

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

Hung Q. Ngo, *Member, IEEE*, Dazhen Pan, and Chunming Qiao, *Member, IEEE*

Abstract—Constructing fast wavelength division multiplexing switches with cheap, integratable components, less power consumption and noise accumulation, 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 nonblocking 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, consume little power, do not accumulate much noise, 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.

Index Terms—Limited-range wavelength converters, non-blocking WDM cross-connects.

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 U.S., 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 10 Gb/s (OC-192) or higher [15]–[17], electronically switching the traffic carried on tens of fibers at each node, or equivalently, thousands of wavelengths at an aggregated throughout of several Terabits

per second, becomes challenging due not only to the high costs of optical–electronic–optical (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 throughput [8], 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), 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 [27], 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 LiNBO₃ 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 semiconductor optical amplifiers (SOAs) and passive InP gate arrays, mainly due to noise accumulation of the SOAs. On the other hand, optical microelectromechanical system (MEMS) switches can be large but switch only at the speed of milliseconds. These and other switching technologies have been described in [14].

In this work, we address the aforementioned challenge by—for the first time—presenting both strictly and rearrangeably nonblocking WDM switch designs using a combination of *arrayed waveguide grating routers* (AWGRs) and *limited range wavelength converters* (LWCs). An AWGR, which has been used for optical packet-switching and as a cross-connect in [11]

Manuscript received March 22, 2004; revised March 4, 2005; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor A. Somani. The work of H. Q. Ngo was supported in part by the National Science Foundation under NSF CAREER Award CCF-0347565 and by Telcordia Technologies Subcontract FA8750-04-C-0249 from the DARPA SRS Program. The work of C. Qiao was supported in part by the National Science Foundation under NSF 0312562. Part of this paper appeared in the Proceedings of IEEE INFOCOM 2004.

The authors are with the Department of Computer Science and Engineering, State University of New York at Buffalo, Amherst, NY 14260 USA (e-mail: hungngo@cse.buffalo.edu; dpan@cse.buffalo.edu; qiao@cse.buffalo.edu).

Digital Object Identifier 10.1109/TNET.2005.863459

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

and [18], is wavelength sensitive, and as such is not the same as a space switch. For instance, a signal carried on wavelength λ_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 λ_j at the first input. 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 λ_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 [5] 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) not only can switch faster but also are less expensive in terms of wavelength conversion cost than any existing nonblocking switch designs based on full-range or wider-range WCs such as FWCs [27]. 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 [23] 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 wavelength converter 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 which offered a lower cost than the SNB counterparts, can also be applied.

The rest of this paper is organized as follows. Section II defines key concepts and notations used throughout the paper. Section III presents and analyzes an RNB construction for one request model, where an input signal is to be routed 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 gets routed to a specific output wavelength on a specific output fiber. Section V discusses three other constructions and analyzes their blocking properties. Section VI com-




Fig. 1. Heterogeneous WDM cross-connect.

pare and contrasts our constructions with other known designs. The reader can refer to Table I at the end of the paper for an overall picture of different designs discussed in this paper. Finally, 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 $f_1 k_1 = f_2 k_2$ (see Fig. 1). The set of input wavelengths need not have any relation with the output counterpart. This kind of WXC was considered in [25] under the name “heterogeneous WXC.”

In this paper, we consider the “homogeneous” 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, 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, the difference is that the output wavelength λ' is also specified.

Note that the (λ, F, F') -model is useful for switching optical packets/bursts synchronously or one batch at a time, as well as for optical circuit-switching in general. The $(\lambda, F, \lambda', F')$ -model is particularly useful for asynchronous switching of optical bursts using JET (and void filling) [22], [28], as well as certain circuit-switching applications requiring specific quality of service (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 the (λ, F, F') -model, a new request (λ, F, F') is said to be *valid* if λ is a free wavelength in input fiber F , and there are at most $k-1$ existing connections at output fiber F' . Under the $(\lambda, F, \lambda', F')$ -model, a new request $(\lambda, F, \lambda', F')$ is *valid* if λ is free in F and λ' is free in F' .

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

A *request frame* under the (λ, F, F') -model is a set of requests such that no two requests are from the same input wavelength/fiber pair, 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 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 2.1: Note that RNB or SNB under the $(\lambda, F, \lambda', F')$ -model implies RNB or SNB under the (λ, F, F') -model, respectively. Also, under the same model, being SNB implies being 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 $b = 1$. 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 [26] for related information.

C. Arrayed Waveguide Grating Routers

Passive AWGRs are attractive optical switching components because they are commercially available and inexpensive, relatively simple to fabricate, and they consume virtually no power [4], [23]. AWGRs also exhibit flat spectral response and low insertion loss. They can be designed to perform multiplexing and demultiplexing simultaneously, and are suitable for a large number of channels. A potential drawback is that they are temperature-sensitive and thus require careful temperature control.

