IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS, VOL. 24, NO. 4, APRIL 2006

Overspill Routing In Optical Networks: A True Hybrid Optical Network Design

Erik Van Breusegem, *Student Member, IEEE*, Jan Cheyns, *Student Member, IEEE*, Davy De Winter,
Didier Colle, *Member, IEEE*, Mario Pickavet, *Member, IEEE*, Filip De Turck, *Member, IEEE*,
and Piet Demeester, *Senior Member, IEEE*

***Abstract*—To efficiently support the highly dynamic traffic patterns of the current Internet in large-scale switches, we propose a new hybrid optical network design: Overspill Routing In Optical Networks (ORION). By taking advantage of the reduced (electronic) processing requirements of all-optical wavelength switching, the electronic bottleneck is relieved. At the same time, ORION achieves a level of statistical multiplexing comparable to the more traditional point to point WDM solutions, circumventing the bandwidth inefficiencies of all-optical wavelength switched networks, caused by dynamic traffic patterns. The result is a true hybrid optical network design, forming a bridge between these two switching concepts. In this paper the generic concept of ORION is described. An example node design, based on current advanced optical technologies, is described in detail. The ORION concept is also evaluated, comparing it with its two composing technologies, optical wavelength switching and point to point WDM, as well as a third, more trivial, hybrid one, through several case studies.**

***Index Terms*—Circuit switching, optical switches, packet switching, wavelength division multiplexing.**

I. INTRODUCTION

IN the past, optical networks opened up a vast amount of bandwidth, with the breakthrough of optical fiber and all-optical amplification. A second boom in bandwidth capacity came with the development of (Dense) Wavelength Division Multiplexing, (D)WDM. As a consequence, pure bandwidth availability is no longer the most important cost determining factor on the technological side. Originally, optical connections were used as high capacity point-to-point interconnections. Data was converted to the electronic domain every hop, and processed at an electronic router. This resulted in highly efficient sharing of the transmission resources. At the same time, however, most of the traffic handled in any router is transit traffic, i.e. still has to travel one or more hops to reach the exit point of the network. In backbone networks, 70% of all traffic being transit is not an exception [1], mainly due to their low meshing degree. Given the current trend that optical transmission capacity grows faster [2] than electronic processing capability, the electronics are considered by many the major bottleneck to resolve. Consequently, several solutions were proposed to reduce the strain on the electronic layer. The most important one certainly is the concept of wavelength paths (or lightpaths). In this technology a wavelength passes transparently through an optical cross-connect (OxC), without the IP router inspecting the carried data. The drawback of this approach is that a wavelength path cannot be re-used by intermediate nodes when it is not used at full capacity. This can result in an increased amount of wavelengths needed to fulfill a set of demands. A logical next step in this evolution are Automatically Switched Optical Networks (ASONs), where the wavelength paths are set up by the control plane, without (explicit) intervention of the network operator. This allows adapting to traffic variations on a medium to large time scale. However, these solutions inherently remain wavelength switched. Although the flexibility is considerably increased, the granularity, in terms of bandwidth and time scale, is still coarse.

The main contribution of this paper is a switch and network architecture proposal, called Overspill Routing In Optical Networks (ORION). ORION succeeds in using the full benefits of the ASON, without sacrificing the high statistical multiplexing gains from the earlier point to point WDM optical networks. In this paper, we will explain ORION applied to IP/MPLS directly over WDM (see [3] for a thorough discussion of IP over WDM), although it is not limited to that setup. Although ORION as a hybrid is quite unique in its design and concept, several other proposals try to achieve similar objectives, through some kind of intermediate form of simple point to point WDM and wavelength switched networks. The closest resemblance ORION has with any other proposal is certainly the one described in [4] and [5]. Others may seem similar, like [6], [7] or [8], but we will show in the next section that they are actually, in our view, an extreme form of more classical switching architectures. The overspill mechanism we will explain is sometimes likened to deflection routing [9], light-trees [10], light-trails [11] or light-frames [12], but we will also illustrate that it is fundamentally different. Finally, the CHEETAH [13] concept comes to mind as related. This is also a true hybrid, but is actually a form of a combined

Manuscript received March 29, 2004; revised October 26, 2005. This work was partly funded by the European commission through their Information Society Technologies (IST) program, more specifically through the projects STOLAS (IST-2000-28557), LASAGNE (FP6 506760), and NOBEL (FP6 507509). The flemish government also funded part of this research through the IWT-GBOU project "Optical Networking and Node Architectures". E. Van Breusegem and D. Colle thank the IWT for its financial support for their PhD and postdoctoral grants. J. Cheyns is Research Assistant with the Flemish Fund for Scientific Research (FWO).

E. Van Breusegem, J. Cheyns, D. Colle, M. Pickavet, F. De Turck, and P. Demeester are with the Department of Information Technology (INTEC), Ghent University, Gent, 9050, Belgium (e-mail: {erik.vanbreusegem, jan.cheyns, didier.colle, pickavet, demeester, filip.deturck}@intec.UGent.be).

D. De Winter is with the Department INWE, Hogeschool Gent, J. Kluyskensstraat 2, Gent, 9000, Belgium (e-mail davy.dewinter@hogent.be).

Digital Object Identifier 10.1109/JSAC-OCN.2006.04002.

> ***"Furthermore, from an overall networking perspective, a hybrid solution combining the merits of fast (optical) circuit switching with those of optical packet switching may offer better cost and performance. Indeed, such a solution may reduce the throughput requirements of packet switches."***
>
> *"Optical Switching Networks: From Circuits to Packets"*, A. Hill and F. Neri,
> Guest editorial in IEEE Comm. Mag., vol. 39, no. 3, March 2001

Fig. 1. Citation from *IEEE Communications Magazine* [15].

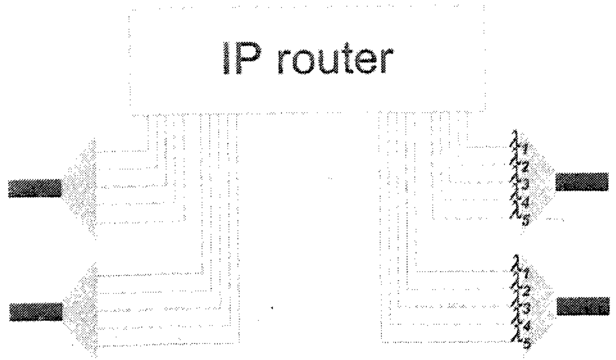

Fig. 2. "Point to Point WDM" conceptually.

architecture (see also Section IV, with two physically distinct networks in parallel. The remainder of the paper is organized as follows. In Section II, we introduce the ORION concept on a functional node and network level. We continue in Section III by presenting a suggested hardware design that can support the functions required. Subsequently, we present in Section IV and Section V an evaluation of ORION, showing detailed case studies and trends. Note that for our evaluations we use a dimensioning approach, in which we calculate a proper amount of wavelengths given a certain demand and topology. The dual problem, where loss and throughput are derived from a certain pre-dimensioned topology is tackled in [14]. That study also compares ORION with an optimized dimensioned network using multi-layer traffic engineering principles. A summary and conclusion in Section VI wraps up the paper.

II. CONCEPTUAL DESIGN

We will start with a short revisit of the basic principles of packet switching, and wavelength switching. But first of all, we would like to draw the attention to the following quote (Fig. 1), which states hybrids may be a very interesting option indeed. We believe our proposal is one of these interesting options.

A. Point to Point WDM

Fig. 2 shows an "optical" node, which actually is nothing more than an IP router connected to (D)WDM fiber. All wavelengths are terminated at this IP router, no optical transparency is available.

In this paradigm, which we term *point to point WDM* in this paper, links never bypass an IP router. As a consequence the wavelengths on the links are still fully shared, so the gains of statistical multiplexing apply. The available bandwidth is used very efficiently, as all wavelengths are constantly accessible. However, this also means that routing information of each packet is examined at each node, and processed again by the IP router. Hence, every router will have a lot of transit traffic to cope with. This results in big IP routers, and expensive forwarding hardware to cope with it. This kind of switching is very commonly deployed today. The paradigm can be further shifted into the optical world [6], when packets are kept optical end-to-end. Their headers with routing information, however, are still electronically processed, as optical processing is not a feasible technology yet. Available in several variants, this switching style is usually referred to as Optical Packet Switching (OPS) or Optical Burst Switching (OBS). Although the strain on the electronic layer can be considerably less in the case of OPS/OBS, the fundamental characteristics have much in common with *point to point WDM*. OPS/OBS itself is mostly still in an experimental stage, with mainly (limited) demonstrators, implementing components like a burst scheduler (e.g. [16]), or limited complete nodes (e.g. [17], [18], [19], or [20]). In *point to point WDM* architectures contention (congestion) has to be solved at intermediate nodes (e.g. for OPS, see [21]). This means that delays and losses can be large and very unpredictable, since they depend on global network load. Furthermore, we need as many line speed interfaces to the electronic domain as there are wavelengths.

