

Performance Analysis of a Hybrid Optical Switch

M. De Leenheer, C. Develder, J. Vermeir, J. Buysse, F. De Turck, B. Dhoedt, and P. Demeester

Dept. of Information Technology

Ghent University - IBBT

Gaston Crommenlaan 8 bus 201, 9000 Gent, Belgium

Email: {marc.deleenheer,chris.develder}@intec.ugent.be

Abstract—To cope with ever-increasing traffic demands in transport networks, all-optical switching is currently perceived as a solution to remove bottlenecks imposed by O/E/O conversions during data transfer. The successful realization of this concept is in large part dependent on the optical switch, which must support a wide range of traffic patterns, while remaining feasible to build both in an economical and practical sense. In this paper, we show a generic design for a hybrid optical switch composed of both slow and fast switching fabrics, and present a performance analysis to provide deeper insight in its behaviour. To this end, we propose and evaluate scheduling algorithms required at the edge of the network to map traffic on the different portions of the core switch, and present a simulation analysis covering a wide range of traffic parameters and switch design choices. These results show the effectiveness of the hybrid switch in catering for short-lived circuits (bursts) by only a limited amount of costly high-speed switching components.

I. INTRODUCTION

Optical networks have a proven track-record in the context of long-haul, point-to-point networking, where large amounts of data are transported in a cost-effective way. However, interest is growing to use optical networks in edge and even access networks (e.g. FTTH), mostly because of the predictable performance of photonic technology (i.e. high bandwidth, low latency). A major issue is O/E/O (optical/electronic/optical) conversions in the network, mostly because the speed of electronic processing can not match the bandwidths currently offered in the form of 40 Gbps and higher. For this reason, current research is focusing on all-optical networking solutions. As of today, it is possible to create all-optical networks through the use of circuit-switched paths, which essentially reserve one or more full wavelengths between end points. For instance, Lambda Grids are a general term to refer to Grid applications making use of wavelengths (i.e. lambdas) to connect high-performance computing sites over an optical network. However, novel applications are appearing which demand a much more fine-grained access to bandwidth capacity, as is demonstrated for instance in consumer Grids [1]. In such a scenario data sizes become smaller, since aggregation of multiple data sources is much harder, and the bandwidth utilization would drop dramatically if full wavelengths were to be used by these applications. Consequently, the network must support reservation and allocation of bandwidth on a sub-wavelength scale. In this paper, we propose a generic hybrid optical switch architecture, which supports both circuits and bursts. We show through simulation analysis of a single node

that this architecture has improved performance over single-technology nodes (e.g. a circuit-only node).

Previous work on hybrid optical switching can be classified in two sections. First, several efforts propose performance models for hybrid optical nodes, focused on achieving accurate and scalable logical (as opposed to physical layer) performance calculation [2], [3], [4], [5]. Further research has targeted the possible improvements by using hybrid optical in contrast to single-technology approaches [6], [7]. However, only recently work has appeared which shows initial studies on the architecture and design of such hybrid optical switches [8].

This work can be considered as performance modeling effort, but is the first to identify and evaluate essential parameters related to data traffic and switch design. Additionally, the results w.r.t. the behaviour of individual nodes can be used as guideline for the design and optimization of optical transport networks, based on hybrid switching nodes. To this end, our work starts from a generic model for a hybrid optical switch, allowing us to make very general assumptions while still drawing important conclusions, valid for a wide range of hybrid optical switch designs. This is achieved by implementing the generic model in a discrete event simulator, and running a series of experiments aimed at analyzing the performance and behaviour of a single switch. The performance analysis is mostly focused on the acceptance probability of the switch; this is a measure for the fraction of traffic which can be successfully switched. As such, it is an important parameter for the optimization of the network's operation as a whole.

The rest of this paper is structured as follows. In Section II we discuss in detail the problems related to construction and performance optimization of an optical switch, and present a generic model for the design of a hybrid optical switch. The following Section III introduces two scheduling algorithms for the assignation of traffic onto different wavelengths, and their influence on the traffic pattern offered to the switch. Extensive simulations of a single hybrid optical switch are shown and discussed in Section IV, while our conclusions are presented in Section V.

