ARCHITECTURES FOR OPS METRO RINGS: COMPARING ACTIVE VERSUS PASSIVE NODES

A dimensioning point of view

Ruth Van Caenegem, Chris Develder, Elise Baert, Didier Colle, Mario Pickavet and Piet Demeester

Dept. of Information technology (INTEC), Ghent University – IMEC, Sint-Pietersnieuwstraat 41, 9000 Gent, Belgium, email: ruth.vancaenegem@intec.ugent.be, tel: +32 9 2649993, fax:+32 9 264 9960

Abstract: Next generation metropolitan area networks (MANs) should provide high bandwidth in a flexible manner: they should efficiently exploit available resources, support multiple traffic types and offer rapid provisioning. Optical Packet Switching (OPS), with its packet-level granularity and hence efficient and flexible bandwidth sharing, fulfils these requirements very well [1]. We discuss the approach and results in comparing three ring architectures (with/without active components allowing for spatial reuse) in terms of resources required.

Key words: WDM, OPS, ILP, Capacity planning

1. INTRODUCTION

The metropolitan area network exchanges information between access networks (interconnected through e.g. an IP router) and the WAN. Consequently, the bandwidth needed in the MAN is much higher than in the access networks, and more fluctuating in time and space than in the WAN. Additionally the MAN has to support heterogeneous services ranging from non time critical (e.g. internet surfing) to quality demanding services (e.g. real time video). In such a dynamic environment packet switching is the most ideal approach. The OPS architecture for a metropolitan area network has to compete with recent, relatively cheap technologies as Ethernet-based optical networks approaching 10 Gbit/sec (standardised in IEEE Std

802.3ae-2002 [4]), or the resilient packet rings (standardisation ongoing IEEE 802.17). Therefore it is appropriate to limit the architectural complexity and hence its cost. Recent examples of OPS MAN architectures propose ring based networks with bufferless nodes [5-10].

In the European DAVID (Data and Voice integration over DWDM) project [2], multiple MAN architectures are compared. Here, we outline the three DAVID metro ring architectures and discuss the two different MAN optical packet add/drop multiplexer (OPADM) designs: a passive one, and an active one.

The goal of this study is to compare the three architectures in terms of amount of resources needed to support a given traffic pattern. The motivation of this study is to find out which architecture is the most cost efficient: the capital expenditure (CAPEX) of a real network after all has a strong relation with the implemented resources.

This paper is organized as follows: in Section 2, we will outline the goal of the benchmarking study we performed. The problem statement, cost measures, as well as the methodology used and the input parameters are summarised in Section 3. The results are discussed in Section 4, and conclusions summarized in Section 5.

2. BACKGROUND

2.1 The DAVID MAN

In the DAVID concept, sketched in Figure 1, the MAN comprises slotted WDM rings collecting traffic from several optical packet add/drop multiplexers (OPADMs). Rings are interconnected by a bufferless hub, which also provides access to a backbone (WAN). The rings constitute a shared medium, requiring a medium access control (MAC) protocol [3] to arbitrate access to the slotted channels. One wavelength, noted λ_c, acts as a dedicated control channel.

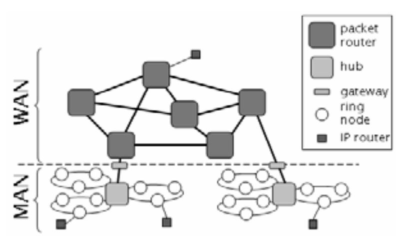

Figure 1. DAVID network architecture

2.2 Node architectures

DAVID proposes two OPADM architectures. The first one is an active node and is depicted in Figure 2 b. The node has an active transit path — hence "active node"—, comprising a waveband demux, with at least for one band a wavelength demux. The Rx and Tx structures realise a tuneable receiver, resp. transmitter, capable of accessing a single wavelength per timeslot. The combination of the Tx, Rx and transit-path structure (excluding waveband demux and mux) form a so-called "babyboard", which is the basic unit that can be added to a Ring Node to increase its capacity. This means that if a Ring Node has to add and/or drop more bandwidth to/from the ring than the bandwidth corresponding to a single wavelength, an extra babyboard needs to be installed (thus resulting in the use of an extra waveband). The active node allows an incoming packet to be erased from the ring, and to replace it with a new one. Because of this erasing capability, there is no need for spectral separation of Rx and Tx signals.

The second one limits the use of advanced optical technologies, choosing commercial and mature ones instead [10]: it uses couplers and off-line filters to minimize physical cascadeability issues. This node (Figure 2 a) has a passive transit path and has at least a single Tx and a single Rx element; it's denoted further as "passive". The granularity with which the node capacity can be increased is a single wavelength, implying the addition of a Tx/Rx for that particular wavelength. Note that for the capacity planning performed here, whether or not these Tx/Rx are tuneable doesn't matter. (We may assume they are fixed and wavelength conversion will only take place at the hub; cf. spectral separation of up- and downstream and thus of Tx and Rx wavelengths.)