An $m \times m$ AWGR on m wavelengths $\{\lambda_0, \dots, \lambda_{m-1}\}$ has a fixed routing pattern defined as follows (see Fig. 2). The inputs and outputs of the AWGR are numbered from 0 to $m - 1$. An input signal carried on wavelength λ_i at input fiber j gets




Fig. 2. An $m \times m$ arrayed waveguide grating router.

routed to output fiber numbered $((i - j) \bmod m)$ on the same wavelength (λ_i).

Remark 2.2: Some authors define the output line to be

$$((j - i) \bmod m)$$

instead. This is not a discrepancy, as we can always renumber the wavelengths and input lines. For example, renaming λ_i by λ_{m-1-i} , and input j by $(m - 1 - j')$, then $(i - j) = (j' - i')$.

Just like a prism, an AWGR allows deflection of an optical beam at various directions, which has the effect of similar to a wavelength “space-switch.” When used in combination with tunable wavelength converters, we can turn an AWGR into an optical cross-bar as shall be seen later. For instance, in Fig. 2, if λ_i were converted to λ_x at the input fiber F_j , then λ_x would be routed to output $F'_{(x-j) \bmod m}$ instead of $F'_{(i-j) \bmod m}$. Additionally, if optical delay lines are used, we can turn this cross-bar into a time/space switch, facilitating optical burst/packet routing.

III. THE (λ, F, F') -REQUEST MODEL

A. Rearrangeably Nonblocking Construction

We use WXC-RNB-1 to denote the construction given in this section. A sample construction of WXC-RNB-1 is shown in Fig. 3. 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 3.1: 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

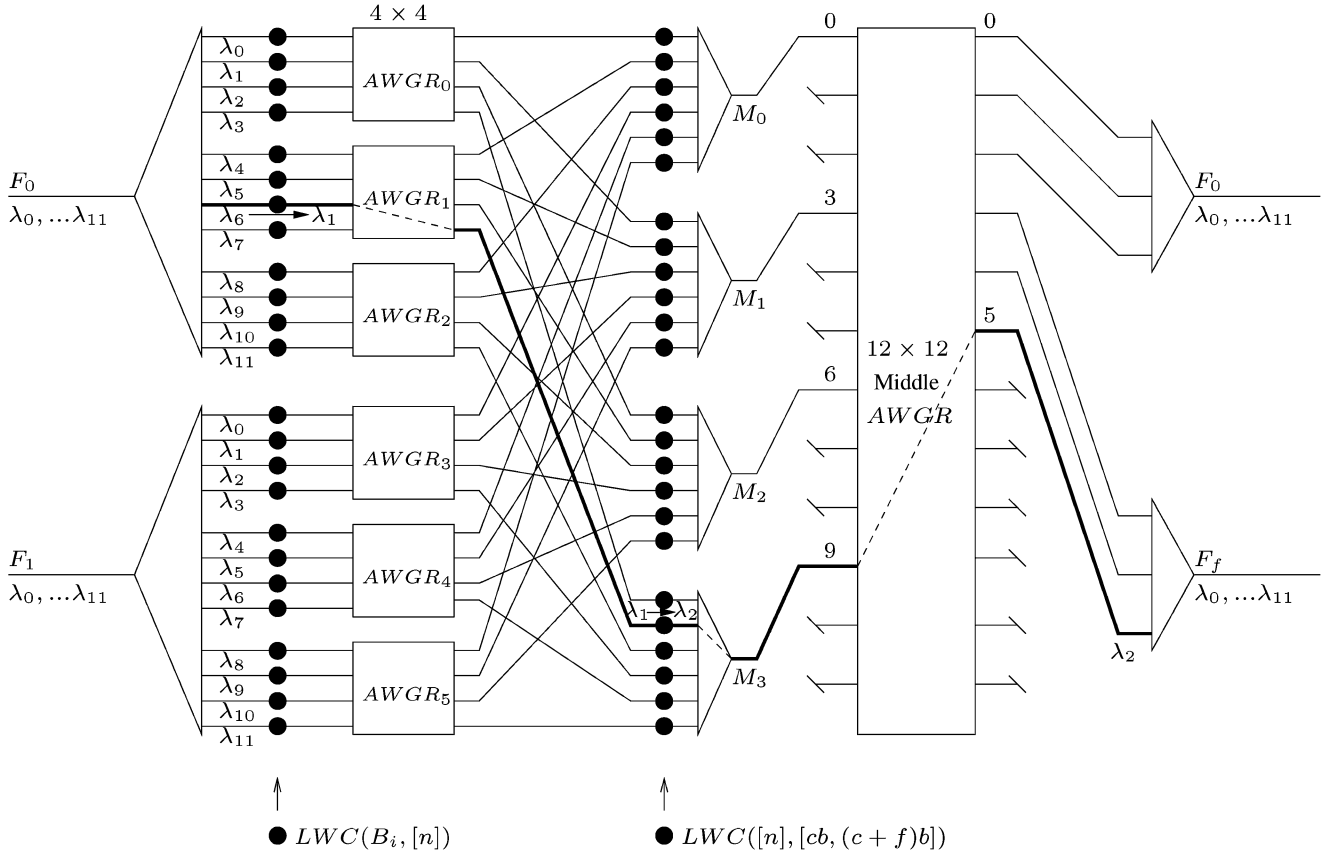


Fig. 3. Example of the WXC-RNB-1 construction for the (λ, F, F') -request model. The parameters are $k = 12$, $n = 4$, $b = 3$, and $f = 2$.

