WDM Cross-connect Architectures with Reduced Complexity

Gordon Wilfong
Bell Labs, Lucent Technologies
Murray Hill, NJ 07974

Benny Mikkelsen
Bell Labs, Lucent Technologies
Crawford Hill, NJ 07733

Chris Doerr
Bell Labs, Lucent Technologies
Crawford Hill, NJ 07733

Martin Zirngibl
Bell Labs, Lucent Technologies
Crawford Hill, NJ 07733

Abstract

With the arrival of wavelength division multiplexed (WDM) systems carrying large numbers of wavelength channels, it appears that new cross-connect architectures that allow large numbers of wavelength channels to cross-connect at the wavelength granularity level are required. We propose a number of such architectures. The different architectures are compared in terms of complexity, cost, loss, crosstalk and filter narrowing. An analysis of the various non-blocking properties of such cross-connects is also described.

1 Introduction

The next generation wavelength division multiplexed (WDM) transmission systems will carry as many as 80 channels or 400 Gb/s total capacity per fiber. Consequently, large optical cross-connects will be needed in the near future to interconnect multiple fiber transmission lines in a central office. As an example, a cross-connect for 8 incoming fibers and 8 outgoing fibers each carrying 80 channels, will need 640 input/output ports each one capable of accepting at least OC-48 rate and eventually OC-192. Current cross-connects are based on electronic switch fabrics; they first time division demultiplex the high incoming rate (for example OC-48) into lower rates (for example OC-1) and then cross-connect at the lower rate, thus they have a very fine granularity. However, with the emergence of WDM systems that carry large numbers of wavelength channels, a new level of cross-connects with wavelength granularity seem to be highly desirable.

In principle, one can build a much larger cross-connect based on smaller cross-connects. However, to maintain strictly non-blocking characteristics the complexity of a cross-connect in practice generally scales like M^2 where M is the number of input/output ports. Thus, a 1000x1000 cross-connect could, in principle, be built out of one hundred 100x100 cross-connects (the size of the largest currently available cross-connect), but from a cost and size point of view, this would not be practical. Optical cross-connects have been demonstrated by routing the channels in the wavelength domain. An example is the MONET cross-connect [1]. However, this type of cross-connect is blocking in the wavelength domain and does not scale well either.

The goal of this work is to investigate new architectures that allow us to cross-connect large numbers of wavelength channels at the wavelength granularity. In its most general architectural form, a strictly non-blocking optical cross-connect with k input fibers and N wavelength channels per fiber consist of 5 stages [2, 3]; demultiplexing, wavelength interchanging/adaptation, space switching, wavelength inter-changing/adaptation, and finally, a multiplexing stage (see Figure 1). The complexity of the cross-connect in Figure 1 will generally scale with $(kN)^2$ if the fabric is based on space

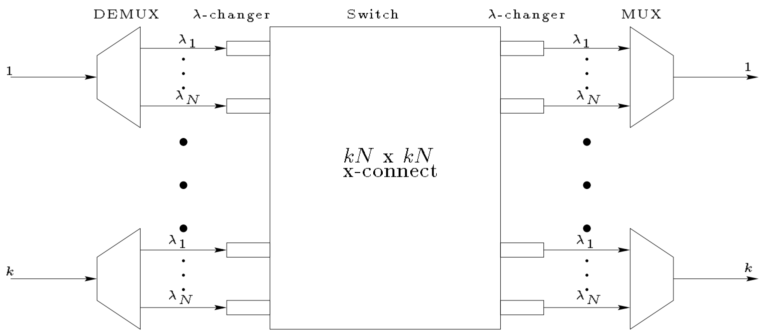

Figure 1: A kNxkN cross-connect connected to k fibers carrying N wavelengths. The kNxkN switching fabric is in general strictly non-blocking.

switches only. In this work we will show three new cross-connect architectures that are based on multi-wavelength modules, which are capable of routing simultaneously N wavelengths. The number of required modules scales only with k^2 or less (i.e., k^2 modules with N complexity). The significant reduction in complexity is traded for a decrease in blocking performance; one of the proposed architectures is strictly non-blocking in the space domain and rearrangeably non-blocking in the wavelength domain, whereas the two others are rearrangeably non-blocking in both the wavelength and space domain. Another very important advantage of the proposed architectures is their utilization of amplifiers. Since the wavelength channels are optically multiplexed in the interconnection fibers, only a small number of optical amplifiers are needed to compensate for the inevitable loss in the interconnection fabric.

We will first describe the elements needed for the proposed cross-connects. We define their functionality and also show some possible implementations. We then show the three cross-connect architectures based on these elements. Then we will compare the different architectures in terms of complexity, cost and important optical properties such as loss, crosstalk and filter narrowing. Finally we analyze the various non-blocking properties of the proposed cross-connects as well as those of two additional designs.

2 Optical cross-connect building blocks

The basic building elements in an optical cross-connect are WDM routers, space switches, optical filters, wavelength interchangers and optical amplifiers. We can construct new key sub-modules by combining these different functional building elements. In the next paragraphs, we will describe the sub-modules being used in our cross-connect architectures and discuss possible implementations of the modules. First, however, we will depict the general layout of the proposed cross-connects.

2.1 Proposed cross-connect architecture

The cross-connects we are proposing are based on the architecture shown in Figure 2. As shown, wavelength interchanging (WI) modules with a single input and output fiber are used at the input to change the wavelength of each of the N incoming channels. The WI modules are then followed by

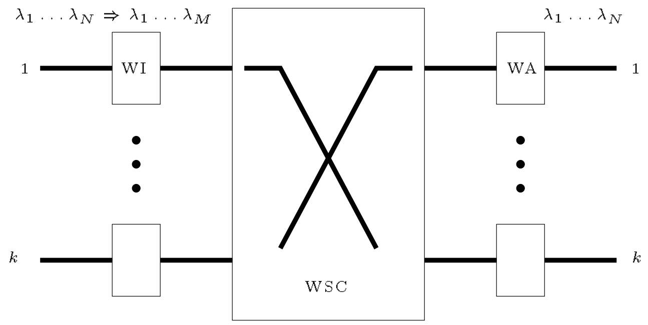

Figure 2: Layout of proposed cross-connects. M wavelength channels are multiplexed in interconnection fibers. If M is different from N then WAs required.

a wavelength selective space switch that route wavelength channels independently from the k inputs to the k outputs. The basic building module for the space switch is discussed in the next paragraph. As also shown in the figure, wavelength adapters (WA) might be needed at the output to allow for a multi-vendor environment. Note, that two channels can not have the same wavelength at the output of the WI module since all channels are multiplexed onto a single fiber. Consequently, our cross-connect is at most rearrangeably non-blocking in the wavelength domain. It can, however, be shown that for two of our architectures, at most $2k$ connections out of the Nk connections have to be rearranged to unblock a demand (see Section 5). In addition, the blocking performance can be improved by having more internal wavelengths than transmission wavelengths (i.e., $M > N$). In its most general case, M can also be less than N if we have more than one output fiber from the WI module.