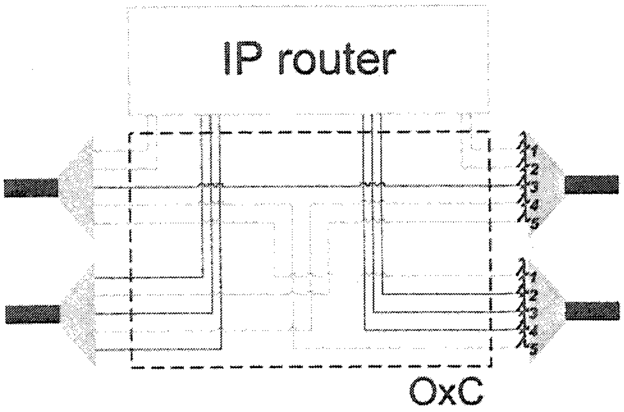

Fig. 3. "Wavelength switched WDM" conceptually.

B. Wavelength Switched WDM

One proposed solution, increasingly being deployed, to counteract the growing electronic processing bottleneck is shown in Fig. 3. Instead of terminating all wavelengths at each hop, some wavelengths bypass the IP router. The optical network uses these wavelength channels transparently to form wavelength paths (e.g. as in [22]) between two nodes in the network. To allow for dynamic configurability an Optical Cross-Connect (OxC) can be used. By changing settings in the OxC, wavelength paths can be set up or torn down. The great advantage of this approach, which we term *wavelength switched WDM* here, is that by constructing these wavelength paths, the electronic layer can be bypassed. As a consequence the amount of transit traffic that has to be handled by this electronic layer can be significantly reduced. Also,

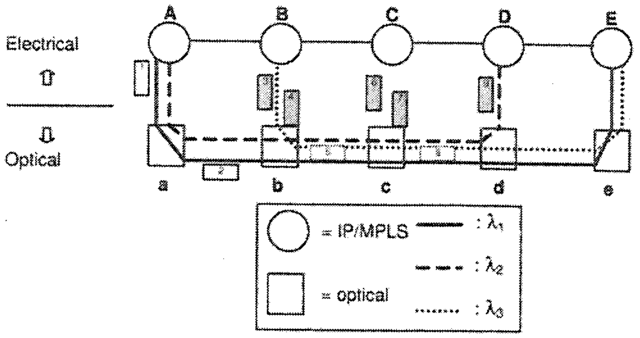

Fig. 4. An example wavelength switched network to demonstrate ORION functionality. Capital letters denote IP/MPLS routers, small letters denote the OxCs. Packet flow is indicated by the numbers in the packets.

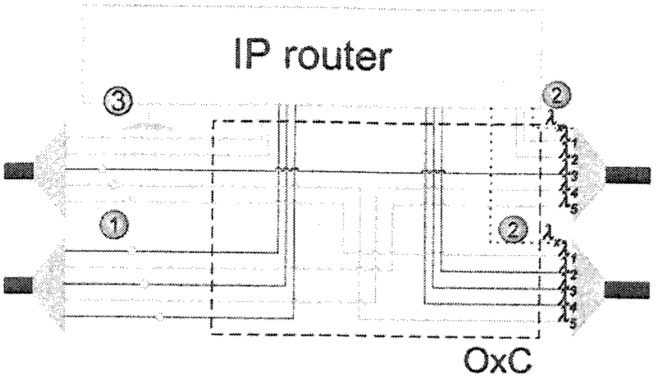

Fig. 5. ORION conceptually.

less interface cards are needed, reducing costs. On the other hand, setting up connections takes at least a round trip time through the network. For large networks this can be quite time-consuming, making it less suited for the transport of bursty or highly variable traffic [8]. Another drawback is the rather large granularity of the circuits especially as bit rates of a single wavelength channel continue to rise. Therefore, setting up and tearing down of wavelength paths is mainly useful for longer term traffic variations. On a shorter time-scale, *wavelength switched WDM* can suffer from severe bandwidth inefficiencies.

C. ORION

Thus, point to point WDM faces the electronic bottleneck, while wavelength switched WDM suffers from inefficient bandwidth usage. What ORION offers is union of the two concepts. We would like to illustrate our concept first on a network level by means of Fig. 4.

The basic idea in ORION is to start from a *wavelength switched WDM* network, where some wavelength paths are established. In a normal wavelength switched network a lightpath passes an intermediate node transparently, i.e. the node cannot access the data in the passing wavelength. In the example of Fig. 4 node D has no access to the IP packets traveling in wavelength λ_1. Only wavelengths starting or terminating in that node can be accessed. Suppose we have the connections A-a-b-c-d-D on λ_2, B-b-c-d-e-E on λ_3 and A-a-b-c-d-e-E on λ_1, established as indicated on Fig. 4. The basic problem with wavelength switching occurs in the following scenario: when the capacity of the lightpath AD on λ_2 is fully used, and there is some temporary overload for this connection, the wavelength λ_1 cannot be used to carry some of its traffic, as it is dedicated for connection AE. It is, however, perfectly possible that the lightpath AE (on λ_1) is only carrying a low traffic load, so that in fact at some moments in time it is empty (again, for clarity in this discussion we assume IP directly over WDM). Thus, while connection AD has a capacity shortage, AE has capacity in abundance. While ordinary wavelength switched WDM networks cannot solve this type of situation instantly, ORION enables the network to use this capacity on λ_1 (normally dedicated to the connection AE) for the connection AD by switching to the so-called overspill mode, which essentially operates the network as if it is a point to point WDM network. The following happens: Packets (belonging to connection AD) are inserted in λ_1 (normally dedicated to connection AE), we call this overspilling, hence the name Overspill Routing In Optical Networks (ORION). Thus, packets inserted in a wavelength different from the lightpath of the connection are called overspill packets. Packets that are sent on their correct lightpath are lightpath packets. When packets of connection AD are inserted in λ_1, node b has to be able to detect these (we will use a special marking, explained later), get them out and let B handle them. Next, they have to be routed further to their destination, which, in this case, is possible by sending them out again as overspill packets on λ_3 (λ_1 may also be an option). At node c they need to get lifted out again and be handed over to C, which does the same. In our case they are sent out again as overspill packets, again on λ_3 (again, λ_1, if free, could be used). The packet flow can be seen on Fig. 4. Packets in gray travel on an ORION interface (see further) between the OxC (b) and the IP/MPLS router (B). The packet flow is indicated by the numbers in the packets. So how will we support this ? A conceptual ORION node is shown in Fig. 5 (we will go in more detail in Section III. The OxC is again present, and allows for the node to operate in wavelength switched WDM mode.

In addition to this functionality, there are "taps" (1) added before the OxC, as well as "drains" (2) at the exit. The taps are capable of temporarily redirecting the established wavelength paths to the IP router. This means that, while packets on these wavelength paths normally bypass the IP router, we reintroduce the capability of local termination. To save interface cards however, the taps are grouped together (3). In this example all are grouped together in one interface, effectively limiting redirection to the IP router to only one packet at a time. At the same time, at the exit, the drains are actually tunable lasers, that can tune to any wavelength supported within the fiber.

What this kind of node concept allows us to do is, first of all, operate it just like a standard wavelength switched WDM node. This means that data inserted in a wavelength follows a predestined path, as with standard wavelength switching. However, it can occur, as explained earlier, that there is not

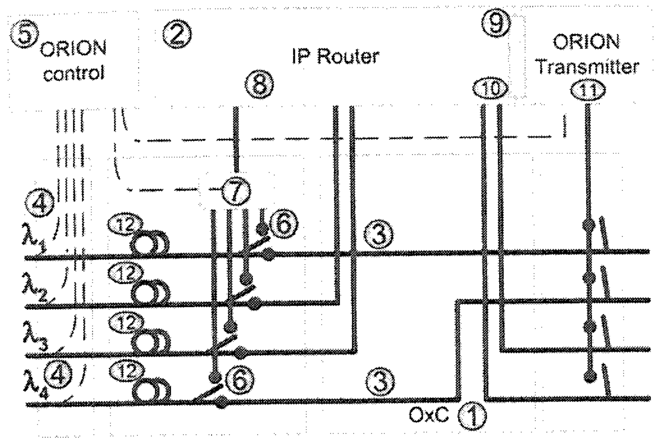

Fig. 6. An example node architecture for the ORION network.

enough capacity on a direct wavelength path. The taps and drains then allow for an operation of the node in a point to point WDM like mode. Changing this mode of operation (between wavelength switched and point to point WDM) will happen on the fly at packet resolution level, by using a special tag for the overspill packet (see Section III), which overrides the standard wavelength switched behavior. We define that specific packet to be in overspill mode from then onwards, as it "spills over" to underused wavelength paths. When the overriding tag is detected, the data is treated as if switched in point to point WDM mode, and is directed through the taps to the IP router. Note the fundamental difference with deflection routing [9] here: deflection routing is not capable of temporary disabling of the wavelength switched behavior. Deflection can only use a concatenation of wavelength paths, possibly forcing data to take a longer physical path. ORION, in contrast, offers full freedom. Note that this is also not a multicast light-tree [10] concept: by default the packet goes through transparently, and only exceptionally it passes through the IP router. No duplication (e.g. drop & continue) of any packet is performed at any time, like in [11] or [12]. At the exit of the node we add generic insertion capability (2), the counterpart of the removal ability at the entry of the node. With the addition of the tunable lasers, any wavelength can be chosen, as opposed to the inflexible wavelength switched approach, where wavelengths that do not terminate locally are inaccessible by the IP router. We will use this ability to send data in overspill mode on wavelength paths that are not fully used. One important problem is that, when sending data on a normally inaccessible wavelength, we need to be careful not to garble up data which may possibly be already present on the now accessible wavelength path. Therefore, we need some form of idle detection. How this concept can be implemented is further clarified in the next section.

