

ROUTING IN AN AWG BASED OPTICAL PACKET SWITCH

Jan Cheyns, Chris Develder, Erik Van Breusegem, Ann Ackaert, Mario Pickavet, Piet Demeester

IBCN group, Department of Information Technology

Ghent University, St.-Pietersnieuwstraat 41, 9000 Ghent, Belgium

Tel. +32 9 267 35 92, Fax. +32 9 267 35 99

{jcheyns, cdvelder, evbreuse, aackaert, pickavet, demeester}@intec.rug.ac.be

Index Terms:

Optical Router, Arrayed Waveguide Grating, Internal blocking

Abstract

For the next generation of the optical internet, focus is now moving from circuit switched networks, which occupy a wavelength continuously regardless of the demand at that time, towards optical packet/burst switching. By only occupying a wavelength when data is to be transmitted, a more efficient utilisation of bandwidth in optical fibres is strived for. As bandwidth in fibres keeps increasing, the bottleneck of the optical network is now moving towards the switching node, since evolution of electronic routers cannot follow the speed of bandwidth increase. Thus a key component in these novel networks is the optical node. Through this node we want to switch traffic very fast and reliable, preferably transparent. Lack of efficient and practically realisable optical memory however makes migration from electronic routers to optical routers a non-straightforward transition. In most optical nodes payload traffic can be switched transparently, whilst control information (e.g. in a header, on a control channel) is still converted to the electronic domain in every node, since optical processing is far from mature. In this paper we present a possible architecture for such a node, combining Array Waveguide Gratings and all-optical tuneable wavelength converters. The concept of this switch is explained and the node is evaluated in terms of loss rate. We will see that an inherent problem of this switch is its internal blocking. This drawback can be greatly overcome by using an intelligent and efficient wavelength assignment algorithm within the node. Simulation of slotted operation will give some numerical results.

1. Introduction

In the past, optical networking experienced a first great boom with the introduction of optical fibre and a second time when optical amplifiers (EDFA) allowed the deployment of (D)WDM. In these optical networks, bandwidth is usually assigned (semi-)statically on a per wavelength basis. Design [1] and survivability [2] of these networks were and are still a field of great interest. However the coarse granularity results in poor bandwidth efficiency. Several techniques are now under study to improve the bandwidth utilisation: Optical Packet Switching (OPS) [3]-[4], Optical Burst Switching (OBS) [5]-[6], Wavelength Routed OBS [7], etc. All of these techniques have the common property that a wavelength is only occupied for limited time, benefiting from statistical multiplexing directly at the optical layer.

The driving force in the success of optical communication up to now was concentrated around the continuous increase of available bandwidth, through different ways: a) hardware evolutions, such as fibre improvements (lower loss, lower dispersion,...) and optical amplifiers in a broad wavelength range b) the use of time and wavelength division multiplexing. In turn the hardware improvements cleared the path for increasing bitrates per wavelengths. Further evolution of optical networks will however need to focus on other parts of the network, the switching nodes. As bandwidth keeps increasing, the use of switches that convert the optical input signal to the electronic domain, perform the switching functionality electronically, and then convert this signal back to the optical domain will soon become very complicated and expensive. This evolution of implementing switching functionality in the optical domain can be seen in circuit switched optical networks [8], where its benefits can help survivability [9]. It seems obvious that this trend will, in the end, lead to optical packet switched networks. For both techniques optically transparent switches are preferred, allowing smooth adaptation to any future changes in bitrate, data format, and so on. Clearly, circuit and packet switched networking also show some important differences in the requirements for switching nodes. The most important issue is the switching time. As in a circuit switched network a wavelength path is set up for a relatively long period, a slow switch can suffice. However, if the holding time of the wavelength decreases, this switching overhead time can lead to inefficiencies. Hence fast switches are required to make packet switching profitable.

This paper will describe a design for such a fast switch for optical packet-based networks, based upon an Arrayed Waveguide Grating (AWG). After a short general introduction on optical switching nodes in Section 2, the

specific AWG node will be described in more detail in Section 3. We will continue in Section 4 by addressing the wavelength assignment problem: how to choose the wavelength to forward incoming packets to? Section 5 will present a performance evaluation of the AWG node, and how to choose a suitable configuration in order to reduce packet loss. The paper will be concluded in Section 6.

2. Generic Node Structure

In Fig. 1, a generic node structure is displayed, with some terminology that will be used throughout this paper. Clearly, three different parts can be seen: the input interface, the switching element and the output interface.

2.1. Input Interface

Several functions are performed at the input interface. If the network is run in slotted operation, the nodes will synchronously switch packets from all their input ports to the outgoing ports. Therefore, one of the first things to do, is synchronisation between the different wavelength channels [10]. Indeed, packets arriving on different wavelengths will not necessarily all be properly aligned in the time domain, and this will certainly be the case for different input fibres. Thermal effects, as well as chromatic dispersion, play an important role here. Getting all packets to start at the same time is done by synchronisation, which can be achieved through e.g. switched FDLs. Of course for asynchronous networking techniques no synchronisation between the different channels is necessary, so the input interface for an asynchronously operated network will be simpler.

Another key action realized by the input interface, is header extraction. If the header is carried on the same wavelength as the data (in-band), extraction can be accomplished by splitting off a small part of the optical power (typically 10%), which is then opto-electronically converted into electronic information. Other techniques use a dedicated wavelength to carry the headers (out-of-band), and in this case extraction can be done by opto-electronic conversion of this wavelength. An alternative for this dedicated control wavelength is the use of an out-of-band control channel per data wavelength, such as Sub Carrier Multiplexing [4], (Differential) Phase/Frequency Shift Keying [11],.... Remark that, regardless which technique is used, the header is always converted into the electronic domain—the reason will become clear in the next paragraph.

