

Nonblocking WDM Switches Based on Arrayed Waveguide Grating and Limited Wavelength Conversion

Hung Q. Ngo
201 Bell Hall,
Computer Science and Engineering,
SUNY at Buffalo,
Amherst, NY 14260, USA.
hungngo@cse.buffalo.edu

Dazhen Pan
201 Bell Hall,
Computer Science and Engineering,
SUNY at Buffalo,
Amherst, NY 14260, USA.
dpan@cse.buffalo.edu

Chunming Qiao
201 Bell Hall,
Computer Science and Engineering,
SUNY at Buffalo,
Amherst, NY 14260, USA.
qiao@cse.buffalo.edu

Abstract—Constructing wavelength division multiplexing (WDM) switches with cheap components and low complexity is an important problem in optical networking.

Typically, there are two request models widely considered. In one model, a connection request asks to go from a wavelength on an input fiber of the WDM switch to a particular wavelength on an output fiber. In the other, a connection only needs to get to a particular output fiber, irrespective of what wavelength it will be on.

In this paper, we give novel constructions of strictly non-blocking and rearrangeably nonblocking WDM switches for both request models using limited range wavelength converters and arrayed waveguide grating routers. We fully analyze their blocking characteristics. Our designs are all relatively simple and easy to be laid out, and are useful for both optical circuit-switching and optical packet/burst switching.

As far as we know, these are the first of such constructions.

I. INTRODUCTION

Despite the recent downturn in the telecommunication industry in general and the optical networking sector in particular, Internet traffic is still doubling every year. In order to meet the ever increasing bandwidth demand from a large number of users in scientific computing and academic communities, as well as in military and other government agencies, there is a renewed interest in photonic switching as evident from several ongoing and planned national-scale projects in the US, Europe, and Asia (see for example, DARPA's recent BAA on Data in the Optical Domain Networks¹).

As the number of wavelengths in a wavelength division multiplexed (WDM) network increases to hundreds or more per fiber, and each wavelength operates at 10Gbps (OC-192) or higher [1]–[3], electronically switching the traffic carried on tens of fibers at each node, or equivalently, thousands of wavelengths at an aggregated throughput of several Terabits per second, becomes challenging due not only to the high costs of optical-electronic-optical or OEO conversion, but also more importantly, to the high power consumption (and heat dissipation) and large footprint (or space consumption). As

the optical device and component technology, and in particular, opto-electronic integration technology matures, photonic switching systems not only can potentially achieve hundreds of Terabits per second or higher ([4]) throughput, but also can be more cost-effective than their electronic counterparts even for applications requiring a lower throughput. Indeed, certain types of photonic switching fabrics such as the so-called WDM cross-connects (WXC's), or dynamic, reconfigurable optical add-drop multiplexers (OADMs), have already been deployed as an economic way to handle a large amount of traffic at the wavelength granularity.

In this paper, we will focus on cost-effective designs of photonic switching fabrics for WDM networks. As pointed out in [5], for cross-connecting wavelengths, that is, switching at the wavelength granularity for wavelength-routed or optical circuit-switched networks, it is not cost-effective to simply demultiplex all incoming f fibers, each having k wavelengths, into $N = f \cdot k$ wavelengths, and then use a photonic, purely space domain switch (or space switch for short) having a large (e.g., $N = 1000$) number of inputs and outputs, along with a stage of N full-range wavelength converters or FWCs (each of which needs to be capable of converting any one of the k wavelengths to any other one of the k wavelengths).

One of the biggest challenges is thus to design cost-effective photonic switching fabrics that can scale in size beyond a hundred of inputs and outputs, and at the same time, switch fast (e.g., tens of nanoseconds or less). For example, while one can achieve sub-nanosecond switching speed using (polarization independent) $LiNbO_3$ couplers based switches, not many such couplers can be integrated into a single module to form a large switch. Similarly, it is difficult to build a large switch with SOAs and passive InP gate arrays, mainly due to noise accumulation of the SOAs. On the other hand, optical MEMS switches can be large but switch only at the speed of milliseconds. These and other switching technologies have been described in [6].

In this work, we address the aforementioned challenge by, for the first time, presenting both strictly and rearrangeably

¹<http://www.darpa.mil/mto/solicitations/index.html#baa0319>

nonblocking WDM switch designs using a combination of *arrayed waveguide grating routers* or AWGRs and *limited range wavelength converters* or LWCs. An AWGR is wavelength sensitive, and as such is not the same as a space switch. For example, a signal carried on wavelength i , $0 \leq i \leq k - 1$, at the first input of a $k \times k$ AWGR will be (statically) routed to say, output i of the AWGR without being able to go to another output, nor a different wavelength at the output i . In order to route the signal to a different output, say $j \neq i$, it has to be converted to wavelength j at input i . The signal may need to be converted again at output j if it is to be carried on a different wavelength (including the original wavelength i). On the other hand, signals from multiple inputs of an AWGR can be routed to the same output (as long as they use different wavelengths), which make it more versatile (functionally more powerful) than a space switch (wherein only one input signal can go to any given output).

AWGRs are also interesting because they can be integrated in a large scale, and unlike a passive star coupler, the signals going through an AWGR are virtually lossless. In addition, with wavelength converters (WCs), fast switching is possible since the switching speed depends only on the speed of wavelength conversion, which is tens of nanoseconds or less even with current technology. Of course, in a $W \times W$ AWGR router, in order to be able to route a signal to any one of the W outputs, fast tunable WCs that can cover the entire spectrum of the W wavelengths are needed. And if W is large, as in [7] where $W = N = f \cdot k$, these WCs become too expensive and perhaps impractical.

Our designs of strictly and rearrangeably nonblocking WDM switches based on AWGRs and LWCs (each of which covers the spectrum of less than k wavelengths) can not only switch faster but also be less expensive than any existing nonblocking switch designs based on full-range or wider-range WCs such as FWCs [5]. As far as we know, these are the first such designs based on multistage construction that are strictly nonblocking (SNB) in both the space and wavelength domains. For example, the design described in [8] is blocking even though it used FWCs. In addition, a trivial extension to a multistage design using FWCs would lead to rearrangeably nonblocking (RNB) at the best. Later, we will also present our RNB designs using a minimal number of LWCs, thus having a lowest cost compared to any existing RNB designs.

Our designs are useful for both optical circuit-switching (or wavelength routing) networks and optical packet/burst switching. More specifically, while our SNB designs are most effective, for most optical circuit-switching networks, as well as optical packet/burst switched networks where optical packets/bursts are switched *synchronously* or *one batch at a time*, our RNB designs offer adequate performance at a lower cost than their strictly nonblocking counterparts.

The rest of the paper is organized as follows. Section II defines key concepts and notations used throughout the paper. Sections III presents and analyzes an RNB construction for one request model, where an input signal wants to get to a specific output fiber while it does not care about the exact wavelength

it gets carried on. Section IV describes and analyzes SNB and RNB constructions for the other request model, where an input signal wants to go to a specific output wavelength on a specific output fiber. Section V discusses three other constructions and analyzes their blocking properties. Section VI compares and contrasts our constructions with other known designs. The reader might want to refer to Table I at the end of the paper to get an overall picture of different designs discussed in this paper. Lastly, Section VII concludes the paper and discusses future works.

II. PRELIMINARIES

A. Basic concepts

A general WDM cross-connect (WXC) consists of f_1 input fibers with k_1 wavelengths on each, and f_2 output fibers with k_2 wavelengths on each, where $f_1 k_1 = f_2 k_2$. The set of input wavelengths need not have any relation with the output counterpart. This kind of WXC's were considered in [9] under the name "heterogeneous WXC's."