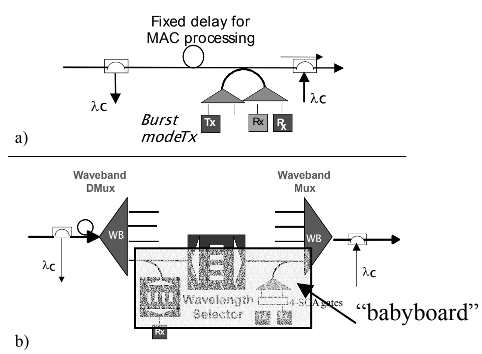

Figure 2. Node architectures a) DOBRN, b) NoSR and SR

2.3 MAN ring architectures

The goal of this paper is to compare the two node architectures used in three ring architectures proposed within the DAVID consortium. The three ring architectures are illustrated in Figure 3, and described below:

- **No spatial reuse (active nodes): NoSR**; this is the original David proposal where every packet sent has to cross the Hub at least once (also all intra-ring traffic). This is a result of a centralised MAC protocol that implies certain traffic streams to cross the ring segment between source and destination nodes twice e.g. B-C in Figure 3. Packets are sent/received on one of the wavelengths in a multislot. A multislot is a set of slots (each contains one packet) transported together in WDM, with a joint header (on a separated control channel).
- **Spatial reuse (active nodes): SR;** where packets do not have to cross the Hub, i.e. intra-ring traffic between nodes belonging to the same ring only has to pass the ring segment between source and destination nodes (assume both nodes use the same waveband). This also allows for spatial reuse: the same wavelength can be re-used, as for D-A and B-C in Figure 3, but implies a more complex distributed MAC protocol.
- **Dual Bus Optical Ring Network (passive nodes): DBORN [10];** A ring is composed of two separated wavebands, where one is used exclusively for upstream “sending”, i.e. putting new packets on the ring, and the other for downstream “receiving”, i.e. taking packets from the ring. This implies all traffic needs to pass through the Hub (wavelength conversion in the Hub).

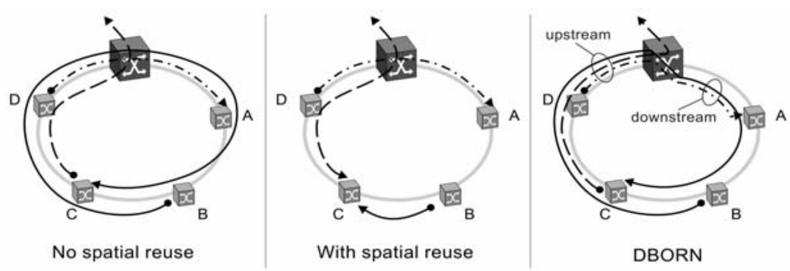

Figure 3. Capacity requirements of the three ring architectures illustrated for three demands.

3. A NETWORK DIMENSIONING POINT OF VIEW

The goal of this study is to compare the three architectures in terms of amount of resources needed to support a given traffic pattern. The motivation of this study is to find out which architecture is the most cost efficient: the capital expenditure (CAPEX) of a real network after all has a strong relation with the implemented resources.

It is clear that from the perspective of the number of Rx/Tx elements alone, DBORN will be cheaper. The only possible "penalty" is that the number of different wavelengths used on a ring will be higher. It can be expected that the SR approach will be the cheapest in terms of number of wavelengths (wavebands) needed on the ring to provide for the demands.

3.1 Problem statement and solution methodology

To accomplish the comparison of active and passive nodes, we solve the following problem:

– **Given:** a set of nodes V (including the Hub) that will form a single MAN (possible multiple rings); a set of (candidate) links between elements v of V; a demand matrix D (where D[i,j] will denote the amount of traffic to be transported from node i to node j).
– **Find:** what nodes will be connected in what rings to what Hub? What is the amount of resources needed?

3.1.1 Cost measures

The cost measures we use are (indicators of) the amount of resources required to deploy a MAN using either one of the proposed architectures. We consider the following measures:

– **Rx/Tx capacity:** for the active node structures (NoSR, SR) this is the number of installed babyboards. For the passive node structure (DBORN), we summed the number of Rx and Tx elements and divided it by two, to allow for a somewhat fair comparison: in the active node, a babyboard contains both a Rx and a Tx which can address the capacity of a single wavelength.
– **Link capacity:** this comprised counting for each physical link the number of wavelengths used, and summing these counts for all links of the MAN. Note that in case of the active node architecture (NoSR, SR), we account for waveband concept: wavelengths are installed per band of four wavelengths and thus we count the complete band in this cost measure (even if the capacity of four wavelengths is not fully used).
– **Number of wavelengths:** here, we count the number of different wavelengths used in each MAN ring. In the multi-ring case, we sum these numbers over all rings. Again, where the waveband concept applies, we count complete bands i.e. increment with 4 for each

addressed band. This number of wavelengths will determine the Hub dimension, since it will amount to its number of i/o ports.

The planning algorithms we have developed, focus on minimising link capacity and number of wavelengths, and only secondarily aim at limiting Rx/Tx capacity. Our results indicated that especially for this first measure, cost could be greatly influenced (i.e. reduced) by making appropriate planning choices.

Note that this dimensioning study is only a single (but quite important) facet of an in-depth assessment of the pros and cons of Active and Passive architectures. This paper therefore is to be complemented with e.g. studies on the architectures' capabilities to deal with dynamic traffic in a network with given amount of resources, as e.g. in [3].

3.1.2 Methodology

We developed a network planning algorithm starting from an ILP-formulation of the planning problem. Yet, the high degrees of freedom hamper reaching optimality within reasonable time. Below we present a sample description of the ILP-formulation for a DBORN-unidirectional ring, and the heuristic approach. The ILP-formulation for other rings differs in constraints. Because of space limitations we cannot give them.

3.1.2.1 ILP-formulation DBORN

In the ILP-formulation as well as the heuristic we use subdemands. A subdemand is a part of the traffic between the source and destination that is transported by one wavelength, and thus cannot exceed the capacity of a wavelength.