2.2 The wavelength selective optical cross-connect (WSC)

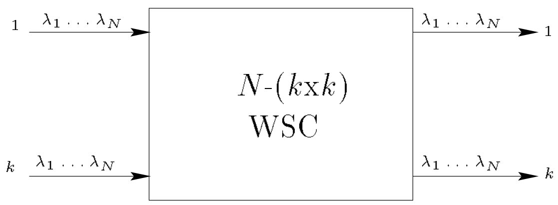

Figure 3: A wavelength selective cross-connect with k incoming fibers each carrying N wavelengths. Each wavelength can be cross-connected amongst k fibers, but not between different channels (no wavelength interchange).

The wavelength selective cross-connect (WSC) is the key sub-module in our three cross-connect configurations to be presented later. The N wavelength kxk WSC has k input ports and k output ports (Figure 3). It consists conceptually of k WDM demultiplexer/multiplexers and of N non-blocking kxk cross-connects. Each of the kxk cross-connects allows cross-connecting one of the wavelength channels on each fiber, say wavelength channel λ_i. Since the kxk WSC (whenever we will leave out the N wavelength notation from now on, it is understood that the particular element can process N wavelength channels) does not allow the wavelength of the individual channels to change, it is blocking in the wavelength domain.

Examples of a 1x1 WSC are shown in Figure 4. A 1x1 WSC is simply a multi-wavelength filter (MWF) that gives individual control for each channel. Each channel can thus be passed or blocked by the MWF. Such a device can be built with tunable fiber gratings (Figure 4a), or back-to-back routers interconnected with an array of semiconductor optical amplifiers (Figure 4b). The latter can be either monolithically integrated in InP [4] or hybridly integrated on a Silica motherboard. Also note that the switching speed of semiconductor optical amplifiers is less than micro-seconds, permitting very fast reconfiguration of the module.

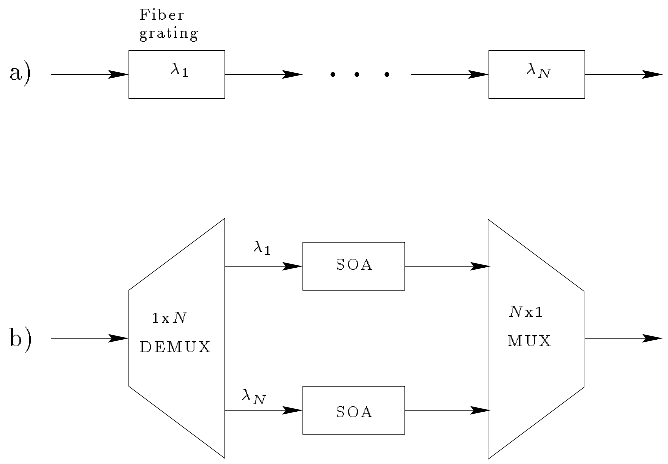

Figure 4: Examples of N wavelength 1x1 WSC which is simply a multi-wavelength filter that can be assembled for instance from tunable fiber gratings a) or two back-to-back optical MUXs interconnected by an array of semiconductor optical amplifiers (SOAs) b).

Examples of a 2x2 WSC, or a wavelength selective optical add-drop multiplexer (OADM) are shown in Figure 5. By simply adding two circulators to the tunable fiber gratings shown in Figure 4a, one gets the equivalent functionality of an OADM. Monolithic 16 wavelength OADM have been recently demonstrated by using two back to back interleave chirped waveguide grating routers [5]. The MONET [1] cross-connect is an example of an 8 wavelength 6x6 WSC.

2.3 The wavelength interchanger

A wavelength interchanger (WI) is, in general, only capable of converting the wavelength of a single channel [6]. Hence, to construct a WI that allows us to arbitrarily change the wavelengths of N incoming channels, we need to demultiplex the N channels and then convert the different channels individually (see Figure 1). The cross-connects we are proposing require WI modules as shown in Figure 6a). The WI has one input fiber and one output fiber. Restrictions are, of course, that no two channels shall have the same wavelength. Such a WI module can be built by broadcasting the N incoming channels to N optical translators, in front of which a tunable filter selects a single channel (see Figure 6b). The translator/interchanger can be an all-optical or an opto-electronic translator. The optical filter can be based on the MWF (1x1 WSC) described previously.

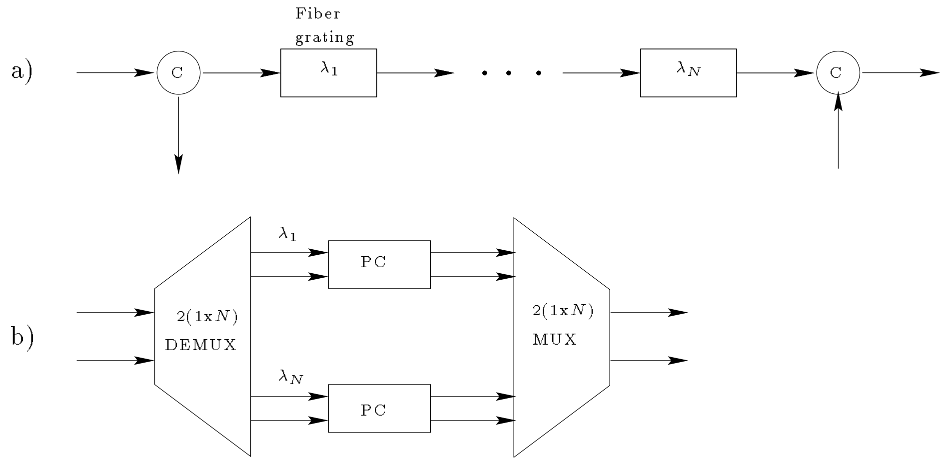

Figure 5: Examples of N wavelength 2x2 WSC based on circulators and tunable fiber gratings a) and on back-to-back chirped optical demuxes and phase controllers (PC) b).