In this paper, we consider a simpler version where each WXC² has f input fibers and f output fibers, each of which can carry a set $\Lambda = \{\lambda_0, \dots, \lambda_{k-1}\}$ of k wavelengths. It should be noted that our constructions can be extended quite straightforwardly to the heterogeneous network case. The restriction was chosen merely for presentation clarity.

Let \mathcal{F} and \mathcal{F}' denote the set of input and output fibers, respectively. In the (λ, F, F') -request model (**model 1**), a connection request is of the form (λ, F, F') , which means that a connection is to be established from wavelength $\lambda \in \Lambda$ of input fiber $F \in \mathcal{F}$ to any free wavelength in output fiber $F' \in \mathcal{F}'$. In the $(\lambda, F, \lambda', F')$ -request model (**model 2**), the difference is that the output wavelength λ' is also specified.

Note that model 1 is useful for switching optical packets/bursts synchronously or one batch at a time, as well as for optical circuit-switching in general. Model 2 is particularly useful for asynchronous switching of optical bursts using JET (and void filling) [10], [11], as well as certain circuit-switching applications requiring specific QoS.

We next define the concepts of strictly nonblocking (SNB) and rearrangeably nonblocking (RNB) for both request models.

Consider a WXC with a few connections already established. Under model 1, a new request (λ, F, F') is said to be *valid* if λ is a free wavelength in fiber F , and there are at most $k - 1$ existing connections at F' . Under model 2, a new request $(\lambda, F, \lambda', F')$ is *valid* if λ is free in F and λ' is free in F' .

A *request frame* under the (λ, F, F') model is a set of requests such that no two requests are from the same wavelength of the same input fiber, and that there are at most k requests to any output fiber.

A *request frame* under the $(\lambda, F, \lambda', F')$ model is a set of requests such that no two requests are from the same input

²We will use the term WXC to refer also to WDM switches where switching speed maybe fast enough for optical packet/burst switching.

wavelength/fiber pair, and no two requests are to the same output wavelength/fiber pair.

The following definitions hold for both request models. A request frame is *realizable* by a WXC if all requests in the frame can be routed simultaneously. A WXC is *rearrangeably nonblocking* iff any request frame is realizable by the WXC. A WXC is *strictly nonblocking* iff a new valid request can always be routed through the WXC without disturbing existing connections.

Remark II.1. Note that RNB or SNB under the model 2 implies RNB or SNB under model 1. Also, under the same model, SNB implies RNB.

B. Commonly used notations

The following notations are used throughout the paper.

Let A and B be two subsets of wavelengths (having at most k wavelengths each). A limited wavelength converter (LWC) capable of converting any wavelength in A to any wavelength in B is denoted by $LWC(A, B)$.

We shall assume that $k = nb$, where n and b are two positive integers for which $n > f$. In practice, the number of wavelengths k per fiber is much larger than the number f of fibers; hence this assumption is practically no restriction. In the worse case, we can always take $b = 1$. Also, our designs also apply to the case where $k = f$ if we use FWCs, and where $k < f$ if wider-range WCs are available.

Implicitly, we agree that $\lambda_i = \lambda_{i \bmod k}$ when $i \geq k$ to avoid writing too many mod. Given natural numbers $i < j$, we use $[i, j]$ to denote the set of $(j - i)$ wavelengths $\{\lambda_i, \dots, \lambda_{j-1}\}$. Also, define $[j] = [0, j]$ for short.

In several of our designs, the wavelength set Λ is partitioned into b bands B_0, \dots, B_{b-1} of size n each. Band B_i consists of the set $[in, (i + 1)n]$ of wavelengths.

In this paper, graph theoretic concepts and notations we use are fairly standard. The reader is referred to [12] for related information.

C. Arrayed Waveguide Grating Routers

Passive arrayed waveguide grating routers (AWGRs) are attractive optical switching components because they are commercially available and inexpensive, relatively simple to fabricate, and they consume no power [8], [13].

A $W \times W$ AWGR on m wavelengths has a fixed routing pattern defined as follows. Number the inputs and outputs from 0 to $W - 1$, and the wavelengths from 0 to $m - 1$. Then, an input signal on wavelength i at input fiber j gets routed to the same wavelength on output fiber $((i - j) \bmod W)$.

Remark II.2. Some authors define the output line to be $((j - i) \bmod W)$ instead. This is not a discrepancy, as we can always re-number the wavelengths and input lines. For example, renaming λ_i by λ_{W-1-i} , and input j by $(W - 1 - j')$, then $(i - j) = (j' - i')$.

III. REARRANGEABLY NONBLOCKING CONSTRUCTION FOR THE (λ, F, F') -REQUEST MODEL

We use WXC-RNB-1 to denote the construction given in this section. A sample construction of WXC-RNB-1 is shown in Figure 1. Each band B_i are inputs to an $n \times n$ AWGR, preceded by one $LWC(B_i, [n])$ for each λ in B_i . The $n \times n$ AWGRs are numbered continuously from $AWGR_0$ to $AWGR_{fb-1}$.

At the second stage of the network, only one $k \times k$ AWGR is used. We refer to this AWGR as the *middle AWGR*. We use only n inputs numbered $0, b, \dots, (n - 1)b$ of the middle AWGR. A multiplexor M_c , $0 \leq c \leq n - 1$, is connected to input numbered cb of the middle AWGR.

For each $c = 0, \dots, n - 1$ and $j = 0, \dots, fb - 1$, output c of $AWGR_j$ is connected to a separate

$$LWC([n], [cb, (c + f)b])$$

and then to the multiplexor M_c .

At the last stage of the construction, there are f multiplexors combining signals to the output fibers. We connect only the first fb output lines of the middle AWGR to the output multiplexors. For each $i = 0, \dots, fb - 1$, output numbered i of the middle AWGR is connected to the multiplexor for output fiber $F_{[i/b]}$.

Remark III.1. In the description of our construction above, we used an AWGR in the middle which can take a WDM link (i.e., a fiber with multiple wavelengths) as one input. This type of AWGR splits the WDM signals carried on a link inside the fabric as if there is a passive splitter inside the fabric (and certainly do not need a wavelength demultiplexer inside). It utilizes the effectively different refractive indices (with respect to wavelengths) created by different waveguide lengths and angles (bending curves) to statically (or passively) route different wavelengths to different output ports.

One can also replace this type of AWGR by several AWGRs which take as inputs fibers with only one wavelength on each. If this type of AWGRs is used, there is no need to have the multiplexors at the second stage any more.

The following lemma is a nice observation which leads to the proof that this construction is rearrangeably nonblocking in the (λ, F, F') -model. We use the standard notation $\mathbb{Z}_p = \{0, \dots, p - 1\}$, for any positive natural number p .

Lemma III.2. Let n, b be natural numbers. Let $G = (U \cup V; E)$ be a bipartite multi-graph, where U and V form the vertex partition, and E is the edge set. Further assume that each vertex in U has degree at most n , and each vertex in V has degree at most $k = nb$.

Then, the edges of G can be colored with colors in the set $\mathbb{Z}_n \times \mathbb{Z}_b$ so that vertices in U are incident to colors with different first coordinates, and vertices in V are incident to different colors.

