biesbroek, Tiago Fioreze, Lisandro Zambenedetti Granville, Aiko Pras

► To cite this version:

Rudolf Biesbroek, Tiago Fioreze, Lisandro Zambenedetti Granville, Aiko Pras. On the Performance of Grooming Strategies for Offloading IP Flows onto Lightpaths in Hybrid Networks. 16th EUNICE/IFIP WG 6.6 Workshop on Networked Services and Applications - Engineering, Control and Management (EUNICE), Jun 2010, Trondheim, Norway. pp.1-10, 10.1007/978-3-642-13971-0_1 . hal-01056496

HAL Id: hal-01056496

<https://inria.hal.science/hal-01056496>

Submitted on 20 Aug 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

On the Performance of Grooming Strategies for Offloading IP Flows onto Lightpaths in Hybrid Networks

Rudolf Biesbroek¹, Tiago Fioreze¹, Lisandro Zambenedetti Granville², Aiko Pras¹

¹ University of Twente, Design and Analysis of Communication Systems (DACS)
Enschede, The Netherlands

² Federal University of Rio Grande do Sul, Institute of Informatics
Porto Alegre, Brazil

Abstract. Hybrid networks take data forwarding decisions at multiple network levels. In order to make an efficient use of hybrid networks, traffic engineering solutions (*e.g.*, routing and data grooming techniques) are commonly employed. Within the specific context of a self-managed hybrid optical and packet switching network, one important aspect to be considered is how to efficiently and autonomically move IP flows from the IP level over lightpaths at the optical level. The more IP traffic moved (offloaded), the better. Based on that, we investigate in this paper different strategies to move IP flows onto lightpaths while observing the percentage of offloaded IP traffic per strategy.

Key words: Grooming strategies, IP flows, lightpaths, ns-2, hybrid networks

1 Introduction

The need for a separation between heavy applications and the normal Internet traffic over a shared network infrastructure has increased the importance of hybrid networks. Through the use of hybrid network infrastructures, backbone networks are able to provide better performance by means of faster delivery and more reliable data transmission. In such a hybrid environment, IP flows can traverse a hybrid network through either a lightpath or a chain of routing decisions. Moving large amounts of data from the IP level to the optical level enables flows to experience faster and more reliable transmissions with optical switching than with traditional IP routing. Moreover, transmitting data flows at the optical level is cheaper than transmitting them at the IP level [11].

In order to configure a hybrid network and create lightpaths for IP flows, a management mechanism is required. Currently, GMPLS signaling and conventional management are important solutions for that [3]. GMPLS coordinates the creation of lightpaths by employing signaling messages that are exchanged by adjacent, intermediate nodes in the path between the source and destination end nodes of a flow [12]. In the conventional management, on the other side, a central manager individually configures each node in the transmission path. Both GMPLS and conventional management rely on human decisions in order to select which flows would remain at the IP level and which other flows should be offloaded to the optical level. As expected, the human intervention turns the whole process slow and error-prone.

Based on the aforementioned state-of-the-art for the management of hybrid networks, it would be interesting to have a decision making process that could be automated in order to minimize human intervention. Having that in mind, a new management approach for hybrid networks is under investigation at the University of Twente, namely self-management of hybrid optical and packet switching networks [7, 9, 8]. One of the main challenges in such an investigation is to find out appropriate lightpath setups in which the available capacity of optical wavelengths is consumed in an optimal manner. For example, through the multiplexing of many flows into a single wavelength. Techniques for that, while considering certain design conditions (*e.g.*, minimum cost), are generally referred as *traffic grooming* [6, 14].

In this context, we pose the following research question to be answered in this paper: *what traffic grooming strategy offloads the highest percentage of IP traffic to the optical level?* Depending on the grooming strategy employed, the percentage of offloaded traffic could differ significantly. At the optical level, each wavelength has a fixed amount of available bandwidth. In most cases, the sum of the offloaded flow rates will not fill the fully available wavelength capacity, leaving some of the capacity unused. Therefore, grooming techniques should strive to minimize the amount of unused capacity, which increases the possible offload percentage.