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 indexes (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 multiplexers 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 integer p .

Lemma 3.2: 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 3.3: 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 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 3.2: 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 [13], 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 [3], [6], [7], 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 3.4: 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

$$U = \{u_0, \dots, u_{fb-1}\}$$

$$V = \{v_0, \dots, v_{f-1}\}.$$

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 3.2]. 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 $(c, d) \in \mathbb{Z}_n \times \mathbb{Z}_b$.

As an example, we will follow the arguments of this proof with the request (λ_6, F_0, F_1) , i.e., $(p, q, j) = (6, 0, 1)$ of the sample design shown in Fig. 3. Suppose this request gets colored $(c, d) = (3, 2)$. Note that $n = 4$, $b = 3$, and $k = 12$.

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 $\lambda_{(c+p) \bmod n} \in [n]$.

In our example, λ_6 will be converted to $\lambda_{(6+3) \bmod 4} = \lambda_1$.

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 $\lambda_{(c+p) \bmod n}$. 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.

In our example, λ_1 will get out on output 3 (numbered from 0) of its AWGR, then go through LWC [4], [9], [15], the multiplexor, and finally to input numbered 9 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 $\lambda_{(c+p) \bmod n}$ 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} .

In our example, λ_1 is converted to $\lambda_{(3+1)3+2} = \lambda_2$ and finally gets to output fiber F_1 as desired.

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

B. Strictly Nonblocking Case

Since the (λ, F, F') -model is less restrictive than the $(\lambda, F, \lambda', F')$ -model, one might think that nonblocking cross-connects under the former are (strictly) cheaper than those under the later. This intuition is not necessarily true. In particular, one can show that we save nothing if we require our cross-connect to be strictly nonblocking. The proof of this fact requires lengthy rigorous settings which go beyond the scope of this paper. The reader is referred to [20] for more details.

The key point is that, as far as SNB under the (λ, F, F') -model is concerned the SNB construction for the $(\lambda, F, \lambda', F')$ -model 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 the three-stage Clos network [2].

A. Rearrangeably Nonblocking Construction

Fig. 4 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 4.1: The WXC-RNB-2 construction described above is rearrangeably nonblocking under the $(\lambda, F, \lambda', F')$ -request model.