Remark III.3. Let (c, d) be a color in the set $\mathbb{Z}_n \times \mathbb{Z}_b$, then c is the first coordinate and d is the second coordinate of this

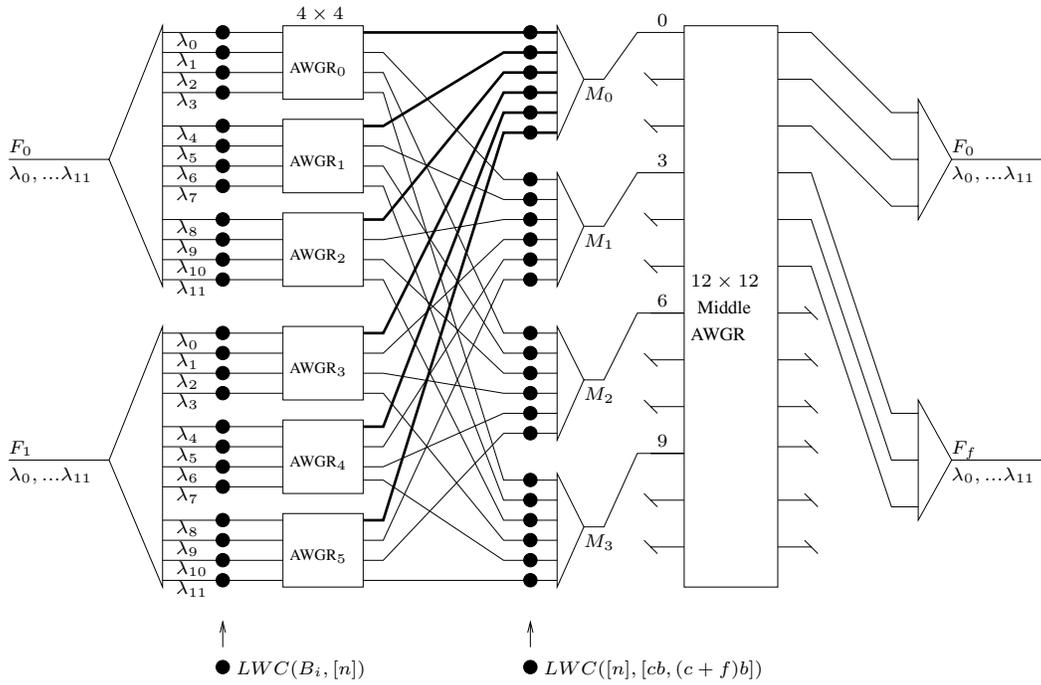


Fig. 1. Example of the WXC-RNB-1 construction for the (λ, F, F') -request model. The parameters are $k = 12, n = 4, b = 3$, and $f = 2$.

color. Two colors in the set $\mathbb{Z}_n \times \mathbb{Z}_b$ are different iff either the first or the second coordinates are different.

Proof of Lemma III.2. Construct a bipartite multi-graph $G' = (U \cup V'; E')$ from G by splitting each vertex $v \in V$ into b vertices $v^{(0)}, \dots, v^{(b-1)} \in V'$ so that each of the $v^{(i)}$ is incident to at most n edges which were previously incident to v . This is certainly possible, since the degree of v in G is at most nb . In this way, G' is a bipartite multi-graph with maximum degree at most n .

By König's line coloring theorem [14], the graph G' can properly be edge-colored with at most n colors. Recall that a proper edge-coloring is a coloring for which all vertices are incident to edges of different colors. A few (fast) polynomial-time coloring algorithms can be found in [15]–[17], for instance.

We use \mathbb{Z}_n for this set of n colors, which shall be the first coordinates of the final coloring. For each $i = 0, \dots, b-1$, all edges incident to the i th copy $v^{(i)}$ of a vertex $v \in V'$ get i as their second color-coordinate.

The coloring (with two-coordinate colors) of G' induces a coloring of G with the desired property. \square

Theorem III.4. *The WXC-RNB-1 construction described above is rearrangeably nonblocking under the (λ, F, F') -request model.*

Proof. Let \mathcal{R} be a request frame for our WXC under the (λ, F, F') -request model. We shall show that all requests in \mathcal{R} can be routed simultaneously through the network.

Recall that we number the input-stage AWGRs from AWGR_0 to AWGR_{fb-1} . Let $G = (U \cup V; E)$ be an $fb \times f$ bipartite multi-graph, i.e. $|U| = fb, |V| = f$, constructed from

\mathcal{R} as follows. Let

$$\begin{aligned} U &= \{u_0, \dots, u_{fb-1}\} \\ V &= \{v_0, \dots, v_{f-1}\}. \end{aligned}$$

There is (a copy of) an edge $(u_i, v_j) \in E$ for each request $(\lambda_p, F_q, F_j) \in \mathcal{R}$ where $i = qb + \lfloor p/n \rfloor$. Basically, there is a vertex u_i for band $B_{i \bmod b}$ on input fiber $F_{\lfloor i/b \rfloor}$.

For any particular pair (p, q) , there can be at most one request from λ_p in input fiber F_q . Hence, vertices in U of the graph G have maximum degree n . For any output fiber F_j , there can be at most k requests to it; hence, vertices in V of G have maximum degree $k = nb$.

Consequently, G can be edge-colored satisfying the conclusions of Lemma III.2. Each request $R \in \mathcal{R}$ thus gets a color in the set $\mathbb{Z}_n \times \mathbb{Z}_b$. We shall use this coloring to establish routes for all requests in \mathcal{R} simultaneously.

Let $(\lambda_p, F_q, F_j) \in \mathcal{R}$ be a request which gets colored $(c, d) \in \mathbb{Z}_n \times \mathbb{Z}_b$.

The basic idea is to route this request from λ_p to the c th output line of its corresponding first-stage AWGR. Next, by construction there are b output lines of the middle AWGR which are connected to the multiplexor on output fiber F_j . The second-stage LWC finishes the job by routing our connection to the d th line of this set of output lines. Lastly, the properties of the coloring ensures that our routing algorithm creates no conflict at the multiplexors as well as at the LWCs.

In order to get to output c of the corresponding input-stage AWGR, the LWC converts λ_p to $\lambda_{(c+(p \bmod n)) \bmod n} \in [n]$. We claim that no output of any input-stage AWGR is used twice. The claim directly follows from the fact that requests coming out of the same band get different colors, thus get

routed through different output line of the band's AWGR. It follows easily that, all the second-stage LWCs are going to be used at most once.

Let us now go back to the request (λ_p, F_q, F_j) which is now on $\lambda_{(c+(p \bmod n)) \bmod n}$. As the connection gets out on output c , it is going through an LWC($[n], [cb, (c+f)b]$) and then to the multiplexor connected to input numbered cb of the middle AWGR.

The b outputs of the middle AWGR connected to F_j are numbered $bj, bj+1, \dots, bj+b-1$. The d th line on this group is thus numbered $bj+d$. For $\lambda_{(c+(p \bmod n)) \bmod n}$ on input cb to get out on output $bj+d$, the wavelength needs to be converted to

$$\lambda_{(bj+d+cb)} = \lambda_{(c+j)b+d} \in [cb, (c+f)b].$$

(Recall that $\lambda_x = \lambda_{x \bmod k}$.) Hence, the request (λ_p, F_q, F_j) is routed to the right destination, so do all requests in \mathcal{R} .

As noted earlier, our routing algorithm does not use any LWC twice, which means there is no conflict at the LWCs. We next confirm that there is also no multiplexor getting the same wavelength more than once.

Consider first the multiplexors at the middle AWGR. The c th multiplexor gets wavelengths of the form $\lambda_{(c+j)b+d}$ for different values of j and d . Recall that $d \in \mathbb{Z}_b$, and $j \in \mathbb{Z}_f$. Suppose there is some conflict at the c th middle multiplexor, then there are two requests to some F_{j_1} and F_{j_2} which get colored (c, d_1) , (c, d_2) such that

$$(c+j_1)b+d_1 = (c+j_2)b+d_2 \pmod{k}.$$

This implies

$$(j_1-j_2)b = d_2-d_1 \pmod{k}.$$

As $0 \leq d_1, d_2 \leq b-1$, it must be the case that $j_1 = j_2 = j$. Moreover, as $k = nb$, we must have $d_1 = d_2 \bmod b$, which means $d_1 = d_2 = d$ also. This contradicts our coloring which says that two requests to the same F_j get different (c, d) pairs.

