

Dynamic Optical Path Allocation in Multi-Layer Traffic Engineering

Marco Ruffini, Donal O'Mahony, Linda Doyle.
 Centre for Telecommunication Value Chain Research,
 University of Dublin,
 Trinity College, Dublin 2, Ireland.

Abstract:

Abstract- Hybrid electro-optical architectures have been recently proposed as cost-effective solution to challenge near-future traffic demands. We believe that IP routing technology can be closely integrated with optical circuit switching using a multi-layer traffic engineering approach.

In this paper we propose a novel network architecture that, by exploiting the flow correlation of Internet packets, redirects the traffic into dynamically created optical paths, providing optical bypass of the IP layer.

The simulation results we provide are based on real traffic traces collected from the pan-European GÉANT network.

I. INTRODUCTION

For over ten years Internet traffic has grown at exponential rate, a trend that will continue as next generation high-bandwidth applications (IPTV, gridnets, etc.) begin to emerge. Deloitte & Touche predicted that traffic will exceed the Internet's capacity as soon as this year. Although quite pessimistic, this gives a clear idea of the challenges network operators are facing.

Although the optical transport layer, with WDM links capable of carrying Terabits per second of data seems to be ready for the challenge, things are different at the IP layer. The routing cost per Gigabit of data is in fact still very high compared to the transport cost and its slow decreasing rate suggests that this technology will not represent a cost-effective solution to the exponentially increasing traffic.

Integrating the IP and optical layers using multi-layer traffic engineering is a very promising approach for near future traffic demands, where IP and wavelength routing are closely integrated to optimize network performances. IP traffic could be engineered to exploit the dynamic switching capability of the optical layer to bypass routing when possible, producing extensive cost saving for the operators.

Besides protocol compatibility issues, a feasible solution should also be compatible with the current Internet infrastructure and organization: a highly distributed and heterogeneous network, connecting different independent domains, each following their own policies.

In this paper we introduce the Optical IP Switching concept, presenting a novel method of forwarding IP traffic through dynamically established optical circuits.

The nodes distributedly operate multilayer traffic engineering by creating dynamical optical paths that adapt the

underlying optical topology to the traffic characteristic encountered at the IP layer. The intrinsic distributed operation makes our solution also compatible with the existing domain-based structure of the Internet.

II. OPTICAL IP SWITCHING CONCEPT

In the Optical IP Switching architecture we propose [1], each IP router constantly analyzes the traffic, by sampling packets at a pre-defined rate. By using the longest prefix matching algorithm to determine the output interface, each packet can be classified depending on incoming interface, outgoing interface and destination prefix. The only information we consider is the packets size.

Each router collects data during an "observation" period before making a decision. At decision time, the router analyzes the statistics collected, considering the total amount of data between each pair of incoming and outgoing interfaces.

If the data examined is over a pre-established "path threshold", the node creates an optical cut-through path (Figure 1), between its upstream neighbour (where the selected incoming channels originates) and its downstream neighbour. The node will also indicate to its upstream neighbour the list of network prefixes to be routed through the newly created optical path. This information will be stored in the routing table of the upstream node as illustrated in Figure 2.

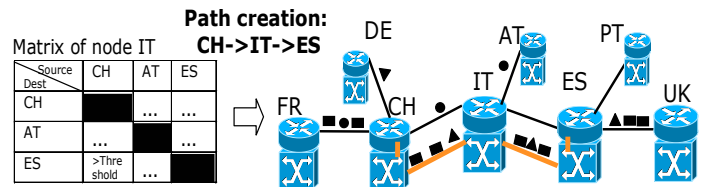


Figure 1 Path creation process

Routing table of node CH, after path creation

Prefix	Next Hop
194.22.1.1/16	Node IT
86.44.1.1/16	Node DE
193.16.1.1/16, 201.23.45.1/24	Cut-Through 1 (to IT)

Figure 2 Routing table of OIS node

We want to emphasize the fact that our cut-through path creation mechanism only involves three nodes initially, while

path extensions involving other nodes can be operated subsequently in a distributed fashion.

The path extension mechanism allows an existing optical cut-through path to be extended towards an upstream or downstream neighbour so that more nodes can avail of the transparent switching (Figure 3).

The mechanism is actuated by an extension algorithm, whose purpose is to select a subset of the prefixes switched by the original paths. Only the packets specified by this subset will be carried by the new extended path, while the remaining data will be routed through the default IP links.