The WI module can also be realized by using fixed filters (i.e., a demultiplexing stage) in front of translators with tunable output wavelengths (see Figure 6c). Note that this approach requires power combiners at the translator output, whereas demultiplexers are possible at the input. This is opposite to the implementation shown in Figure 6b. Other possible implementations of the WI module rely on switches placed in between a demultiplexing stage and an array of translators with fixed output wavelengths. The NxN space switch can be an optical switch or an electronic cross-connect (see Figure 6d).

No matter what approach we use, the WI is a fairly expensive module consisting of at least N active elements like transmitters and filters, and its use has to be minimized as much as possible. In the architectures shown below, we will need k WI modules.

3 Proposed cross-connect architectures

In the following paragraphs, we will introduce the three different cross-connect architectures. The general approach in the three architectures is to combine k N-wavelength WI modules with a network of N-wavelength kxk WSCs. We show in Section 5 that a cross-connect based on these two elements can be made strictly non-blocking in the space domain and rearrangeably non-blocking in the wavelength domain. It should be noted that wavelength adapters (WAs) are not shown at the output ports in the layouts for our proposed cross-connects, however, they can easily be included to allow for multi-vendor environment.

3.1 WI-broadcast&select (Figure 7)

The output fiber of each WI is split k-times and filtered by k^2 1x1 WSCs (MWFs). Note that the space fabric for each wavelength is strictly non-blocking. It is also able to perform broadcasting. It is furthermore modular in the sense that one extra input can be accommodated by adding one extra MWF at each output. Only one stage of filtering is needed and only one stage of amplifiers for k=8. The coherent crosstalk accumulates by a factor equal to k, i.e., 9 dB for k=8.

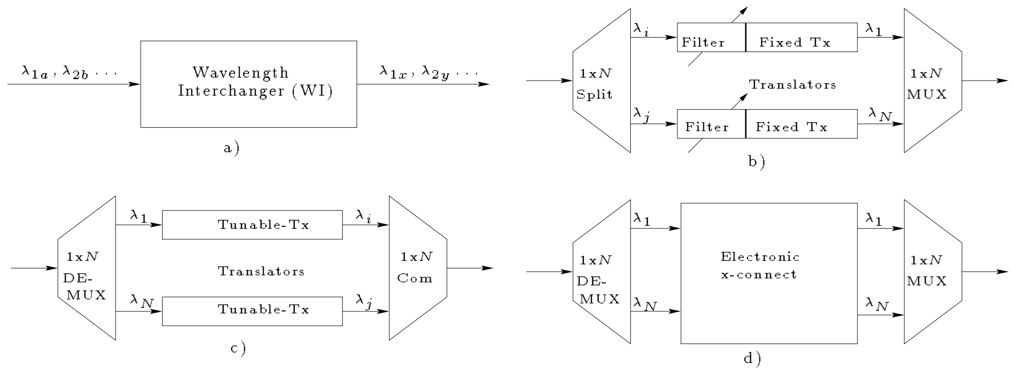

Figure 6: a) Wavelength interchanger swaps channels and wavelengths. It could be built with tunable optical filters followed by translators with fixed wavelength transmitters b) or from fixed filters and tunable lasers c) or with an electronic space fabric d).

3.2 WI-Beneš (Figure 8)

The k WIs are followed by a Beneš architecture of N-wavelength 2x2 WSCs (OADMs). For k=8, we need 5 stages of OADMs or a total of 20 to have a rearrangeably non-blocking cross-connect in the space domain. Assuming a loss of 10 dB per OADM, the total loss through the fabric is 50 dB, thus two stages or a total of 16 optical amplifiers would be needed. The coherent crosstalk accumulates by a factor equal to the number of stages, i.e., 5 or 7 dB for k=8.

3.3 Beneš/2-WI-Beneš/2 (Figure 9)

In this case, the WIs are put in the middle of the Beneš architecture at the expense of one additional stage. Unlike the previous architecture, this network gives us the flexibility of choosing a specific output wavelength. However, due to the extra stage the cross-connect requires slightly more optical amplification and the coherent crosstalk accumulates by a factor of 6 or 8 dB for k=8.

4 Comparison of the three cross-connects

In Table 1, we compare the three architectures in terms of number of modules, fiber interconnections, number of stages, loss, number of OAs. As an example, we choose k=8. All architectures require k WIs.

	# WI	# 2x2 WSC	# 1x1 WSC	# fibers	# stages	Loss	# OAs	x-talk
WI-Be	k	$k(2\log k - 1)/2$	-	4(# WSC)	$2\log k - 1$	10(# stages)	k/3(# stages)	$2\log k - 1$
$k = 8$	8	20	-	80	5	50 dB	16	7dB
Be/2	k	$k(2\log k)/2$	-	4(# WSC)	$2\log k$	10(# stages)	k/3(# stages)	$2\log k$
$k = 8$	8	24	-	96	6	60 dB	16	8dB
BroadC	k	-	k^2	4(# WSC)	1	$10 + 2k(dB)$	k/30(Loss)	k
$k = 8$	8	-	64	256	1	30 dB	8	9dB

Table 1: Comparison between the different architectures. 10 dB loss per WSC and 30 dB gain per optical amplifier (OA) are assumed.

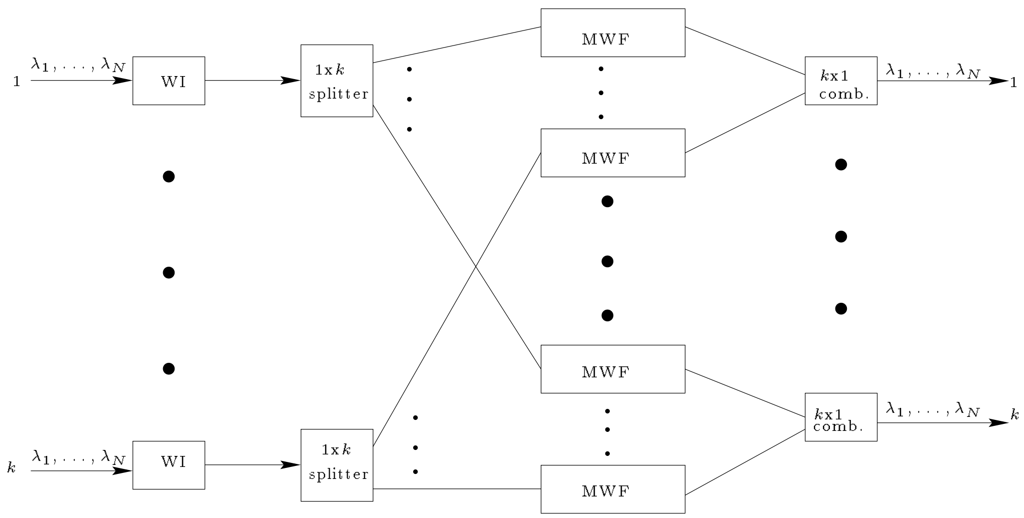

Figure 7: NkxNk cross-connect based on broadcast&select network with 1x1 WSCs (MWF).