Lastly, consider the multiplexors at the last stage. The j th multiplexors get wavelengths of the form $\lambda_{(c+j)b+d}$ for different values of c and d . A similar argument as above completes the proof. We omit the details. \square

Remark III.5. One might wonder if this idea can be used to construct an SNB WXC for the same request model. We do have one such construction similar to the one described above, i.e. only two stages of wavelength converters and AWGRs are used. Unfortunately, the number of AWGRs and LWCs is too large to be practical. This is explained by a somewhat surprising theorem shown in Ngo [18], which basically says that a WXC is SNB in the (λ, F, F') -request model if and only if it is SNB in the $(\lambda, F, \lambda', F')$ -request model.

As far as SNB under the (λ, F, F') -request model is concerned, the construction shown in Section IV-B is sufficient.

IV. THE $(\lambda, F, \lambda', F')$ -REQUEST MODEL

The constructions in this section are motivated from the idea of three-stage Clos network [19]. Readers who are familiar with the Clos networks will undoubtedly recognize the parallels.

A. Rearrangeably nonblocking construction

Figure 2 shows an illustration of this construction, which we refer to as WXC-RNB-2.

The first stage of the construction is identical to the WXC-RNB-1 construction described in Section III.

The second stage consists of n AWGRs of size $fb \times fb$, numbered from 0 to $n-1$. The i th output of the j th first-stage AWGR is connected to an LWC($[n], [fb]$), and then to the j th input of the i th middle-stage AWGR.

The last stage, as before, consists of one multiplexor for each output fiber. We number the multiplexors from 0 to $f-1$, as usual. The connection patterns from the middle-stage AWGRs to the multiplexors are identical. The i th output of a middle stage multiplexor is connected to an LWC($[fb], B_{i \bmod b}$), and then to the multiplexor numbered $\lfloor i/b \rfloor$.

Theorem IV.1. *The WXC-RNB-2 construction described above is rearrangeably nonblocking under the $(\lambda, F, \lambda', F')$ -request model.*

Proof. Let \mathcal{R} be a request frame for our WXC under the $(\lambda, F, \lambda', F')$ -request model. We shall show that all requests in \mathcal{R} can be routed simultaneously through the network.

Recall that we number the input-stage AWGRs from AWGR₀ to AWGR _{$fb-1$} . Let $G = (U \cup V; E)$ be an $fb \times fb$ bipartite multi-graph, i.e. $|U| = |V| = fb$, constructed from \mathcal{R} as follows. Let $U = \{u_0, \dots, u_{fb-1}\}$ and $V = \{v_0, \dots, v_{fb-1}\}$. There is an edge $(u_i, v_j) \in E$ for each request $(\lambda_r, F_s, \lambda_p, F_q) \in \mathcal{R}$ where

$$\begin{aligned} i &= sb + \lfloor r/n \rfloor \\ j &= qb + \lfloor p/n \rfloor. \end{aligned}$$

For any particular pair (p, q) , there can be at most one request from λ_p of input fiber F_q , and there can be at most one request to λ_p of output fiber F_q . This conclusion comes from the fact that \mathcal{R} is a request frame.

Thus, it is straightforward that the maximum degree of G is at most n . Intuitively, there can be at most n requests from any band on fiber s , and at most n requests to any band on fiber q .

By König's line coloring theorem [14], we can properly edge-color G with at most n colors. Hence, each request in \mathcal{R} gets a color between 0 and $n-1$, such that requests coming out of the same input band and fiber or to the same output band and fiber get different colors.

In the rest of the proof, we shall use this coloring to route requests in \mathcal{R} . The idea is to route request colored c to the middle-stage AWGR numbered c .

Consider a request $(\lambda_r, F_s, \lambda_p, F_q) \in \mathcal{R}$, which gets colored c . We tune the LWC($B_{\lfloor r/n \rfloor}, [n]$), which is connected to λ_r

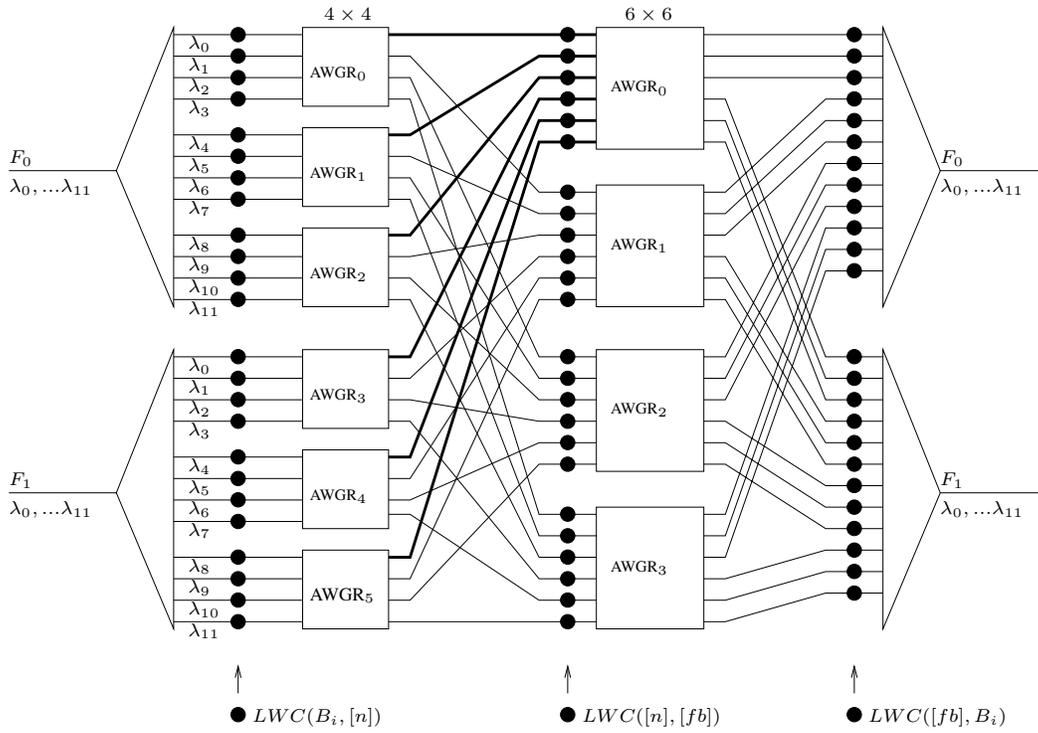


Fig. 2. An example of the WXC-RNB-2 construction for the $(\lambda, F, \lambda', F')$ -request model. The parameters are $k = 12, n = 4, b = 3$, and $f = 2$.

