Feasibility of Flow-Based Optical Provisioning in GÈANT

Marco Ruffini, Donal O'Mahony, Linda Doyle
CTVR, Universiy of Dublin, Trinity College
ruffinim@tcd.ie

Abstract: We present a novel network architecture that establishes dynamic optical paths by analyzing local network traffic. We demonstrate its feasibility and effectiveness through a network analysis using real traces collected from the pan-European GÈANT network.

OCIS codes: (060.4250) Networks; (060.4510) Optical communications

1. Introduction

Traditionally the difference between telephone and data networks has been quite sharply demarcated. Since a telephone call is characterized by a set up time that is negligible compared with the call duration, it makes sense to allocate it into an end-to-end circuit. On the other hand, data transmission has to be routed on a packet-by-packet basis to be cost effective. Historically, telephone calls were injected into circuits at the edge of the network, while data traffic was sent to an IP router towards a packet switched architecture. Today this natural separation has been blurred by the convergence of many heterogeneous applications into the Internet network. In particular voice and multimedia stream applications are frequently routed on a packet-by-packet basis and this has become a clear trend (e.g. new e-science applications, IPTV).

From a network perspective, routing such streams of traffic is very inefficient, as the strong route correlation of the packets is not exploited: the extreme granularity of a packet switched architecture is both unnecessary and unwise.

Flow-based switching [1] was proposed as a solution to improve the efficiency of the IP network under these circumstances. The main purpose is to locate and redirect packets with high route correlation (i.e. those originating a large traffic flow) towards a dedicated circuit dynamically allocated within the same network. The capability of using packet or virtual-circuit switching as appropriate to forward the encountered traffic is the key to restore high efficiency in the Internet network. Besides, this approach would allow quality of service differentiation, a feature more and more network operators seem to be interested in.

We have extended the flow-based switching concept down to the optical layer proposing a network architecture called Optical IP Switching (OIS) [2]. An OIS node is an IP router with optical switching capabilities, able to identify and aggregate large traffic flows, allocating them into dynamically created optical cut-through paths. The foundation of flow-based switching lies in the co-existence of both small and large traffic flows within the same network. This characteristics has been thoroughly investigated in literature, and interestingly it was proved that Internet traffic is characterised by a heavy-tailed distribution [3], [4]. As a consequence most of the flows encountered in the network are "mice" made up of a small number of packets; on the other hand the few remaining "elephant" flows carry most of the network traffic. This characteristic perfectly matches the idea of switching large flows at the optical level, where optical paths might be limited in number but suited to carry a large amount of data traffic. However proving the existence of "mice" and "elephant" flows on Internet links is not sufficient to support the idea of a flow based optical switching architecture. In fact it is very important also to investigate the propagation of these flows through the network, their inclination towards aggregation and the time variation of their characteristics.

In this paper we apply our flow-based optical switching model to the pan-European GÈANT network, using real data traces. Our results show how the time-variation of "elephant" flows influences the switching capabilities required of the network nodes.

2. Optical IP Switching (OIS) architecture

The idea behind Optical IP Switching is to adapt the underlying physical topology to the traffic flows encountered at the IP layer. The main feature of the system is that the decision making process is completely distributed, based only on local observation of the traffic. An OIS node monitors the traffic by sampling IP packets at a certain rate, using mechanisms similar to those used by Netflow. The goal is to identify elephant flows, aggregate them on the basis of their upstream and downstream direction, and evaluate the feasibility of switching those flows into a dedicated optical cut-through path, established between the selected upstream and downstream neighbours. The advantage for the middle node in such a path is that it does not need to route those flows, which are instead injected by the upstream neighbour into the new transparent cut-through path. A path extension process then allows upstream and downstream neighbours to extend the optical path to their own neighbours, availing of the advantages of transparent

©OSA 1-55752-830-6

switching. The node in this case selects from the existing path the flows directed towards a particular downstream neighbour: if this subset is large enough the path is extended downstream carrying only the subset of flows. A similar process is required for upstream extension.

The main advantage of this distributed decision mechanism is that every node can autonomously evaluate the convenience of switching or routing a flow aggregate, depending on its available electrical and optical resources. This makes the concept very suitable for inter-domain networking and more generally for highly heterogeneous network, as it allows every node to make its own traffic analysis and optimization. From a network perspective this approach creates an opportunity for very high cost savings [5]. It simplifies the network signalling and management, and increases the routing efficiency providing automatic adaptation of the network resources to the traffic needs.