III. DETAILED NODE DESIGN

To illustrate that we can achieve the required functionality for ORION without having to terminate each and every packet at an intermediate node, we (functionally) designed an example switch. We just describe it here, but we actually built a working ORION network, although the optics were emulated to reduce deployment costs [24]. Fig. 6 depicts the design of a single node in a four wavelength network. No assumption here is made whether the incoming wavelengths belong to the same fiber or not. Also, for simplicity in the discussion, we assume the wavelengths can be converted to each other without problems. However, we would like to stress this is no requirement of ORION.

Fig. 6 shows a central optical cross-connect (1), which has four wavelengths coming in, of which two, λ_2 and λ_3 are terminated towards the electrical IP router (2), and the two others, λ_1 and λ_4, pass through transparently (3). Each packet is assumed to be intensity modulated, while at the same time there exists a possibility to modulate orthogonally a label at slower speed than the line rate. An example of this orthogonal labeling approach, using FSK or DPSK, can be found in [17]. A very important characteristic is, from a network perspective, that this orthogonal signal is at a slower line speed, which makes the electronics that have to process it cheaper.

As a way of detecting whether or not a packet is in overspill mode, i.e. it is on a wavelength passing through but it does not belong there, the orthogonal header contains a label indicating this status. Before the packet enters the switch we split some of the power at (4) and read this label (5). If an overspill packet is detected the 1x2 fast optical switches (6) are set up to lead the packet towards the electrical IP router. Note that in the figure we group together (7) all wavelengths to save interface cards. It reflects the expectation that packets in overspill will form the exception rather than occur very frequently. It however also implies that we need some contention resolution for the case that both pass-through wavelengths contain overspill packets at the same time, since we can only receive one overspill packet at a time. Deciding on how many wavelengths should be grouped together in this way is a design parameter which depends on how many overspill packets there can be expected simultaneously. In general, we reduce the range of wavelengths where 1x2 switches are installed, M, to a smaller range N. The best implementation approach and the exact dimensioning of both M and N remain topics for further study.

After the 1x2 switch is set, the overspill packet can be received through a wide-band receiver (8). The electrical IP router can then decide (9) to either send it via the normal regime (wavelength switched) with fixed (or slow tunable) lasers (10), either again in overspill mode (11). In standard ORION a packet always remains in overspill mode (once overspill, always overspill). If, however, we allow overspill packets to enter a (under-used) wavelength path originating at an intermediate node leading to the correct destination, we use the term ORION with wavelength path re-entry [23]. Note that there are no 1x2 switches needed on wavelengths λ_2 and λ_3, since these wavelengths are received by the IP router anyway. They can, however, be useful when rearranging the network (in the ASON way). After such a rearrangement, λ_2 and λ_3 may become part of a wavelength path, and thus no longer terminate locally, at which point the overspill capability on these wavelengths can be put to good use.

The reading and detection, together with deciding on eventual contention resolution and the setting of the 1x2 switches of course takes some time. Therefore, we provide for some delay lines (12). Note that all wavelengths can be equipped with delay lines. This allows keeping all wavelengths multiplexed in the fiber until after the delay line. The orthogonal signal detection would then require an extra demultiplexer (hence leading to a trade-off between having one extra multiplexer or several delay lines per fiber).

The required delay lines also have a second function if we want to insert packets in overspill mode on the two pass-through wavelengths. The same way we used an orthogonal label to identify overspill packets, we can also apply it for detection of idle periods on the wavelengths. Each normal packet (i.e. all packets not in overspill mode) can be fitted with two labels, one at its end and one at its beginning, identifying the exact position of the underlying (intensity modulated) packet on the line. Alternatively power detection instead of orthogonal labels could be used to assess whether or not a packet is present on the pass-through wavelengths, or only a label at the beginning of the packet which also indicates its length. Which option is more appropriate (from a feasibility point of view) is not clear yet. In both cases, however, we thus acquire knowledge of the start and end of a normal (lightpath) packet.

The orthogonal label thus can have a double function:

- It identifies overspill packets
- It identifies the start and end of occupied periods on the line

The information about availability is passed through to the ORION control module (5), whose function now also is to provide the electrical edge router with a wavelength on which an overspill packet can be sent, should the provided capacity be insufficient.

In order to be able to access all outgoing wavelengths without having to install a transmitter (and interface card) for each wavelength, a fast tunable laser (which we term ORION transmitter) can be employed (11). Again, the shown schematic implies that only one overspill packet can be sent at a time, which can, of course, be changed by adding more tunable lasers.

The advantages of ORION are already reflected partly in the node design. First of all, the electronics controlling the overspill part need to work only at the speed of the orthogonally labeled signals, which can be much lower than the actual line rate (e.g. 622 Mb/s for a 10Gb/s system). Secondly, packets which are inserted on a direct wavelength path are left untouched until they reach their destination. This means that no expensive interface cards at the line rate are needed for pass-through traffic, and that the delay of the packets throughout the network can be guaranteed, and are minimal. Once they enter the direct wavelength path, their delivery is assured (apart from physical corruption). And finally, since we only rely on information obtained via the orthogonal label, this part of the node is totally oblivious to the underlying line encoding or data rate. A drawback of this scheme is that all packets need an orthogonal label, which will make the transmit infrastructure more expensive. Possible solutions, as mentioned earlier, may be to use power detection instead of orthogonal markers to observe the line. That way, orthogonal labels are only needed for overspill packets. Finally, note that the IP router (2) in this design can be replaced by an OPS/OBS switch. The advantages of ORION still apply: smaller OPS/OBS switches (less traffic seen, which means less processing) and a high utilization rate, resulting in potentially cheaper networks.

IV. DIMENSIONING

A. *Evaluated Technologies*

Through case studies we compared standard ORION (without wavelength path re-entry) with three other technology types. We compared them using a double metric: on one hand the amount of wavelengths needed in the network to fulfill demands, and on the other hand the amount of traffic the IP routers will have to handle. The types we compared were :

- ORION : our new proposed concept
- point to point WDM : Sometimes referred to as fully packet switched, or link-by-link grooming, this is one of the two technology types composing ORION. Since full statistical multiplexing is possible, all demands over a link can be combined, resulting in a minimal amount of required wavelengths. However, a maximum amount of packets need to be switched in the IP routers.
- wavelength switched WDM : Also referred to as "end-to-end wavelength paths", or no grooming, this is the other technology type ORION consists of. Since there is no statistical multiplexing between demands, having either a different source or destination node (as explained earlier), no demands can be combined, which results in the highest amount of wavelengths needed to route them. In essence, all demands on a link are dimensioned for separately. On the other hand, since there is no transit traffic anywhere, IP routers only have to route demands originating or terminating here. Thus, a minimum amount of traffic will go through the IP routers. Note that we assumed full wavelength conversion throughout all of our evaluations, as well as no restrictions for possible wavelength paths (e.g. length).
- Combined : As a means to compare ORION with another hybrid architecture, we also evaluated a wavelength switched WDM network combined with a parallel point to point WDM network. Unlike ORION however, these two networks are completely separate, and cannot use each others spare capacity. Incoming traffic at the edge of the network uses the wavelength switched WDM network if capacity on a direct wavelength path is available, otherwise it is sent over the parallel point to point WDM network. Like standard ORION (also termed "once ORION, always ORION"), once traffic enters the point to point WDM network (*overspill mode* in case of ORION) it does not re-enter the wavelength switched WDM network. Note that the combined architecture can be considered a form of grooming. The CHEETAH [13] concept, although slightly different, can also be considered as a form of combined.