Table 1. Constants of the ILP-formulation for unidirectional rings

name	value	Description
d_{ij}	$[0, +\infty)$	demand from node i to node j
sd_{ijs}	$[0, BW_\lambda]$	subdemands from node i to node j
fw_{ijk}	1	link k used to receive d_{ij}
	0	link k not used to receive d_{ij}
fv_{ijk}	1	link k used to send d_{ij}
	0	link k not used to send d_{ij}
n	1 or 4	number of wavelengths per waveband
BW_λ	2.5 or 10 Gbit/s	bandwidth of one wavelength
L		SR and NoSR: total number of wavebands per ring DBORN: number of wavelengths.

name	value	Description
K	[4, 8]	number of nodes (without hub) per ring in SR- or NoSR-ring
	[4, 16]	number of nodes (without hub) per ring in DBORN-ring

Table 2. Variables of the ILP-formulation for unidirectional rings

name	value	Description
w_{ijsb}	1	waveband b used to receive subdemand sd_{ijs}
	0	waveband b not used to receive subdemand sd_{ijs}
v_{ijsb}	1	waveband b used to send subdemand sd_{ijs}
	0	waveband b not used to send subdemand sd_{ijs}

The numbering of the links is done as follows:
- link k is the link between node k and node k+1
- link 0 is the link between the hub and node 0
- link K is the link between node K-1 and the hub

Definition of fv and fw:

$$fv_{ijk} = \begin{cases} 1 & i < k \leq K \qquad i,j = 0\ldots K \\ 0 & \forall \text{ other combinations of } i,j,k \end{cases}$$

$$fw_{ijk} = \begin{cases} 1 & 0 \leq k \leq j \qquad i = 0\ldots K \qquad j = 0\ldots K\text{-}1 \\ 0 & \forall \text{ other combinations of } i,j,k \end{cases}$$

Constraints:

We will first elucidate our notation.The subscripted $\forall$ followed by indices means that for the logical operation we consider these indices in their whole range. All the others remain constant. E.g. $XOR_{\forall b}(v_{ijsb}) = XOR(v_{ijs1}, v_{ijs2}, v_{ijs3}, v_{ijs4}, v_{ijs5})$.

1. All subdemands should be routed over the ring and thus use exactly one waveband.

$XOR_{\forall b}(v_{ijsb})$ b=1…L, $\forall$i,j=0…K, $\forall$s=1…S_{max} where $d_{ijs}>0$
$XOR_{\forall b}(w_{ijsb})$ b=1…L, $\forall$i,j=0…K, $\forall$s=1…S_{max} where $d_{ijs}>0$

2. Don't exceed bandwidth per waveband per link $\forall$b=1…L $\forall$k=0…K

$$\sum_{i=0}^{K}\sum_{j=0}^{K}\sum_{s=0}^{S_{max}} sd_{ijs}.fw_{ijk}.w_{ijsb} \leq BW_{\lambda} \qquad \sum_{i=0}^{K}\sum_{j=0}^{K}\sum_{s=0}^{S_{max}} sd_{ijs}.fv_{ijk}.v_{ijsb} \leq BW_{\lambda} \qquad (1)$$

3. Flow-conservation For each node k counts: The incoming traffic + the traffic transmitted by the node = the outgoing traffic + the traffic received by the node. ∀k= 0…K-1.

$$\left(\sum_{b=1}^{L}\sum_{i=0}^{K}\sum_{j=0}^{K}\sum_{s=0}^{S_{\max}} w_{ijsb}.sd_{ijs}.fw_{ijk} + \sum_{b=1}^{L}\sum_{i=0}^{K}\sum_{j=0}^{K}\sum_{s=0}^{S_{\max}} v_{ijsb}.sd_{ijs}.fv_{ijk}\right) + \sum_{b=1}^{L}\sum_{j=0}^{K}\sum_{s=0}^{S_{\max}} v_{kjsb}.sd_{kjs} = \sum_{b=1}^{L}\sum_{j=0}^{K}\sum_{s=0}^{S_{\max}} w_{iksb}.sd_{iks} + \left(\sum_{b=1}^{L}\sum_{i=0}^{K}\sum_{j=0}^{K}\sum_{s=0}^{S_{\max}} w_{ijsb}.sd_{ijs}.fw_{ijk+1} + \sum_{b=1}^{L}\sum_{i=0}^{K}\sum_{j=0}^{K}\sum_{s=0}^{S_{\max}} v_{ijsb}.sd_{ijs}.fv_{ijk+1}\right) \quad (2)$$

Similar for the hub:

$$\left(\sum_{b=1}^{L}\sum_{i=0}^{K}\sum_{j=0}^{K}\sum_{s=0}^{S_{\max}} w_{ijsb}.sd_{ijs}.fw_{ijK} + \sum_{b=1}^{L}\sum_{i=0}^{K}\sum_{j=0}^{K}\sum_{s=0}^{S_{\max}} v_{ijsb}.sd_{ijs}.fv_{ijK}\right) + \sum_{b=1}^{L}\sum_{j=0}^{K}\sum_{s=0}^{S_{\max}} w_{Kjsb}.sd_{Kjs} = \sum_{b=1}^{L}\sum_{i=0}^{K}\sum_{s=0}^{S_{\max}} v_{iKsb}.sd_{i1s} + \left(\sum_{b=1}^{L}\sum_{i=0}^{K}\sum_{j=0}^{K}\sum_{s=0}^{S_{\max}} w_{ijsb}.sd_{ijs}.fw_{ij0} + \sum_{b=1}^{L}\sum_{i=0}^{K}\sum_{j=0}^{K}\sum_{s=0}^{S_{\max}} v_{ijsb}.sd_{ijs}.fv_{ij0}\right) \quad (3)$$

Objective: We minimize the amount of wavelengths.