In terms of number of WSC modules, the WI-Beneš has the best performance. The Beneš/2-WI-Beneš/2 has slightly more modules but it gives us output wavelength flexibility. However it also has the most stages resulting in the highest loss and filter narrowing. The WI-broadcast&select architecture is clearly best from a filtering narrowing, optical loss and blocking perspective at the expense of the largest number of modules needed and more accumulated crosstalk. The comparison is, however, not completely straightforward, because the WI-broadcast&select architecture needs only 1x1 WSCs (MWF) versus 2x2 WSCs (OADMs) for the Beneš type networks (the complexity of one 2x2 WSC corresponds roughly to two 1x1 WSCs).

The WSCs are allowed to be polarization dependent in the WI-Beneš and the WI-broadcast&select architectures, since the WI modules at the input of the cross-connects perform polarization conversion as well (WSCs with polarization maintaining fiber pigtails are however required).

The three architectures are all rearrangeably non-blocking in the wavelength domain as shown in Section 5. Furthermore, it can be shown that at most $2k$ connections out of the Nk connections have to be rearranged to unblock a demand for the WI-Beneš and the WI-broadcast&select architectures. The number of required connections that have to be rearranged for the Beneš/2-WI-Beneš/2 to unblock a demand is more complicated to predict. Though, the wavelength blocking can be reduced for all three architectures by having more internal wavelengths than transmission wavelengths. This, however, requires that the cross-connect be equipped with wavelength adapters at the output, which might be required anyway.

For all three architectures, the cost of the cross-connect is expected to be dominated by the cost of the optical translator/converter units in the WI modules, which are inevitable in any cross-connect architecture. Therefore we believe our architectures are very attractive for the realization of practical and cost effective cross-connects.

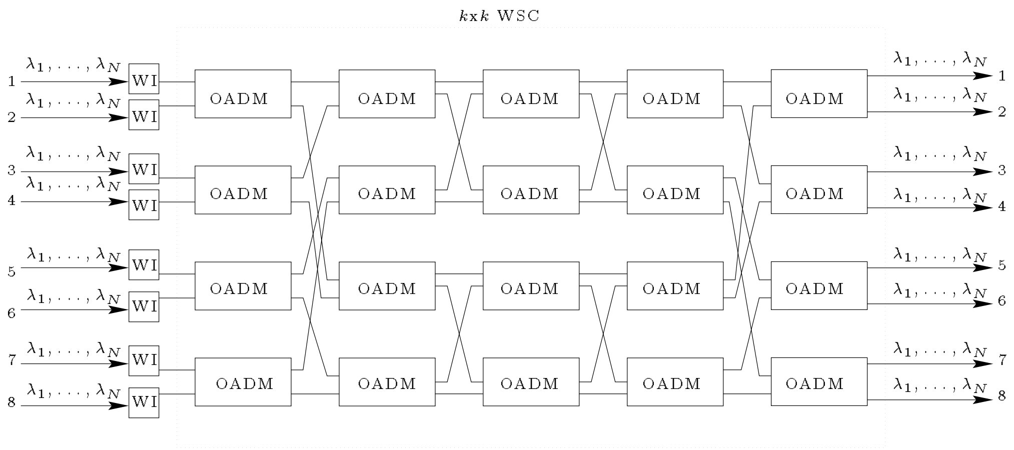

Figure 8: NkxNk cross-connect based on k WIs and a kxk-WSC composed of a Beneš architecture of N wavelength 2x2 WSCs (OADMs).

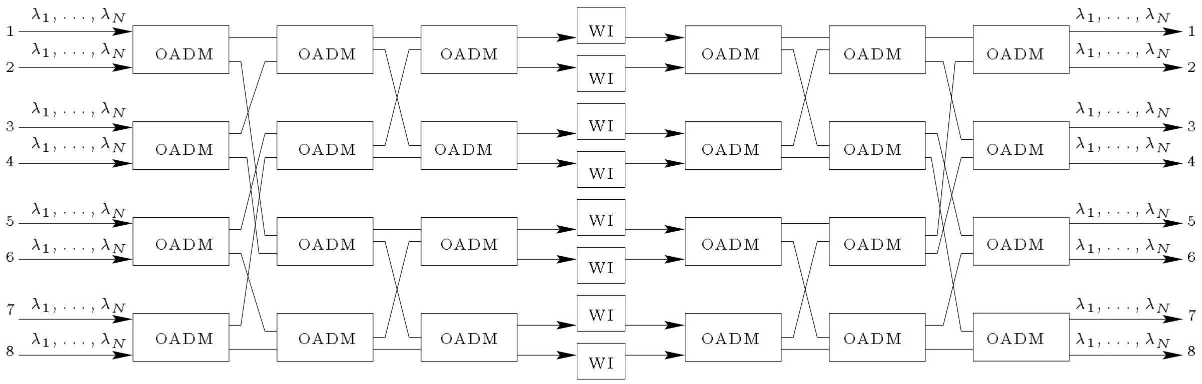

Figure 9: NkxNk cross-connect based on two 1/2-Beneš kxk-WSCs and an array of WIs in between.

5 Non-blocking characteristics

In this section we study the non-blocking properties of the WDM cross-connects introduced in Section 3 with some additional ones based on the Cantor architecture [7].

There is, of course, a vast literature on the non-blocking properties of various space domain cross-connects (see for example [8]. A demand in that setting means a request to find a route between a particular input line and a particular output line that does not interfere with other routed demands. A space domain cross-connect can be pathwise rearrangeably non-blocking (meaning that any set of demands can be routed through the cross-connect but after routing a set of demands any additional demands that need to be routed may require some of the previously routed demands to be re-routed) or it can be pathwise strictly non-blocking (meaning that not only can any set of demands be routed through the cross-connect but any additional demands that need to be routed will never require the re-routing of the previously routed demands).

In contrast, in the WDM setting, a demand is a request to route the channel with a given wavelength from a given input line to a specified output line. Along the route a demand may change its wavelength at places only where wavelength interchangers are available. Thus for a given demand, not only must a route be determined but also a wavelength must be assigned to the demand for each link in the route. Therefore in the WDM case, there are two types of non-blocking characteristics of interest; namely, the pathwise non-blocking characteristics as in the space domain case as well as non-blocking characteristics in the wavelength domain. That is, when an additional demand is initiated, the previously routed demands may require their routes to be changed, their wavelength assignments to be changed or both. We consider the two types of non-blocking characteristics of a number of WDM cross-connect designs in this section.