In this paper we evaluate the performance of some grooming strategies. These strategies have the purpose of grooming many IP flows, regardless the granularity of the IP flows (*e.g.*, bit-rate), over the available lightpaths. The list of strategies that we investigate here is inspired by an earlier research on strategies and related algorithms for achieving dynamic routing of data flows for global path-provisioning [13]. Whereas the authors of the previous research have investigated the blocking probability while observing different offloading strategies to accommodate LSPs on established lightpaths, we observe the percentage of IP traffic that can be offloaded to the optical level.

Through the use of simulation we evaluate the performance of grooming strategies while observing the percentage of traffic that is offloaded by each one of them. For that, we employ three different strategies: *dedicated*, *spreading*, and *packing*. As a side research, we also observe the energy consumption of each strategy. In order to do that, we look at the number of in-use wavelengths while accommodating the offered flows that need to be offloaded to the optical level.

The remainder of this paper is structured as follows. In Section 2 we review the current status in the field of traffic grooming in hybrid networks. In Section 3 we describe our simulation model and present a network topology used in our evaluation scenarios. In Section 4 we discuss the simulation results and finally, in Section 5, we close this paper with final remarks and perspectives for future work.

2 Related Work

Our research is inspired by the research performed by Sabella *et al.* [13] who focus on a solution for an online-routing function, which allows the network to promptly react to traffic changes. The authors have proposed a strategy and related algorithms to achieve dynamic routing of data flows. To accommodate new traffic requests, they have proposed the use of two algorithms: (i) a routing algorithm, to find a route for the re-

quested traffic, and (ii) a grooming algorithm, to assign for any link of the route the traffic to an optical channel. Looking at the latter, the authors concluded that, by choosing the right grooming strategy, a reduction from two up to about four time of refused bandwidth for a network load of 70% and 55%, respectively, can be achieved. Moreover, the authors have argued that the gain of the proposed strategy (*packing* strategy) is greater when the average granularity of LSP's is coarser, and have remarked that this gain tends to diminish when the network becomes uniformly congested.

An important difference between our work and of Sabella *et al.* [13] is the goal of the grooming function; where Sabella *et al.* have aimed to maximize non-blocking probability when multiplexing LSP into the given wavelengths, we aim to achieve maximum percentage of offloading IP traffic when sending high amount of traffic reaching up to 100% of the total bandwidth.

Drummond *et al.* [5] carried out a similar investigation through the use of simulation. They showed that the NS-2 simulator, combined with the OWns package, is able to simulate grooming capabilities of IP flows into wavelengths. Despite the use of simulation to observe different grooming strategies, Drummond *et al.*, did not consider the performance of such strategies as we consider in this paper.

Operational research aims at providing analytic methods to structure and understand complex situations as well as to use this understanding to improve and predict a system's behavior [10]. Based on that, our work is aligned with the operational research field, since we aim at formulating a model that enables us to analyze and understand the behavior of our system by means of simulation. The result of our simulation enables us to analyze the system behavior regarding to the formulated method, which leads to the best performing system (*e.g.*, highest offload percentage).

3 Simulation Model

In this section we describe the model we use to simulate the offloading strategies considered in this paper. We then present: a network topology (subsection 3.1), the flow handling (*i.e.*, starting, offloading, and termination) (subsection 3.2), the evaluated criteria (subsection 3.3), the scenarios (subsection 3.4), and simulation tools (subsection 3.5). This simulation model enables the evaluation of the performance of our system in terms of percentage of offloaded IP traffic.

3.1 Topology

Our topology (Figure 1) consists of two routers being logically connected via an OC-192 link, which actually comprises of eight OC-24 links. The unidirectional transmitted data (IP flows) is sent from Router 1 to Router 2. These IP flows vary in rate between 1 Mbps to 500 Mbps. In order to avoid traffic loss, the total bandwidth of the transmitted flows is limited to 9.952 Gbps (the equivalent of an OC-192 link). All the flows generated by Router 1 are initiated at the IP level and they stay at such a level until the offload procedure moves them to the optical level.