2.2.Switching Matrix

This is the core element of the switching node: here the actual switching from one port to another is performed. The switch controller configures the switching matrix. Based on the information in the (now electronically available) headers, a decision is made. Here we see the main reason why the header is still converted into the electronic domain. The header can undergo complex processing in the switch controller, using the mature electronic technology. These complex operations will be the main causes of delay within a switching node. This processing delay calls for a similar delay of the payload information, as we need to wait for the result of the header processing to know which actions have to be taken to correctly switch the data. This delay can be implemented using switched Fibre Delay Lines [12], however an elegant proposition to avoid their use is the JET technique in OBS, where the header information is simply sent ahead in time in respect to the payload [5]. Electronic header processing will probably still be used for quite some time, as pure optical processing [12] is still a very immature (although fast evolving) technology, which makes it practically unavailable at this moment. This part of the switch is the most critical for switching speed. The time between the controller having calculated a new configuration for the switching matrix and the moment when the switch is stable in the new state must be small. So once the controller has passed on the correct settings of all necessary parameters, the reaction time for the switching mechanism has to be as small as possible.

Crucial for this element is also its transparency, as was explained above. The switch should just guide the light containing the payload information from input to output, without using any information on bitrate, data coding technique...

2.3.Output Interface

The output interface's main function is to rewrite header information for the packet. Notice that the content of the header is very likely to be different before and after the switch (e.g. label swapping operations), so the old header will have to be replaced with a new one. Additionally, the output stage is used to "clean up" the optical signals. Signals on different output ports can have different power levels, which is solved by equalisation. The power at the output of the switching element can also be (too) low, which is overcome by using a regenerator. Moreover, if a 2R all-optical regenerator is used, the Signal-to-Noise Ratio (SNR) can also be sufficiently improved [14]. Again all elements in this stage should preferably be transparent.

3. An AWG based node

In this section we will introduce an AWG based switching node. Our attention will go to switching matrix and output interface. The input interface is not within the scope of this paper.

3.1. An AWG based Switching Matrix

3.1.1. The Array Waveguide Grating

An Array Waveguide Grating (AWG) is a diffractive element, and as such it doesn't contain any active devices, making it a very reliable component. It is sometimes also referred to as a PHASed ARray (phasar). We will shortly describe the functionality of an AWG, without giving any physical details (these can be found in e.g. [15]).

The dashed box in Fig. 2 shows a black box model of an 8x8 AWG. Depending on the wavelength at input port i , the light comes out at output port j . This allows us to write out a simple matrix M , where the input ports are written vertically (index i) and the output ports horizontally (index j). An element M_{ji} in this matrix is the wavelength, which will come out on port j , if it entered the device at input port i . All these wavelengths are equally spaced by $\Delta\lambda$. An AWG has a cyclic nature: if $M_{ji}=\lambda_a$, there exists a λ_{cyc} for which $\lambda_b=\lambda_a+k.\lambda_{cyc}$ ($k \in \mathbf{Z}$, λ_{cyc} is called the period) shows the same behaviour as λ_a . In other words, if λ_b enters on port i , it will also exit on port j , just as it was the case for λ_a . It is possible to design the AWG in such a way that for an $N \times N$ AWG, the period is $N \cdot \Delta\lambda$. Putting all of this together, we come to a table like the one shown in Fig. 3. In this figure another property of the AWG can be seen. Let $\lambda_i=\lambda_0+i.\Delta\lambda$. If from a certain input port λ_i goes to output port k , then λ_{i+1} will go to output port $(k+1) \bmod N$. This is a physical constraint that has an impact on the available configuration options, since it limits the degree of freedom to a large extent. We will come back on this issue in section 3.2.

As an example two configurations of the AWG are shown on both Fig. 2 and Fig. 3, in order to clarify this notation further. The signal at input port 4 can be sent to output port 4 if it were on λ_0 , (dotted line) while it would end up at output port 5 if the wavelength is λ_1 (full line). For clarity, during the rest of our qualitative discussions, we will continue working with this example of an 8x8 AWG, and a switch based upon it. Generalisation to an AWG (and switch) with higher dimensions, such as the one used to produce the performance results, is straightforward.

3.1.2. Tuneable Wavelength Converters

With the necessary notations introduced, we will now elaborate on the node concept itself. As can be seen in Fig. 3, the signal on input port 4 can be sent to output port 4 if it is on λ_0 and to output port 5, if it is on λ_1 . So if we can manipulate the wavelength of a signal, this feature can be used in a switching matrix. In Fig. 2, the switch concept is further explained with the example of two input fibres and two output fibres (in this study we will only consider symmetric switches, i.e. the number of incoming and outgoing fibres are the same), each carrying 4 wavelengths. In a first step the signals are demultiplexed into separate wavelength channels. Each of these input channels is connected to an input port of the AWG via an all-optical Tuneable Wavelength Converter (TWC) [16]. By setting the TWC we can control at which output port the signal comes out. Note that it is allowed for a certain output port to contain more than one signal simultaneously, provided that these signals are on different wavelengths (e.g. in the table of Fig. 3, both input ports 0 and 5 can reach output port 0 at the same time, as they will do so using different wavelengths, λ_0 resp. λ_3). We will come back to this feature in Section 5.1. Now all that has to be done is to combine the output ports into the different output fibres (two in our example). This is not a trivial matter, as will become clear from the following discussion.