5.1 Definitions

Let k be the number of incoming (and outgoing) fibers of the cross-connect and let N be the number of wavelengths on each fiber. A *demand* d is a triple (a, b, λ) with the meaning that the channel on input line a using wavelength λ is to be routed to output line b. A *demand set* is a set of demands that satisfies the following conditions:

(i) for all incoming fibers a the number of demands of the form (a, x, y) for any x and y is no more than N

(ii) for all outgoing fibers b the number of demands of the form (x, b, y) for any x and y is no more than N

(iii) there are not demands of the form (a, b_1, λ) and (a, b_2, λ) where b_1 and b_2 might or might not be the same

Each component of the various cross-connect designs will be one of the following: a wavelength interchanger (WI), an N-wavelength kxk wavelength selective cross-connect (kxk N-WSC), a splitter or a combiner.

A *single channel* version of a cross-connect is the resulting network when there is only one wavelength (that is, a space domain cross-connect).

A demand set D is said to be *satisfied* by a cross-connect C if each demand can be simultaneously routed from its required input line to its required output line and assigned a wavelength from the set $\lambda_1, \ldots, \lambda_N$ over each link in the route so that

(i) for each demand $d = (a, b, \lambda_i)$ the wavelength assigned to it over the first link of its route is λ_i,

(ii) no two demands are assigned the same wavelength on any given link and

(iii) for each demand, the wavelength assigned to the demand over consecutive links in its route are the same unless the common node of the links contains a WI.

We are only interested in cross-connects that can satisfy every possible demand set. The various cross-connects of interest will be differentiated based on how much work must be performed to update the routes and wavelength assignments of a given demand set when an additional demand is added to the demand set.

Let C be a cross-connect and D be any demand set satisfied by C. Consider any new demand d such that $D' = \{d\} \bigcup D$ is a demand set. Then C is said to be *wavelength rearrangeably non-blocking* if for any such D and d, the demand set D' can be satisfied with no changes to the routes followed by the demands of D but possibly with changes to the wavelengths assigned to them. We say that C is *rearrangeably non-blocking* if D' can be satisfied, however demands of D might have to be assigned a new route as well as a new wavelength. Cross connect C is *pathwise rearrangeably non-blocking* if for any such D and d, D' can be satisfied without disturbing the wavelengths of the demands of D but might require rerouting of the demands in D.

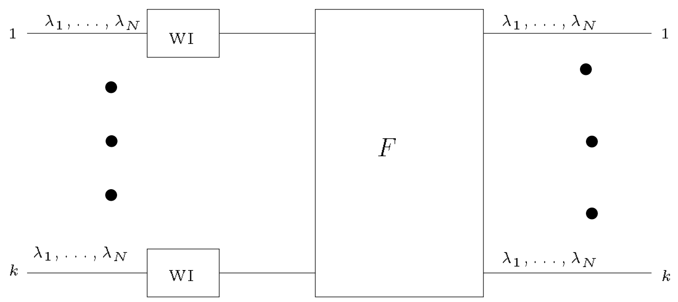

Figure 10: Standard design cross-connect.

Consider any cross-connect C that has the basic form as shown in Figure 10. The fabric F can be anything that does not contain any device to change the wavelength of any channel. That is, after the wavelength interchanger WI on each input line, no demand can have its wavelength changed. Thus a wavelength assignment for a route of a given demand can be thought of as just the single wavelength that is assigned to the demand over every link after the WI. Then we say that an assignment of wavelengths to demands is *valid* if no demands from the same input line or to the same output line are assigned the same wavelength. Then in order to satisfy a demand set D we need a valid wavelength assignment and a route from the appropriate WI to the appropriate output line so that no two demands that are assigned the same wavelength share a common link within F. Suppose F has the property that for any demand set D and any valid wavelength assignment there is always a routing of the demands within F so that no two demands with the same wavelength meet within F. Such a cross-connect C is said to be a *standard design cross-connect.*

5.2 Cross-connects under consideration

We study the blocking characteristics of several different designs for cross-connects. We consider the WI-broadcast&select, the WI-Beneš and the Beneš/2-WI-Beneš/2 cross-connects as described in Section 3 as well as two architectures based on the Cantor network [7]. The two cross-connects based on the Cantor network were not discussed in the earlier sections as increased complexity was judged to render them unlikely to be of practical interest. However they are included in this non-blocking analysis for completeness.

The WI-Beneš, the WI-broadcast&select and the WI-Cantor (described below) are standard design cross-connects whose fabrics are respectively, a Beneš network, a broadcast&select network and a Cantor network.

The fabric of a WI-Cantor cross-connect is made up of $\log k$ Beneš networks, see Figure 11. It turns out that this design has similar blocking characteristics as the WI-broadcast&select cross-connect but essentially requires $O(k(\log k)^2)$ basic elements rather than the k^2 required by the WI-broadcast&select cross-connect. It should be noted however that the basic elements of the Cantor design are more complex than those in the broadcast&select architecture.

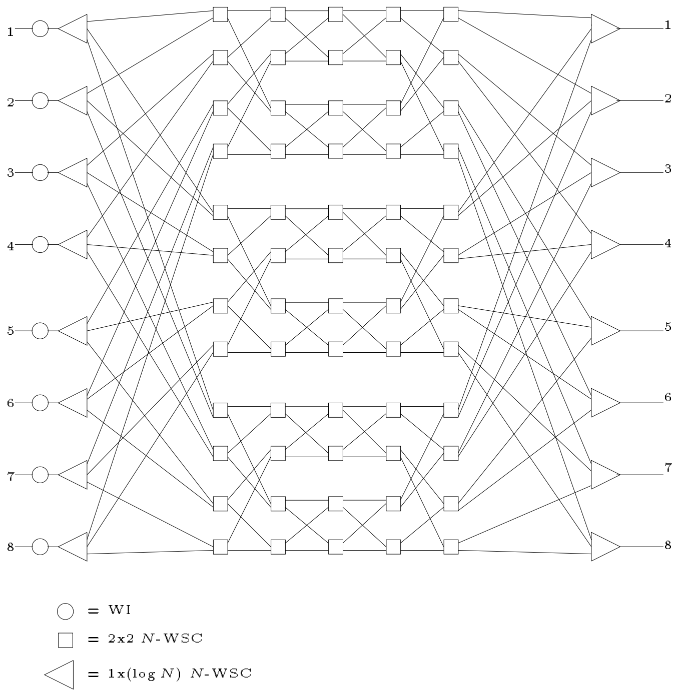

Figure 11: WI-Cantor cross-connect for $k = 8$.