The extension algorithm plays an important role in the trade-off between length of the optical path and amount of data carried by the path. An optical cut-through path can aggregate together only packets sharing a common path. When an existing path is extended, statistically, only a subset of the original packets will share the new longer path, diminishing the amount of data transported by the optical channel and consequently the channel efficiency. On the other hand however longer cut-through paths increase the number of transparent hops, enhancing the cost-saving potentials of optical switching. From this perspective, the extension algorithm has the task to optimize the cost-efficiency problem delineated by this trade-off.

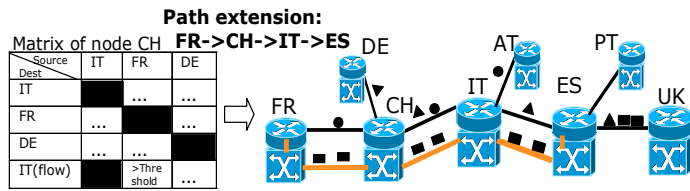


Figure 3 Path extension mechanism

III. RESULTS

We have taken the pan-European GÉANT network as reference model for our simulations, using empirical data collected from the access points. The GÉANT dataset appeared to be especially suited for our case, as it complements the traffic traces with the BGP routes collected from the border routers. Data are made available by researchers from the Computing Science and Engineering dept. at the University of Louvain-la-Neuve, who also provide C-BGP [2], a network simulator capable of reconstructing the BGP network from the routes included in the dataset.

Figure 4 and 5 show the results of our simulations. The x axis reports the traffic size: ‘x1’ denotes the traffic as collected from the original traces, which is then multiplied by progressively increasing factors.

On the y axis we report the percentage of the total traffic that could be switched into optical paths bypassing the IP layer (values are averaged over all nodes).

The results are reported for different interface rates (1 and 10Gbps). The difference between the two curves originates from the different “path threshold” values adopted (equal to 10% of the interface rate). Using 1Gbps channel rate, we see that the switching capacity increases initially with traffic up to

the “x8” value, before starting to decrease. The curve decreases when, due to the increase in traffic size, some of the flow aggregates become bigger than the channel capacity, and cannot be fitted into a dynamic optical path. The use of 10Gbps interfaces shifts the phenomenon towards higher levels of traffic.

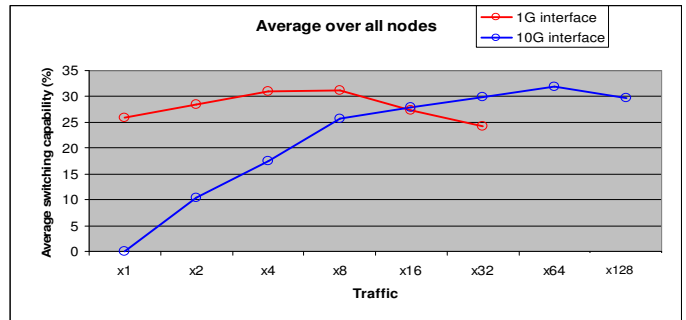


Figure 4 Simulation result averaged over all nodes

Figure 4 showed the switched traffic averaged over all the nodes. However we noticed during our simulations that the distribution of switched traffic is not uniform: while some nodes are capable of switch up to 90% of traffic, others do not switch any data. In Figure 5 we present the simulation taking in consideration only the nodes switching more than 10% of the traffic. As we can see the difference is noticeable, with averages of switched traffic well over 50%.

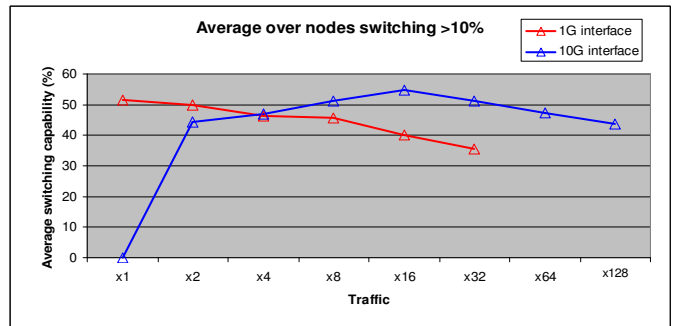


Figure 5 Simulation results averaged over nodes switching more than 10% of traffic

REFERENCES

- [1] M. Ruffini, D. O’Mahony, L. Doyle. “A Testbed Demonstrating Optical IP Switching (OIS) in Disaggregated Network Architectures” Proceedings of IEEE Tridentcom 2006, Barcelona, Spain, 1-3 March, 2006.
- [2] B. Quoitin, S. Uhlig. “Modeling the Routing of an Autonomous System with C-BGP”. IEEE Networks Magazine, Nov/Dec. 2005