

Evaluation of ORION in predimensioned networks

Erik Van Breusegem, Bart Puype , Didier Colle, Jan Cheyns, Mario Pickavet, and Piet Demeester

Department of Information Technology (INTEC)
Ghent University - IMEC
Sint-Pietersnieuwstraat 41, B-9000 Gent, Belgium.
Tel.: +32 9 33 14938, Fax: +32 9 33 14899
E-mail: Erik.VanBreusegem@intec.UGent.be

Abstract. The emerging wavelength switched networks reduce the strain on packet forwarding. Unfortunately, that solution is not really efficient on a bandwidth level, and is not ideally suited for bursty traffic. Packet switched solutions, whether electronic or optical, can use statistical multiplexing to cope with bursty traffic and yield better bandwidth efficiency. In earlier work we presented a novel network concept, Overspill Routing In Optical Networks (ORION), that can combine these two worlds, withholding their advantages. In this document we further evaluate the benefits of this architecture, focusing on ORION over a predimensioned groomed topology. This is a latex adaption to the word template.

Keywords: Circuit Switching, Optical Switching, Packet Switching, Wavelength Division Multiplexing

1 Introduction

In the past optical networks opened up a vast amount of bandwidth, with the breakthrough of optical fiber and all-optical amplification. A second boom in bandwidth capacity came with the development of (Dense) Wavelength Division Multiplexing, (D)WDM. As a consequence, pure bandwidth availability is no longer the most important cost determining factor on the technological side. Originally, optical connections were used as high capacity point-to-point interconnections. Data was converted to the electronic domain every hop, and processed at an electronic router. This resulted in highly efficient sharing of the transmission resources. At the same time, however, most of the traffic handled in any router is transit traffic, i.e. still has to travel one or more hops to reach the exit point of the network. 70% of all traffic being transit is not an exception [1], which is of course mainly due to the low meshing degree of backbone networks. Given the current trend that optical transmission capacity grows faster than electronic processing capability, the electronics will become a major bottleneck. Consequently, several solutions were proposed to reduce the strain on the electronic layer. The most important one certainly is the concept of wavelength paths (or lightpaths). In this technology a wavelength passes

transparently through an optical cross-connect (OXC), without the IP router inspecting the carried data. The drawback of this approach is that a wavelength path cannot be re-used by intermediate nodes when it is not used at full capacity. This can result in an increased amount of wavelengths needed to fulfill a set of demands. A logical next step in this evolution are Automatically Switched Optical Networks (ASONs), where the wavelength paths are set up by the control plane, without (explicit) intervention of the network operator. This allows adapting to traffic variations on a medium to large time scale (several minutes to even hours or days). However, these solutions inherently remain wavelength switched. Although the flexibility is considerably increased, the granularity, in terms of bandwidth and time scale, is still coarse.

The main contribution of this paper lies in the more sophisticated evaluation of our switching concept, Overspill Routing in Optical Networks (ORION), leading to several new insights. In the earlier papers we started with a certain network demand and then dimensioned a network, utilizing one of four switching technologies. While this evaluation gave valuable insights, it also had a few shortcomings. Most notably, a very simple routing scheme was used (shortest path only), wavelength paths were setup only end-to-end, and a wavelength path could only be used by one demand at a time. Also, no loss evaluation was possible since the network was dimensioned for a certain demand with a certain confidence level. Thus, the actual performance could not be observed. This paper addresses these issues, by turning the evaluation upside down. We start with an already dimensioned (groomed) network for a given demand (pattern), and then evaluate the benefits of ORION using loss and throughput calculations. But first, we start with shortly revisiting our recently proposed switch and network architecture proposal, which we believe succeeds in using the full benefits of ASON, without sacrificing the high statistical multiplexing gains from earlier point to point WDM (packet switched) optical networks. This introduction is mostly a short version of the concept as described in [2], which also contains a good summary and references to other ORION work.