$$\text{minimize} \quad \sum_{k=0}^{K}\sum_{b=1}^{L}\left[OR_{\forall j,i,s}\left(fw_{ijk}.w_{ijsb}\right)\right] + \sum_{k=0}^{K}\sum_{b=1}^{L}\left[OR_{\forall j,i,s}\left(fv_{ijk}.v_{ijsb}\right)\right] \quad (4)$$

i,j=0…K s=0…$S_{\max}$

Cost measures:

$$\text{Rx/Tx capacity} = \sum_{i=0}^{K}\sum_{b=1}^{L}\left(OR_{\forall j,s}\left(v_{ijsb}\right) + OR_{\forall j,s}\left(w_{ijsb}\right)\right) \quad (5)$$

j=0…K s=0…$S_{\max}$

$$\text{Link capacity} = \sum_{k=0}^{K}\sum_{b=1}^{L}\sum_{j=0}^{K}\sum_{i=0}^{K}\sum_{s=0}^{S_{\max}}\left[OR\left(fw_{ijk}.w_{ijsb}\right) + OR\left(fv_{ijk}.v_{ijsb}\right)\right] \quad (6)$$

$$\text{nr. of used wavelengths} = \sum_{b=1}^{L}\left[OR_{\forall j,i,s}\left(v_{ijsb}\right)+OR_{\forall j,i,s}\left(w_{ijsb}\right)\right] \quad (7)$$

$i,j=0\ldots K \qquad s=0\ldots S_{max}$

3.1.2.2 Heuristic approach for DBORN

Table 3. Symbols used in the single ring heuristics - unidirectional

name	value(s)	description
i	[0, K]	index number of a source node (K=total number of nodes)
j	[0, K]	index number of a destination node
s	[0, S_{max}]	index number of a subdemand
d_{ij}	[0,+∞)	demand from node i to node j
$d'_{i,hub}$	[0,+∞)	all demand from node i to the hub
$d'_{hub,j}$	[0,+∞)	all demand from the hub to node j
$sd'_{i,hub,s}$	[0, BW_λ]	subdemand from $d'_{i,hub}$
$sd'_{hub,j,s}$	[0, BW_λ]	subdemand from $d'_{hub,j}$
f_{ijk}	1	link k used for demand d_{ij}
	0	link k not used for demand d_{ij}
n	1 or 4	number of wavelengths per waveband
BW_λ	2.5 or 10 Gb/s	bandwidth of one wavelength
L		SR and NoSR: total number of wavebands per ring DBORN: total number of wavelengths Max(up, down)
K	[4, 8]	number of ring nodes in SR- or NoSR-ring
	[4, 16]	number of ring nodes in DBORN-ring

The algorithm for a DBORN ring goes as follows:

1. The demand matrix D is transformed in a new demand matrix D', defined as follows: d'_{ihub} is the total amount of traffic transmitted by node i and crossing or dropped at the hub, d'_{hubj} is the total amount of traffic, crossing or added by the hub, received by node j. The new D' is in the remaining of the algorithm used as THE demand matrix.
2. The new demands are split up so that as few as possible wavelengths are installed (i.e. subdemands are given as much as possible the maximal capacity BWλ. The filling of the wavelengths will be the easiest with as much as possible subdemands of maximal size. The next steps 3 and 4 are independently executed for all nodes.
3. Determine which wavelength is used for which subdemand (sd'ihubs). When a wavelength doesn't meet one of the conditions (i) not used or (ii) enough free capacity, we go on to the next wavelength. Once all subdemands originating a node are assigned a wavelength, we go on with step 4.

4. For all in-groups (sd'hubjs) wavelengths are being chosen. As the wavelengths used for upstream (out-groups) differ from the wavelengths used for downstream (in-groups), this step in the algorithm can be executed independently from the previous step.
5. Calculation of the resources used.

3.1.2.3 Multi-ring problem

For the multi-ring problem, we use a heuristic tabu-search approach. Given an input topology with candidate links, we solved the original planning problem in phases:

– **Phase I** – ring generation: find all possible rings that can be formed in the given topology. Rings may consist of nodes which are all connected through direct physical links (i.e. part of the given topology), but also multiple physical hop connections are allowed. In the latter case, there are “pass-through” nodes, which are part of the physical route followed by a ring, but where no babyboards (or Rx/Tx in case of DBORN) will be installed.
– **Phase II** – ring selection: find a suitable subset of all possible rings that contains all nodes, and where each ring contains the Hub. For this phase, we used a tabu-search heuristic.
– **Phase III** – hub identification: we repeat phase II for multiple hub choices, and calculate the required amount of resources for each of the rings that are part of the candidate solution. For the capacity calculation phase, we resort to our single ring planning heuristic and we apply to each individual ring. We select the hub leading to the smallest overall cost (see previous section 3.1.1 for info on the cost measures used).

3.1.3 Traffic matrices

As outlined a crucial input to the planning/dimensioning problem is constituted by the traffic matrices. In our analysis, we used following traffic matrices:

– **Uniform:** in this case, we will set D[i,j]=d for each node pair (i,j), and investigate the evolution of the cost for increasing d, compared for the three cases (NoSR, SR, DBORN). This means that the traffic coming from each node will be equally spread over all other nodes in the ring.
– **Server:** here, we will assume all traffic comes from/goes to a single node, i.e. D[h,i]=D[i,h]=d1 for all i≠h, and all other demands D[i,j]=d2 (d2<d1 and i,j≠h). Again we will compare the three ring architectures for increasing d1 and d2.

– **Neighbour:** here, we assume there is only traffic between neighbouring nodes. We consider (i) the case of having traffic between a node and its predecessor seen in the direction in the ring (D[i,i-1]=d, other D[i,j]=0) as well (ii) as the case of traffic between each node and its immediate successor in the direction of the ring (D[i,i+1]=d, other D[i,j]=0). Note that the distinction between (i) and (ii) will only be visible when unidirectional rings are used.
– **Random:** for all i,j: D[i,j]=random(0,2d) where random returns a value uniformly distributed over [0,2d] with d the mean. Again we will compare the three ring architectures for increasing d.
– **DAVID:** this is the traffic matrix proposed in the DAVID project and, what's more it is proved to be very realistic, i.e. the traffic pattern is similar to that of ADSL concentration points of Telefonica in the environment of Madrid (see further).