3.2. Output interface

First, if we look back at Fig. 2 and suppose that every input fibre has the same number of wavelengths (4 in the example), it is clear that only a limited number of wavelengths are in use inside the AWG (instead of all N). Consequentially the table is now no longer completely filled. For example, the wavelengths λ_4 to λ_7 (indicated in grey) are no longer present, only the entries in black remain in the table. In Fig. 4 some tables for possible ways of connecting the output ports of the AWG-switch are shown. As discussed in Section 3.1.1, the rows of the table correspond to the output ports. Since certain output ports are joined together in one output fibre, we can group the corresponding rows of the table (output ports) going to the same output fibre together and create one table per outgoing fibre. As an example, in Fig. 4(a), the table for O_0 is shown for the configuration of Fig. 2. One last transformation will now result in a more compact table notation, with no more empty cells. We let the rows λ_k represent the wavelengths, and the columns the output ports o . Now the element at position (λ_k, o) is the input port that will be routed to output port o by converting it to λ_k , as shown in Fig. 4(b) and Fig. 4(c). Due to physical

properties of the AWG, as already mentioned in 3.1.1, we can see that the rows of the table are not completely independent. To be more specific, once the first row of the table is set, the rest of the rows automatically follow from this. Note that a table of this form is valid for one output fibre. So the node is specified by as much tables as there are output fibres. It should also be clear that we are not concerned about the exact output port the incoming packet is switched to, the only concern is for the packet to arrive in the correct output fibre.

We wish to ask the reader's attention for the terminology used here, and further in the discussion. When input/output *port* is mentioned, this is the input/output waveguide of the AWG. These are not to be confused with the input and output fibres. So in fact an input port corresponds to a demultiplexed signal of an input fibre, it is thus a wavelength channel. An output port however can carry more than one wavelength, as already discussed above. So we can state that an output port carries a number of wavelengths of its corresponding output fibre.

4. The Switch Controller: Wavelength assignment

Now that the concept of the node is clear, we can move on to the operation of the node. At a certain moment in time a header will arrive, indicating what to do with the associated payload. First, the header is read and processed to determine the output fibre. Using the above-developed table for that specific output fibre, we need to find out to which wavelength the signal has to be converted by the TWC in order to allow switching to that correct output fibre. Thus, a wavelength assignment algorithm is needed. We will try to find an optimal and efficient algorithm, considering slotted operation, with fixed size packets. It should be stressed here that for both the following algorithms the wavelength assignment is done for each output fibre independently from each other, allowing for parallelism in a practical implementation. This assumes that each packet has a unique output fibre (e.g. determined by a table lookup, if the packet headers would carry MPLS-like information). This assumption excludes routing mechanisms where a packet could choose from several different output fibres, we do not consider this case here.

We will start by describing a heuristic approach in 4.1, followed by a solution based on a known problem in graph theory in 4.2.

4.1. Heuristic

Several heuristics were studied, of which we will present one here: LUW-LFP: Least Used Wavelength – Least Flexible Port. Since slotted operation is considered, we know which input ports are active at the start time of every slot. We now have a set of active input ports that need to get a wavelength assigned to them. Every input port of the AWG can be converted to a number of wavelengths that can perform switching to the correct output fibre, the issue is now to make an intelligent decision in order to get the highest possible throughput. First we will illustrate the two driving ideas behind this heuristic:

- The first idea (LUW) is that if a wavelength can potentially be used for a large number of input ports, we will wait as long as possible to assign this wavelength to a certain input port. Indeed, when choosing a wavelength that is shared by many other ports, this could affect all of these ports; by assigning a wavelength that can only be used for a small number of ports, we ensure that only a small number of ports now have one less possible wavelength. In other words wavelengths that can only be used for a small number of active input ports will be chosen first. An example might clarify this even more. Suppose input port 1 can choose a wavelength from the set $\{\lambda_1, \lambda_2\}$, input port 2 from $\{\lambda_1, \lambda_2\}$ and input port 3 from $\{\lambda_1, \lambda_3\}$. Starting by assigning λ_1 to an input port is no good idea since assigning it to input port three results in the unnecessary blocking of either input port 1 or 2. On the other hand starting with λ_3 is a good way of working, as it has no consequence for the other input ports.
- The other idea (LFP) is that if an input port has only a small number of possible wavelengths, we prefer to assign a wavelength to this input port first. For an input port with lots of possible wavelengths the assignment of the wavelength is less critical, so it can be postponed. Again we clarify this with an example. Suppose input port 1 can be routed using a wavelength from the set $\{\lambda_0, \lambda_1, \lambda_3\}$, while input port 2 can only use λ_2 . It is obvious that we assign λ_2 to input port 2 first. If we were to make a choice for input port 1 first there is a risk that λ_2 would be chosen, resulting in an unnecessary blocking of port 2.

Now that the ideas behind the algorithm are explained, the general description follows. Each of the input ports get a wavelength assigned to them, one after the other. The order in which the input ports are being dealt with

is part of the algorithm. LFP will determine the order in which ports get a wavelength assigned to them, while LUW will result in the wavelength assignment itself.