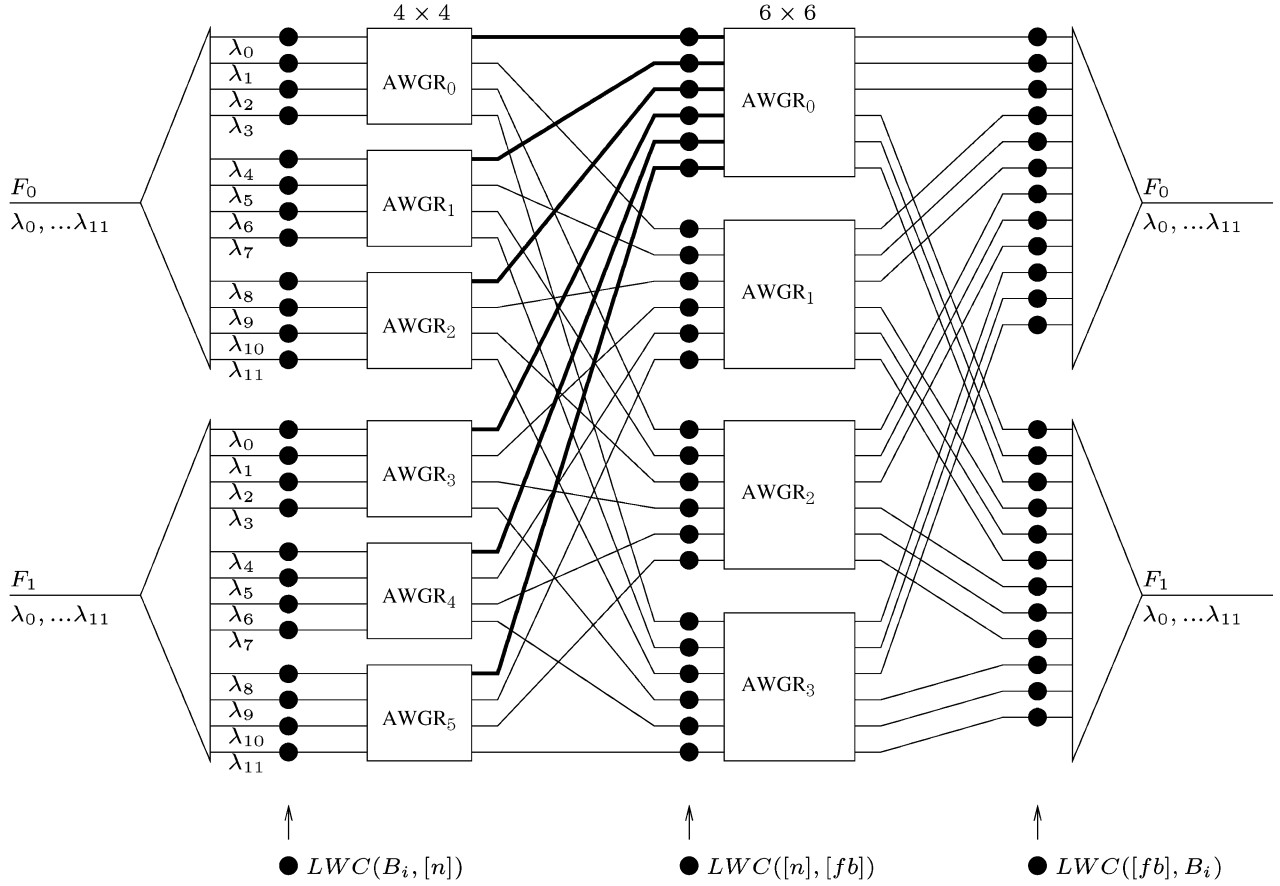


Fig. 4. Example of the WXC-RNB-2 construction for the $(\lambda, F, \lambda', F')$ -request model. The parameters are $k = 12$, $n = 4$, $b = 3$, and $f = 2$.

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 AWGR_0 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

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

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 [13], 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 $\text{LWC}(B_{\lfloor r/n \rfloor}, [n])$, which is connected to λ_r 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

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

to

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

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

Lastly, this output line is connected to an $\text{LWC}([fb], B_{\lfloor p/n \rfloor})$, 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

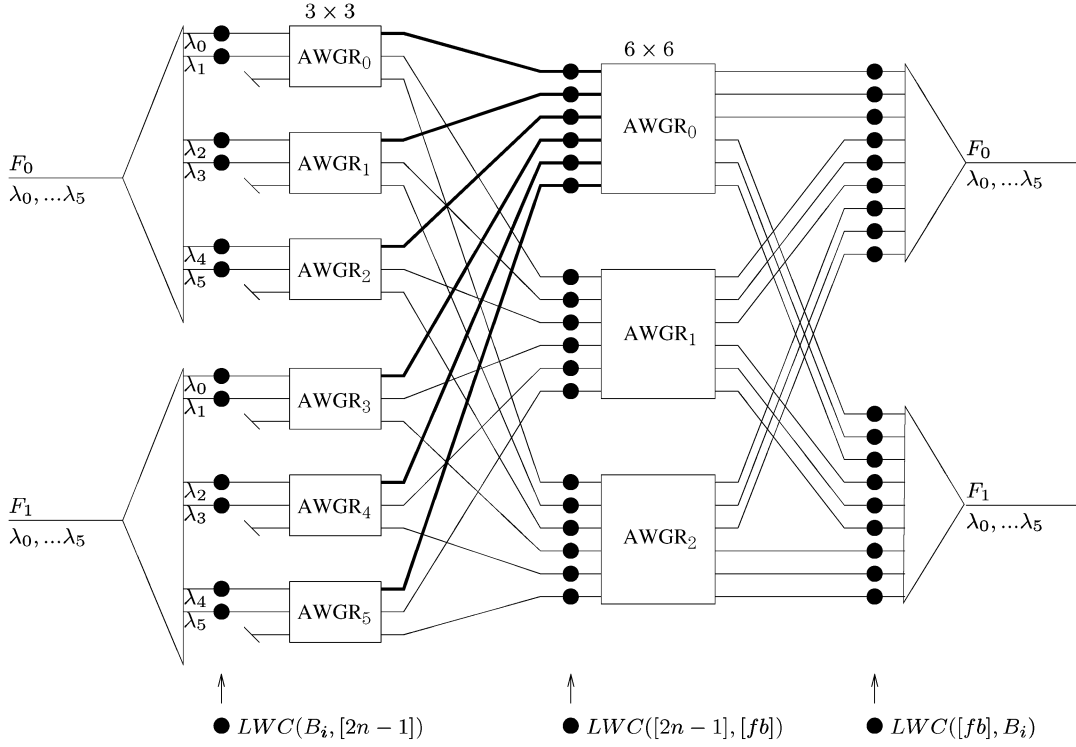


Fig. 5. Example of the WXC-SNB-2 construction for the $(\lambda, F, \lambda', F')$ -request model. The parameters are $k = 6$, $n = 2$, $b = 3$, and $f = 2$.