B. *Dimensioning Method*

1) Moving from a fixed demand to a distribution: Often, when wavelength switched WDM and point to point WDM are compared, static traffic profiles are used [1]. A certain value is assigned for each node pair, reflecting how much traffic

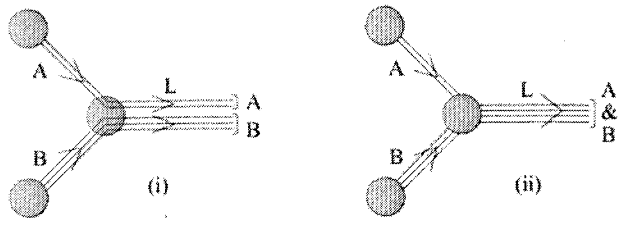

Fig. 7. Example of separate dimensioning (i), as in wavelength switched WDM, and grouped dimensioning, as in point to point WDM.

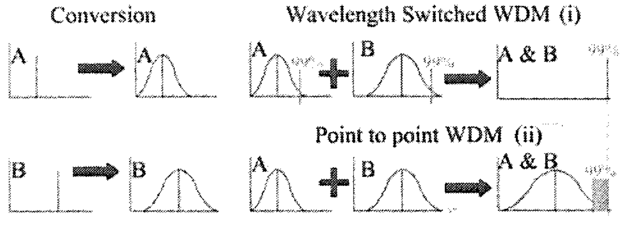

Fig. 8. Illustration of the conversion to distributions instead of crisp demands (left), and an example calculation used in the dimensioning (right).

K = the network we are dimensioning

$\mathbf{L_i}$ = Link i in K

$\mathbf{W(L_i,S)}$ = #wavelengths installed on L_i in a network S

$\mathbf{D(L_i,S)}$ = collection of demands in S carried by L_i

$\mathbf{P_\alpha(D)}$ = percentile x (in %) of demand D

CEIL(x, y) = x rounded up to the nearest unit of y

$\mathbf{C_\lambda}$ = capacity of one wavelength

$\mathbf{T(N_i)}$ = size of IP router N_i (in Gb/s) in K

$\mathbf{D_{end}(N_i,S)}$ = collection of demands leaving network S at IP router N_i

$\mathbf{D_{start}(N_i,S)}$ = collection of demands entering network S at IP router N_i

$\mathbf{D_{transit}(N_i,S)}$ = collection of demands in network S passing IP router N_i

Fig. 9. Definitions used in the dimensioning algorithm

is requested from one node in the network to another. Based on these values necessary provisions are then calculated. In order to be able to do a fair calculation, however, a static demand matrix is not sufficient. These type of calculations may underestimate the statistical multiplexing (if any) effect when operating in a shared environment, like our point to point WDM, resulting in an overestimation of required capacity.

For example, suppose there is a demand A of 20 Gb/s, and a demand B of 30 Gb/s, travelling over the same link L (fiber), as in Fig. 7 (left). In a typical calculation, supposing 10 Gb/s wavelengths, this would mean that demand A got 2 wavelengths over this fiber, while demand B gets 3. In reality however, these demands will fluctuate, which can be more accurately modeled by distributions, which at the same time takes the statistical multiplexing effect into account. Suppose we want to be sure that in 99% of the cases both traffic streams A and B will not encounter a bottleneck on this particular individual link. We now have two different cases : On the one hand, the same situation as in Fig. 7 (left): demands A and B do not share their capacity (as in wavelength switched WDM). This means that of each distribution the 99% percentile must taken individually, a process shown in Fig. 8 (top). In Fig. 7 (right), on the other hand, A and B share their capacity (point to point WDM), which means it is sufficient to take the 99% percentile of their combined distribution, as in Fig. 8 (bottom). This results in a capacity saving, the gray area shown in the lower right corner of Fig. 8.

This method of working did, however, present us with a problem, since in cases IIa-IIc, the provided traffic demand forecast provided only crisp (i.e. fixed) values, so we had to adapt them to distributions as used in our model. Since we had no information available concerning traffic variability in a network of such scale, we based ourselves upon measurements of two smaller, public networks. The first one is Belnet [31], a Belgian public network, the second Uninett [32], the Norwegian public network. The general validity of our calculations is hence directly tied to these networks as being representative (in terms of user behavior) for most networks. Since the results of both networks are about the same, we believe this is the case. We used them to construct an equation to transform the static demand to a normal distribution. We took the value itself as the mean μ, and the standard deviation (σ) was calculated as $\sigma = 0.52 * (\mu^{0.914})$. For example, if the traffic demand matrix predicts a demand of Hamburg to Berlin of 37,1 Gb/s for 2008, we will model this traffic as Gaussian with a mean of 37,1 Gb/s, and a standard deviation of $0.52 * (37.1^{0.914}) = 12{,}9$ Gb/s.

2) Algorithm: In this paragraph we describe, from a high-level viewpoint, the algorithms used to calculate our results in more detail. The (Java) source code can be obtained upon request to the authors. All percentile calculations, needed to go from distributions back to crisp numbers, were obtained through sampling, with a 99% statistical confidence that the obtained number is within a 0.1% interval [33]. We used the COLT library for this statistical support, a software distribution providing an infrastructure for scalable scientific and technical computing in Java [34]. For the network logic we resorted to the TRS library, developed at our group [25]. Some additional custom classes provided for generic distribution math and special purpose logic. In order to express our calculation method a little more formally and consistent, we first define some symbols, shown in Fig. 9.

Using these definitions, the formula used to calculate the amount of wavelengths for a given link, as well as for determining the amount of traffic in an IP router, becomes quite simple, as shown in Fig. 10. We assumed shortest path routing, with a weight of 1 for each link. In essence, it is the application of the techniques described in the previous paragraph. In the case of point to point WDM (Fig. 10 - left), there are no optical bypasses. This means that for the IP router dimensioning, transit traffic needs to be taken into account, while in wavelength switched WDM (Fig. 10 - right), no transit traffic is present. The other important difference is that, when calculating the amount of installed wavelengths for *point to point WDM*

$$W(L_i,K) = \mathrm{CEIL}\,(P_{99}(\overset{D \in D(L_i,K)}{\Sigma} D), C_\lambda)$$

$$T(N_i) = \mathrm{CEIL}\,(P_{99}(\,D_{start}(N_i,K) + D_{end}(N_i\,,K) + D_{transit}(N_i\,,K)), C_\lambda)$$

wavelength switched WDM:

$$W(L_i,K) = \overset{D \in D(L_i,K)}{\Sigma}\ \mathrm{CEIL}\,(\,P_{99}(D), C_\lambda)$$

$$T(N_i) = \mathrm{CEIL}\,(\,P_{99}(D_{start}(N_i\,,K) + D_{end}(N_i\,,K)\,), C_\lambda)$$

Fig. 10. Expressions for the calculation of installed wavelengths and IP router traffic for the two baic technology types

```
ORION / combined common code
        $K_1$ = dimensioned network from point to point WDM algorithm
        initialize O, the list of open demands, to all demands
        initialize R, the list of remaining demands, to { }
        while ( O != { } ) {
                D = highest demand $\in$ O
                attempt wavelength path setup in $K_1$ for D, along shortest path
                        => succes: confirm path & reduce D by $C_\lambda$
                        => failure: remove D from O, add D to R
        }
```

Fig. 11. First portion of the pseudo-code for calculating the amount of wavelengths installed per link. This portion is common for ORION and combined. Note that this indeed means that before this pseudocode K_1 needs to be obtained by executing the point to point WDM calculation algorithm.

```
combined only code:
        wavelength switched network
                = network $K_1$ from previous step
        point to point WDM network $K_2$
                = result of point to point WDM algorithm
                  with remaining demands R
        $W(L_i,K) = W(L_i,K_1) + W(L_i,K_2)$, with:
                • wavelength path fraction = $W(L_i,K_1)$
                • point to point WDM fraction = $W(L_i,K_2)$
ORION only code:
        $W(L_i,K) = W(L_i,K_1)$, with:
                • wavelength path fraction = $W(L_i,K_1)$ - $W(L_i,K_2)$
                • overspill (point to point WDM mode) fraction = $W(L_i,K_2)$
```

Fig. 12. Second portion of the pseudo-code for calculating wavelength usage, for combined and ORION, calculating the K_2 network. The ORION only code assumes calculation of K_1 and K_2 as in the combined code.

a link, for point to point WDM, the demand distributions are combined before returning to a crisp value. This is in contrast with the wavelength switched WDM calculation, which first returns to crisp values for each demand, after which the summation of all demands over a certain link is taken.

Since ORION and combined are both hybrids, they have similar methods for calculating wavelength usage. The algorithm for IP router traffic is even identical. The wavelength usage calculation is more complicated compared to the previous two technology types, since ORION and combined are actually a combination of them. As a result we actually need to dimension two networks, K_1 and K_2, before we are able to derive the final result. In case of combined these networks are directly used in the end result, whereas for ORION the K_2 network is actually "physically embedded" in the basic network K_1. The first portion of the algorithm is in common, as shown in Fig. 11. Basically, the idea is to first determine how many wavelengths there are needed at a bare minimum to fulfill demands, as in point to point WDM. In a next step this point to point WDM network, K_1, is converted to a wavelength switched WDM network, setting up as much paths as possible using simple shortest path calculations. Then the second portion of the algorithm is executed, which differs for ORION and combined as shown in Fig. 12.