Each of the active input ports i have a set of possible free wavelengths S_i to arrive in its destination fibre — a wavelength is free as long as no input port has this wavelength assigned to it. For each free wavelength λ_j we can calculate N_j , the number of input ports that can use λ_j to arrive at the correct output fibre. We now select the free wavelengths with the lowest N_j , and call this set the Least Used Wavelengths set (LUW set). For each of the free wavelengths λ_k in this LUW set we create an input port set I_k . This input port set contains the active ports that can be correctly routed by the free wavelength λ_k and that haven't had a wavelength assigned to them yet. For all input ports p in this set I_k , we then calculate the respective flexibility f , which is simply the number of free wavelengths that can route the input port p correctly. Now, the port with lowest flexibility is chosen, and where more than one choice is possible, a random decision is made. The corresponding free wavelength from the LUW set is then assigned to this input port—this wavelength is now no longer free. Again if more than one wavelength is possible, a random decision is made. Now the algorithm can restart with the updated knowledge (the number of ports that have had a wavelength assigned to them is increased by 1), in order to try and assign a wavelength to every active input port.

4.2. Maximum Matching

In this section we will first show how the problem of assigning the wavelengths in the AWG can be converted to a graph problem. Once this graph representation is available, a known optimal and efficient algorithm can be used to do the assignment.

A graph can be drawn for every output fibre, just like there was a table per output fibre. Converting the problem to a graph is quite simple. We can consider the wavelengths to be vertices, just as the input ports. The edges of the graph are then possible wavelength assignments, which can be read from the tables. In Fig. 5, we show this graph for the table in Fig. 4(b).

A known problem in graph theory is that of matchings, more specific Maximum Matchings. A matching of graph G is a subgraph of G that consists of pairwise non-adjacent edges. For a bipartite graph (vertices can be split up in two classes, vertices in the same class have no edges between them), the interpretation of a matching is pretty straightforward: a vertex of the first class (A) is matched to a vertex of the second class (B). A clarifying example

might be A and B being men and women, edges meaning two persons (a male and a female) knowing each other. Finding the largest numbers of possible marriages is then a maximum matching problem. In our case the two classes are active input ports on the one hand, and wavelengths on the other hand. If two vertices are matched, we can consider the wavelength assigned to the input port. The definition of a matching that the two edges have to be non-adjacent is indeed a valid condition for our problem: an input port can only have one wavelength assigned to it and a wavelength (in a single output fibre) can only be assigned to one input port. The graph constructed above is indeed bipartite, as no edges exist between two nodes, nor are there edges between two wavelengths. A maximum matching is the matching with the highest number of edges. In our case this corresponds to the maximum number of input ports having a wavelength assigned to it, so the highest switch throughput.

For Maximum Matching an efficient algorithm exists [17]. The complexity of this algorithm is of the order $O(pq)$, p being the order and q the size of the graph. In our case this leads to a complexity of $O(W^3)$, W being the number of wavelengths per fibre.

4.3.Evaluation

In this section we will briefly compare the two methods explained above. As a reference we will also add results for a completely random wavelength assignment: in random order, each input port in turn is assigned a wavelength from its set of possible wavelengths (reduced with previously assigned wavelengths) using a uniform distribution for these wavelengths. In Fig. 6 and Fig. 7 we show some results for a 3 fibre (3 input and 3 output fibres) system for a varying number of wavelengths per fibre.

In Fig. 6, a comparison in blocking probability is illustrated, for a Poisson traffic load of 0.8. Clearly both the heuristic and the maximum matching have about the same performance, and are significant improvements over the totally random wavelength assignment. Fig. 7 shows the time needed to perform the simulations for increasing number of wavelengths. We see that although the heuristic gives good results, it is a lot slower. The maximum matching algorithm on the other hand is only a little bit slower than the random case. So the improved performance over the random assignment comes only at a marginal cost. In the following performance study of the switching node, we will continue using this Maximum Matching algorithm.

5. Node Performance

5.1. Internal blocking

Have a closer look at the table in Fig. 4(b), and suppose we want to switch input ports 1,3 and 5 to output fibre O_0 . The table learns that only converting these input ports to either λ_1 or λ_3 can do this. So we have a problem, since three ports are contending for only two wavelengths, resulting in the loss of one of the input ports. Since this loss is not due to contention (i.e. there are no more packets destined for the output fibre than this fibre carries wavelengths), this is a problem of internal blocking of the switch. The question now rises whether this internal blocking is the same for all output configurations, or not. In Fig. 4(c) another table is shown, this time the first 4 output ports are put grouped in output fibre O_0 : input port 4 isn't even in the table, meaning that a signal on input port 4 can never be switched to output fibre O_0 . So it is clear that the way the output ports are combined into an output fibre is very important, and has a severe impact on the performance of the node.

Since internal blocking occurs due to the fact that it is not allowed for two packets to be on the same wavelength in the same output fibre, it might seem a good idea (from a performance point of view) to add Fixed Output Wavelength Converters at each output port of the AWG, before multiplexing them into the output fibre. However this does not solve the problem, on the contrary, blocking gets worse. Indeed, if no output converters are used, it is possible for an output port to contain more than one wavelength, as was discussed in 3.2. This is no longer possible when each output port is equipped with a wavelength converter. This loss in flexibility can not be overcome by the flexibility gain due to the extra wavelength converters. As an example have a look at Fig. 4 b, and suppose ports 1,2,3 and 4 need to go to output fibre O_0 . If no converters are present, this can be realised by making the following wavelength assignments to these input ports: input port 1: λ_1 , input port 2: λ_2 , input port 3: λ_3 and input port 4: λ_0 , so output port 4 carries two signals but on a different wavelength. However with converters one of the input ports will suffer loss, since one output port can only carry one signal, and there are only three output ports possible for the 4 active input ports.