II. PROBLEM STATEMENT AND SWITCH DESIGN

The basic function of an optical switch (also referred to as OXC or Optical Cross Connect) is straightforward: it must create a path between an input and an output port for each incoming data packet. The decision which output port a data packet should be directed to is usually made in a control unit

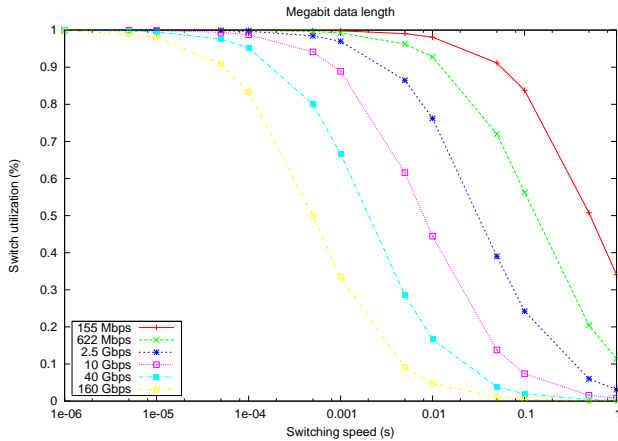


Fig. 1. Upper bound for utilization of an optical switch for different switch speeds and bandwidths (data size is 10 MB)

available at each optical switch. This unit receives control information from each data transfer, which can be a reservation packet long in advance for circuit switching, or a header prepended to the actual data for packet switching. In this paper, we assume data is sent in bursts, i.e. OBS or Optical Burst Switching, and control information is sent ahead of the actual data on a separate control plane (i.e. out-of-band signaling). The time between the control packet and the actual data transfer is denoted by T_{offset} , and is the time available to the switch to reconfigure its internal cross-connections. Each switching fabric (see [9] for current technologies) is limited by its switching speed T_{switch} , and thus a data burst can only be switched successfully if $T_{switch} < T_{offset}$. A related parameter is the switch utilization, and is bounded by:

$$\frac{T_{data}}{T_{data} + T_{switch}}$$

An illustration of this can be found in Figure 1, which shows the maximum utilization of an optical switch as a function of varying switching speeds. The data transferred has a size of 80 Mbit (10 MB), and the experiment is repeated for different link speeds. If we take, for instance, a switch speed of 10 ms (a representative value for MEMS-based switches), we see that the switch utilization is 76% for a 2.5 Gbps link speed. This value drops to below 20% for 40 Gbps link speeds, and the situation clearly becomes worse for even higher bandwidths. Obviously, the same argument holds for a fixed bandwidth and decreasing data sizes. The example shows that, to support very long data transfers (i.e. circuits), slow switching speeds are usually sufficient to obtain a high switch utilization, even for very high speed link rates. However, for smaller data transfers (burst or even packet sizes), high speed switching fabrics are required to achieve acceptable throughput in optical switching nodes. As current and emerging applications generate data according to very diverse distributions (both the data sizes and the instants of time at which the data is created), the idea arose to integrate multiple types of switching fabric into a single optical switch. This concept is generally referred to as

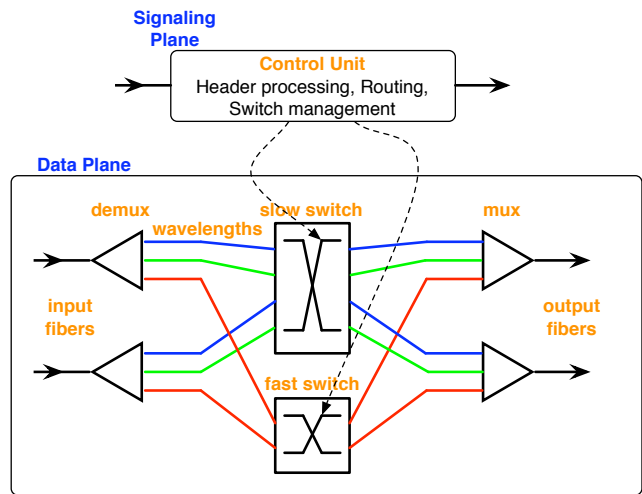


Fig. 2. Generic model for a hybrid optical switch

hybrid optical switching, and becomes essential if a single, unified data plane needs to support a wide range of users and applications.