In the case of combined, which consists of a wavelength switched WDM and a separate point to point WDM network, the converted network K_1 is used as wavelength switched WDM network, and a point to point WDM network, K_2, is constructed to fulfill the unserved leftover demand from the first portion. Note that, like we do in the ORION calculation as well, we fix a degree of freedom here: we can always add more wavelengths to K_1, reducing IP router strain, or setup less wavelength paths to reduce wavelengths used. There are two reasons we use the dimensioned capacity for point to point WDM as resource pool for setting up wavelengths : One is the fact that we will at least need this amount of wavelengths anyway in the network. The other is that ORION always sets up as many paths as possible with the given resources, and using this basis makes it easier (and fairer) to compare with combined. The second portion of ORION differs from combined, since ORION actually uses spare capacity within the wavelength switched WDM network. This means that ORION can be thought of as having the point to point WDM network, K_2, physically embedded within the wavelength switched WDM network K_1. As a final note we would like to remark that although the presented evaluation algorithms and method may be rather simple, other, more complex simulation [27] and evaluation [14] efforts confirm the results we obtained.

V. EVALUATION

A. *Cases*

In this paper we compared these four technology types using two different topologies in four case studies. The first topology, shown in Fig. 13, is only used in case I. It is a simple 4x4 grid - "Manhattan street" - topology.

The second topology, shown in Fig. 14, is used in the three other presented case studies, IIa through IIc: It is the more elaborate reference core network topology, developed within the COST266 [28] and IST-LION [29] projects. It describes a hypothetical pan-European network, with corresponding demand forecasts for 2008. Both the topology and traffic demands are available at [30]. These cases, IIa to IIc, study different aspects of the ORION behavior: Case IIa, is a detailed discussion where the network is more realistic than in

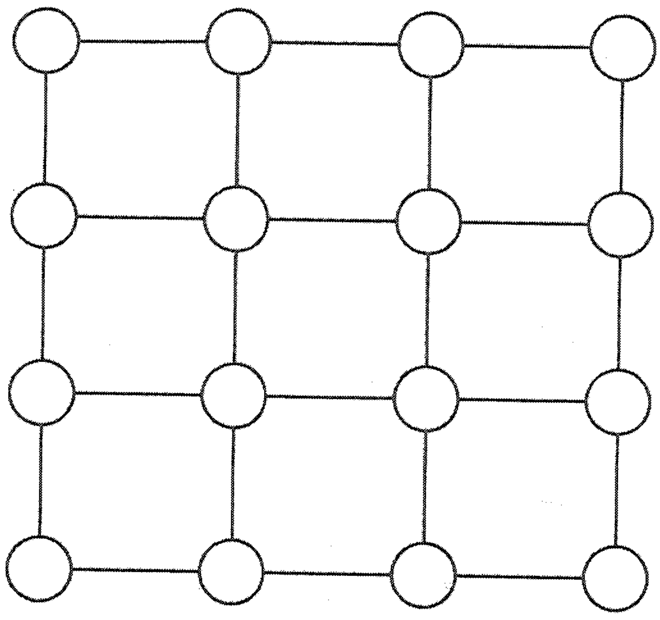

Fig. 13. The grid topology of the first case study. Note that there are no wrap-arounds at the edges.

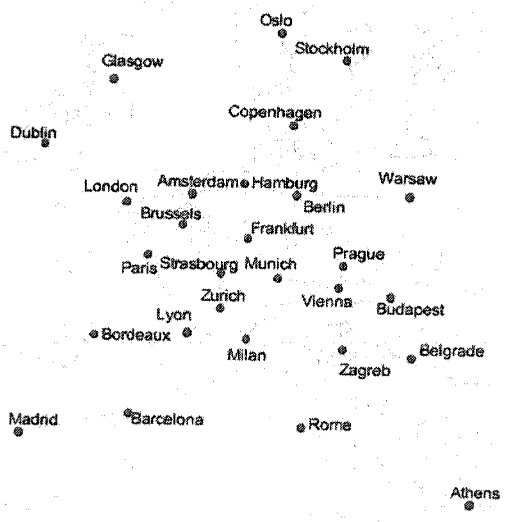

Fig. 14. The reference network used in case studies IIa-c.

case I (the grid topology); Case IIb, shows how the technology types react when traffic patterns exhibit increasing dynamism. The last case study, Case IIc, shows how the technology types compare when the wavelength capacity is increased. For every case we dimensioned each network four times. Each time we used a different network technology (cf. Section IV-A, and took into account its specific characteristics.

B. Case I: Grid Topology

In this first case we evaluated a simple "Manhattan street" topology of 16 nodes (4x4 grid). All demands were the same, having a normal distribution with mean (μ) of 16 Gb/s, and standard deviation (σ) of 7,63 Gb/s. Fig. 15 shows, for each of the four technology types, the average amount of traffic an IP router sees (left bars), and the average amount of wavelengths (right bars - 10Gb/s per wavelength) that would be installed on a link. Note that in the case of ORION and combined the right bars consist of two components. For combined the two components indicate how much wavelengths each separate network it consists of consumes. For ORION, since bandwidth is assigned ad hoc, the interpretation is slightly different : the "packet fraction" indicates how much traffic at most is expected to be in overspill (¿ 99% of the time). The total is the amount of wavelengths installed in the network.

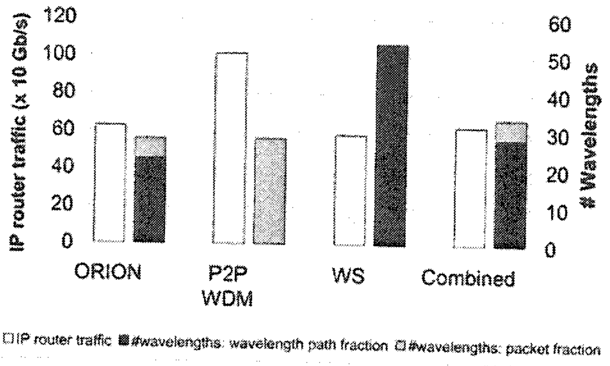

Fig. 15. Case study results of the grid topology, per techonology type : average amount of traffic seen by IP routers (left bars - left axis) and average number of wavelengths required in the network to fulfull demand (right bars - right axis).

When looking at average IP router load, it is very clear that point to point WDM has to handle a lot more than the other three alternatives. The wavelength switched WDM solution naturally performs the best, but more interestingly, the combined technology and ORION are very close to this lower limit. When looking at average wavelengths needed in the network, the right bars, we see the opposite (again, less is better). The wavelength switched WDM architecture uses a lot more wavelengths than the others. ORION uses the same amount of wavelengths as the packet switched case. We also see an important feature of ORION illustrated: only a small fraction is in overspill (grey area). At the same time, ORION has to switch only a little more packets than wavelength switching. Combined approaches ORION in terms of wavelength usage, but uses slightly more wavelengths. Both hybrids however, are remarkably balanced in resource usage when compared to the two more extreme approaches. Therefore, it is quite clear that these type of hybrid architectures make most (economical) sense when bandwidth and switching costs are both significant cost driving factors.

C. Case IIa: NRS Network: Detailed

A more realistic case study is based on Fig. 14, with a corresponding traffic demand forecast for 2008 [30]. This case study also assumes 10 Gb/s per wavelength installed. As stated earlier, only end-to-end direct wavelength paths are set up, to enable us to compare two extreme network architectural options, point to point WDM and wavelength switched WDM respectively. Fig. 16 shows the results summarized, the same way as with the discussed grid topology.

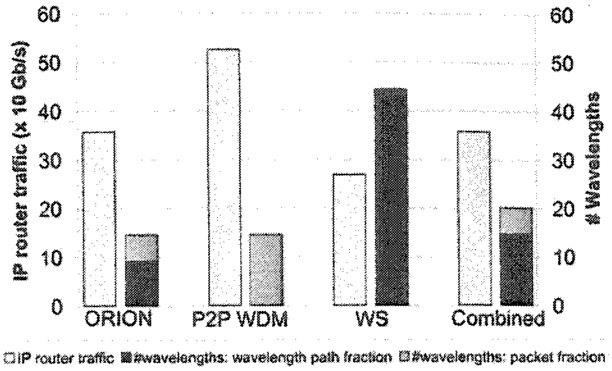

Fig. 16. Case study results of the nrs network, per techonology type : average amount of traffic seen by IP routers (left bars - left axis) and average number of wavelengths required in the network to fulfull demand (right bars - right axis).