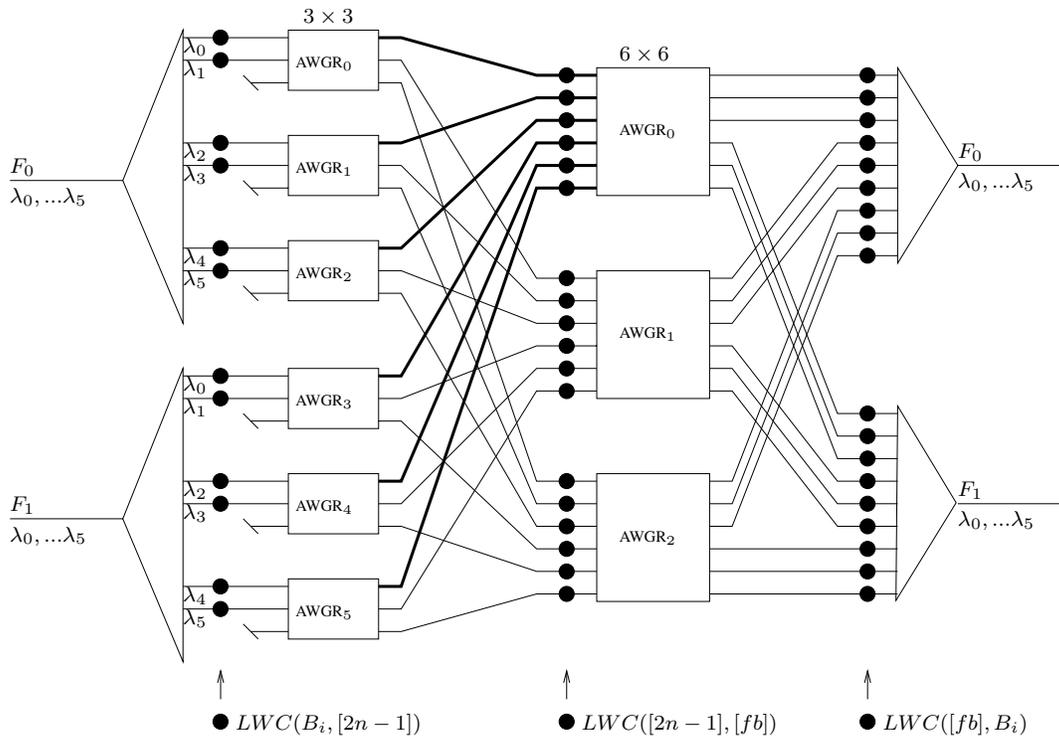


Fig. 3. An example of the WXC-SNB-2 construction for the $(\lambda, F, \lambda', F')$ -request model. The parameters are $k = 6, n = 2, b = 3$, and $f = 2$.

at input fiber F_s , so that it converts λ_r to $\lambda_{(c+(r \bmod n)) \bmod n}$ and hence it will get out on output c of its input-stage AWGR.

This signal shall get to the c th middle-stage AWGR at its $(sb + \lfloor r/n \rfloor)$ input line. At this point, the corresponding LWC is tuned to convert

$$\lambda_{(c+(r \bmod n)) \bmod n} \in [fb]$$

to

$$\lambda_{((qb+\lfloor p/n \rfloor)+(sb+\lfloor r/n \rfloor)) \bmod fb} \in [fb],$$

which shall get routed to output numbered $qb + \lfloor p/n \rfloor$ of this AWGR.

Lastly, this output line is connected to an LWC($[fb], B_{\lfloor p/n \rfloor}$), then multiplexed to output fiber F_q . This last LWC shall convert the previous wavelength to λ_p as desired.

To complete the proof, we need to show that we did not use any LWC twice.

As the requests from the same input band and fiber get different colors, no two requests shall be routed through the same output line of any input-stage AWGR. This means we did not use any second-stage LWC twice.

Similarly, the requests to the same output band and fiber get different colors, they will come to the band from different middle-stage AWGRs. \square

B. Strictly nonblocking construction

Figure 3 shows an illustration of the construction, called WXC-SNB-2.

At each input fiber, each line in band B_i is connected to an LWC($B_i, [2n-1]$) and then to one of the first n inputs of a $(2n-1) \times (2n-1)$ AWGR. Thus, the first stage consists of fb AWGRs of larger size than those in the RNB construction.

At the second stage, there are $(2n-1)$ AWGRs of size $fb \times fb$. The i th output of the j th first-stage AWGR is connected to an LWC($[2n-1], [fb]$), and then to the j th input of the i th middle-stage AWGR.

The last stage has one multiplexer for each output fiber. The connection patterns from middle-stage AWGRs to the multiplexers are identical. The i th output of a middle-stage AWGR is connected to an LWC($[fb], B_{i \bmod b}$), and then to the multiplexer on output fiber $F_{\lfloor i/b \rfloor}$.

Theorem IV.2. *The WXC-SNB-2 construction described above is strictly nonblocking under the $(\lambda, F, \lambda', F')$ -request model.*

Proof. Suppose our WXC has had a few connections set up. Let $(\lambda_i, F_j, \lambda_p, F_q)$ be a new valid connection request, i.e. λ_i and λ_p are free wavelengths on input fiber j and output fiber q , respectively. We want to show that we can find a route for this request through the WXC without disturbing the routes of existing connections.

After being demultiplexed at input fiber j , λ_i comes to LWC($B_{\lfloor i/n \rfloor}, [2n-1]$), and then connected to the $(i \bmod n)$ th input of the first-stage AWGR numbered $jb + \lfloor i/n \rfloor$.

As there can be at most $n-1$ existing connections which come from band $B_{\lfloor i/n \rfloor}$ of input fiber F_j , there must be at least

n outputs of AWGR $_{jb+\lfloor i/n \rfloor}$ which carry no signal. With the help of the LWC($B_{\lfloor i/n \rfloor}, [2n-1]$), the signal on λ_i can get out on any of these free output lines. Let \mathcal{A} be a set of n middle-stage AWGRs to which n of these free lines are connected to. Then, a signal on λ_i can get to any AWGR in \mathcal{A} . However, λ_i shall be converted to a wavelength in $[2n-1]$.

To this end, notice that λ_p belongs to band $B_{\lfloor p/n \rfloor}$ of output fiber F_q . By construction, each AWGR in the set \mathcal{A} has one output line connected to an LWC($[fb], B_{\lfloor p/n \rfloor}$), and then to the multiplexer on output fiber q . Let \mathcal{L} be this set of output lines. Since there can only be at most $n-1$ existing connections to wavelengths in band $B_{\lfloor p/n \rfloor}$ of output fiber F_q , and since $|\mathcal{A}| = n$, at least one line $l \in \mathcal{L}$ carries no signal.

Suppose l corresponds to the AWGR $A \in \mathcal{A}$. We can now let the signal on the original λ_i get to AWGR A , at which it is converted by the LWC($[2n-1], [fb]$) to some wavelength in $[fb]$ in order to get out at l . Lastly, the LWC($[fb], B_{\lfloor p/n \rfloor}$) will convert the previous wavelength to λ_p . \square

V. ANALYSES OF SOME OTHER CONSTRUCTIONS

The design given in Ramamirtham-Turner [8] was not RNB, even in the less restrictive (λ, F, F') -request model. The Ramamirtham-Turner design essentially looks like the one in Figure 4, except that there are no wavelength converters at the second stage. One can also change the patterns of connecting

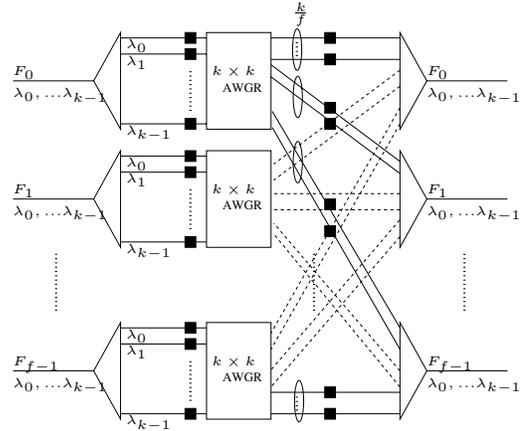


Fig. 4. A one-stage construction that is generally blocking, but should be more powerful than the one in [8].

the AWGRs to the output fibers. This design has the advantage of being simple and has RNB throughput up to 87% of an RNB design under the (λ, F, F') -request model.

A few natural questions arise. For instance, one might wonder if adding the second stage of FWCs, as shown in Figure 4, would yield an RNB design, or even SNB.

In this section, we analyze the design shown in Figure 4 and another “natural” design shown in Figure 5, which has one additional stage of $k \times k$ AWGRs and FWCs. The analyses motivates another interesting construction discussed in Section V-C.