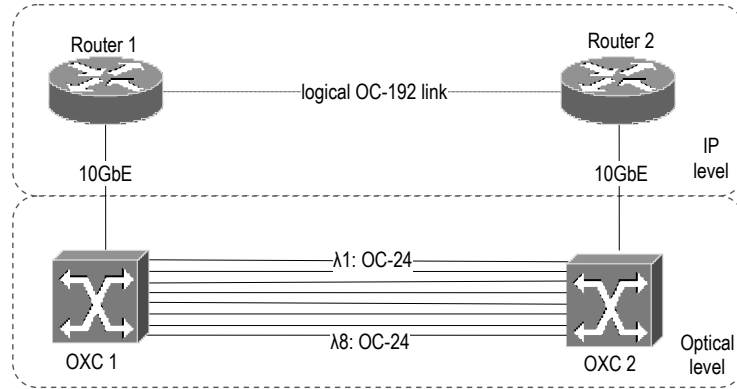


Fig. 1. Our simulation topology.

3.2 Start and Termination of Flows

The start and termination of flows is regulated as described in Figure 2. When an arrival event occurs, a new flow is started and the next flow arrival event is scheduled.

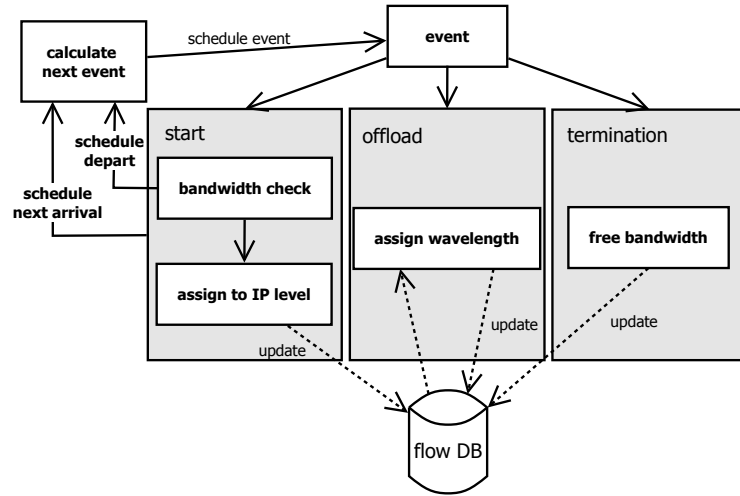


Fig. 2. Sketch of the flow handling (start, offload, and termination) during the simulation.

For the inter arrival of flows we assumed a negative exponential distribution ($\lambda = 0.1234175251$). For the termination of a flow, we assume a Weibull distribution ($\lambda = 0.0190299$, $k = 0.17494315$). These assumptions are based on the analysis performed on the IP data collected from the University of Twente network.

Upon starting the flow, a bandwidth check is performed to ensure the available bandwidth in the network and to prevent packet loss. The total available capacity C^A is found by taking the used bandwidth C_i^U per wavelength i , and the used bandwidth C_{IP}^U on the IP level, together with the total link capacity L_c ($=9.952$ Gbps).

$$C^A = L_c - \left(\sum C_i^U + C_{IP}^U \right) \quad (1)$$

If enough bandwidth has been determined, a depart event is scheduled and the flow is assigned to the IP level. The offload event is triggered in a fixed interval of one second. This event will cause the offload process of moving IP flows from the IP level to the optical level. It is important to mention that the order of offloading the flows is determined according to the rate of the flow: *flows with the highest rates are offloaded first*. When a flow has been offloaded and assigned to a wavelength, it will stay on this wavelength until it terminates. When a departure event is triggered, the associated flow will be terminated. At the moment of its termination, a flow can reside at the IP level or at the optical level. In both cases, the bandwidth associated with the flow will be released. In case this event will cause a wavelength to become empty, this wavelength will be torn down in order to save energy.