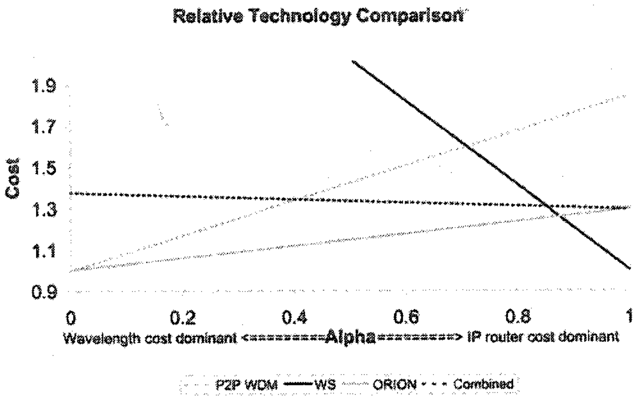

Fig. 17. A crude comparison between four technology types, depending on relative cost.

The same trend shows up: ORION and combined strike a balance between its two composing architectures, point to point WDM and wavelength switched WDM. Another fact, not shown on the figure, is that there are some hot spots in the network, where a lot of transit traffic passes. These hot spots result in a very large difference between the packet and wavelength switched case. For example, in Berlin, the packet switched network needs to process more than three times the IP traffic the wavelength switched based solution sees. These hot spots also show up in other studies [1]. However, ORION, having to handle the same amount of packets as the combined case, shows a significant improvement, leaning close to the wavelength switched WDM case. At hot spot Berlin (worst case) the increase in IP router strain is about 80% (instead of over 300% experienced by point to point WDM!) when compared to wavelength switched WDM.

Fig. 17 presents these results from a crude economical standpoint. Shown is the total network cost, in function of a factor alpha (α), which indicates the relative importance in the cost of respectively the amount of traffic a node needs to switch, and the amount of wavelengths that need to be provisioned in the network. The cost is defined as :

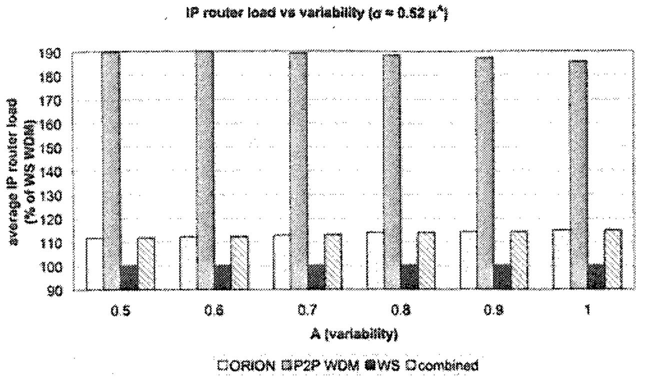

Fig. 18. Impact of increasing variability on IP router strain for the four different technology types: Y-axis expresses load as a % of the wavelength switched case.

$$\alpha * \frac{averageIProuterLoad}{averageIPRouterLoadWavelengthSwitched} + (1-\alpha) * \frac{averageWavelengthsUsed}{averageWavelengthsUsedP2PWDM}$$

Thus, α=0 indicates that the only deciding factor is the amount of wavelengths, while α=1 indicates that the cost of handling packets will be the sole cost factor. The cost for transmission resources furthermore is normalized to the minimal amount of wavelengths (averaged out) required (i.e. the wavelength cost of the point to point WDM architecture is 1), while the cost for processing resources is normalized to the minimal average amount required in the nodes (i.e. the node cost of wavelength switched WDM is 1). As clearly can be observed from the figure, ORION comes out as the most cost-effective option for almost all values of α, being consistently better than the combined architecture. Only when the amount of packets handled is very dominating in the total cost ($\alpha >$ 0.9) does the wavelength switched WDM option come out cheaper. Note the interesting cross-over point around α=0.7, where the cost of both "traditional" options is equal.

The most important conclusion however is the trend that ORION is the most cost-effective architecture, unless the handling of traffic becomes the dominant cost factor. Therefore it seems an architecture like ORION is mostly suited as a transition architecture for the medium term. In the long term, if the current traffic growth continues while it retains its characteristics, wavelength path based solutions will probably become the best solution for the backbone (granted, this is assuming bandwidth growth mainly comes from an increased amount of wavelengths, not increasing bandwidth per wavelength). The ORION architecture then can migrate gradually towards the user (in the metropolitan network), until even there wavelength services become predominant. Note that this is, however, a crude comparison, since much depends on the cost of an actual ORION node. It does however, illustrate the potential usefulness of ORION.

D. Case IIb: NRS Network: Impact of Increasing Dynamism

In order to generalize the obtained case study results, we also studied what happens when traffic varies more (or

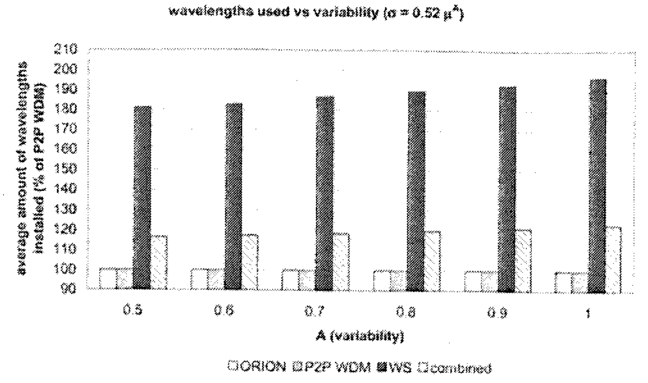

Fig. 19. Average number of wavelengths per link in function of variability, for the four technology types:. Y-axis expresses wavelength usage as a % of the point to point WDM case.

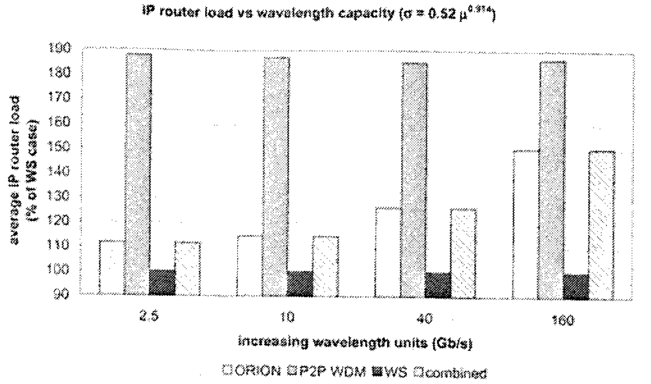

Fig. 20. Impact of increasing wavelength capacities, under the same traffic conditions.

less). We did this by changing the equation, $\sigma = 0.52 * (\mu^{0.914})$, used to convert our crisp demands to distributions. More specifically, we changed the exponent, introducing a new variable A (variability), indicating how strongly traffic varies through time. Fig. 18 shows the results of the IP router dimensioning when this A is gradually increased, effectively increasing standard deviation while the average demands are kept constant. The range of A shown is what reasonably can be expected in a typical network, again based upon the Uninett and Belnet measurements. To show growth trends more clearly, the results for each value of A are expressed as a percentage of the wavelength switched WDM case, which has the smallest IP router load.

The two hybrids (ORION and combined) increase slightly compared to wavelength switched WDM. This can be explained by the fact that more traffic will experience overspill as variability increases, thus increasing the amount of packets seen by IP routers. In effect, the increased variability causes ORION to behave more like a point to point WDM architecture. On the other hand, the gap between wavelength switched WDM and point to point WDM is somewhat reduced. This can be explained by the fact that more demands pass the IP router in point to point WDM, enabling point to point WDM to make better use of statistical multiplexing, offsetting the high amount of transit traffic somewhat. The advantage of ORION compared to point to point WDM however, remains very large. Fig. 19 shows the other side of the dimensioning, the average amount of wavelengths installed per link, now normalized to point to point WDM, which uses the smallest amount of wavelengths. Clearly the wavelength switched WDM approach shows the least tolerance to highly variable traffic patterns, causing it to overdimension even more (nearly +20%). The combined architecture suffers significantly less (+16%), but is still beaten by the flexibility of the point to point WDM approach, and ORION. Note that ORION in this case is better than combined, whereas when looking from an IP router load standpoint, they perform the same. The most important conclusion from these two trend graphs regarding the hybrids is, however, that they can withstand variability with regard to transmission resource usage, but are slightly more vulnerable on the IP layer. In other words, as long as variability stays within the current traffic trends, ORION offers significant advantages. If, however, variability is extremely high (which can happen in current metro or access networks) or very low (well below traffic patterns of current backbone networks), hybrids may not be the perfect tool for the job.