3. Analysis of the GÈANT dataset

GÈANT is a pan-European data communication network interconnecting different National Research and Education Networks in Europe, serving over 3500 research and education institutions. The network topology includes 23 nodes, connected by links with capacities ranging from 155 Mbps to 10 Gbps.

We have taken the GÈANT network as a model for our simulation study, using empirical data directly collected from GÈANT nodes. Since our goal is the study of how traffic flows are routed in a real network and the possibility of aggregating them into a dedicated optical path, we could not base our analysis on isolated traffic traces. The GÈANT 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[6], a network simulator capable of reconstructing the BGP network from the routes included in the dataset. The tool embeds a clustering algorithm that reduces the number of BGP entries by more than two orders of magnitude, making it possible to simulate a real network scenario. The traffic traces, collected using Netflow with sampling rate of 1/1000, are summarized depending on their source/destination prefix and only the total number of bytes over a 15 minutes period is provided. This has the two-fold effect of saving storage space for the data files while keeping the traces anonymous. The disadvantage is that information about the precise timing of the flows is lost: a condition that however does not influence our study, since the mechanism that creates optical cut-though paths averages the observed traffic over a period of some minutes.

C-BGP proved extremely valuable, as it reconstructed the routing path of the traffic traces from the source node, where they were collected, towards their destination. We have fed these routing paths into our simulator, which grouped traces with a common path into simulated optical cut-through paths. Only “elephant” flows with average data rate higher than a certain threshold (100 Kbps) were taken in consideration. The minimum average rate over which a flow aggregate could trigger the creation of a dedicated cut-through path was set to 100 Mbps (equal to 10% of the bandwidth of Gigabit Ethernet links). The traffic traces we have considered were differently distributed in time: three of them were within one hour interval, three were on adjacent days and the rest approximately one week apart from each other. We have selected traces collected during weekdays, between 3.30 and 4.30 p.m., in order to analyse peak hours.

4. Discussion of the results

We report on the analysis of traffic propagation from the ingress towards the egress of the GÈANT network. Although this study focuses on the feasibility of the Optical IP Switching architecture, the results we provide can be considered as a general demonstration of flow aggregation capability in GÈANT. The algorithm we use to accomplish this task considers one trace at a time and does not try to optimize the percentage of traffic that is switched other than routed: the optimization issue is part of our future work.

Fig 1 shows the amount of traffic switched in optical cut-through path versus the remaining routed traffic, for each node in different time periods. It is interesting to observe that this ratio, at each node, has a mild variation over time. Analysis of the number of optical ports, transmitters and receiver used at each node, support this observation, showing similar quasi-static behaviour. This implies that the amount of optical resources required at each node is reasonably constant over time: a primary requirement for the practical implementation of any optical network. On more detailed examination of our traces we noticed that nodes with higher ratio of switched traffic (AT, DE2,HU) are mainly routing data between GÈANT nodes, while others with little or no switched traffic (DE1, SI, CH) are dealing with data coming from or directed to external networks.

In some cases (most notably the last 3 traces for DE2, and last 5 for AT), there is a sudden change in the amount of switched vs. routed traffic. These changes appear to be step-like, as the values remain stable after the change. By observing the traces leading to this behaviour we could directly correlate the traffic increase at DE2 with the decrease at CH, and conclude that probably the change was a consequence of an internal re-engineering of the

GÈANT network connections. The variation observed at the AT node instead is simply due to a general drop of traffic towards AT (mainly coming from DE1 and CH).