Global network optimization not only depends on efficiency and utilization, but also on the feasibility to offer this technology in a cost-effective and practical way. Current optical switching technologies offer a broad range of switching speeds, but faster switching speeds generally have two distinct disadvantages: cost and scalability. For instance, micro-electromechanical switches (MEMS) have a typical switching time in the millisecond range, while it is technologically feasible to produce port counts of for instance 1000x1000. In contrast, Semiconductor Optical Amplifier technology (SOA) can only scale up to 32x32 port counts at very high cost, but at the same time can achieve switching speeds in the nanosecond range. Hence, cost-effectiveness is an important driver for hybrid optical switch designs requiring only a limited amount of expensive fast switching components. As our results show (e.g. Section IV-D), even a minimal amount of fast switching fabrics can achieve considerable improvements in network performance.

Figure 2 shows a generic design of a hybrid optical switch, which is composed of two separate switching fabrics. The signaling plane, depicted out-of-band in the figure, informs the Control Unit of the imminent arrival of a data burst. If optical signaling is used, additional O/E and E/O conversions are necessary before and after the Control Unit.

A final note is related to the practical realization of the switch, where several architectural choices remain an open research challenge. For instance, a sequential design (where the fast switching fabric is cascaded behind the slow fabric) such as the one presented in [8], allows reconfiguration of the fast wavelengths, at the expense of an increase in dimensionality of the slow switch. The design depicted in Figure 2 places the two switching fabrics in parallel, and as such loses its reconfigurability for a slightly smaller slow switching matrix.

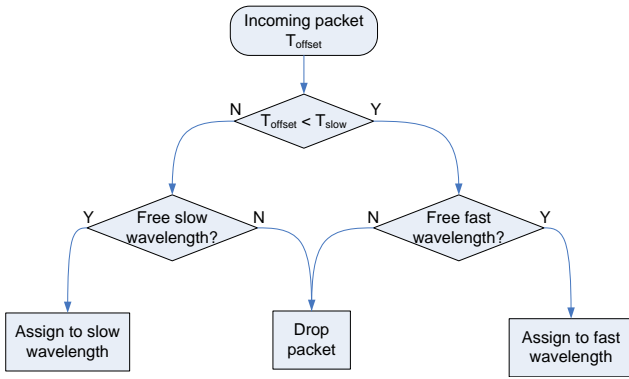


Fig. 3. *Simple* scheduler for hybrid optical switches




Fig. 4. *Greedy* scheduler for hybrid optical switches

III. SCHEDULING ALGORITHMS

In case hybrid switching is used, an additional algorithm is required to map generated traffic on either slow or fast wavelengths. This algorithm will be executed at the network's edge, thus before entering the all-optical data transport network. We propose two algorithms: *simple* (Figure 3) and *greedy* (Figure 4) scheduling. The *simple* scheduling algorithm works by looking up the requested switching time of the burst, and if it is possible to switch slow (i.e. offset between header and data (T_{offset}) is larger than the switching time of the slow switch T_{slow}), the burst is assigned to a slow wavelength in case one is free, otherwise the burst is dropped. In case fast switching is required, a fast wavelength is used if one is free. *Greedy* scheduling also allows slow packets on the fast wavelengths in case no slow wavelength is available. It uses available bandwidth more aggressively, at the risk of assigning valuable fast wavelengths to slow data bursts. The other option, where fast packets are allowed on the slow wavelengths was not implemented because of its obvious non-optimal use of wavelength capacity; fast packets on the slow wavelengths cannot be switched in most cases, and additionally consume capacity required for slow packets.