3.1.4 Topologies

For the multi-ring planning case, we started from several topologies with various connectivity properties. The topology given as an input to our planning problem consists of nodes, and acceptable physical links between them. Thus, when there is a link between node N1 and node N2 in a given input topology, this means the solution to the ring planning problem is allowed to contain rings with a direct link between node N1 and N2. When there is no link between N1 and another node N3, this means that when the solution needs a connection needs to have a connection N1-N3, this will comprise multiple spans (crossing pass-through nodes, see phase II of the multi-ring planning algorithm in 3.1.2).

The topologies used to address the multi-ring problem were the following:
– **Mesh:** fully meshed network, where every link between any two nodes is allowed. In this case, there are no restrictions on the physical ring topologies.
– **Star:** this is a star-like topology, where one "central" node has a link to every other node, while the others have only links to two neighbouring nodes and the central node.
– **3Links:** like the Mesh case, this is a symmetrical topology, but not all links are allowed: each node has three direct links to neighbouring nodes.
– **Rings:** this is a topology consisting of three physical rings, all passing through a central Hub node.

The purpose of using these different physical topologies was to qualify the impact of physical topology constraints on the outcome of the ring planning problem.

4. RESULTS

4.1 Single ring

In this section we will summarize our results for the single ring planning case. We will subsequently address four questions: (1) is it advantageous to exploit two counter-rotating rings, and use both (under failure-free conditions)?; (2) does the active node architecture make sense when spatial reuse is not exploited (i.e. does NoSR make sense, or is it always outperformed by SR)?, (3) what are the benefits of the passive DBORN node structure from a dimensioning point of view?, (4) does the waveband concept make sense?

Figure 4 and Figure 5 summarize some results of single ring dimensioning for Uniform and Random traffic patterns. For the sake of brevity, we have omitted similar graphs for other traffic patterns. Clearly, with an increasing demand (recall section 3.1.3 for definition of the demand d), we need extra resources.

With respect to Rx/Tx capacity, we find that the Passive DBORN architecture performs best. This was to be expected: since Rx and Tx spectrum are completely decoupled, there is no conflict between up- and downstream traffic in the wavelength domain. This type of conflict does lead to a slightly higher Rx/Tx capacity requirement with the active architecture for NoSR and SR. In addition, the cost corresponding with a single wavelength in the passive node structure may be lower, since the node architecture is less complex.

The graphs plotting the required amount of wavelengths show that the price paid for the minimal Rx/Tx capacity achieved by DBORN is an increase in the number of wavelengths used, which roughly doubles (compared to NoSR) because of the spectral separation of Rx and Tx. When the number of wavelengths used is an issue (i.e. limited number of wavelengths compared to the traffic volume to carry, and/or either technological or cost-related constraints on the Hub dimensions), the active node structure is most profitable and by using the spatial reuse concept (SR) this number of required wavelengths can be reduced to the minimum, as will be discussed further in more detail.

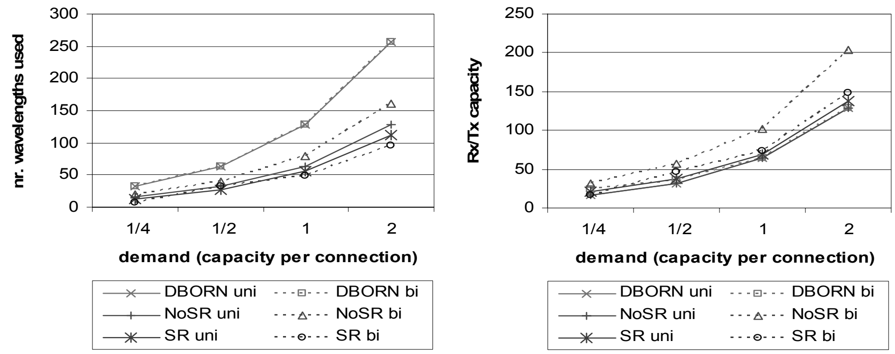

(a) number of wavelengths used (b) Rx/Tx capacity

Figure 4. Sample results for Uniform demand (d is quantified in number of wavelengths)

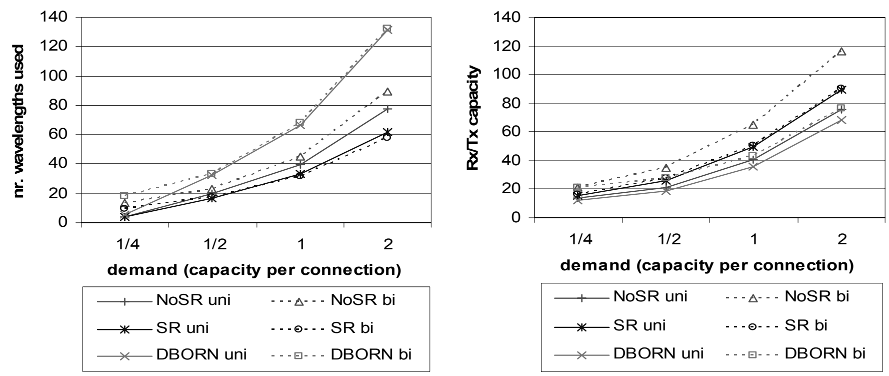

(a) number of wavelengths used (b) Rx/Tx capacity

Figure 5. Sample results for Random demand (D[i,j]=random(0,2d) for all node pairs; d is quantified in number of wavelengths)

4.1.1 Bidirectional vs. unidirectional rings

A first question we promised to address was whether the deployment of bidirectional rings (i.e. using two counter-rotating rings also under failure-free conditions) is useful. The results presented earlier contained for each architecture (NoSR, SR, DBORN) a curve labeled “uni” and “bi”. The Uni case was where traffic only flows in a single direction. In the bi(directional ring) case, we assume babyboards can access either a ring where traffic flows clockwise, or the one with counter-clockwise rotating traffic. Thus, when a babyboard is installed, a decision has to be made on which of the two rings it is inserted. In our planning algorithm, we start by placing babyboards on only one of the rings for all traffic flows whose shortest path between source and destination lies along this ring. Subsequently, we try to fill up any unused bandwidth with remaining flows. The remaining traffic then is put on the other ring.