3.3 Evaluation Criteria

During the simulation, flows are generated and, whenever possible, offloaded to the optical level. The evaluation of the simulation is done with the goal of finding the best performing offload strategy (*e.g.*, the strategy that has the highest percentage of offloaded IP traffic). Thus, we take the amount of traffic that resides at the IP level and compare it with the amount of traffic at the optical level. We use the percentage of offloaded traffic as a measurement to determine the performance of the offloading strategy.

$$\text{offloaded}(\%) = \frac{\sum C_i^U}{C_{total}^U} \times 100 \quad (2)$$

Where C_{total}^U is formulated as:

$$C_{total}^U = \sum C_i^U + C_{IP}^U \quad (3)$$

We also evaluate the energy consumption at the optical level by monitoring the number of wavelengths used during our simulation. The power consumption values of each optical element is depicted in Figure 3. The 8 x 1 Gbps transponders represent the 8 x OC-24 lightpaths connecting the OXC with the WDM terminal, and a 10 Gbps transponder connecting the WDM terminal with a demux. This demux connects to its counter-part, with amplifiers in between. Then, all the optical elements aforementioned repeat themselves in inverted order until OXC2. It is important to highlight that the transponders are switched on and off on-demand. They are automatically switched on when there is data to be transmitted and switched off when there is no data transmission (to save energy). The minimum energy consumption (*e.g.*, no data is sent) comprises of the energy consumption of the OXCs + the WDM terminals + 10GTx/Rx packs + WDMs + amplifiers on both sides. The amount of energy consumption when data

is transmitted is the minimum amount of energy consumption *plus* the corresponding link's transponders of the in-use wavelengths.

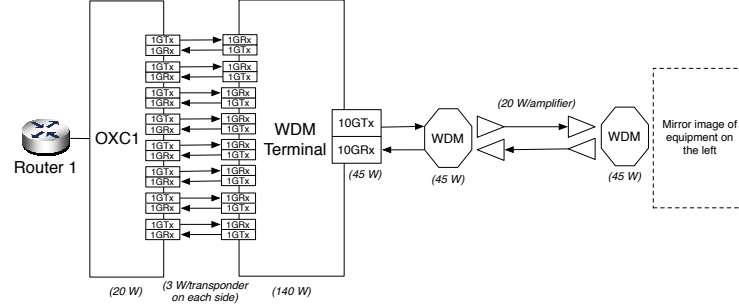


Fig. 3. Energy consumption per element (watt) at the optical level.

3.4 Scenarios

When IP flows are offloaded to the optical level, and more than one wavelength is able to serve the offloaded flow, a criteria is needed to select the proper wavelength. In our attempt to find the best performing strategy for offloading IP flows from the IP level to the optical level, three different strategies are considered: *dedicated*, *spreading*, and *packing*. For the sake of performance evaluation we also investigate the effects when offloading only the biggest flow per offloading event, instead of offloading all possible IP flows. Furthermore, we look into the effects of allowing more coarser flows by limiting the maximum flow rate to the equivalent of an OC-24 link (*i.e.*, 1.244 Gbps).

Dedicated: This strategy offloads one flow per unused wavelength. In practice, this means that a maximum of eight flows (one per wavelength) are offloaded to the optical level. The other remaining flows stay at the IP level until one of the offloaded flows ends, releasing thus the wavelength. When a new offload event occurs, the biggest flow at the IP level is offloaded, and subsequently, dedicatedly assigned to one wavelength.

Spreading: This strategy aims for an equally distributed utilization (in terms of bandwidth) of the wavelengths. By choosing the least loaded wavelength for assigning the offloaded flow, an attempt is made to equally divide the flows over the available wavelengths. This offloading algorithm chooses the wavelength with the most bandwidth capacity available to serve the flow to be offloaded.

Packing: Opposite to the *spreading* strategy, this strategy chooses the wavelength which is most loaded, *i.e.*, by checking all the available wavelengths and selecting the most loaded one able to serve the requested flow. Since the other wavelengths are kept

less loaded, the chance to serve a flow with a large bandwidth request increases. By using this strategy an attempt is made for maximizing the probability to find a wavelength that is able to serve a flow with a large bandwidth demand.

Offloading only the biggest flow: While non-offloaded flows are kept at the IP level due to not enough bandwidth available, all other flows at the IP level are offloaded to the optical one. For the sake of our performance evaluation, we also observed the performance when we offload only the flow with the highest rate at the IP level. If this flow cannot be offloaded, all flows are kept at the IP level until the next offload event is triggered, serving the highest possible priority to the biggest flow.