2 ORION in a Nutshell

This paragraph explains how ORION works at a network level, which is sufficient for the discussion in the remainder of the paper. In ORION, we start from a wavelength switched WDM network, where some wavelength paths are established. In a normal wavelength switched network a lightpath passes some node transparently, i.e. the node cannot access the data in the passing wavelength. In the example of Fig. 1 node D has no access to the IP packets travelling in wavelength λ_1 . Only wavelengths starting or terminating in that node can be accessed. Suppose we have the connections A-a-b-c-d-D on λ_2 , B-b-c-d-e-E on λ_3 and A-a-b-c-d-e-E on λ_1 , established as indicated on Fig. 1. The basic problem with wavelength switching occurs in the following scenario: when the capacity of the lightpath AD is fully used, and there is some temporary overload for this connection, the wavelength λ_1 cannot be used to carry some of its traffic, as it is dedicated for connection AE. It is, however, perfectly possible that the lightpath AE (on λ_1) is only carrying a low traffic load, so that in fact at some moments in time it is empty (we focus on packets directly over WDM here). Thus, while connection AD has a capacity shortage, AE has capacity in abundance.

these design choices are mainly to clearly separate the two modes of operation in an ORION network - point to point WDM like for overspill, and wavelength switched for lightpath packets - and to avoid further complicating the design (needlessly).

The overspill mechanism is meant as a *temporary solution* for a *transient* overload scenario. If the experienced imbalance in the network remains for a longer period of time, wavelength paths should be rearranged to accommodate for the changing traffic pattern. Typically this would mean either an intervention from the operator or a dynamic rearrangement protocol. Thus, overspill is meant as a solution for the time frame not handled by these slower mechanisms. This would mean overspill could work in a time frame of seconds up to even hours or days, depending on the underlying rearrangement technology and the severity of the imbalance. The overspill mechanism itself could even function to postpone unnecessary tearing down and construction of wavelength paths, thereby enhancing network stability. Note that the overspill mechanism can also work *during path setup*, reusing the otherwise wasted time during setup or teardown of a wavelength path.

3 Evaluation Methodology

In this document we want to compare three switching technologies:

1. An optimized optical network, with some established lightpaths.
2. That same optimized optical network, with ORION enabled IP routers.
3. The same optical network as in 1, but without lightpaths. This means the optical links are only used as interconnects between IP routers. In other words, no OXCs are used, and the lightpaths are thus “converted” to a series of terminated wavelengths. It can also be thought of as replacing each OXC connection with an electronic interface to the IP router.

Designing such an optimized optical network is a big challenge in itself, and not the focus of our study. In essence, one needs to establish wavelength paths that are neither underused, nor overused, and this under changing demands over relatively long time scales. This type of dimensioning is commonly called multilayer traffic engineering, and is a research topic in itself. Therefore we relied on existing work, described in [3]. In essence, we took the dimensioned network and its associated demands as a starting point. From those, we looked at loss and throughput, for changing demands. It is important to note that the dimensioning is conservative, in that it strongly avoids a highly loaded link, by applying a specific cost function (see [3] for more details). This is to incorporate short term fluctuations of the given (static) demands. Since ORION is designed to actually deal with such fluctuations we introduced them, by increasing the demands from the original dimensioning. Thus, each evaluation is roughly done in the following way:

1. Firstly, a dimensioned network for a certain traffic profile and physical topology is obtained, using the methodology described in [3].
2. A point to point network is generated from the dimensioned network. As stated, this simply involves converting the lightpaths to simple fiber interconnects. Thus, no IP routers are bypassed any longer, but the same amount of bandwidth resources is used.
3. Next, starting from the demand set used in step 1, a new demand set is created.

demands and possible routings. Note that we always work directionally i.e. a link from A to B can have different load, capacity and thus throughput if traversed from B to A.

Instead we adopted an iterative approach, as follows:

1. Initialize the throughput estimate for each link at 100%.
2. Take a list view of all the links in the network.
3. Select the next link in the list view.
4. Using the current throughput estimates, calculate how much demand arrives at this link.
5. Update, if necessary, the current throughput for this link, by adjusting it so that no more than the $\max(\text{link capacity, demand})$ travels over the link.
6. Go to step 3 unless all links are visited.
7. Go to step 2 unless there is no significant difference with the previous iteration.