For DBORN, we find that bidirectional rings do not make much sense. Both Rx/Tx capacity and number of wavelengths used cannot be diminished by deploying counter-rotating rings. On the contrary: we note an increase in Rx/Tx capacity and a slight increase in number of wavelengths used. (Note that for the Uniform traffic pattern we see no differences, but this is because the demands are all of the same size and can fill integer multiples of wavelengths.)

In case of NoSR, using an active node structure without exploiting its spatial reuse capability, we find that deploying bidirectional rings considerably increases Rx/Tx capacity and to a lesser extent the amount of wavelengths used. The only advantage is that the total link capacity addressed may be slightly reduced as the majority of the traffic can follow the shortest path[1]. However, since all traffic has to pass through the Hub, the links directly connected to the Hub form the bottleneck which can not substantially be alleviated by providing two counter-rotating rings.

In the SR case, where the spatial reuse concept is exploited, bidirectional rings can be advantageous. While the advantages in case of more or less symmetrical traffic conditions (as Uniform and Random) are not that pronounced, they do become clear when the proportion of intra-ring traffic rises (e.g. the Neighbour pattern). Allowing to choose the shortest path indeed does increase the opportunities to fill in unused capacity through spatial reuse.

We conclude[2] that from perspective of Rx/Tx capacity or number of wavelengths, only in a spatial reuse concept (SR), bidirectional rings should be deployed. However, when link capacity (see 3.1.1. for definition) dominates overall CAPEX, counter-rotating rings may be useful.

[1] The following of the shortest path is the main reason for the increase in Rx/Tx capacity, since it forces nodes to have a babyboard on both rings, while from an Rx/Tx capacity point of view this is not always strictly necessary.

[2] Note that these conclusions consider error-free conditions only, it is clear that for protection purposes a counter-rotating ring always is useful.

4.1.2 The spatial reuse concept

The active node structure allows spatial reuse. The results presented above indicate that even when it is not exploited (NoSR), the active structure can be useful to limit the number of wavelengths used (and therefore link capacity). From a Rx/Tx point of view, it does not perform worse than the DBORN case — but not better either.

When the number of wavelengths available on a fibre becomes a limiting factor, the passive node structure should clearly be replaced by an active one. In this case, to exploit the available bandwidth most efficiently, the Spatial Reuse concept can be of considerable help: SR can reduce the amount of necessary wavelengths with about 25% compared to NoSR for even fairly symmetrical traffic patterns such as Random.

4.1.3 Passive node structure

The advantage of the passive node structure used in DBORN is of course its architectural simplicity, allowing for instance longer ring circumference. From the dimensioning results, we find that it also leads to the minimal Rx/Tx capacity requirements. However, the number of wavelengths used in the ring is considerably higher, and consequently also the link capacity (i.e. amount of wavelengths used on all the links). This considerable increase may call for deployment of more efficient SR architectures when the traffic volume to be carried in the MAN rises to node-to-node demands close to the order of a complete wavelength or higher.

4.1.4 The waveband concept

A peculiar aspect of the active node architecture used either in NoSR or SR is the Waveband concept: a babyboard is able to access n=4 wavelengths constituting a band, but only one at a time (i.e. in a particular timeslot): the corresponding Rx/Tx capacity is that of a single wavelength. It is reasonable to question the concept, and compare it with a band-less concept (or the case n=1). From an architectural viewpoint, this means the band mux/demux can be omitted and fewer switching components are required in the transit path. Here, we discuss the effects of abandoning the band concept on the required resources to provide for a given traffic demand.

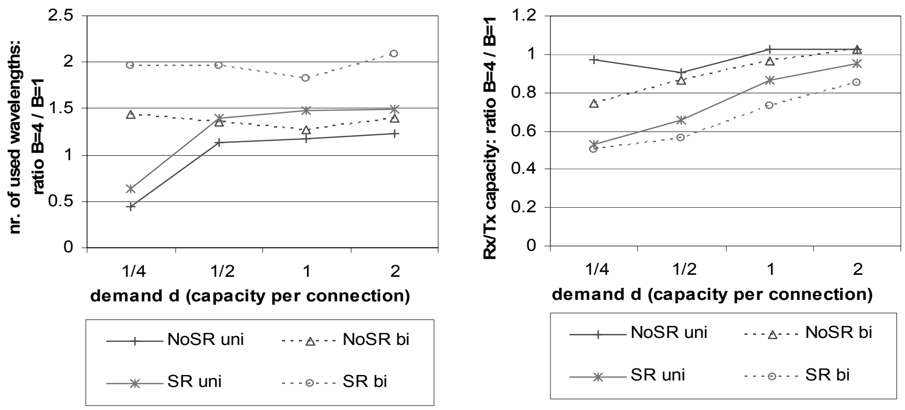

(a) number of wavelengths used

(b) Rx/Tx capacity

Figure 6. Sample results for Random demand (D[i,j]=random(0,2d) for all node pairs). The values plotted are ratios of the numbers for the case with bands (B=4) divided by the case without bands (B=1).