Remark V.1. In most analyses given in this section, we ignore the issue of divisibility for presentation clarity. In the figures,

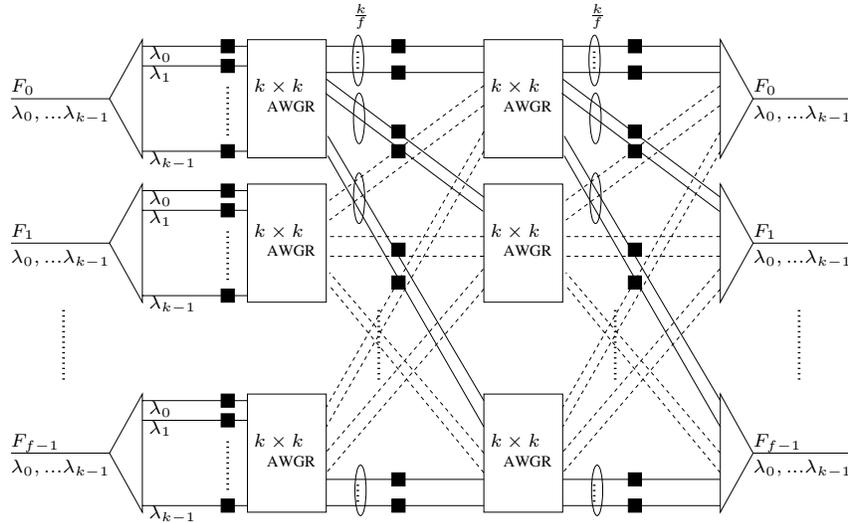


Fig. 5. A two stage construction using FWCs that is RNB-2 but not SNB-1.

the black squares represent FWCs.

A. A one stage construction

We analyze the construction shown in Figure 4 in this section. We ignore the trivial case when $f = 1$.

Theorem V.2. *The construction shown in Figure 4 is not rearrangeably nonblocking under the $(\lambda, F, \lambda', F')$ -request model, no matter how we connect the AWGRs to the output multiplexors.*

Proof. No matter what the connection pattern is, there must be an output multiplexor to which the first AWGR has $\leq k/f$ outputs connected. Without loss of generality, assume there are $\leq k/f$ links from the first AWGR to the first multiplexor M_0 .

We first show that there exists a request frame \mathcal{R} in which there are $(f-1)\frac{k}{f} + 1$ requests of the form $(\lambda_i, F_0, \lambda_j, F_0)$, so that converting λ_i to λ_j would direct a signal on input λ_i to an output link **not** connected to M_0 .

To form such a request frame, we need a one on one matching of size $(f-1)\frac{k}{f} + 1$ between the λ_i and their candidates. We verify P. Hall's matching condition ([20]) to show the claim.

Let $\bar{\Lambda}$ be any subset of size $(f-1)\frac{k}{f} + 1$ of $\Lambda = \{\lambda_0, \dots, \lambda_{k-1}\}$. Note that, for each $\lambda_i \in \bar{\Lambda}$ there are at least $(f-1)\frac{k}{f}$ candidate wavelengths for λ_j , one for each output line not connected to M_0 . Moreover, two different λ_i have two different candidate sets. Thus, any subset of size ≥ 2 of $\bar{\Lambda}$ has at least $1 + (f-1)\frac{k}{f}$ total candidates. A subset of size 1 certainly has at least 1 candidate. Consequently, P. Hall's matching theorem ([20]) shows the existence of such a request frame.

To realize a request in such an \mathcal{R} , each λ_i has to use a second stage FWC. However, the number of requests $(f-1)\frac{k}{f} + 1$ is strictly more than the number of available FWCs, which is at most $\frac{k}{f}$. Hence, this construction is not RNB. \square

Theorem V.3. *The construction shown in Figure 4 is not strictly nonblocking under the (λ, F, F') -request model, no matter how we connect the AWGRs to the output multiplexors.*

Proof. Without loss of generality, assume there is a set L of $\leq k/f$ links from the first AWGR to the first multiplexor M_0 .

Let S be the set of wavelengths to which λ_0 has to be converted to, in order to get out on any link in L . Noting that $k/f \leq k-1$, it is easy to construct a state of the network in which λ_0 is not part of any current request, and each FWC on each link in L is used to convert an incoming wavelength to separate wavelength in S .

The request (λ_0, F_0, F_0) is now blocked. \square

Lastly, we show that even in the less restrictive (λ, F, F') -request model, this design is generally not RNB. Unlike the previous two theorems, we only prove the following theorem for the particular connection pattern shown in Figure 4, where there are two AWGRs which have the first k/f outputs connected to M_0 .

Remark V.4. Much less restrictive forms of the following theorem can also be shown, at the expense of clarity.

Theorem V.5. *Consider the construction shown in Figure 4. Suppose there is an output multiplexor to which two AWGRs have their first k/f outputs connected to. Then, when $f \geq 6$ the construction is not rearrangeably nonblocking under the (λ, F, F') -request model.*

Proof. Without loss of generality, we assume the two AWGRs are $AWGR_0$ and $AWGR_1$ and the multiplexor is M_0 .

Let $\bar{\Lambda} = \{\lambda_0, \dots, \lambda_{k/2-1}\}$. Consider a request frame $\mathcal{R} = \mathcal{R}_0 \cup \mathcal{R}_1$, in which \mathcal{R}_i contains $k/2$ requests from wavelengths in $\bar{\Lambda}$ of input- F_i to output- F_0 , $i = 0, 1$. The main idea is still to show that there are not enough FWCs to avoid conflicts at M_0 .

For each $\lambda_i \in \bar{\Lambda}$, let $W(\lambda_i)$ be the set of wavelengths λ_i has to be converted to, in order to get out to M_0 . Note that $|W(\lambda_i)| = k/f$, and, due to the circular shift nature of AWGRs,

$$\left| \bigcup_{0 \leq i \leq k/2-1} W(\lambda_i) \right| = k/2 + k/f - 1.$$

(We ignore the issue of divisibility for clarity.) For each $i = 0, 1$, the requests in \mathcal{R}_i can only use at most k/f second-stage FWCs, leaving the other $k/2 - k/f$ requests not being able to convert after getting out of the AWGR.

Consequently, when

$$k/2 - k/f > \frac{1}{2}(k/2 + k/f - 1),$$

or $k/2 > 3k/f - 1$, there must be some conflict between requests in \mathcal{R}_0 and requests in \mathcal{R}_1 .

We can ensure $k/2 > 3k/f - 1$ whenever $f \geq 6$. \square

B. A two-stage construction

The one-stage construction in the previous section is fairly restrictive. In this section, we analyze the two stage construction shown in Figure 5 and its variations.

Lemma V.6. *Let f, k be positive integers where $f|k$. Let $G = (U \cup V; E)$ be a k -regular bipartite multi-graph, where U and V form the vertex partition, and E is the edge set.*

Then, the edges of G can be colored with f colors such that each vertex is incident to exactly k/f edges of any particular color.

Proof. Splitting each vertex $w \in U \cup V$ into k/f copies of degree f each, we obtain an f -regular graph G' . König's theorem [14] ensures that G' is f -edge colorable. (Again, coloring algorithms can be found in [15]–[17].) This coloring of G' induces a desired coloring of G . \square

Theorem V.7. *The construction shown in Figure 5 is rearrangeably nonblocking under the $(\lambda, F, \lambda', F')$ -request model. The connection patterns could be arbitrary, as long as they are evenly divided into groups of size k/f at each stage.*

Proof. Without loss of generality, consider a full request frame \mathcal{R} of size fk . If all full request frames are realizable, then all request frames are realizable.

Construct an $f \times f$ bipartite multi-graph $G = (U \cup V, E)$, where

$$\begin{aligned} U &= \{u_0, \dots, u_{f-1}\}, \\ V &= \{v_0, \dots, v_{f-1}\}, \end{aligned}$$

and there is an edge (u_i, v_j) for each request $(\lambda_p, F_i, \lambda_q, F_j) \in \mathcal{R}$.