Then we consider two non-standard design cross-connects. The first is the Beneš/2-WI-Beneš/2 cross-connect described in Section 3. Finally we modify the WI-Cantor cross-connect much as we modified the WI-Beneš cross-connect by moving the wavelength interchangers into the middle of the fabric, see Figure 12. Unfortunately, this results in having to use $k \log k$ wavelength interchangers as opposed to the k required in the other designs. However as will be shown, this design has the desirable property that it is strictly non-blocking (in both the wavelength and path senses). As in the previous example, this is not a standard design. We call this cross-connect a *Cantor/2-WI-Cantor/2 cross-connect.*

5.3 Non-blocking characteristics in standard design cross-connects

We begin by showing that in a standard design there are inherent rearrangement properties in the wavelength assignment domain. Then we consider how the blocking properties (in the space domain) of the fabric F affect the blocking characteristics of the cross-connect. We first require some definitions and a result from graph theory.

A *bipartite multigraph* $G = (V, U, E)$ is a graph with two sets of nodes U and V and with a collection of edges E with the property that each edge $e = \{u, v\}$ has one endpoint in U and the other in V. G is a multigraph since we allow multiple edges with the same endpoints. The *degree* of a node w in G is the number of edges in E having w as one of its endpoints. An *edge coloring* of G is an assignment of a color to each edge of E so that no edges sharing a common endpoint are assigned the same color. A *minimum edge coloring* of G is an edge coloring that uses the smallest number of colors possible.

Lemma 5.1 *Let G be a bipartite multigraph and m be the maximum degree of any node of G. Then a minimum edge coloring of G uses exactly m colors.*

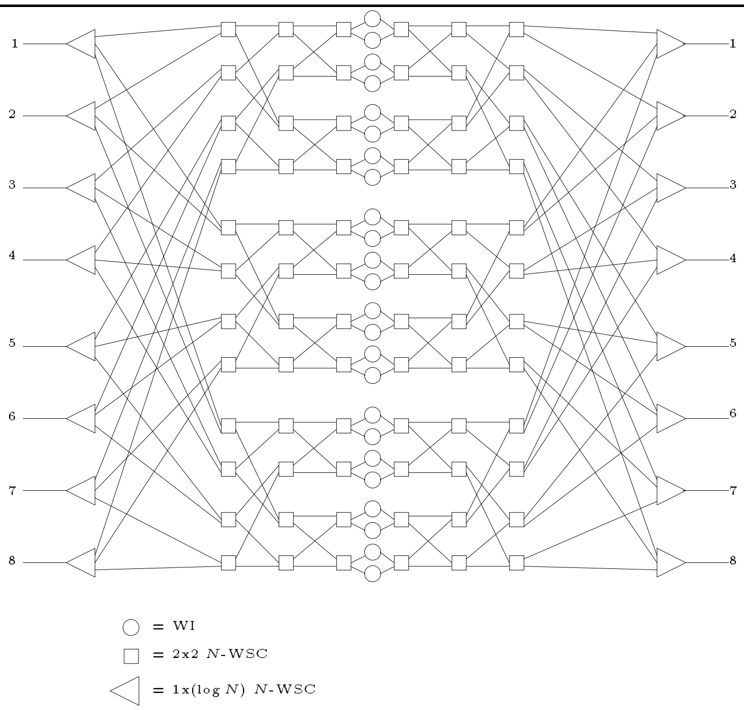

Figure 12: Cantor/2-WI-Cantor/2 cross-connect for $k = 8$.

Proof: See [9]. ∎

We wish to show that any standard design cross-connect is wavelength rearrangeably non-blocking. To accomplish this we use the result in Lemma 5.1.

Lemma 5.2 *Any standard design cross-connect is wavelength rearrangeably non-blocking.*

Proof: In order to show that a standard design cross-connect is wavelength rearrangeably non-blocking it suffices to show that for any demand set D there is some assignment of wavelengths to the demands so that no demands from the same input line or to the same output line are given the same wavelength. To accomplish this we use the result of Lemma 5.1.

We define a bipartite multigraph $G = (V, U, E)$ as follows. Let $U = \{u_1, u_2, \ldots, u_k\}$ and $V = \{v_1, v_2, \ldots, v_k\}$ where u_i represents the i^{th} input line and v_j represents the j^{th} output line. For every demand $d = (i, j, \lambda)$ in demand set D define an edge $\{u_i, v_j\}$ in E (and these are the only edges in E). Note that there might be multiple edges $\{u_i, v_j\}$ in E since there might be demands (u_i, v_j, λ) for multiple values of λ. By Lemma 5.1, we know that we can color the edges of E using m colors and no two edges with a common endpoint have the same color where m is the maximum degree of any node of G. If N is the number of wavelengths on each incoming fiber, then there can be no more than N demands requesting channels from any input line and no more than N demands to any output line. Hence $m \leq N$. Therefore we have an assignment of $m \leq N$ wavelengths (colors) to the demands (edges) so that no demands coming from the same input line or going to the same output line are assigned the same wavelength. ∎

There is a particularly simple constructive proof of Lemma 5.1 due to Vizing [10] that iteratively colors the edges of the graph G. It has the property that when the next edge is to be colored, the number of previously colored edges that might have to be reassigned a new color is restricted to those with one of two colors. Vizing's construction is based on an "alternating path" argument as follows.

Let m be the maximum degree of any node of G and so we have m colors available. By induction we assume that so far we have used no more than m colors. Suppose the new edge is from node u to node v. If there is some color not already assigned to any edge with endpoint u and not assigned to any edge with endpoint v then assign it to edge $\{u, v\}$. Otherwise, since at most $m-1$ edges incident to u are colored there must be some color, say c_1 that is available at u. Similarly, there must be some color, say c_2 that is available at v. From v construct the longest path P whose first edge has color c_1 and whose subsequent edges alternate between c_2 and c_1. If w is the last node on P then it is easy to see that $w \neq u, v$ since G is bipartite and by the definition of c_1 and c_2. For each edge e along P, swap the color of e from c_1 to c_2 or vice versa. Then color c_1 is available at both u and v and so we can assign color c_1 to edge $\{u, v\}$. Notice that using this construction, only those edges with colors c_1 or c_2 (and were edges actually in the alternating path constructed) had their colors changed by the need to color the new edge.

In terms of cross-connects, the interpretation of this result is that when a new demand is created in a standard design cross-connect, assigning a wavelength to that demand can be achieved while causing only those already existing demands that have been assigned one of two wavelengths, say λ_1 and λ_2, to have to be assigned different wavelengths.