To evaluate the influence of the scheduling algorithm on the traffic offered to the switch, we implemented a scenario where generated traffic is composed of 50% fast and 50% slow bursts, and the wavelengths consists of 5 slow and 5 fast wavelengths.




Fig. 5. Influence on generated traffic for different scheduling algorithms

The resulting Figure 5 shows the load generated at the switch itself, i.e. after scheduling is performed. First, we see that even without scheduling it is not possible to generate 100% load, since the arrivals are generated independently with both interarrival times and data lengths as exponential distributions. Secondly, as expected, *greedy* scheduling generates higher loads at the switch than *simple* scheduling.

IV. SINGLE NODE SIMULATIONS

The following sections present simulation results of a single hybrid optical switch, which consists of 2 input and 2 output fibers, with each fiber carrying 10 wavelengths. No wavelength conversion is present in any of the switch designs. Unless noted otherwise, the hybrid switch supports 5 slow and 5 fast wavelengths, and each incoming data burst has a 50% probability of choosing the first output fiber (see the discussion on traffic directionality in Section IV-B for more information). The bandwidth of each wavelength is 10 Gbps, and traffic is generated according to a Poisson process with an average interarrival time of 15 ms. Data lengths follow an exponential distribution, with a varying average to establish the switch's load. The offset times between control packet and data are modeled as a 2-phase hyperexponential distribution, with probability density function (pdf) $f = p_{slow} \cdot f_{slow} + p_{fast} \cdot f_{fast}$. Unless otherwise noted, $p_{slow} = .8$ and $p_{fast} = .2$, and the pdf of the slow (resp. fast) traffic is an exponential distribution with average 100 ms (resp. 10 ns). The slow switching fabric has a speed $T_{switch} = 1$ ms, while the fast switching speed is $T_{switch} = 1$ ns. These values are representative for a MEMS-based (resp. SOA-based) switch. For both slow and fast-only designs, a generated data burst is mapped on the first available wavelength; if no free wavelength is found, the burst is discarded. For the hybrid design, the scheduling algorithms as discussed in Section III are used.

A. Comparison between different switch architectures

Figure 6 shows the acceptance probability for different switch designs. For all architectures, the generated load consists of 50% fast and 50% slow traffic (i.e. $p_{slow} = p_{fast} =$




Fig. 6. Architectural comparison: acceptance probability vs generated load




Fig. 8. Influence of directionality of traffic




Fig. 7. Architectural comparison: acceptance probability vs offered load




Fig. 9. Influence of fraction of fast/slow traffic (before scheduling)

.5). For lower loads, the hybrid designs have the best performance, while higher loads allow the fast-only design to have the highest acceptance probability. Also note that *greedy* scheduling always performs at least as good as the *simple* scheduling approach. The reason for the behaviour of fast-only vs hybrid can be seen in Figure 7, which shows the acceptance probability of bursts that are actually offered to the switch (i.e. after the scheduling process). The hybrid switch achieves the highest performance for all offered loads. This implies that for lower generated loads, the scheduling algorithm does not drop a significant amount of packets, but the assignment to different wavelengths allow the hybrid switch to have better performance. For higher loads, packets are dropped by the scheduler before entering the switch, and this amounts to more dropped packets than the fast-only drops internally.

B. Influence of directionality of traffic

Previously we have assumed that traffic has a 50% probability of choosing each output fiber. Now we investigate other options, i.e. where traffic has a 10%, 30% or 50% probability of choosing the first output fiber (with corresponding probabilities for the other output). The generated network load is

composed of 20% fast traffic and 80% slow. Figure 8 shows the details of this traffic directionality for all switch designs (for clarity, hybrid is shown only for *simple* scheduling). Clearly, having a high probability of choosing the same output port for two consecutive burst (case 10% - 90%) results in the highest acceptance probability. For this specific case, the slow-only is even able to outperform the hybrid approach, although when switching gains importance (i.e. directionality values closer to 50%) this advantage quickly disappears. In a final remark we confirm that, although not shown, *greedy* scheduling outperforms *simple* scheduling consistently as before.