3.5 Simulation Tool

For the simulation of flows over an optical connection, we made use of the NS-2 simulation tool [1]. This tool does not simulate WDM networks [5] [4] without an additional module such as Optical WDM Network Simulator (OWns) [2]. Some features needed for our simulation, such as grooming multiple flows into the same wavelength, are not available in the OWns module. Just like NS-2, the OWns module is based on open source, which allowed us to code this function.

4 Simulation results

In this section we discuss the simulation results. Our simulations have been carried out in a time interval of 200,000 seconds (roughly two days).




Fig. 4. Offload percentage per grooming strategy. Flow rates vary from 1 Mbps to 500 Mbps.

As expected, the performance of the *dedicated* strategy is considerably lower than the *spreading* and *packing* strategies (Figure 4). This can be explained by the fact that only one flow per wavelength can be assigned, leaving significant amount of bandwidth unused. In our scenario we have used flow rates between 1 to 500 Mbps. Since one OC-24 link has a total bandwidth capacity of 1.244 Gbps, the maximum possible offloading percentage is about 40% when the maximum link capacity is in use. Obviously, this is not realistic in practice. As Figure 4 shows, the offloaded percentage lies between 20% and 30%. This number can increase up to 50% when the granularity of the flows are coarser (*e.g.*, flow rates between 1 Mbps to 1.244 Gbps). We also observed the performance when offloading only the biggest flow, as described in subsection 3.4. In this case, the simulation results show no significant differences for the *dedicated* strategy because the flows are offloaded almost in the same order.

The *spreading* and *packing* strategies perform better than the *dedicated* strategy. By allowing multiple flows per wavelength, the used capacity per wavelength increases significantly. As Figure 4 shows, both *spreading* and *packing* perform roughly between 90% and 100%. There are some moments where the *spreading* strategy show better results than the *packing* strategy, but on average the *packing* strategy (96.7%) slightly outperformed the *spreading* strategy (96.4%). The main reason for the better performing *packing* strategy is because it aims for maximizing the probability to serve a flow with a high bandwidth demand as described in section 3.4.




Fig. 5. Offload percentage per grooming strategy. Flow rates vary from 1 Mbps to 1.244 Gbps, using the ‘offload biggest flow only’ alternative.

Unlike the *dedicated* strategy, the *packing* and *spreading* strategies perform a little less efficient when the flows are coarser (*i.e.*, higher variance in flow rate). The time a flow needs to reside at the IP level grows when the flow rate is higher, resulting in

lower offload percentages. When rates vary between 1 Mbps to 1.244 Gbps, the *packing* strategy shows better performance (94.7%) in comparison to the *spreading* one (93.9%).

Looking at the ‘offload biggest flow only’ alternative, strong dropping is observed (Figure 5). This is caused by flows with a high bandwidth demand residing on the IP level due to insufficient available bandwidth on the optical level. When one of the wavelengths is able to accommodate the ‘waiting’ flow, a quick recovery of the offload percentage takes place. On average we can see that the *packing* strategy outperforms the *spreading* strategy when the ‘offload biggest flow only’ alternative is applied.

4.1 Energy Consumption

By observing the energy consumption of the different strategies, we can conclude that due to the small amount of time that one or more wavelengths are idle, there is no significant difference in energy consumption among the considered strategies. We do see however a small difference in the *packing* strategy when compared with the other ones. This difference arises mainly at the start of the simulation, when the total used bandwidth has not reached the near maximum transmission rate yet. When the network load is close to maximum, all eight wavelengths are in use and, independently from the used strategy, the energy consumption is quite similar. Table 1 summarizes the percentage of time that all wavelengths were in use.

Strategy	Percentage
Packing	99.82%
Spreading	99.97%
Dedicated	99.98%

Table 1. The percentage of time with all eight wavelengths in use.

5 Conclusions and Future Work

Through use of simulation we have showed the performance behavior of three different offloading strategies, *i.e.*, *dedicated*, *packing*, and *spreading*. We also observed the effects when using coarser flow rates, and we looked into an offload variant to serve flows with the biggest rate first. We conclude that, regardless the granularity of the flow rates, the *packing* strategy is superior to its *spreading* variant. At the cost of about 1%, the alternative ‘offload biggest flow only’ can provide the highest possible offload priority to the biggest flows.