5.2. Simulation of slotted operation

In order to get an idea of the performance differences for different table configurations, we ran simulations for each of these configurations, using the above-mentioned Maximum Matching Wavelength assignment algorithm. Two types of traffic were generated. The destination distribution is assumed uniform, so the probability for a packet going to an output fibre is the same for all output fibres (namely $1/F$, if F denotes the number of output fibres). The arrival distribution of packets was assumed to be Poisson. Secondly, since we are confronted with internal blocking, we made a distinction between contention-free and normal traffic. In normal traffic it is possible that the number of packets destined to the same output fibre is higher than the number of available wavelengths, this is called contention. In contention-free traffic the maximum number of packets going to the same output fibre was limited to the number of wavelengths in the system, so contention can not occur. For a non-blocking switch loss probability with contention-free traffic would be zero.

5.3. Tables

We will make another hypothesis, namely that the number of wavelengths (on a single fibre) should be a multiple of the number of fibres. If this condition is not fulfilled, it will result in unfairness between the different input ports, one port having a higher loss probability than another.

The description of some table configurations that were studied follows. Let F denote the number fibres and W the number of wavelengths, with the output ports numbered from 0 to $F \cdot W - 1$.

- **T1:** output fibre i contains ports $i \cdot W + k$, $k=0, \dots, W-1$, this is simply grouping the output ports ordered, like was done in Fig. 4(c). As discussed above this is certainly not a good configuration, as some ports can't be switched to a certain output fibre.
- **T2:** output fibre i contains ports $i \cdot W/F + kW + 1$, with $k = 0 \dots F-1$ and $i = 0 \dots W/F-1$
- **T3:** output fibre i contains ports $i + k \cdot F$ with $k=0 \dots W-1$
- **T4:** is a more irregular variant of T3. We will start by explaining how output fibre 0 is configured, it contains output ports $p \cdot W$ and $2 \cdot F - 1 + k \cdot F + p \cdot W$ with $k = 0 \dots W/F-2$ and $p = 0 \dots F-1$. The other fibres are composed by taking the remaining output ports in the natural order and assigning them to the remaining output fibres in a

round robin way. So if we number the remaining output ports $j=0, 1, \dots, (F-1) \cdot W - 1$. Then fibre i contains the output ports for which $j \bmod (F-1) = i-1$.

5.4. Results

Note that for all of the studies discussed in the sections below, buffering (e.g. via switched FDL's) was not included in the node. The impact of buffers would be to reduce the blocking probability. This blocking could be caused either by contention or by internal blocking. However the most important study in this presentation is the comparison with a non-blocking switch. If the internal blocking increases, this increases the need for optical buffers, which we want to avoid as much as possible.

5.4.1. Different Table Configurations

We will now present and discuss some results. In order to have sufficient distinction between the different table configurations we can no longer work with the simple example of a 2 fibre and 4-wavelength system—indeed, for this simple case, e.g. strategies T2 and T3 coincide. Thus, in order to study a more realistic node configuration, we choose 3 fibres and a 12-wavelength system. First of all we will show the difference between table configurations. This can be seen in Fig. 8 for non-contending traffic. The result is quite remarkable, as several orders of magnitude of difference can be seen. Using the T1 configuration, the loss seems to be independent of the load. This is indeed correct. Let's return to the T1 example with 2 fibres and 4 wavelengths, corresponding to Fig. 4 c, where we see that input port 4 can never reach O_0 , the same goes for input port 0 destined for O_1 . Since all ports are equally likely to carry traffic, and destinations are also assumed uniform, this fact alone leads to a loss of $1/8 \cdot 1/2 + 1/8 \cdot 1/2 = 1/8$. Writing out the complete table for the case of 3 fibres each carrying 12 wavelengths, leads in the same way to input ports 0, 12 and 24 being unable to reach two of the three output fibres, while all other input ports have one output fibre they can't reach. This leads to a loss probability of $3/36 \cdot 2/3 + 33/36 \cdot 1/3 = 13/36$. These calculations are underestimates of the actual loss, however loss is so high due to this effect that extra loss due to contention or internal blocking has no visible effect on the overall performance. The most irregular table T4 results in the lowest blocking probability. Intuitively this might be logical, as in a very regular table structure several ports tend to have the exact same (sub-)set of possible wavelengths, so they behave exactly the same and are thus very

strongly coupled. If this coupling is weaker, better wavelength choices can be made for the different input ports, resulting in a lower loss probability. The coupling is now in fact smeared out more randomly over the input ports.

The same four tables are compared in Fig. 9, but now traffic is allowed to be contending. As was expected the blocking probability is now much higher, especially for higher loads, where contention is more likely to occur. We see that for high loads the blocking probability using T4 and allowing contention, is several orders of magnitude higher than when contention was excluded. This can be interpreted as the internal blocking having almost no impact on the total blocking probability, i.e. loss due to contention is dominating loss due to internal blocking. T2 also shows an overall loss behaviour close to the one of the good table T4, although with non-contending traffic there was a clear performance difference, so we still prefer the T4 configuration. For the tables T1 and T3 the contrary holds. As loss rates are already very high using contention-free traffic, the internal blocking has a severe detrimental effect on the switching node's performance, since it is dominating the overall loss probability.

5.4.2. Comparison with non-blocking switch

As a consequence of the results from the previous paragraph, the following simulations consider only the T4 configuration.

In Fig. 9, showing the blocking probability for a switch with 3 fibres (in and out), each containing 12 wavelengths, allowing contention in the traffic, the behaviour of an internally non-blocking switch was added. It can be very clearly seen that the behaviour of the AWG-based switch combined with the Maximum Matching Wavelength Assignment algorithms lies very close to that of an non-blocking switch. As comparison, the behaviour of the AWG-based switch combined with a random wavelength assignment was also shown on the same graph, indicating again the importance of a good wavelength assignment algorithm.