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

Fig. 5 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 multiplexor for each output fiber. The connection patterns from middle-stage AWGRs to the multiplexors 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 multiplexor on output fiber $F_{\lfloor i/b \rfloor}$.

Theorem 4.2: 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

$$LWC\left(B_{\lfloor \frac{i}{n} \rfloor}, [2n-1]\right),$$

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 multiplexor 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 OTHER CONSTRUCTIONS

The design given in [23] was not RNB, even in the less restrictive (λ, F, F') -request model. The Ramamirtham–Turner design essentially looks like the one in Fig. 6, except that there are no wavelength converters at the second stage. Their design

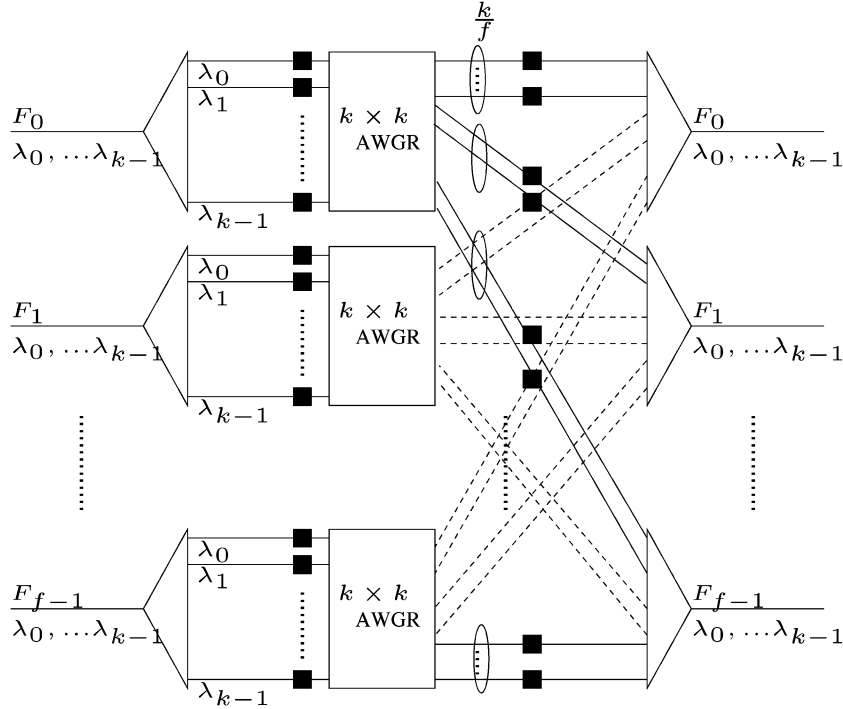


Fig. 6. One-stage construction that is generally blocking, but should be more powerful than the one in [23].

has several advantages: it is physically simple to be laid out, it has only one stage (implying less power consumption, power attenuation), it also has RNB throughput up to 87% of an RNB design under the (λ, F, F') -request model [23] (this result was obtained through simulation for bursts arriving according to an exponential distribution).

A few natural questions arise from the kind of construction in Fig. 6. (Note again that, the design in the figure has one additional stage of wavelength converters as compared to the Ramamirtham–Turner design.) For instance, one might wonder if the second stage of FWCs yields an RNB design, or even SNB. If the second stage does not induce nonblockingness, would the blocking behavior be any better than the 87% obtained in [23]. We answer the former question here, leaving the latter for future works.

In this section, we analyze the design shown in Fig. 6 and another “natural” design shown in Fig. 7, which has one additional stage of $k \times k$ AWGRs and FWCs. The analyses motivates another interesting construction discussed in Section V-C. Some of the constructions we analyze in this section are blocking. Their analyses are included for pedagogical reasons, for one thing. They also are simple and symmetrical designs with few stages, which are important features when we physically realize them. Thus, even when they are blocking, if they have good probabilistic blocking behavior, then they can be practically useful. (The last issue is left for future works.)

Remark 5.1: In most analyses given in this section, we ignore the issue of divisibility for presentation clarity. In the figures, the black squares represent FWCs.

A. One-Stage Construction

We analyze the construction shown in Fig. 6 in this section. We ignore the trivial case when $f = 1$.

Theorem 5.2: The construction shown in Fig. 6 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 ($AWGR_0$) to the first multiplexor (M_0). Let $X = \{x_1, \dots, x_m\}$ be the outputs of $AWGR_0$ connected to these links. Then $m \leq k/f$.

Consider a request of the form $(\lambda_i, F_0, \lambda_j, F_0)$. In order to realize this request, λ_i has to be converted to some $\lambda_{i'}$ such that $((i' - i) \bmod k) \in X$. If $i' \neq j$, then a second-stage FWC needs to be used to convert $\lambda_{i'}$ to λ_j . From $AWGR_0$ to M_0 there are $\leq k/f$ FWCs. Hence, the existence of $> k/f$ requests $(\lambda_i, F_0, \lambda_j, F_0)$ which require the use of a second-stage FWC would imply that the construction is not RNB.