Notice that once wavelengths have been assigned to the demands then we can think of F as actually N independent single channel copies of F, one for each wavelength, since demands of different wavelengths can not "block" each other within F. Thus for example, consider the WI-Beneš cross-connect, the standard design cross-connect based on a Beneš network. When a new demand is given, it may be the case that some demands with one of two colors, say λ_1 and λ_2, have to have their colors swapped. Then since a one channel Beneš network is pathwise rearrangeably non-blocking [11], all the demands of colors λ_1 and λ_2 may potentially have to have their routes changed in order to satisfy all the demands.

Note that in the case of the WI-broadcast&select cross-connect, when a new demand is introduced only those demands involved in the alternating path might have to be adjusted (and only by changing their wavelengths) since no routes through F need be disturbed. For a WI-Cantor cross-connect, when a new demand is added only those demands involved in the alternating path might need their wavelengths and hence their routes modified.

Therefore we conclude the following:

Theorem 5.3 *The WI-Beneš, a standard design cross-connect whose fabric F is a Beneš network, is a rearrangeably non-blocking cross-connect and the addition of a new demand causes rearrangement (of wavelengths and routes) of at most demands having one of two wavelengths.*

Theorem 5.4 *The WI-broadcast&select cross-connect, a standard design cross-connect with a broadcast&select fabric F, is wavelength rearrangeably non-blocking and the addition of a new demand causes changes in the wavelengths of at most demands having one of two wavelengths.*

Theorem 5.5 *The WI-Cantor cross-connect, a standard design cross-connect whose fabric F is a Cantor network, is rearrangeably non-blocking and the addition of a new demand causes changes in the wavelengths of at most demands having one of two wavelengths and only those demands that have their wavelength changed might also require their routes to be altered.*

5.4 Non-blocking characteristics in the non-standard designs

The Beneš/2-WI-Beneš/2 cross-connect with the WIs in the middle is not a standard design and so the analysis of the previous section does not apply. We will show that, like the standard design WI-Beneš cross-connect, this cross-connect is rearrangeably non-blocking in the pathwise sense but without the additional property of the standard design cross-connect that rearrangements are restricted to demands with one of two wavelengths. Since the Beneš/2-WI-Beneš/2 cross-connect has the property that a connection never has to have its wavelength changed it is wavelength strictly non-blocking.

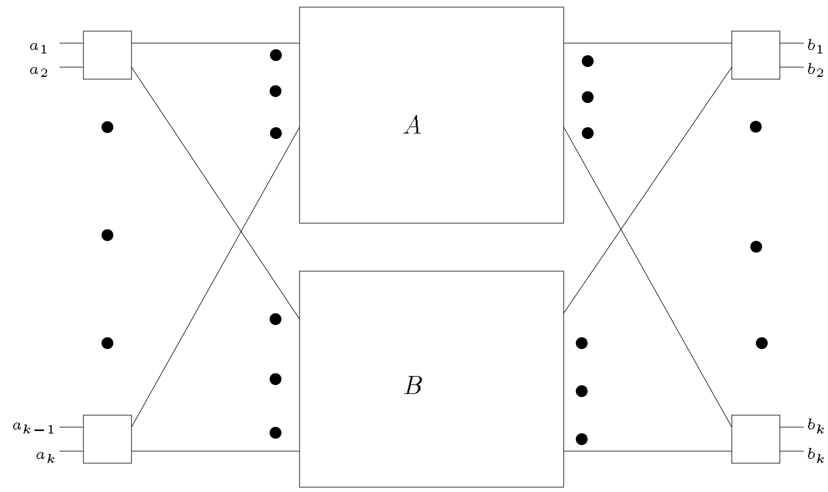

Figure 13: Illustrating looping algorithm.

Theorem 5.6 *The Beneš/2-WI-Beneš/2 cross-connect is pathwise rearrangeably non-blocking.*

Proof: The proof uses a multi-color version of the well-known "looping algorithm" as described in [12] for a single channel Beneš network. The looping algorithm is a simple recursive scheme that starts at the outermost stages and then removes the leftmost and rightmost stages of the network and recurses on the remaining smaller networks. It works as follows. Let $a_1, a_2, \ldots, a_k$ be the input lines and $b_1, b_2, \ldots, b_k$ be the output lines as shown in Figure 13. We assume that there is a demand from every input line (and hence a demand to each output line). Demands are just pairs of the form (a_i, b_j) with the meaning that we want to connect input line a_i with output line b_j. If there is a demand between input line a_i and output line b_j then define $\pi(a_i) = b_j$ and $\pi(b_j) = a_i$. Let $\bar{i} = i - 1$ if $i = 2^j$ for some $j > 0$ and $\bar{i} = i + 1$ otherwise. Thus for example, a_i and $a_{\bar{i}}$ are the two input lines that share a leftmost 2x2 N-WSC in the network. The network shown in Figure 13 is based on the recursive definition of a Beneš network and so A and B are both $k/2$ x $k/2$ Beneš networks. We now describe how to route the demands at the outermost level and the same rule is applied to the inner networks A and B recursively. Choose some demand $(a_i, \pi(a_i))$ and route it from a_i through A to $\pi(a_i) = b_j$. Then route $b_{\bar{j}}$ through B to $\pi(b_{\bar{j}})$. If $\pi(b_{\bar{j}}) = a_{\bar{i}}$ then this "cycle" is done and we arbitrarily choose an unrouted demand and continue as above. Otherwise we route demand $(\pi(b_{\bar{j}}), \pi(\pi(b_{\bar{j}})))$ through A and repeat the process until we end up at $a_{\bar{i}}$. This clearly routes every demand without conflict.

In the case where there are multiple wavelengths each demand is thought of as a 4-tuple $(i, \lambda_u, j, \lambda_v)$ meaning that we want to connect the channel with wavelength λ_u on input line a_i with the channel with wavelength λ_v on the output line b_j. (Notice that this means that the wavelength assignments to every demand are fixed and so there is no notion of wavelength rearrangeability.) Then we essentially do the same thing as in the single channel case except we "match" wavelengths as we construct a routing. That is, suppose we start with demand $(i, \lambda_u, j, \lambda_v)$ which we route through A and on the left side it has wavelength λ_u and on the right side it has wavelength λ_v. Thus the next demand we route is the one of the form $(k, \lambda_w, \bar{j}, \lambda_v)$ for some k and w and we route this one through B. Then as before we would route demand $(\bar{k}, \lambda_w, r, \lambda_s)$ for some r and s through A (unless $\bar{k} = i$ and $\lambda_w = \lambda_u$). Again it is obvious that this method routes all demands so that no two with the same initial wavelength (or final wavelength) are routed on the same link (i.e. if they share the same initial (or final) 2x2 N-WSC and they are of the same initial (or final) wavelength then one is routed through A and the other through B). Thus for a given demand set we can always find a routing to satisfy the demand set but the addition of a new demand might potentially require the re-routing of all previously routed demands. ■

Thus while a change in the demand set caused limited rearrangements in the case of a standard design cross-connect (i.e. rearrangements limited to demands of one of two wavelengths), in the case of the Beneš/2-WI-Beneš/2 cross-connect a single demand change might require re-routing of all of the previously routed demands (although no changes to wavelengths are ever made).