For the above described example, this would work as follows:

- Initialize (relative) throughput of L_1 and L_4 at 1 (=100%)
- Looking at L_1 , A has an estimated demand of 4 Gb/s here, and B 8 Gb/s (since no loss is assumed at L_4). This gives a total demand of 12 Gb/s, and hence a loss rate of 0.167. The throughput at link L_1 is adjusted to 0.833.
- Looking at L_4 , A has an estimated demand of $4 * 0.833 = 3.332$ Gb/s here, and B 8 Gb/s. Total demand thus is 11.332 Gb/s. The new throughput estimate for L_4 is 0.882.
- Throughputs have changed significantly, so the previous steps are repeated so that :
 - L_1 becomes $10 / (4 + 0.882 * 8) = 0.905$, thus L_4 becomes $10 / (8 + 0.905 * 4) = 0.861$
 - L_1 becomes $10 / (4 + 0.861 * 8) = 0.918$, thus L_4 becomes $10 / (8 + 0.918 * 4) = 0.857$
 - L_1 becomes $10 / (4 + 0.857 * 8) = 0.921$, thus L_4 becomes $10 / (8 + 0.921 * 4) = 0.856$
 - etc., until enough precision is reached.

In the above example the throughput estimates converge monotonically towards the actual values ($L_1 = 0.922144$ and $L_4 = 0.855536$, when solved using algebra). This is not always the case. For example, we sometimes observed that a certain link is initially found as having a big loss, but in the next iteration has no loss at all. Those kind of big jumps are usually caused by big bottleneck links which are updated late in an iteration. Throughout all evaluations we made however, the solution always converged (usually within 10 iterations to a precision over 10^{-8}). Thus, while convergence was a major concern when implementing this method of calculating loss, in practice we found no problems.

Another concern is whether the converged solution is the right one. More specifically, our concern was that the algorithm could converge to different throughput values depending on the start values. Clearly, in a real network there is only one solution under the assumed conditions. Although we have no theoretical proof, the algebraic equations (inequalities) we solved for some simple scenarios consistently agreed with the algorithm. Also, we tried several different start values for the algorithm, and they all converged to the same value. Therefore, we are confident in its correctness.

3.2 Determining Overspill

With the technique described in the previous paragraph we can determine throughput and loss statistics for an arbitrary network of any size, with or without wavelength paths,

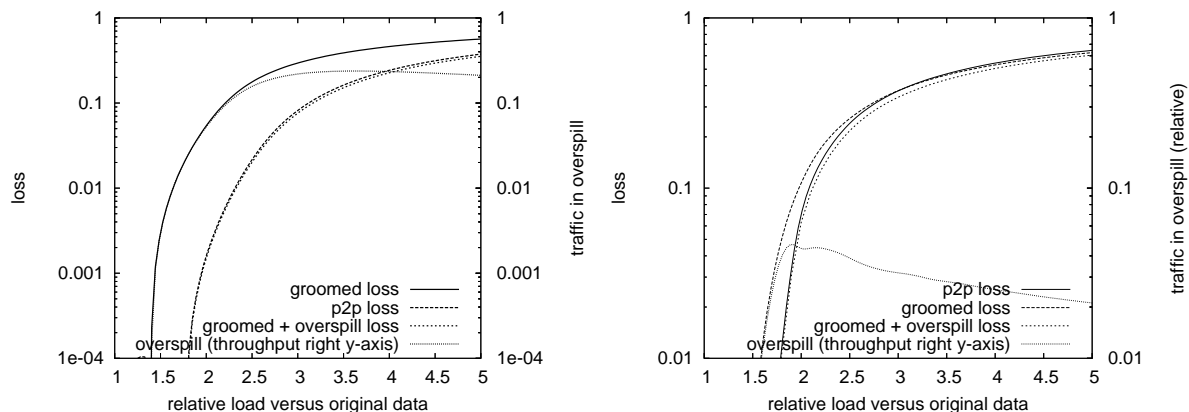


Fig. 4. Results : A typical case of overspill usage in a network (left), and a strange special case (right)

network. From the throughput figures we can compute for each original demand how much traffic gets through, how much traffic gets through in overspill (and where it enters the overspill mode), and how much traffic is lost (and where).