As a last graph Fig. 10 shows the blocking probability with increasing number of wavelengths. As we expect, a higher number of wavelengths results in a lower loss probability.

6. Conclusion

In conclusion, we have studied an AWG-based switching node. Results of simulations showed that the internal blocking of such a node could be well solved by using an intelligent wavelength assignment algorithm. Such an algorithm was proposed, and verified that it is indeed improving performance. It is based upon the translation of

the wavelength assignment problem to a graph problem, combined with the known algorithm for finding a Maximum Matching. Further the importance of the right configuration of the output was stressed. Again simulation results resulted in the proposition of a possible, well performing architecture. Important parameter in this output port configuration is the non-regular structure of these output port combinations.

Acknowledgements

This work was partly funded by the IST project STOLAS, under contract number IST-2000-28557 and the IWT-GBOU project 010058 “Optical Networking and Node Architectures”.

Both Jan Cheyns and Chris Develder are Research Assistants of the Fund for Scientific Research - Flanders (Belgium) (F.W.O. - Vlaanderen). Erik Van Breusegem would like to thank the IWT for its financial support.

References

- [1] N. Wauters, P. Demeester, *Design of the optical path layer in multiwavelength cross-connected networks*, IEEE J. on Sel. Areas in Comm., Vol. 14, Nr. 5, June 1996, pp. 881-892
- [2] D. Colle, S. De Maesschalck, C. Develder, P. Van Heuven, A. Groebbens, J. Cheyns, I. Lievens, M. Pickavet, P. Lagasse, P. Demeester, *Data-centric optical networks and their survivability*, IEEE J. on Sel. Areas in Comm., Vol 20, Nr. 1, January 2002, pp 6 –20
- [3] P. Gambini, M. Renaud, C. Guillemot, F. Callegati, I. Andonovic, B. Bostica, D. Chiaroni, G. Corazza, S. L. Danielsen, P. Gravey, P. B. Hansen, M. Henry, C. Janz, A. Kloch, R. Krahenbuhl, C. Raffaelli, M. Schilling, A. Talneau, L. Zucchelli, *Transparent optical packet switching: network architecture and demonstrators in the KEOPS project*, IEEE J. on Sel. Areas in Comm., Vol. 16, Nr. 7, Sept. 1998, pp. 1245 –1259
- [4] M. O’Mahony, K. M. Guild, D. Hunter, I. Andanovic, I. White, R. Penty, *An optical packet switched network (WASPNET)-concept and realisation*, Optical Networks Mag., Vol. 2, Nr. 6, November/December 2001, pp. 46-53
- [5] C. Qiao, M. Yoo, *Optical Burst Switching (OBS) - A New Paradigm for an Optical Internet*, J. High Speed Networks (JHSN) on WDM Networks, Vol. 8, No. 1, 1999, pp. 69-84
- [6] M. Yoo, C. Qiao, S. Dixit, *QoS Performance of Optical Burst Switching in IP-over-WDM Networks*, IEEE J. on Sel. Areas in Comm., Vol. 18, Nr. 10, October 2000, pp. 2062-2071

- [7] M. Düser, P. Bayvel, *Performance of a Dynamically Wavelength-Routed Optical Burst Switched Network*, IEEE Phot. Techn. Lett., February 2002, to appear
- [8] Giles C. R., Aksyuk V., Barber B., Ruel R., Stulz L., Bishop D., *A silicon MEMS optical switch attenuator and its use in lightwave subsystems*, IEEE J. of Sel. Topics in Quantumel. Vol. 5, Nr. 1, January-February 1999, pp. 18-25
- [9] S. De Maesschalck, D. Colle, A. Groebbens, C. Develder, I. Lievens, P. Lagasse, M. Pickavet, P. Demeester, F. Saluta, M. Quagliotti, *Intelligent Optical Networking for Multilayer Survivability*, IEEE Commun. Mag., Vol. 40, Nr. 1 January 2002, pp. 42-49
- [10] L. Zuchelli, M. Burzio, P. Gambini, *New solutions for optical packet delineation and synchronisation in optical packet switched networks*, Proc. Europ. Conf. On Opt. Comm. 1996, Oslo, 1996, vol. 3, pp. 301-304
- [11] T. Koonen, G. Morthier, J. Jennen, H. de Waardt, P. Demeester, *Optical Packet Routing in IP-over-WDM networks deploying two-level optical labelling*, Proc. of Europ. Conf. On Opt. Comm. 2001, Amsterdam, Vol. 4, October 2001, pp. 608-609
- [12] S. Yao, B. Mukherjee, S. Dixit, *Advances in Photonic Packet-Switching: An Overview*, IEEE Comm. Mag., Vol. 38, Nr. 2, February 2000, pp. 84-93.
- [13] K. Stubkjaer, *Semiconductor optical amplifier-based all-optical gates for high speed optical processing*, IEEE J. of Sel. Topics in Quantumel. Vol. 6, Nr. 6, November/December 2000, pp. 1428-1435
- [14] D. Cotter, A. D. Ellis, *Asynchronous Digital Regeneration and Networks*, J. Lightwave Techn., Vol. 16, No. 12, December 1998, pp. 2068-2080.
- [15] M.K. Smit, C. van Dam, *Phasar-based WDM devices: principles, design and applications*, IEEE J. Sel. Topics Quant. Electron, Vol. 2, 236-250, 1996
- [16] D. Wolfson, A. Kloch, T. Fjelde, C. Janz, B. Dagens, M. Renaud, *40 Gb/s all-optical wavelength conversion, regeneration and demultiplexing in an SOA-based all-active Mach-Zehnder interferometer*, IEEE Phot. Techn. Lett., Vol. 12, 2000, pp. 332-334
- [17] G. Chartrand, O. P. Oellermann, *Applied and Algorithmic Graph Theory*, Mc Graw-Hill, New York, 1975, pp.161-204