Clearly G is k -regular and can be colored according to Lemma V.6. Let \mathcal{Z}_f be the set of colors. We separately route k/f requests from F_i which were colored c to the c th second-stage AWGR. Each of these requests can be routed on a separate link to the AWGR. As there are also f/k requests

colored c to output fiber F_j , the FWCs at the c th AWGR can be used to distribute k/f requests on f/k links to each output fiber. The last stage of FWCs finishes the job by converting their incoming wavelengths to the desired wavelengths. \square

Remark V.8. The details of the above proof can be done in a similar fashion as in Theorem III.4. We were brief in this proof due to the similarity and the space limit. The proof implicitly contains a routing algorithm.

Theorem V.9. *The construction is shown in Figure 5 is not strictly nonblocking under the (λ, F, F') -request model.*

Proof. The proof of this fact is quite simple, yet fairly tedious to present formally. Inspection for the case $f = 2$ should give the reader a good idea of how a blocking network state can be constructed. We omit the details for this version of the paper. \square

C. A rearrangeably nonblocking construction for $f = 2$ under the (λ, F, F') -request model

In this section, we present an interesting and very simple construction, albeit limited to $f = 2$. We believe that this construction contains a good idea to be explored further.

The construction is shown in Figure 6. It is self-explanatory. All wavelength converters are FWCs.

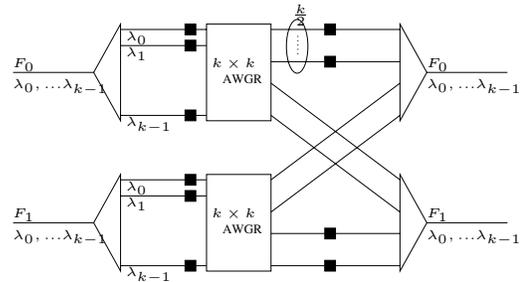


Fig. 6. A two fiber construction that is RNB-1 using less than 2 stages of FWCs.

Theorem V.10. *The construction for $f = 2$ shown in Figure 6 is rearrangeably nonblocking under the (λ, F, F') -request model.*

Proof. Consider a fixed output fiber F_j , $j = 0, 1$. We first claim that any request frame \mathcal{S} of at most $k/2$ requests of the form (λ_p, F_q, F_j) , $p \in \mathbb{Z}_k$, $q \in \mathbb{Z}_2$, is realizable without using the second-stage converters. In fact, we can route all requests in \mathcal{S} in a greedy manner. Consider a state of the cross-connect where there are $< k/2$ connections already established to F_j . Let (λ_p, F_q, F_j) be a new request to F_j . The input FWC for λ_p allows $k/2$ different choice for λ_p to get out on F_j . Each of these choices requires λ_p to be converted to a different wavelength. As there are currently $< k/2$ wavelengths at the multiplexor for F_j , there is always one available choice for λ_p .

Next, we show that any request frame \mathcal{S}_j of requests of the form (λ_p, F_q, F_j) , for a fixed j , is realizable. This fact will complete the proof.

Let \mathcal{S}_{ij} be the subset of requests of \mathcal{S}_j initiated from fiber F_i . Without loss of generality, assume $j = 0$ and $|\mathcal{S}_j| = k$.

Suppose $|\mathcal{S}_{10}| \geq k/2$, then all requests in \mathcal{S}_{10} can be routed through the first output of the second AWGR. (The first output is connected to output fiber F_0 .) For the set \mathcal{S}_{00} we can route them through separate outputs of the first AWGR. These outputs are connected to FWCs, which shall convert the requests to the free wavelengths. Note that this is possible since $|\mathcal{S}_j| = k$ and there are k available wavelengths.

On the other hand, suppose $|\mathcal{S}_{10}| < k/2$. This set \mathcal{S}_{10} along with $k/2 - |\mathcal{S}_{10}|$ requests from \mathcal{S}_{00} can be routed greedily. The other $k/2$ requests from \mathcal{S}_{00} make full use of $k/2$ wavelength converters. \square

Remark V.11. The above proof implicitly yields a routing algorithm, whose details we omit.

VI. COMPARISONS WITH KNOWN CONSTRUCTIONS

A. The (λ, F, F') -request model

Wilfong et al. [5] proposed several rearrangeably nonblocking architectures, and also a strictly nonblocking construction under the (λ, F, F') -model. Their architectures did not make use of AWGRs. The basic components of their designs were multiplexors and demultiplexors, 2×2 wavelength selective cross-connects (WSC) or optical add-drop multiplexors (OADM), And a special component named wavelength interchanger (WI) which is capable of permuting the wavelengths on its input fiber to the output. Essentially, a WI is equivalent to k full-range wavelength converters (FWCs).

Their rearrangeably nonblocking designs include the *WI-Beneš* and the *WI-Cantor* cross-connects. The *WI-Beneš* design needs about $f \lg f$ OADMs and f WIs, which are equivalent to kf FWCs. The *WI-Cantor* design needs about $f(\lg f)^2$ OADMs and f WIs, which are equivalent to kf FWCs.

Their strictly nonblocking design is the *Cantor/2-Beneš-Cantor/2* (CBC) cross-connect. This architecture needs about $f(\lg f)^2$ OADMs and $f \lg f$ WIs, which are equivalent to $kf \lg f$ FWCs.

All the above three networks have $O(\lg f)$ number of stages.

Ramamirtham and Turner [8] proposed an architecture using AWGRs, which is not rearrangeably nonblocking. However, their simulations show that their construction can achieve a throughput of 87% of a nonblocking switch. This construction has essentially one stage, uses $f k \times k$ AWGRs, and kf FWCs.

As compared to the constructions above, our 2-stage RNB construction in Section III uses $2fk$ LWCs, $fb n \times n$ AWGRs, and one $k \times k$ AWGR.

It is quite difficult to compare these designs which use different types of optical components. Number-wise, however, our construction for this request model is definitely competitive.

B. The $(\lambda, F, \lambda', F')$ -request model

Rasala and Wilfong [9], [21] described a strictly nonblocking construction using the so-called WDM split cross-connects, which consists of two WSCs a number of WIs in

the middle. If the number of WIs is $2f - 1$, then the split cross-connect is strictly nonblocking. On the other hand, although not mentioned in their papers, it is quite easy to show that n WIs are necessary and sufficient for a split cross-connect to be rearrangeably nonblocking.

In summary, a SNB split cross-connect requires two $(2f - 1) \times f$ WSCs and about $(2f - 1)k$ FWCs. A RNB split cross-connect requires two $f \times f$ WSCs and about fk FWCs.

Our SNB construction from Section IV uses $fb n \times n$ -AWGRs, $2n - 1$ $fb \times fb$ -AWGRs, and a total of

$$kf + 2(2n - 1)fb = 5kf - 2fb$$

LWCs of various kinds. In the worst case when $b = 1$, we need $(5k - 2)f$ LWCs.

Our RNB construction from Section IV requires $fb n \times n$ -AWGRs, n $fb \times fb$ -AWGRs, and a total of $3kf$ LWCs of various kinds.

Again, it is quite difficult to compare these different constructions as we used more limited wavelength converters than their full wavelength converters. On the other hand, the AWGRs are much cheaper than the WSCs. One of our future works is to find a good cost model to compare these different constructions.

Table I compares all constructions in this paper with the known constructions. Entries marked with “-” means NO, and entries marked with “?” are open questions.

VII. CONCLUSIONS AND FUTURE WORKS

We have given a number of novel constructions of rearrangeably nonblocking (RNB) and strictly nonblocking (SNB) WDM cross-connects (or switches) under two different request models. The optical components we used are multiple stages of limited-range wavelength converters (LWCs) and small to medium sized arrayed waveguide grating routers (AWGRs). Our designs are all relatively simple and easy to be laid out, and are useful for both optical circuit-switching and optical packet/burst switching.