We also observed the energy consumption of the optical level, depending on the used strategy. Independently from the employed strategy, transmission rates are near maximum, therefore the percentage of time that one or more wavelengths are not in use is negligible. This has a direct impact on the energy consumption because the in-use wavelengths determine the difference in the percentage of consumed energy. For this reason we conclude that there is no significant difference in power consumption.

In our work we assumed that all flows have constant rates and do not change over time. As future work, it lies in the extension of this work to investigate the performance behavior when flows rates do changes over time. Although we concluded that there is no significant difference in power consumption, it might also be interesting to investigate if this conclusion holds when the network load is *not* near maximum capacity, increasing the chance that one or more wavelengths are not in use and therefore using less energy.

Acknowledgments. This research work has been supported by the EC IST-EMANICS Network of Excellence (#26854). Special thanks to Roeland Nuijts (SURFnet) and Cees de Laat (UvA) for their valuable contribution to this paper.

References

1. Ns-2 network simulator, <http://www.isi.edu/nsnam/ns/>
2. Optical WDM network simulator (OWns), <http://www.eecs.wsu.edu/~dawn/software/ownns.html>
3. Bernstein, G., Rajagopalan, B., Saha, D.: Optical Network Control: Architecture, Protocols, and Standards. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA (2003)
4. Chittenden, A.: Extending OWns to include protection functionality. Master's thesis, University of Pretoria (2005)
5. Drummond, A.C., da Silva, R.T.R.: IP over WDM Module for the NS-2 Simulator. In: Communications Workshops, 2008. ICC Workshops '08. IEEE International Conference on. pp. 207–211. Beijing, (May 2008)
6. Dutta, R., Rouskas, G.N.: Traffic Grooming in WDM Networks: Past and Future. Tech. rep., Raleigh, NC, USA (2002)
7. Fioreze, T., Pras, A.: Using self-management for establishing light paths in optical networks: an overview. In: Proceedings of the 12th Open European Summer School. pp. 17–20. Institut für Kommunikationsnetze und Rechnersysteme, Universität Stuttgart, Stuttgart, Germany (September 2006)
8. Fioreze, T., Pras, A.: Self-management of lambda-connections in optical networks. In: Proceedings of the 1st International Conference on Autonomous Infrastructure, Management and Security (AIMS 2007) Student Workshop. Lecture Notes in Computer Science, vol. 4543, pp. 212–215. Springer Verlag, Berlin (June 2007)
9. Fioreze, T., van de, R.M., Pras, A.: An architecture for the self-management of lambda-connections in hybrid networks. In: 13th EUNICE Open European Summer School and IFIP TC6.6 Workshop on Dependable and Adaptable Networks and Services. Lecture Notes in Computer Science 4606, vol. 4606, pp. 141–148. Springer Verlag, Germany (2007)
10. Hillier, F.S., Lieberman, G.J., Hillier, F., Lieberman, G.: Introduction to Operations Research. McGraw-Hill Science/Engineering/Math (July 2004)
11. de Laat, C., Radius, E., Wallace, S.: The rationale of the current optical networking initiatives. Future Gener. Comput. Syst. 19(6), 999–1008 (2003)
12. Mannie, E.: Generalized Multi-Protocol Label Switching (GMPLS) Architecture. RFC 3945 (Proposed Standard) (Oct 2004), <http://www.ietf.org/rfc/rfc3945.txt>
13. Sabella, R., Iovanna, P., Oriolo, G., D'Aprile, P.: Strategy for dynamic routing and grooming of data flows into lightpaths in new generation network based on the GMPLS paradigm. PHOTONIC NETWORK COMMUNICATIONS 7(2), 131–144 (MAR 2004)
14. Wang, J., Cho, W., Vemuri, V.R., Mukherjee, B.: Improved approaches for cost-effective traffic grooming in wdm ring networks: Ilp formulations and single-hop and multihop connections. J. Lightwave Technol. 19(11), 1645 (2001)