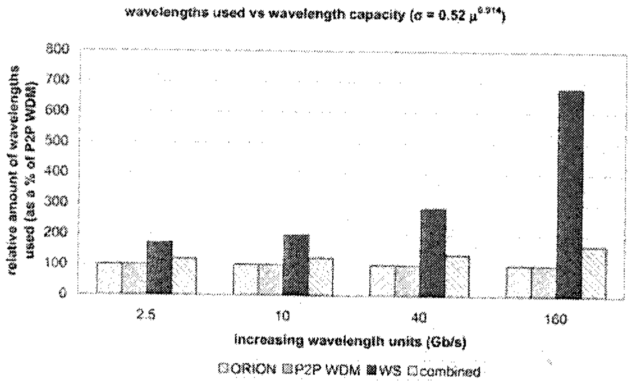

Fig. 21. The other side of the dimensioning: amount of wavelengths used with increasing wavelength capacities. Y-axis is in percentage relative to point to point WDM.

E. Case IIc: NRS Network: Impact of Wavelength Capacity

With the ever increasing bitrates per wavelength, a natural question would be how technologies respond when demands become smaller compared to the wavelength capacity. To determine this we fixed all parameters like in the previous case, but now only changed the bitrate of a wavelength. Fig. 20 shows the impact on the IP router load, again expressed as a percentage of the wavelength switched WDM case. Although traffic patterns do not change, remarkable trends show. While point to point WDM and wavelength switched WDM remain largely unaffected, the combined architecture and ORION rise very significantly in terms of IP router strain (up to more than 50%), moving fast towards the point to point WDM. This is explained by the fewer wavelength paths that can be set up. As a consequence, more traffic needs to be transported in overspill. In fact, in case of 160 Gb/s wavelengths, some demands did not even have a single wavelength path, and all their traffic went in overspill through the network. This actually misuses the overspill mechanism, since it is not meant to transport a large fraction of the network traffic.

As in the previous case, we also show, in Fig. 21, the impact on the average amount of wavelengths deployed per link. Again, to provide for a better comparison between the four technology types, results for different bitrates are normalized to the technology using the least amount (point to point WDM). As the graph shows, the wavelength switched WDM architecture wastes enormous amount of wavelengths with increasing wavelength capacities (up to nearly 7 times as many wavelengths). This is easily explained since most demands occupy only a fraction of a wavelengths capacity, and wavelength switched WDM only sets up direct wavelength paths. Thus even the smallest demands have their own wavelength path. ORION avoids this by not setting up a path at all for most demands, but the downside is, as shown in Fig. 20, that a lot of traffic is in overspill. The same reasoning holds for the combined architecture. If the wavelength capacities are rather small (< 40 Gb/s), ORION seems the best option. At the other end of the spectrum (> 40 Gb/s) only point to point WDM currently (as in, for the given traffic forecast) seems to make sense amongst the four discussed types.

VI. CONCLUSION

In the previous paragraphs we described and evaluated a new network concept, ORION. Its main advantage is the effective combination of point to point WDM, and wavelength switched WDM concepts. The result is a very flexible architecture, capable of using a minimal amount of bandwidth while maintaining a relatively low amount of IP router traffic, relieving the upcoming electronic forwarding bottleneck. A wide variety of cases was studied, from which it is also clear that ORION is not the best option in all cases. Indeed, the design itself assumes that both link bandwidth and forwarding capacity of IP routers (whether they are optical or not) are at least within the same order of magnitude regarding cost. If one of these factors becomes totally dominant, such as bandwidth was in the past, ORION will not prove to be most effective. Therefore, we see ORION mainly as an architecture that may prove very useful as an intermediate stage in the evolution of optical networks. The extent of the application of the concept will of course ultimately depend on the cost of implementation of an ORION node, when compared to other alternatives. To this end, we already performed some further studies. One is a more sophisticated comparison of ORION with, amongst others, an architecture where arbitrary grooming is allowed [14], this time based on loss performance and throughput of a pre-dimensioned network. Also detailed packet level simulations [26] and even an emulation platform [24] were developed, which all support the general conclusions presented here. But a lot needs to be done : e.g. some variants of ORION, like the wavelength path re-entry option, remains to be evaluated. As mentioned earlier, in this variant packets launched in overspill are allowed to re-enter a wavelength path originating at any intermediate hop going to the same destination. This variant requires a more complex control mechanism, discussed in [23], but reduces traffic seen by the IP router even further. And last but not least there are several hardware constraints and challenges which we hope to address in the near future.

ACKNOWLEDGEMENT

The authors would like to thank Uninett for providing the traffic measurements, and Belnet for making their MRTG graphs publicly available.

REFERENCES

[1] S. De Maesschalck *et al.*, "Pan-European optical transport networks: an availability-based comparison," *Photonic Network Commun.*, vol. 5, no. 3, pp 203-225, May 2003.

[2] P. Green, "Progress in optical networking," *IEEE Commun. Mag.*, vol. 39, no. 1, pp. 54-61, Jan. 2001.

[3] N. Ghani, S. Dixit, and T. S. Wang, "On IP-over-WDM integration," *IEEE Commun. Mag.*, vol. 38, no. 3, pp. 72-84, Mar. 2000.

[4] S. Bjornstad, D. R. Hjelme, and N. Stol "A highly efficient optical packet switching node design supporting guaranteed service," in *Proc. ECOC 2003*, Mo4.4.5, pp. 110, 111.

[5] S. Bjornstad, M. Nord, and D. R. Hjelme "QoS differentiation and header/payload separation in optical packet switching using polarisation multiplexing," in *Proc ECOC 2003*, Mo3.4.6, pp. 28, 29.

[6] P. Gambini, *et al.*, "Transparent optical packet switching: network architecture and demonstrators in the KEOPS project," *IEEE J. Select. Areas Commun.*, vol. 16, no. 7, pp. 1245-1259, Sept. 1998.

[7] C. Qiao and M. Yoo, "Optical burst switching: a new paradigm for an optical internet," *J. High Speed Networks*, special issue on optical networking, vol. 8, no. 1, pp. 36-44, Jan. 2000.

[8] M. Dueser and P. Bayvel, "Analysis of a dynamically wavelength-routed optical burst switched network architecture," *J. Lightwave Technol.*, vol. 20, no. 4, pp. 574-586, Apr. 2002.

[9] F. Borgonovo, *Deflection Routing, Routing in Communications Networks.* Hertfordshire, UK: Prentice Hall International (UK) Ltd., 1995.

[10] L. H. Sahasrabuddhe and B. Mukherje. "Light-trees: optical multicasting for improved performance in wavelength-routed networks," *IEEE Commun. Mag.*, vol. 37, no. 2, pp. 67-73, Feb. 1999,

[11] A. Gumaste and I. Chlamtac, "Light-trails: a solution to IP centric communication at the optical domain," in *Proc. IEEE High Performance Switching and Routing HPSR.*

[12] A. Gumaste, I. Chlamtac, and J. P. Jue, "Light-frames: a pragmatic framework for optical packet transport," in *Proc. 2004 IEEE International Conference on Communications*, vol. 3, pp. 1537-1542.

[13] M. Veeraraghavan *et al.*, "CHEETAH: circuit-switched high-speed end-to-end transport architecture," in *Proc. OptiComm 2003*, pp. 214-225.

[14] E. Van Breusegem, B. Puype , D. Colle, J. Cheyns, M. Pickavet, and P. Demeester, "Evaluation of ORION in predimensioned networks," in *Proc. 19th International Teletraffic Congress 2005.*

[15] A. Hill and F. Neri, "Optical switching networks: from circuits to packets," guest editorial in *Commun. Mag.*, vol. 39, no. 3, Mar. 2001.

[16] S. Junghans and C. M. Gauger, "Resource reservation in optical burst switching:architectures and realizations for reservation modules," in *Proc. OptiComm 2003*, pp. 409-413.

[17] K. Vlachos *et al.*, "An optical IM/FSK coding technique for the implementation of a label-controlled arrayed waveguide packet router," *J. Lightwave Technol.*, vol. 21, no. 11, pp. 2617-2628, Nov. 2003.

[18] J. S. Turner, "Terabit burst switching," *J. High Speed Networks*, vol. 8, no. 1, pp. 3-16, Jan. 1999.

[19] M. O'Mahony, K. M. Guild, D. Hunter, I. Andanovic, I. White, and R. Penty, "An optical packet switched network (WASPNET): concept and realization," *Optical Networks Mag.*, vol. 2, no. 6, pp. 46-53, Nov./Dec. 2001.

[20] L. Dittmann (ed.), *et al.*, "The European IST Project DAVID: a viable approach towards optical packet switching," *J. Select. Areas Commun.*, special issue on high-performance optical/electronic switches/routers for high-speed Internet, vol. 21, no. 7, pp. 1026-1040, Sept. 2003.

[21] S. Yao, S. J. B. Yoo, B. Mukherjee, and S. Dixit, "All-optical packet-switched networks: a study of contention-resolution schemes in an irregular mesh network with variable-sized packets," in *Proc. OPTICOMM 2000.*

[22] R. Ramaswami and K. N. Sivarajan, "Design of logical topologies for wavelength-routed optical networks," *IEEE J. Select. Areas Commun.*, vol. 14, no. 5, pp. 840-851, June 1996.