There is no prior multistage SNB or RNB designs based on AWGRs and certainly not LWCs. For example, the design proposed in [8] which used a single stage of AWGRs is blocking, in addition to using full-range wavelength converters (FWCs). Other known designs have used wavelength selective cross-connects and FWCs [5].

We have not considered a nonblocking degree called wide-sense nonblocking [22] (WSNB), which is less restrictive than strictly, yet more powerful than rearrangeably nonblocking. One particular reason is that even in the classical switching networks, there are relatively few results on WSNB. The reader is refer to [23], [24] for some recent WSNB works.

Last but not least, developing a good cost, complexity and performance model taking into consideration the switching speed, the amplification needed, the signal-to-noise ratio, and the integrability, for the purpose of evaluating and comparing various WDM cross-connects, especially those constructed from wavelength converters and AWGRs is an interesting research topic.

TABLE I

TABULATED COMPARISONS OF DIFFERENT CONSTRUCTIONS. THE SECOND HALF OF THE TABLE CONTAINS OUR CONSTRUCTIONS, REFERRED TO BY THE SECTION NUMBERS THEY WERE ANALYZED. NOTE THAT $k = nb$. WE USE $p \otimes q$ TO MEAN p AWGRS OF SIZE q . RNB/SNB-1 MEANS RNB/SNB UNDER THE (λ, F, F') -MODEL. RNB/SNB-2 IS FOR THE $(\lambda, F, \lambda', F')$ -MODEL.

	SNB-1	RNB-1	SNB-2	RNB-2	#FWCs	#LWCs	#AWGRs	#WSCs	#OADMs
WI-Beneš [5]	-	Yes	-	-	kf	-	-	-	$f \lg f$
WI-Cantor [5]	-	Yes	-	-	kf	-	-	-	$f(\lg f)^2$
CBC [5]	Yes	Yes	-	?	$kf \lg f$	-	-	-	$f(\lg f)^2$
Ref. [8]	-	-	-	-	fk	-	$f \otimes k$	-	-
Ref. [9], [21]	Yes	Yes	Yes	Yes	$(2f-1)k$	-	-	2 of dim. $f \times (2f-1)$	-
Section III	-	Yes	-	-	-	$2fk$	$fb \otimes n$ $1 \otimes k$	-	-
Section IV-A	-	Yes	-	Yes	-	$3fk$	$fb \otimes n$ $n \otimes fb$	-	-
Section IV-B	Yes	Yes	Yes	Yes	-	$5fk-2fb$	$fb \otimes n$ $2n-1 \otimes fb$	-	-
Section V-A	-	-	-	-	$2fk$	-	$f \otimes k$	-	-
Section V-B	-	Yes	-	Yes	$3fk$	-	$2f \otimes k$	-	-
Section V-C ($f = 2$)	-	Yes	-	-	$3k$	-	$2 \otimes k$	-	-

REFERENCES

- [1] Lucent Technologies Press Release, "Lucent Technologies unveils ultra-high-capacity optical system; Time Warner Telecom first to announce it will deploy the system," 2001, <http://www.lucent.com/press/0101/010117.nsa.html>.
- [2] Lucent Technologies Press Release, "Lucent Technologies engineer and scientists set new fiber optic transmission record," 2002, <http://www.lucent.com/press/0302/020322.bla.html>.
- [3] Lucent Technologies Website, "What is dense wave division multiplexing (DWDM)," 2002, <http://www.bell-labs.com/technology/lightwave/dwdm.html>.
- [4] Zhigang Jing H. Jonathan Chao, Kung-Li Deng, "A petabit photonic switch (p3s)," in *Proceedings of IEEE Infocom'03*, San Francisco, USA., April 2003.
- [5] G. Wilfong, B. Mikkelsen, C. Doerr, and M. Zirngibl, "WDM cross-connect architectures with reduced complexity," *Journal of Lightwave Technology*, vol. 17, no. 10, pp. 1732–1741, Oct 1999.
- [6] L.Y. Lin, E.L. Goldstein, L.M. Lunardi, and R.W. Tkach, "Optical crossconnects for high-capacity lightwave networks," *Journal of High Speed Networks*, vol. 8, no. 1, pp. 17–34, 1999.
- [7] S. J. B. Yoo et al., "High-performance optical-label switching packet routers and smart edge routers for the next generation internet," *IEEE JSAC Special issue on High-Performance Optical-Electronic Switches-Routers for High-Speed Internet*, To appear.
- [8] Jeyashanker Ramamirtham and Jonathan S. Turner, "Design of wavelength converting switches for optical burst switching," in *Proceedings of the 21st Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, 2002, vol. 2, pp. 1162–1171, IEEE.
- [9] April Rasala and Gordon Wilfong, "Strictly non-blocking WDM cross-connects for heterogeneous networks," in *Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing (STOC'2000, Portland, OR)*, New York, 2000, pp. 513–524, ACM.
- [10] Jinhui Xu, Chunming Qiao, Jikai Li, and Guang Xu, "Efficient channel scheduling algorithms in optical burst switching networks," in *Proceedings of IEEE Infocom'03*, San Francisco, USA., April 2003.
- [11] C. Qiao and M. Yoo, "Optical burst switching (obs) - a new paradigm for an optical internet," *Journal of High Speed Networks*, vol. 8, no. 1, pp. 69–84, 1999.
- [12] Douglas B. West, *Introduction to graph theory*, Prentice Hall Inc., Upper Saddle River, NJ, 1996.
- [13] C. Dragone, "An $n \times n$ optical multiplexor using a planar arrangement of two star couplers," *IEEE Photonic Technology Letters*, vol. 6, pp. 812–815, 1991.
- [14] D. König, "Über graphen und ihre anwendung auf determinantentheorie und mengenlehre," *Math. Ann.*, vol. 77, pp. 453–465, 1916.
- [15] Richard Cole and John Hopcroft, "On edge coloring bipartite graphs," *SIAM J. Comput.*, vol. 11, no. 3, pp. 540–546, 1982.
- [16] Harold N. Gabow, "Using Euler partitions to edge color bipartite multigraphs," *Internat. J. Comput. Information Sci.*, vol. 5, no. 4, pp. 345–355, 1976.
- [17] Harold N. Gabow and Oded Kariv, "Algorithms for edge coloring bipartite graphs and multigraphs," *SIAM J. Comput.*, vol. 11, no. 1, pp. 117–129, 1982.
- [18] Hung Q. Ngo, "WDM switching networks and $[w, f]$ -connectors," 2004, Submitted to STOC 2004.
- [19] C. Clos, "A study of non-blocking switching networks," *Bell System Tech. J.*, vol. 32, pp. 406–424, 1953.
- [20] Phillip Hall, "On representatives of subsets," *J. London Math. Soc.*, vol. 10, pp. 26–30, 1936.
- [21] April Rasala and Gordon Wilfong, "Strictly non-blocking WDM cross-connects," in *Proceedings of the Eleventh Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'2000, San Francisco, CA)*, New York, 2000, pp. 606–615, ACM.
- [22] V. E. Beneš, *Mathematical theory of connecting networks and telephone traffic*, Academic Press, New York, 1965, Mathematics in Science and Engineering, Vol. 17.
- [23] Penny E. Haxell, April Rasala, Gordon T. Wilfong, and Peter Winkler, "Wide-sense nonblocking wdm cross-connects," in *Algorithms—ESA '03 (Rome, Italy)*, vol. 2461 of *Lecture Notes in Comput. Sci.*, pp. 538–549, Springer, Berlin, 2002.
- [24] Frank K. Hwang, *The mathematical theory of nonblocking switching networks*, World Scientific Publishing Co. Inc., River Edge, NJ, 1998.