Figure 6 plots the ratio of the required resources needed when using the waveband concept, divided by the amount of resources needed when abandoning the band concept (n=1). When no spatial reuse is deployed (NoSR), the use of the band concept leads to slightly lower Rx/Tx requirements, but the number of wavelengths increases (except for small demands).The increase in number of wavelengths needed (cf. ratio > 1) stems from the restriction of only accessing the capacity of a single wavelength per band per node: not all n=4 wavelengths comprising a waveband can be completely filled, resulting in the installation of unused bandwidth. Still, since there is a quite large total demand, and all those demands have to cross the Hub, the penalty of imposing a band concept is relatively limited (esp. compared to the SR ring, see further): all demands share the links adjacent to the Hub, and therefore the bands can be filled rather efficiently. For small demands and unidirectional rings, we notice even an improvement: with the waveband concept fewer wavelengths are required. The slight advantage of having bands when demands are small stems from the sub-optimality of our heuristic planning algorithm: we start installing babyboards for demands in a particular order (depending on the order in which nodes are numbered in the demand matrix D), and this means that a particular wavelength assignment in the band-less case (B=1) sometimes forces us to install extra wavelengths because a wavelength is already filled somewhere along the ring, even when the add/drop capacity in a particular node is not yet fully used. (When demands are small, an “unlucky” wavelength assignment when using a waveband concept less often leads to instalment of an extra board, because of the bigger granularity of bands compared to wavelengths.) Note that this same effect is the cause of the lower Rx/Tx capacity requirements of the band concept.

When spatial reuse is adopted (SR), and particularly in the bidirectional case, it becomes far more difficult to efficiently fill the bands. This is reflected in the considerable increase in number of wavelengths: the case n=4 needs about twice as much wavelengths than the band-less case B=1.

4.1.5 The DAVID test-case

To assess the validity of our conclusions, we applied our dimensioning algorithms to the traffic scenario developed in the DAVID Project. We have considered 16 nodes per ring, but four different types of nodes: server nodes, which are peculiar nodes bigger than others and generating more traffic than they receive, and three types of “regular” nodes (i.e. receiving more traffic than they emit) with different weights (big/medium/small indicating the contribution of the node to the ring traffic). We focused on a ring network with one server node (Se), 2 big nodes (B), 4 medium nodes (M) and 9 small nodes (s) with a distribution shown in the table 4.

Table 4. Node repartition for the ring network

Node	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
Type	B	s	M	Se	s	s	B	M	s	s	M	s	M	s	s	s

Knowing the overall structure of the ring network, and in order to fit approximately into the limit in terms of capacity for DAVID (1,28Tbit/s Optical Hub) in case of multiple rings, we considered the following mean capacity per ring: 20G, 40G and 80G.

In addition to that, we have fixed then the ratio between the upstream and the downstream traffic in the network and the weight of each node type on the ring. This is summarized in table 5:

Table 5. Traffic assumption per node

	Se	**B**	**M**	**s**	**total**
up	20.00%	3.20%	1.60%	0.80%	40.0%
down	2.40%	8.40%	4.80%	2.40%	60.0%

We have plotted the results (Figure 7) for SR with bands, SR without the band concept (B=1, “SR”), and DBORN.

From a Rx/Tx perspective, we conclude that the DBORN case, with its passive node structure, and in a unidirectional ring configuration leads to the cheapest solution. For the passive node structure, a bidirectional ring concept only makes sense to limit the link capacity. For large demands (the 80G and 160G case), we notice that we even may gain a few wavelengths to deploy in the ring (thus reducing the required Hub I/O ports).