Now consider the Cantor/2-WI-Cantor/2 cross-connect. The proof that the single channel Cantor network is strictly non-blocking (see [8], [13]) has the following basic form. It is shown that no matter what demands are already routed in a Cantor network, any free input can be routed to more than half of the middle elements of the network and by symmetry, any free output can be reached from more than half of the middle elements. Thus for any free input and output pair there must be some middle element that can be reached by both of them and hence there is an available route between them. We take this proof and show that it can be easily modified to prove that the Cantor/2-WI-Cantor/2 cross-connect is strictly non-blocking for the case of multiple channels (wavelengths). Suppose we wish to route a channel with wavelength λ_u from input line a_i to the channel with wavelength λ_v on output line b_j. Then the proof of the strictly non-blocking property of the single channel Cantor network immediately shows that more than half of the wavelength interchangers can be reached from the input line a_i avoiding conflicts with other previously routed demands that have wavelength λ_u on their input line. Similarly, the proof shows that there are more than half of the wavelength interchangers that can be reached from b_j without collisions with already routed demands whose output channel has wavelength λ_v. Therefore there is some common wavelength interchanger that they both can reach and hence there is a route available. Therefore we conclude:

Theorem 5.7 *The Cantor/2-WI-Cantor/2 cross-connect is strictly non-blocking in both the wavelength sense and the path sense.*

5.5 Summary of non-blocking properties

We have seen that any standard design cross-connect, such as the WI-Beneš cross-connect, the WI-broadcast&select cross-connect and the WI-Cantor cross-connect, is wavelength rearrangeably non-blocking. The standard design cross-connects have the added property that the addition of a new demand requires only that previous demands assigned one of two wavelengths might have to have their wavelengths changed in order to accommodate the new demand. While the WI-Beneš cross-connect might require all demands having one of two wavelengths to be re-routed, the WI-Cantor cross-connect requires re-routing of only those demands with one of two wavelengths that actually have their wavelength changed and the WI-broadcast&select cross-connect never requires any demand to be re-routed.

The Beneš/2-WI-Beneš/2 cross-connect is pathwise rearrangeably non-blocking. That is, when a new demand is added, all previously routed demands might require rerouting but no change in wavelength assignment is ever necessary.

The Cantor/2-WI-Cantor/2 cross-connect has the strongest non-blocking property in that it is strictly non-blocking in both the wavelength and space domains. This power comes at the cost of using $k \log k$ WIs as opposed to the k WIs used in the other cross-connects. The Cantor/2-WI-Cantor/2 and the Beneš/2-WI-Beneš/2 cross-connects both have the additional property that they allow for the specification of the output wavelength. Note that this property could be achieved by any standard design cross-connect by placing additional WIs on each output line.

6 Summary

We have proposed a number of new optical cross-connect architectures based on multi-wavelength modules, which are capable of simultaneously routing N wavelengths. Compared to strictly non-blocking cross-connects, the architectures have a significantly reduced complexity and cost, which is traded for an increased wavelength blocking; one of the proposed architectures is strictly non-blocking in the space domain and rearrangeably non-blocking in the wavelength domain, and another of the practical designs is rearrangeably non-blocking in both the wavelength and space domain. It can, however, be shown that at most $2k$ connections out of Nk connections have to be rearranged to unblock a demand in either case. The third design is rearrangeably non-blocking in the space domain and never requires rearranging in the wavelength domain. Another very important advantage of the three proposed architectures is their utilization of amplifiers. Since the wavelength channels are optically multiplexed in the interconnection fibers, only a small number of optical amplifiers are needed to compensate for the inevitable loss in the interconnection fabric.

References

[1] R.S. Vodhanel et al., "National-scale WDM networking demonstration by the MONET consortium," in *Proc. of OFC'97*, Dallas, Feb. 1997, vol. PD27.

[2] E. Karasan and E. Goldstein, "Optical restoration at the wavelength-multiplex section level in WDM mesh networks," in *Proc. of OFC'98*, San Jose, Feb. 1998, vol. ThJ5.

[3] R.W. Tkach et al., "Fundamental limits of optical transparency," in *Proc. of OFC'98*, San Jose, Feb. 1998, vol. WJ1.

[4] M. Zirngibl et al., "Digitally tunable channel dropping filter/equalizer based waveguide grating router and optical amplifier integration," *IEEE Photon. Lett.*, vol. 6, no. 4, pp. 513–515, April 1994.

[5] C. Doerr, "Proposed wdm cross connect using a planar arrangement of waveguide grating routers and phase shifters," *IEEE Photon. Lett.*, vol. 10, no. 4, pp. 528–530, April 1998.

[6] B. Mikkelsen et al., "Wavelength conversion device," in *Proc. of OFC'96*, San Jose, Feb. 1996, vol. WG1.

[7] D. Cantor, "On construction of nonblocking switching networks," in *Proc. of Symposium on Computer-Communications Networks and Teletraffic*, 1972, pp. 253–255.

[8] J. Y. Hui, *Switching and Traffic Theory for Integrated Broadband Networks*, Kluwer Academic Publishers, Norwell, MA, 1990.

[9] C. Berge, *Graphs and Hypergraphs*, North Holland, Amsterdam, 1973.

[10] O. Ore, *The Four Color Problem*, Academic Press, New York, NY, 1967.

[11] V. E. Beněs, *Mathematical Theory of Connecting Networks and Telephone Traffic*, Academic Press, New York, NY, 1935.

[12] D. C. Opferman and N. T. Tsao-Wu, "On a class of rearrangeable switching networks, Part I: Control algorithm," *The Bell System Technical Journal*, vol. 5, no. 50, pp. 1579–1600, 1971.

[13] Masson, Gingher, and Nakamura, "A sampler of circuit switching networks," *IEEE Computer*, vol. 5, pp. 32–48, June 1979.