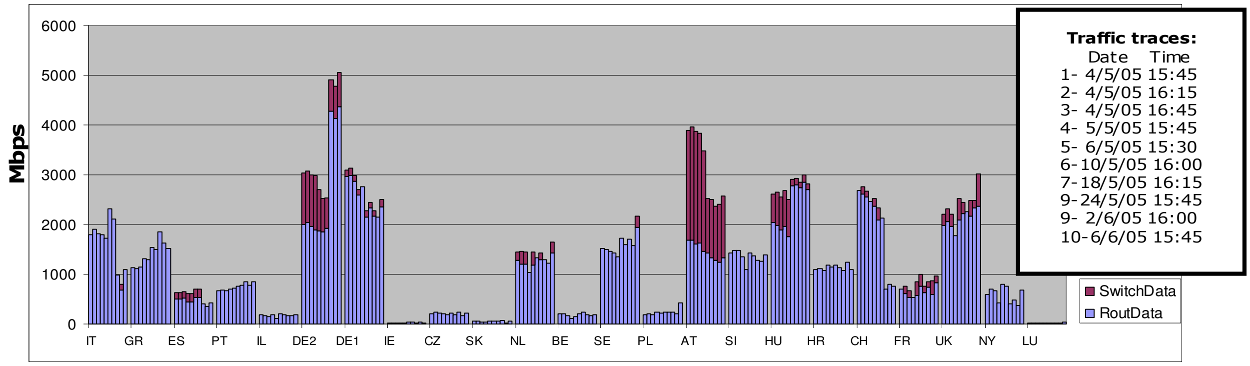

Fig 1: Switched vs. routed traffic over time

Figure 2 shows the correlation of the optical paths over time, reporting the number of occurrences of the same path in the traces under examination. Different figures compare different sets of traces. Figure 2.A for example considers three traces collected within one hour period (traces 1, 2 and 3), where 75 % of the optical paths remain unchanged. If we extend the observation period, comparing traces collected on three adjacent days instead (fig 2.B, traces 1, 4, 5), we notice that the number of paths unchanged decreases down to 55%. Finally in fig 2.C, where we extend the correlation analysis to all the 10 traces, we see that only 6% of them are static over a month time period. This type of behaviour shows that although the ratio of switched versus routed traffic (reported in fig. 1) is subjected to mild variations over time, the structure of underlying flow aggregates is much more dynamic. This variability can be best handled by an architecture that, like Optical IP Switching, is capable of dynamically adjusting its optical connection to the actual traffic needs.

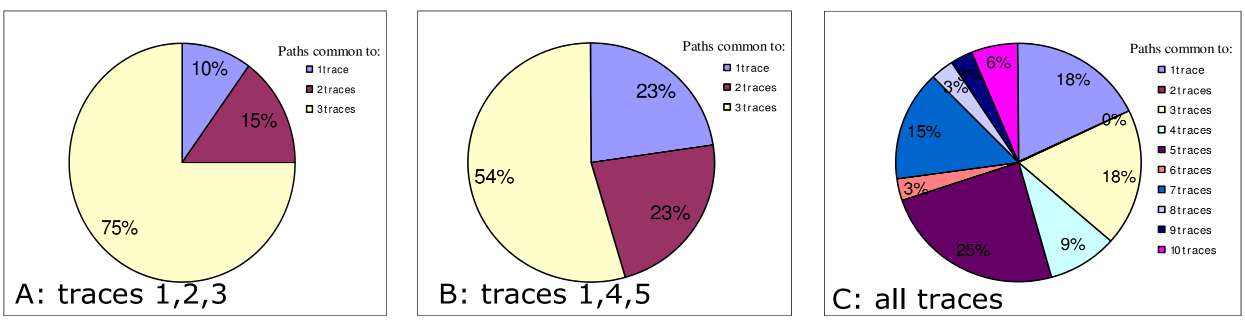

Fig 2: Correlation of optical paths over time

4. Conclusions

In this paper we have introduced a novel optical architecture that increases the routing efficiency by aggregating suitable traffic flows into dedicated optical cut-through paths. We have used real data collected from the pan-European GÈANT network to study the aggregation characteristics of traffic flows. The results we have obtained show the feasibility of our approach and its advantages over static allocation of optical paths. We are currently investigating new path creation/extension algorithms to increase the ratio of switched versus routed traffic.

5. References

[1] P. Newman, G. Minshall, T. L. Lyon.. "IP Switching-ATM Under IP". IEEE/ACM Transactions on Networking, Vol. 6 , No 2 , Apr. 1998.
[2] 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.
[3] N. Brownlee, KC Claffy. "Understanding Internet Traffic Streams: Dragonflies and Tortoises". IEEE Communications Magazine, Oct. 2002.
[4] W. Fang, L. Peterson. "Inter-AS Traffic Patterns and Their Implications". Proceedings of the IEEE Globecomm, Brazil. Dec. 1999.
[5] M. Ruffini, D. O'Mahony, L. Doyle. "A cost analysis of Optical IP Switching in new generation optical networks". Proceedings of Photonics in Switching conference 2006, Heraklion, Greece, 16-18 Oct. 2006.
[6] B. Quoitin, S. Uhlig. "Modeling the Routing of an Autonomous System with C-BGP". IEEE Networks Magazine, Nov/Dec. 2005.