For each $i \in \mathbb{Z}_k$, define

$$C_i = \{j \in \mathbb{Z}_k \mid ((j - i) \bmod k) \notin X\}.$$

Then, for any $j \in C_i$, the request $(\lambda_i, F_0, \lambda_j, F_0)$ requires the use of a second-stage FWC. Note that $|C_i| \geq (f - 1)k/f$, $\forall i \in \mathbb{Z}_k$.

Let Y be any subset of size $k/f + 1$ of \mathbb{Z}_k . We claim that the set collection $\mathcal{C} = \{C_y \mid y \in Y\}$ has a system of distinct representatives by verifying Hall’s condition [9] that the union of any m members of \mathcal{C} has size at least m . The condition holds trivially when $m \leq 1$. When $m \geq 2$, since $C_y \neq C_{y'}$ whenever $y \neq y'$, and since $|C_y| \geq (f - 1)k/f$, the union of m members of \mathcal{C} has cardinality at least $(f - 1)k/f + 1$. Moreover, $(f - 1)k/f + 1 \geq k/f + 1 \geq m$. Hence, Hall’s condition holds for $m \geq 2$.




Fig. 7. Two-stage construction using FWCs that is RNB-2 but not SNB-1.

Now, suppose j_i represents C_i , $i \in Y$, in some particular system of distinct representatives of \mathcal{C} . Then, the set

$$\{(\lambda_i, F_0, \lambda_{j_i}, F_0) | i \in Y\}$$

is a request frame of size $k/f + 1$, in which each request requires a second-stage FWC to be realized. This completes the proof. \square

Theorem 5.3: The construction shown in Fig. 6 is not strictly nonblocking under the (λ, F, F') -request model, no matter how we connect the AWGRs to the output multiplexers.

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 Fig. 6, where there are two AWGRs which have the first k/f outputs connected to M_0 .

Theorem 5.4: Consider the construction shown in Fig. 6. Suppose there is an output multiplexor to which two AWGRs have their first k/f outputs connected. 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 AWGR₀ and AWGR₁ 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 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_j has to be converted to, in order to get out to M_0 , namely

$$W(\lambda_i) = \left\{ \lambda_j | 0 \leq ((j - i) \bmod k) \leq \frac{k}{2} - 1 \right\}.$$

Note that $|W(\lambda_i)| = k/f$, and, due to the circular shift nature of AWGRs,

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

(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

$$\frac{k}{2} - \frac{k}{f} > \frac{1}{2} \left(\frac{k}{2} + \frac{k}{f} - 1 \right)$$

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

Remark 5.5: Much less restrictive forms of the previous theorem can also be shown, at the expense of clarity. Particularly, we do not need to require the first k/f outputs from the two AWGRs to be connected to the same multiplexor, any k/f outputs from them would be sufficient. It is an interesting open problem to find a connection pattern for $f < 6$ which could make the design RNB. Section V-C discusses the case when $f = 2$. It turns out that in this case we need only half as many converters in the second stage for the design to be RNB.

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 Fig. 7 and its variations.




Fig. 8. Two-fiber construction that is RNB-1 using less than two stages of FWCs.

Lemma 5.6: Let f, k be positive integers where f is divisible by k . Let $G = (U \cup V; E)$ be a k -regular bipartite multi-graph, where (U, V) is the vertex bipartition, 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' . Kőnig's theorem [13] ensures that G' is f -edge colorable. (Again, coloring algorithms can be found in [3], [6], and [7].) This coloring of G' induces a desired coloring of G . \square

Theorem 5.7: The construction shown in Fig. 7 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 5.6. Let \mathbb{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 5.8: The details of the above proof can be done in a similar fashion as in Theorem 3.4. The proof implicitly contains a routing algorithm.

Theorem 5.9: The construction shown in Fig. 7 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. \square

C. 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 (the pigeonhole principle based routing idea) to be explored further.

The construction is shown in Fig. 8. It is self-explanatory. All wavelength converters are FWCs.