[23] J. Cheyns, E. Van Breusegem, D. De Winter, D. Colle, M. Pickavet, and P. Demeester, "Controlling LSPs in an ORION network," in *Proc. Broadnets 2004, First Annual International Conference on Broadband Networks*, pp. 74-81.

[24] E. Van Breusegem, D. De Winter, J. Cheyns, D. Colle, M. Pickavet, P. Demeester, and J. Moreau, "Emulation of overspill routing in optical networks," in *Proc. ECOC 2005*.

[25] TRS (Telecom Research Software), a Java-library developed in the IBCN-group, intended to be used in telecom-research to speed up the development of tools and applications; more information at http://www.ibcn.intec.ugent.be/projects/INTERNAL/TRS/index.html

[26] J. Cheyns, E. Van Breusegem, D. De Winter, D. Colle, M. Pickavet, and P. Demeester, "A detailed look on ORION: traffic simulations on a node level," invited at NOC 2004, pp. 474-481, July 2004.

[27] K. Christodoulopoulos, E. Van Breusegem, K. Vlachos, M. Pickavet, E. Varvarigos, and and D. Colle, "Network simulation of overspill routing in optical networks," submitted for ICC 2006.

[28] R. Inkret, A. Kuchar, and B. Mikac, "Advanced Infrastructure for photonic networks," extended final report of COST Action 266, Zagreb, Croatia, Faculty of Electrical Engineering and Computing, University of Zagreb, 2003.

[29] B. Puype, Q. Yan, D. Colle, S. De Maesschalck, I. Lievens, M. Pickavet, and P. Demeester, "Multi-layer traffic engineering in data-centric optical networks: illustration of concepts and benefits (IST LION)," in *Proc. ONDM 2003, 7th IFIP Working Conference on Optical Network Design & Modelling*, vol. 1, pp. 211-226, Feb. 2003.

[30] Available Online at: http://ibcn.atlantis.rug.ac.be/projects/IST/NRS/index.html

[31] BELNET,the Belgian national research network for education, research and public services, available online at http://www.belnet.be

[32] Uninett, The Norwegian national research network for education, research and public services, available online at http://www.uninett.no

[33] K. Pawlikowski, H.-D. J. Jeong, and J.-S. R. Lee, "On credibility of simulation studies of telecommunication networks," *IEEE Commun. Mag.*, vol. 40, no. 1, pp. 132-139, Jan. 2002.

[34] The COLT (Collections Tuned) distribution, Open Source Libraries for High Performance Scientific and Technical Computing in Java, available online at http://hoschek.home.cern.ch/hoschek/colt/

Erik Van Breusegem received M. Sc. degree in computer science from the Ghent University in 2001. Since then, he has been working at the same university as researcher in the department of Information Technology (INTEC), pursuing a Ph.D degree. He is part of the research group INTEC Broadband Communication Networks (IBCN) headed by prof. Piet Demeester. In the context of his Ph.D thesis, he was awarded a scholarship from the Institute for the Promotion of Innovation through Science and Technology in Flanders (IWT-Vlaanderen) in the period 2001-2005. His research deals with optical networks in general, with a focus on hybrid optical network designs. He was involved in this context in several IST projects (STOLAS, NOBEL, LASAGNE) and in the COST-action 266. His work has been published in more than 30 scientific publications in international conferences and journals.

Jan Cheyns (jan.cheyns@intec.UGent.be) received an M. Sc. degree in electrical engineering (option micro- and opto-electronics) from the Ghent University in 1998. One year later he also received an M. Sc. degree in physics engineering. In 1999 he joined the Optical Materials and Devices group of the Department of Information Technology of the same university working on GaN materials. In 2000 he was granted a Fund for Scientific Research of Flanders (FWO-Vlaanderen) scholarship, in order to prepare a Ph.D., moving to the Broadband Communications Network Group. He obrtained this Ph. D. in 2005, on packet and burst switched optical networks and hybrid forms of optical networking. He has been involved in several European IST-projects such as DAVID, STOLAS, LASAGNE, NOBEL. Currently he is employed at COMSOF.

Davy De Winter received M. Sc. degree in computer science from the Ghent University in 2003. Since 2004, he has been working at the Hogeschool Gent as researcher at the INWE-department in collaboration with the department of Information Technology (INTEC) at the Ghent University (in the Broadband Communications Networks (IBCN) group headed by prof. Piet Demeester). During the first year of research he was involved in the IST STOLAS project in the context of hybrid optical network designs. Since the start of 2005, his research focuses on thinclient architectures and protocols, and content-distribution in the internet. He is currently also involved in several IBBT-projects (funded by the Flemish governement) (E-paper and QoE ISBO).

Didier Colle received a M. Sc. degree in electrotechnical engineering (option: communications) from the Ghent University in 1997. Since then, he has been working at the same university as researcher in the department of Information Technology (INTEC). He is part of the research group INTEC Broadband Communication Networks (IBCN) headed by prof: Piet Demeester. His research lead to a Ph.D degree in February 2002. He was granted a postdoctoral scholarship from the Institute for the Promotion of Innovation through Science and Technology in Flanders (IWT-Vlaanderen) in the period 2003-2004. His research deals with design and planning of communication networks. This work is focussing on optical transport networks, to support the next-generation Internet. Up till now, he has actively been involved in several IST projects (LION, OPTIMIST, DAVID, STOLAS, NOBEL and LASAGNE), in the COST-action 266 and 291 and in the ITEA/IWT TBONES project. His work has been published in more than 100 scientific publications in international conferences and journals.

Mario Pickavet received a Master of Science degree and a Doctor of Electrical Engineering degree, specialized in telecommunications, from Ghent University in 1996 and 1999, respectively. Since 2000, he has been a full-time professor at the same university. His research interests are related to broadband communication networks (i.e., IP, MPLS, WDM, SDH, ATM) and include resilience mechanisms, design, and long-term planning of core and access networks. In this context, he was and currently is involved in European IST projects (i.e. LION, DAVID, STOLAS, ePhoton/One, LASAGNE, NOBEL) on IP over WDM next generation networks. He has published a number of international publications on these subject, both in leading journals (e.g., IEEE Journal on Selected Areas in Communications and IEEE Communications Magazine) and proceedings of conferences. Together with Jean-Philippe Vasseur (Cisco) and Piet Demeester (Ghent University), he published the book Network Recovery book in the Morgan Kaufmann series in networking.

Filip De Turck received his M.Sc. degree in Electronic Engineering from the Ghent University, Belgium, in June 1997. In May 2002, he obtained the Ph.D. degree in Electronic Engineering from the same university. From October 1997 to September 2001, Filip De Turck was research assistant with the Fund for Scientific Research-Flanders, Belgium (F.W.O.-V.). At the moment, he is a part-time professor and a post-doctoral fellow of the F.W.O.-V., affiliated with the Department of Information Technology of the Ghent University. Filip De Turck is author or co-author of approximately 100 papers published in international journals or in the proceedings of international conferences. His main research interests include scalable software architectures for telecommunication network and service management, performance evaluation and optimization of routing, admission control and traffic management in telecommunication systems.

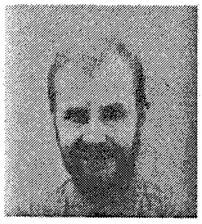

Piet Demeester finished his PhD thesis at the Department of Information Technology (INTEC) at the Ghent University in 1988. At the same department he became group leader of the activities on Metal Organic Vapour Phase Epitaxial growth for optoelectronic components. In 1992 he started a new research group on Broadband Communication Networks. The research in this field resulted in already more than 300 publications. In this research domain he was and is a member of several programme committees of international conferences, such as: ICCCN, the International Conference on Telecommunication Systems, OFC, ICC, and ECOC. He was Chairman of DRCN'98. In 2001 he was chairman of the Technical Programme Committee ECOC'01. He was Guest Editor of three special issues of the IEEE Communications Magazine. He is also a member of the Editorial Board of the Journals "Optical Networks Magazine" and "Photonic Network Communications". He was a member of several national and international PhD thesis commissions. Piet Demeester is a member of IEEE (Senior Member), ACM and KVIV. His current research interests include: multilayer networks, Quality of Service (QoS) in IP-networks, mobile networks, access networks, grid computing, distributed software, network and service management and applications (supported by FWO-Vlaanderen, the BOF of the Ghent University, the IWT and the European Commission). Piet Demeester is currently full-time professor at the Ghent University, where he is teaching courses in Communication Networks. He has also been teaching in different international courses.

IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS

A PUBLICATION OF THE IEEE COMMUNICATIONS SOCIETY

APRIL 2006 VOLUME 24 NUMBER 4 ISACEM (ISSN 0733-8716)

PART II

OPTICAL COMMUNICATIONS & NETWORKING

IEEE