C. Influence of fraction of fast/slow traffic

Previously we assumed generated traffic consists of 50% slow and 50% fast traffic. We now investigate the effects of having more (80%) or less (20%) fast traffic. The hybrid switch consists of 5 slow wavelengths and 5 fast wavelength per fiber. Figures 9 and 10 shows the expected high sensitivity to fast traffic of the slow-only switch design. The fast-only design is completely insensitive to the varying fraction of fast traffic. The hybrid design shows the highest performance for the case where the number of fast/slow wavelengths

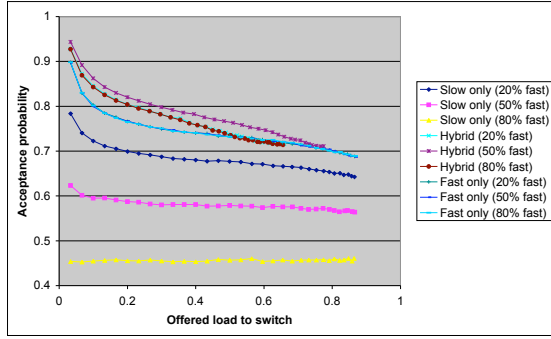


Fig. 10. Influence of fraction of fast/slow traffic (after scheduling)

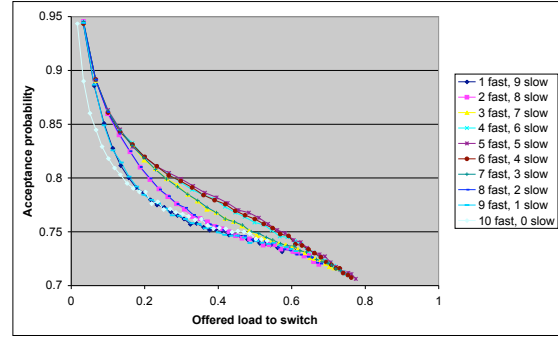


Fig. 12. Influence of fraction fast/slow wavelengths (after *simple* scheduling)

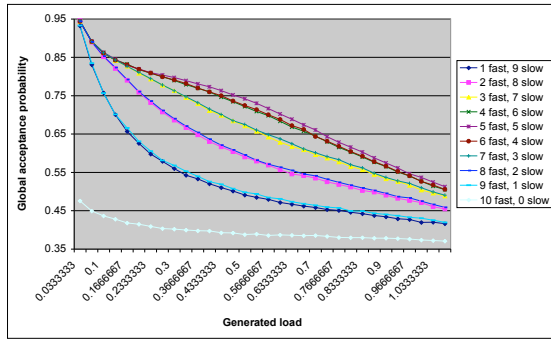


Fig. 11. Influence of fraction fast/slow wavelengths (before *simple* scheduling)




Fig. 13. Influence of fraction fast/slow wavelengths (before *greedy* scheduling)

corresponds to the fraction of fast/slow traffic. Also for the hybrid design, the case with the lowest fraction of fast traffic (20%) performs slightly better than the one with the highest fraction (80%). Finally note that for low loads, the hybrid scheme outperforms even the fast-only switch. The slow-only architecture seems only viable for low fractions of fast traffic, since only then there are (high) load regions where the hybrid does not perform better.

D. Influence of fraction fast/slow wavelengths

In the following series of experiments, we show the influence of the ratio of fast/slow wavelengths, for a fixed ratio of fast/slow traffic. This latter ratio is held constant throughout the following simulations, and equal to .5. Obviously these results show the performance of the hybrid switch design, and we discuss the performance of the *simple* scheduling separately from the *greedy* scheduling.