Figure Captions

Fig. 1 A generic node structure

Fig. 2 Schematic overview of the switch build up

Fig. 3 An AWG table representation

Fig. 4 a: table for one specific output fibre (O_0 in Fig. 2)

b: compact notation for the table in a

c: compact notation for another output port configuration (T1)

Fig. 5 Graph representation of the wavelength assignment problem, in order to translate the problem to a Maximum Matching problem

Fig. 6 Performance comparison of the different wavelength assignment algorithms, using a three fibre system, and varying the number of wavelengths per fibre

Fig. 7 Comparison of the time duration of the different algorithms for a three fibre system, varying the number of wavelengths per fibre

Fig. 8 Comparison of different tables, using a 3 fibre 12 wavelength system, with non-contending traffic, a non-blocking switch would have no loss, T4 has best performance

Fig. 9 Comparison of different tables, using a 3 fibre, 12 wavelength system, with contending traffic, T4 performance close to ideal behaviour

Fig. 10 Performance improves for increasing number of wavelengths, using a three fibre system

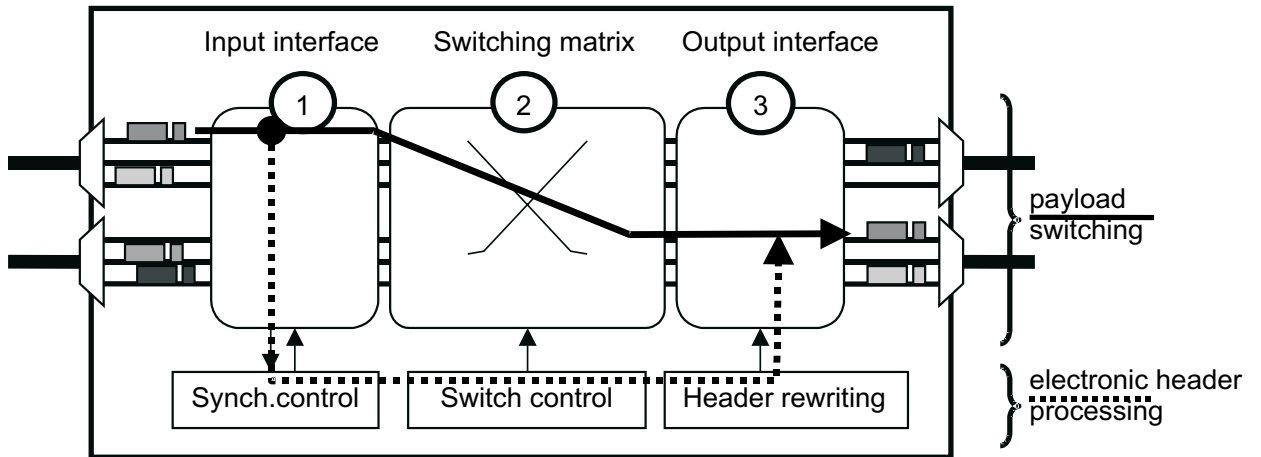


Fig. 1 A generic node structure

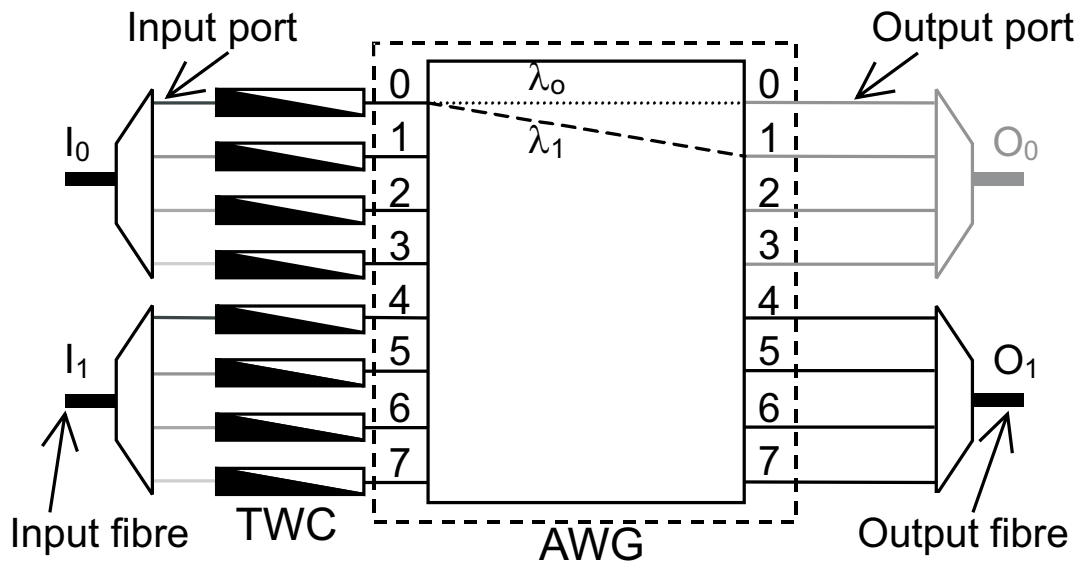


Fig. 2 Schematic overview of the switch build up

		Input ports							
		0	1	2	3	4	5	6	7
Output ports	0	λ_0	λ_7	λ_6	λ_5	λ_4	λ_3	λ_2	λ_1
	1	λ_1	λ_0	λ_7	λ_6	λ_5	λ_4	λ_3	λ_2
	2	λ_2	λ_1	λ_0	λ_7	λ_6	λ_5	λ_4	λ_3
	3	λ_3	λ_2	λ_1	λ_0	λ_7	λ_6	λ_5	λ_4
	4	λ_4	λ_3	λ_2	λ_1	λ_0	λ_7	λ_6	λ_5