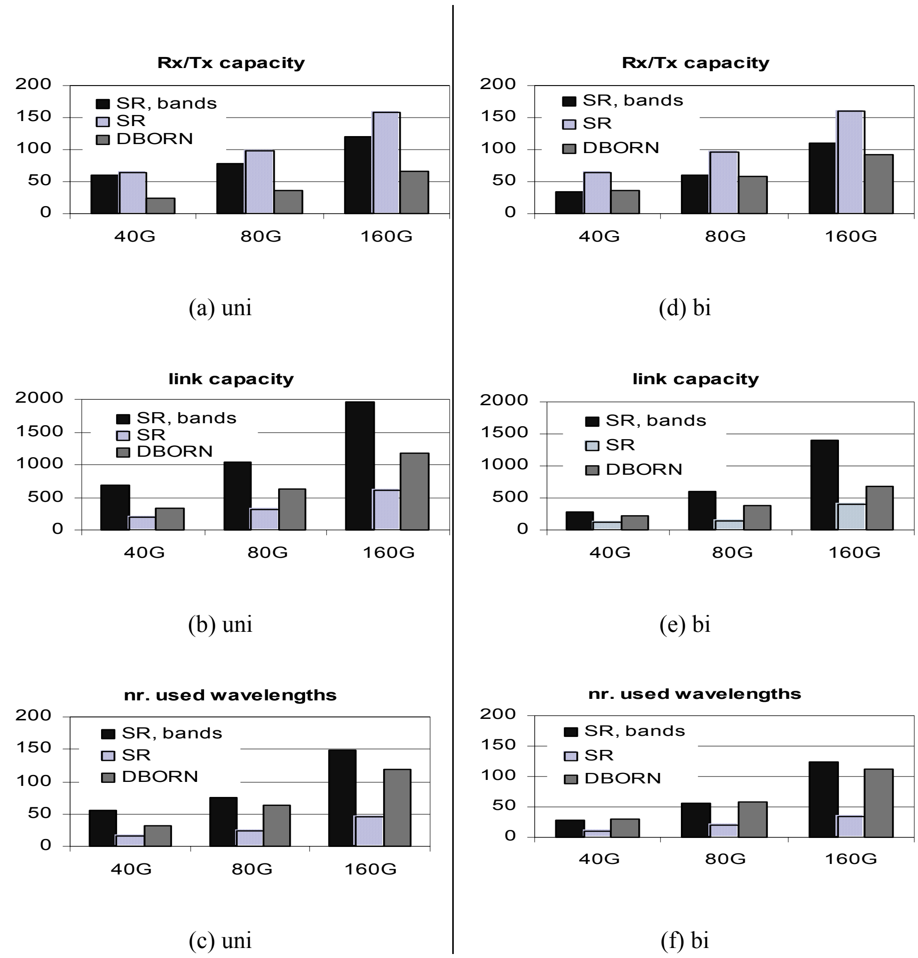

(a) uni (d) bi

(b) uni (e) bi

(c) uni (f) bi

Figure 7. Results for the DAVID test-case: (a-c) with unidirectional rings, (d-f) for bi-directional rings.

The active node structure exploiting spatial reuse in a band concept ("SR, bands") surprisingly does not lead to lower link capacities or even number of wavelengths. This stems from the fact that the band concept hampers the advantages that could be gained from spatial reuse, especially in unidirectional rings. Indeed, when we discard the band concept and rather use a per-wavelength concept (“SR”), the gain in both link capacity and the number of wavelengths required is quite substantial (about a factor 1/3).

4.2 Multiple rings

Our results for the multiple ring approaches basically confirm the conclusions from the single ring studies. Since they do not shed new light on the issues arising in comparing the architectures from a dimensioning point of view, we have chosen not to include a detailed discussion here.

The Hub appeared to be a node that has a central position in as well as the networktopology as the traffic matrix. Due to space limitations it is not possible to discuss this in detail.

5. CONCLUSIONS

We have compared the active versus passive node structures from a capacity dimensioning point of view, and addressed the issues of spatial reuse, the waveband concept and the use of bidirectional rings.

- **Bidirectional versus unidirectional rings:** When the cost of installing Rx/Tx in the nodes (babyboards in case of the active node structure for NoSR and SR) dominates the overall cost, the use of bidirectional rings (for error-free conditions) is not advisable. However, to reduce the number of wavelengths used (and consequently, the link capacities), the use of bidirectional rings is crucial. This reduction is most effective when spatial reuse is allowed.
- **NoSR, SR or DBORN:** When only Rx/Tx costs are important, the active node structures of NoSR and SR are not justifiable. Clearly, when the amount of intra-ring traffic is small (compared to inter-ring traffic or traffic between MAN and WAN, which also crosses the Hub), there is no better alternative than DBORN. However, when there is a considerable amount of intra-ring traffic (and therefore plenty of spatial reuse opportunities), the additional complexity of introducing SR with active node structures should be considered. Especially when traffic volumes increase up to the point where the number of available wavelengths could become a limitation, the active node with its spatial reuse capability will become a key factor in efficiently exploiting the available bandwidth.
- **The waveband concept:** The current concept of having active nodes with spatial reuse where babyboards can address a complete band of four wavelengths, but only have an Rx/Tx capacity corresponding to a single wavelength does not seem to make much sense from a dimensioning

point of view. (It may however be advantageous from a MAC point of view, and its flexibility to deal with traffic fluctuations.) The restricted addressable capacity of the babyboard hampers an efficient use of the available capacity on the fibres. When abandoning the waveband concept, we note a significant reduction in the amount of wavelengths needed to provide a given demand.

– **The choice of the Hub** is influenced by both the network topology and the traffic matrix.

ACKNOWLEDGEMENTS

This work has been supported by the European Commission through the IST-project DAVID (IST-1999-11387), and by the Flemish Government through the IWT GBOU-project "Optical Networking and Node Architectures". C. Develder is supported as a Research Assistant of the Fund for Scientific Research – Flanders (F.W.O.–Vl.), Belgium. D. Colle is supported as a Post-Doctoral Research Assistant of the IWT

REFERENCES

[1] S. Yao, et al., "All-optical packet switching for metropolitan area networks: opportunities and challenges", IEEE Comm. Mag., 39(3), pp. 142–148

[2] L. Dittmann, et al., "The European IST project DAVID: a viable approach towards optical packet Switching", IEEE J. Selected Areas in Communications, Vol. 21, no 7, Sept.2003, pp.1026-1040.

[3] A. Bianco, et al., "Access control protocols for interconnected WDM rings in the DAVID metro network", in Proc. IWDC 2001.

[4] D. G. Cunningham, "The status of the 10-Gigabit Ethernet standard", Proc. European Conf. on Optical Commun. (ECOC) 2001

[5] A. Carena, et al., "RINGO: A demonstrator ofWDM optical packet network on a ring topology", Proc Working Conf. On Optical Netw. Design and Modelling (ONDM) 2002

[6] D. Dey et al., "FLAMINGO: a packet-switched IP-over-WDM all-optical MAN", Proc. European Conf. on Optical Commun. (ECOC) 2001

[7] K. Shrikhande, et al., "HORNET: A packet-over-WDM multiple access metropolitan area ring network", IEEE J. Selected Areas in Commun., vol 18, no 10, Oct. 2000, pp. 2004-2016

[8] A. Stavdas, et al., "IST-DAVID: Concept presentation and physical layer modelling of the metropolitan area network", IEEE .J. Lightwave Techn., vol 21, no. 2, Feb. 2003, pp. 372-383

[9] A. Takada, et al., "Architecture of ultrafast optical packet switching ring network", IEEE J. Lightwave Techn., vol 20, no. 12, Dec. 2002, pp. 2306-2315

[10] N. Le Sauze, et al., "A novel low cost optical packet metropolitan ring architecture", Proc. European Conf. on Optical Commun. (ECOC) 2001