1) *Simple scheduling*: The best performance is achieved by the hybrid switch with a number of fast/slow ports that corresponds to the fast/slow traffic ratio. We also see that cases for corresponding fast/slow port counts (e.g. 8 fast, 2 slow and

2 fast, 8 slow) show very similar behaviour. In these cases, the design with the highest number of fast ports has slightly better performance. The acceptance probability of packets which are actually offered to the switch shows the best performance for the combination of fast/slow ports which corresponds to the ratio of fast and slow traffic (i.e. 50% fast, 50% slow). Performance degrades as the number of fast/slow ports differs more from the ratio of fast and slow traffic.

2) *Greedy scheduling*: The *greedy* scheduling method clearly makes better use of the available capacity, by allowing slow bursts access to the fast wavelengths. In contrast to the *simple* scheduling algorithm, the corresponding fast/slow wavelength counts (e.g. 8 fast, 2 slow and 2 fast, 8 slow) do not have similar performance. Instead, higher number of fast wavelength counts are used more effectively by the *greedy* scheduling and ultimately lead to higher performance.

V. CONCLUSION

In this paper we presented a generic model for a hybrid optical switch, to support all-optical switching in future transport networks. Through simulation analysis, we showed the




Fig. 14. Influence of fraction fast/slow wavelengths (after *greedy* scheduling)

possible performance improvements of the hybrid switch over a wide range of traffic and switch parameters, in comparison to more traditional single-technology switch designs. Also, we introduced two scheduling algorithms required for integrating hybrid switches in a network, and showed their influence on the traffic offered to the switch and its respective effect on the switch performance.

ACKNOWLEDGMENT

This work was partially supported by the FP6 IP Phosphorus project, and the FP6 NoE ePhoton/ONE+ project. M. De Leenheer is funded by the IWT through a Ph.D. grant, and C. Develder is supported by the FWO through a post-doc grant.

REFERENCES

- [1] M. De Leenheer, P. Thysebaert, B. Volckaert, F. De Turck, B. Dhoedt, P. Demeester, D. Simeonidou, R. Nejabati, G. Zervas, D. Klonidis, and M.J. O'Mahony, *A View on Enabling Consumer Oriented Grids through Optical Burst Switching*, IEEE Communications Magazine, 44(3):124-131, Mar 2006
- [2] H.L. Vu, A. Zalesky, E.W.M. Wong, Z. Rosberg, S.M.H. Bilgrami, M. Zukerman, and R.S. Tucker, *Scalable Performance Evaluation of a Hybrid Optical Switch*, Journal of Lightwave Technology, 23(10):2961-2973, Oct 2005
- [3] C.T. Chou, F. Safaei, P. Boustead, and I. Ouveysi, *A Hybrid Optical Network Architecture Consisting of Optical Cross Connects and Optical Burst Switches*, Proc. 12th Int. Conf. on Computer Communications and Networks (ICCCN), pp. 53-58, Oct 2003
- [4] M. Zukerman, E.W.M. Wong, Z. Rosberg, G.M. Lee, and H.L. Vu, *On Teletraffic Applications to OBS*, IEEE Communications Letters, 8(2):116-118, Feb 2004
- [5] E.W.M. Wong, and M. Zukerman, *Analysis of an Optical Hybrid Switch*, IEEE Communications Letters, 10(2): 108-110, Feb 2006
- [6] B. Chen, and J. Wang, *Hybrid Switching and P-Routing for Optical Burst Switching Networks*, IEEE Journal on Selected Areas in Communications, 21(7):1071-1080, Sep 2003
- [7] C. Xin, C. Qiao, Y. Ye, and S. Dixit, *A Hybrid Optical Switching Approach*, Proc. IEEE Globecom, pp. 3808-3812, San Francisco, CA, USA, Dec 2003
- [8] D. Simeonidou, G. Zervas, and R. Nejabati, *Design Considerations for Photonic Routers Supporting Application-driven Bandwidth Reservations at Sub-wavelength Granularity*, Proc. Int. Workshop on Optical Burst/Package Switching (WOBS), San Jose, CA, USA, Oct 2006
- [9] S. J. Ben Yoo, *Optical Packet and Burst Switching Technologies for the Future Photonic Internet*, Journal of Lightwave Technology, 24(12):4468-4492, Dec 2006