	5	λ_5	λ_4	λ_3	λ_2	λ_1	λ_0	λ_7	λ_6
	6	λ_6	λ_5	λ_4	λ_3	λ_2	λ_1	λ_0	λ_7
	7	λ_7	λ_6	λ_5	λ_4	λ_3	λ_2	λ_1	λ_0

Fig. 3 An AWG table representation

a)

		Input ports							
		0	1	2	3	4	5	6	7
Output	0	λ_0					λ_3	λ_2	λ_1
	2	λ_2	λ_1	λ_0					λ_3
	4		λ_3	λ_2	λ_1	λ_0			
	6				λ_3	λ_2	λ_1	λ_0	

b)

		Output ports			
		0	2	4	6
Lambda	λ_0	0	2	4	6
	λ_1	7	1	3	5
	λ_2	6	0	2	4
	λ_3	5	7	1	3

c)

		Output ports			
		0	1	2	3
Lambda	λ_0	0	1	2	3
	λ_1	7	0	1	2
	λ_2	6	7	0	1
	λ_3	5	6	7	0

Fig. 4 a: table for one specific output fibre (O_0 in Fig. 2)

b: compact notation for the table in a

c: compact notation for another output port configuration (T1)




Fig. 5 Graph representation of the wavelength assignment problem, in order to translate the problem to a Maximum Matching problem

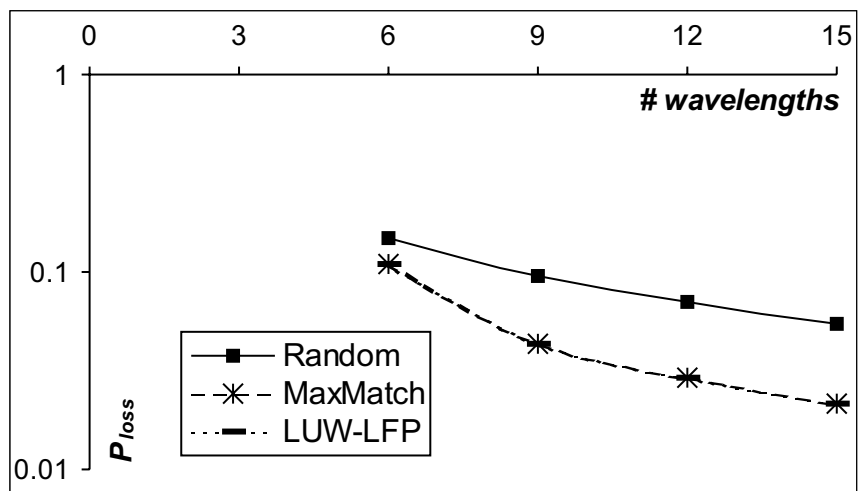


Fig. 6 Performance comparison of the different wavelength assignment algorithms, using a three fibre system, and varying the number of wavelengths per fibre

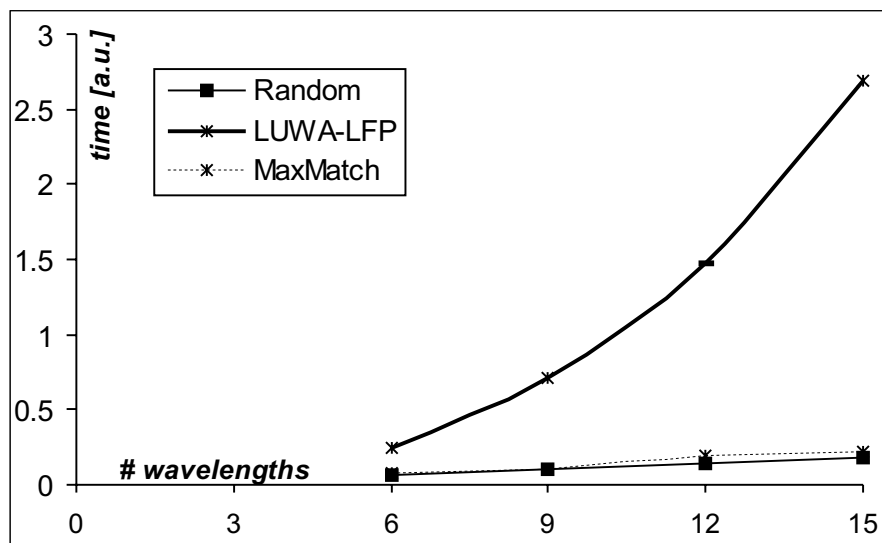


Fig. 7 Comparison of the time duration of the different algorithms for a three fibre system, varying the number of wavelengths per fibre




Fig. 8 Comparison of different tables, using a 3 fibre 12 wavelength system, with non-contending traffic, a non-blocking switch would have no loss, T4 has best performance




Fig. 9 Comparison of different tables, using a 3 fibre, 12 wavelength system, with contending traffic, T4 performance close to ideal behaviour




Fig. 10 Performance improves for increasing number of wavelengths, using a three fibre system