Theorem 5.10: The construction for $f = 2$ shown in Fig. 8 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 choices 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

TABLE I
TABULATED COMPARISONS OF DIFFERENT CONSTRUCTIONS

	SNB-1	RNB-1	SNB-2	RNB-2	#FWCs	#LWCs	#AWGRs	#WSCs	#OADMs
WI-Beneš [27]	-	Yes	-	-	kf	-	-	-	$f \lg f$
WI-Cantor [27]	-	Yes	-	-	kf	-	-	-	$f(\lg f)^2$
CBC [27]	Yes	Yes	-	?	$kf \lg f$	-	-	-	$f(\lg f)^2$
Ref. [23]	-	-	-	-	fk	-	$f \otimes k$	-	-
Ref. [24,25]	Yes	Yes	Yes	Yes	$(2f-1)k$	-	-	2 of dim. $f \times (2f-1)$	-
Section 3	-	Yes	-	-	-	$2fk$	$fb \otimes n$ $1 \otimes k$	-	-
Section 4.1	-	Yes	-	Yes	-	$3fk$	$fb \otimes n$ $n \otimes fb$	-	-
Section 4.2	Yes	Yes	Yes	Yes	-	$5fk-2fb$	$fb \otimes n$ $2n-1 \otimes fb$	-	-
Section 5.1	-	-	-	-	$2fk$	-	$f \otimes k$	-	-
Section 5.2	-	Yes	-	Yes	$3fk$	-	$2f \otimes k$	-	-
Section 5.3 ($f = 2$)	-	Yes	-	-	$3k$	-	$2 \otimes k$	-	-

The second half of the table contains our constructions, referred to by the section numbers in which 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.

Remark 5.11: 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.* [27] 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 (WSCs) or optical add-drop multiplexors (OADMs), and a special component called the 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 [23] 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, and uses $fk \times k$ AWGRs and fk FWCs.

As compared to the constructions above, our 2-stage RNB construction in Section III uses $2fk$ LWCs, $fbn \times n$ 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 [24], [25] described a strictly nonblocking construction using the so-called WDM split cross-connects, which consists of two WSCs and 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 $fbn \times n$ -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 $fbn \times n$ -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 “-” mean 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 [23] 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 [27].

We have not considered a nonblocking degree called wide-sense nonblocking [1] (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 referred to [10] and [12] 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. Initial works along this direction can be seen in [19] and [20].

ACKNOWLEDGMENT

The authors would like to thank the anonymous reviewers for many constructive comments which have helped improve the presentation of this paper.

REFERENCES

- [1] V. E. Beneš, *Mathematical Theory of Connecting Networks and Telephone Traffic*. New York: Academic Press, 1965, vol. 17, Mathematics in Science and Engineering.
- [2] C. Clos, "A study of nonblocking switching networks," *Bell Syst. Tech. J.*, vol. 32, pp. 406–424, 1953.
- [3] R. Cole and J. Hopcroft, "On edge coloring bipartite graphs," *SIAM J. Comput.*, vol. 11, no. 3, pp. 540–546, 1982.
- [4] C. Dragone, "An $N \times N$ optical multiplexor using a planar arrangement of two star couplers," *IEEE Photon. Technol. Lett.*, vol. 3, no. 9, pp. 812–815, Sep. 1991.
- [5] S. J. B. Yoo *et al.*, "High-performance optical-label switching packet routers and smart edge routers for the next generation internet," *IEEE J. Select. Areas Commun.*, vol. 21, no. 7, pp. 1041–1051, Sep. 2003.
- [6] H. N. Gabow, "Using Euler partitions to edge color bipartite multigraphs," *Int. J. Comput. Inf. Sci.*, vol. 5, no. 4, pp. 345–355, 1976.
- [7] H. N. Gabow and O. Kariv, "Algorithms for edge coloring bipartite graphs and multigraphs," *SIAM J. Comput.*, vol. 11, no. 1, pp. 117–129, 1982.
- [8] H. J. Chao, K. L. Deng, and Z. Jing, "A petabit photonic packet switch (P^3S)," in *Proc. IEEE INFOCOM*, San Francisco, CA, 2003, pp. 775–785.