4 Simulation Results

In this section we discuss the two most interesting results we currently obtained from experimenting with the setup as described in the previous paragraphs. In each case, we compare the three options described in section 3. The first one is a general scenario, in which we averaged out the results of 1000 different predimensioned topologies and corresponding demands. In each of those networks, all demands were then increased proportionally, creating a global network spike. This gives a general feel of how ORION behaves. The second result is a more detailed look at one specific predimensioned network with corresponding demand, taken out of the set of 1000 for the general case. They were both performed on a hypothetical pan-european topology [3], [5].

4.1 General Overload Scenario, averaged out over 1000 runs

Fig. 4 (left) shows the general picture. It shows the average loss - defined as total traffic lost divided by total traffic offered - of a groomed topology, that same groomed topology with overspill enabled, and an equivalent point to point topology (equal wavelength resources). This is in function of an increasing demand, relative to the demand that the network originally was dimensioned for. The right Y-axis at the same time shows how much traffic is sent through the overspill mechanism.

The first observation is quite trivial: given a certain demand, in this case at around a 40% traffic increase, the groomed topology gives in first, and shows rapidly increasing loss. The p2p architecture holds out longer, until traffic increases over 70% from our starting point. This is a simple confirmation, and a well known characteristic. Due to higher multiplexing capabilities, the p2p architecture has more flexibility, at the cost of (much) more IP traffic processing. It achieves up to several orders of magnitude less loss.

caused by the fact that in severely overloaded networks a lot of traffic survives few hops, but almost no traffic survives the complete path due to collisions with other traffic. In a wavelength switched solution there a lot less hops at the logical IP layer, meaning that the overall chance to complete the end to end trajectory can be higher, if looking at the total useful network throughput. The end result is that in wavelength switching a small portion of traffic gets through, while most is lost at the very beginning. In effect, this creates an unfair situation which strangely improves total network throughput.

Note that these kind of situations are rare in most well-designed networks (we observed it only in one network in the batch of 1000), and that the difference observed (as you can see in Fig. 4 (right)) is rather small. It may however be useful to try to exploit this behavior in highly variable network environments.

5 Conclusions

In this paper we used an evaluation methodology, allowing us to assess network performance (loss and throughput) of our recently developed ORION concept in a more convincing way. This allowed using more sophisticated multilayer traffic engineered networks. As expected, the difference in performance with or without ORION is not spectacular ($\sim 20\%$ on average in the observed cases), but still very significant. In addition, an interesting effect was discovered in which a point to point switched WDM network can be outperformed by an equivalent wavelength switched network. This same effect also makes a wavelength switched network with ORION enabled outperform slightly a comparable point to point switched WDM network. Thus, the hybrid nature of ORION makes the resulting network more than just the sum of its parts.

6 Acknowledgements

This work was partly funded by the European commission through projects IST-LASAGNE (FP6 506760) and IST-NOBEL (FP6 507509). The Flemish government also funded through the IWT-GBOU project "Optical Networking and Node Architectures", and through a PhD and postdoctoral grant to E. Van Breusegem and D. Colle, respectively.

References

1. S. De Maesschalck et al., "Pan-European Optical Transport Networks: an Availability-based Comparison", *Photonic Network Comm.*, Vol. 5, No. 3, May 2003, pp. 203-225.
2. E. Van Breusegem et al. , "A broad view on overspill routing in optical networks: a real synthesis of packet and circuit switching ?", *Optical Switching and Networking* (inaugural issue), Vol. 1, no.1, pp. 51-64, (2004)
3. B. Puype et al., "Multi-layer traffic engineering in data-centric optical networks - illustration of concepts and benefits (IST LION)", *Proc. of ONDM 2003*, Budapest, Hungary, vol. 1, February 3-5, 2003, pp. 211-226
4. J. Cheyngs et al., "Controlling LSPs in an ORION network", *Broadnets 2004*, San Jose (California, USA), October 25-29 2004.
5. Online available at: <http://ibcn.atlantis.rug.ac.be/projects/IST/NRS/index.html>.