- [9] P. Hall, "On representatives of subsets," *J. London Math. Soc.*, vol. 10, pp. 26–30, 1936.
- [10] P. E. Haxell, A. Rasala, G. T. Wilfong, and P. Winkler, "Wide-sense nonblocking WDM cross-connects," in *Algorithms—ESA '03 (Rome, Italy)*. Berlin, Germany: Springer, 2002, vol. 2461, Lecture Notes in Comput. Sci., pp. 538–549.
- [11] D. K. Hunter, M. H. M. Nizam, M. C. Cia, I. Andonovic, K. M. Guild, A. Tzanakaki, M. J. O'Mahony, L. D. Bainbridge, M. F. C. Stephens, R. V. Penty, and I. H. White, "WASPNET: a wavelength switched packet network," *IEEE Commun. Mag.*, vol. 37, no. 3, pp. 120–129, Mar. 1999.
- [12] F. K. Hwang, *The Mathematical Theory of Nonblocking Switching Networks*. River Edge, NJ: World Scientific, 1998.
- [13] D. König, "Über graphen und ihre anwendung auf determinantentheorie und mengenlehre," *Math. Ann.*, vol. 77, pp. 453–465, 1916.
- [14] L. Lin, E. Goldstein, L. Lunardi, and R. Tkach, "Optical crossconnects for high-capacity lightwave networks," *J. High Speed Netw.*, vol. 8, no. 1, pp. 17–34, 1999.
- [15] Lucent Technologies Press Release. (2001) Lucent Technologies Unveils Untra-High-Capacity Optical System; Time Warner Telecom First to Announce it Will Deploy the System. [Online]. Available: <http://www.lucent.com/press/0101/010117.nsa.html>
- [16] Lucent Technologies Press Release. (2002) Lucent Technologies Engineer and Scientists Set New Fiber Optic Transmission Record. [Online]. Available: <http://www.lucent.com/press/0302/020322.bla.html>
- [17] Lucent Technologies Website. (2002) What is Dense Wave Division Multiplexing (DWDM). [Online]. Available: <http://www.bell-labs.com/technology/lightwave/dwdm.html>
- [18] M. Maier, *Metropolitan Area WDM Networks: An AWG-Based Approach*. New York: Springer, 2003.
- [19] H. Q. Ngo, "Multiwavelength distribution networks," in *Proc. IEEE Workshop on High Performance Switching and Routing (HPSR 2004)*, Phoenix, AZ, 2004, pp. 186–190.
- [20] H. Q. Ngo, "WDM switching networks and [w,f]-connectors," *SIAM J. Comput.*, submitted for publication.
- [21] H. Q. Ngo, D. Pan, and C. Qiao, "Nonblocking WDM switches based on arrayed waveguide grating and limited wavelength conversion," in *Proc. IEEE INFOCOM*, Hong Kong, 2004, pp. 1352–1362.
- [22] C. Qiao and M. Yoo, "Optical burst switching (OBS)—a new paradigm for an optical internet," *J. High Speed Netw.*, vol. 8, no. 1, pp. 69–84, 1999.
- [23] J. Ramamirtham and J. S. Turner, "Design of wavelength converting switches for optical burst switching," in *Proc. IEEE INFOCOM*, vol. 2, 2002, pp. 1162–1171.
- [24] A. Rasala and G. Wilfong, "Strictly nonblocking WDM cross-connects," in *Proc. 11th Annu. ACM-SIAM Symp. Discrete Algorithms (SODA'2000)*, San Francisco, CA, 2000, pp. 606–615.
- [25] A. Rasala and G. Wilfong, "Strictly nonblocking WDM cross-connects for heterogeneous networks," in *Proc. 32nd Annu. ACM Symp. Theory of Computing (STOC'2000)*, Portland, OR, 2000, pp. 513–524.
- [26] D. B. West, *Introduction to Graph Theory*. Upper Saddle River, NJ: Prentice-Hall, 1996.
- [27] G. Wilfong, B. Mikkelsen, C. Doerr, and M. Zirngibl, "WDM cross-connect architectures with reduced complexity," *J. Lightw. Technol.*, vol. 17, no. 10, pp. 1732–1741, Oct. 1999.
- [28] J. Xu, C. Qiao, J. Li, and G. Xu, "Efficient channel scheduling algorithms in optical burst switching networks," in *Proc. IEEE INFOCOM*, San Francisco, CA, Apr. 2003.



Hung Q. Ngo (M'04) received the B.S. degree from Ho Chi Minh City University of Technology, Viet Nam, and the M.S. degree in mathematics and the Ph.D. degree in computer science from the University of Minnesota, Minneapolis.


He is currently an Assistant Professor in the Department of Computer Science and Engineering, State University of New York at Buffalo. His main research interests include algorithms and networking.

Dr. Ngo has been a member of the Association for Computing Machinery (ACM) since 1998.



Dazhen Pan received the B.S. degree in computer science from Nankai University, China, in 1999. He is currently working toward the Ph.D. degree in computer science and engineering at the State University of New York at Buffalo.

His research is on complexity, constructions, and routing algorithms of optical switching networks.



Chunming Qiao (S'89–M'92) is a full Professor at the State University of New York at Buffalo, where he directs the Lab for Advanced Network Design, Analysis, and Research (LANDER), which conducts cutting-edge research work on optical networks, wireless networks, survivable networks, and TCP/IP technologies. He has over 12 years of academic and industrial experience in optical networks. He has published more than 100 papers in leading technical journals and conference proceedings, and is recognized for his pioneering research on optical internet and,

in particular, the optical burst switching (OBS) paradigm. His work on integrated cellular and ad hoc networking systems (iCAR) is also internationally acclaimed.

Dr. Qiao was the IEEE Communication Society's Editor-at-Large for optical networking and computing until recently, when he became an Area Editor of *IEEE Communications Magazine* for the new Optical Communications Supplement. He is also an editor of several other journals and magazines, including IEEE/ACM TRANSACTIONS ON NETWORKING, and has been a guest editor for the IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS and other publications. He has chaired and co-chaired many conferences and workshops on optical communications and networking, including the Opticomm 2002. He is also the founder and chair of the Technical Group on Optical Networks sponsored by SPIE, and a Vice Chair of the IEEE Technical Committee on Gigabit Networking.