Optical Switching Networks with Minimum Number of Limited Range Wavelength Converters

Hung Q. Ngo
Computer Science and Engineering,
State University of New York at Buffalo

Dazhen Pan
Computer Science and Engineering,
State University of New York at Buffalo

Yuanyuan Yang
Electrical and Computer Engineering,
State University of New York at Stony Brook

***Abstract*— We study the problem of determining the minimum number of limited range wavelength converters needed to construct strictly, wide-sense, and rearrangeably nonblocking optical cross-connects for both unicast and multicast traffic patterns. We give the exact formula to compute this number for rearrangeably and wide-sense nonblocking cross-connects under both the unicast and multicast cases. We also give optimal cross-connect constructions with respect to the number of limited-range wavelength converters.**

Keywords: Wavelength-division-multiplexing (WDM), optical switching networks, cross-connects, limited range wavelength conversion.

I. Introduction

Wavelength division multiplexing (WDM) is a key technique to exploit the huge bandwidth of optics. As the number of wavelengths in a WDM network increases to hundreds per fiber and each wavelength operates at the rate of 10Gbps (OC-192) or higher [1], optical communication has become a promising networking choice to meet ever-increasing demands on bandwidth from emerging bandwidth-intensive networking and computing applications, such as data browsing in the world wide web, multimedia conferencing, e-commerce, and video on demand services. The next generation of the Internet is expected to employ WDM-based optical backbones [2].

A *WDM optical switching network* (also referred to as *WDM cross-connect* or WXC) provides interconnections between a group of input fiber links and a group of output fiber links with each fiber link carrying multiple wavelength channels. It not only can provide many more connections than a traditional electronic switching network, but also can offer much richer communication patterns (e.g. unicast and multicast communication patterns under different connection models [3, 4]). Such an optical switching network can be used to serve as an optical cross-connect (OXC) in a wide-area communication network or to provide high-speed interconnections among a group of processors in a parallel and distributed computing system.

Dr. Hung Q. Ngo is supported in part by NSF CAREER Award CCF-0347565. Dr. Yuanyuan Yang is supported in part by NSF grants CCR-0073085 and CCR-0207999.

In this paper, we consider supporting two typical communication patterns, *unicast* (or *permutation*) and *multicast* in a WDM switching network. A unicast communication pattern is a one-to-one mapping between input wavelengths and output wavelengths of a WDM switching network, while a multicast communication pattern is a one-to-many mapping between them.

We will also consider WDM switching networks with different nonblocking capabilities, such as *strictly nonblocking* (SNB), *wide-sense nonblocking* (WSNB), and *rearrangeably nonblocking* (RNB). In an SNB network, any legal connection request can be arbitrarily realized without any disturbance to existing connections. In a WSNB network, a proper routing strategy must be adopted in realizing connection requests to guarantee nonblockingness. In an RNB network, any legal connection request can be realized by permitting the rearrangement of on-going connections in the network.

The major challenge in designing WDM optical switching networks is how to provide maximum connectivity at high speed while keeping minimum hardware cost. To meet the challenge, it is required to keep data in optical domain all the way from its source to destination. One reason is that optical switching is much faster than electrical switching. For example, Lucent's all-optical switch LambdaRouter [5] can transmit at 10 trillion bits per second, while today's fastest electrical switches can reach only about 160 billion bits per second. Furthermore, all-optical switching eliminates the need for costly conversions between optical and electronic signals (so-called O/E/O conversions). As the optical technology matures, photonic switching systems not only can potentially achieve higher throughput [6], but also can be more cost-effective than their electronic counterparts even for applications requiring a lower throughout.

In all-optical switching, either the wavelength on which the data is sent and received must be the same, or all-optical wavelength converters are needed to convert the signals on one input wavelength to another output wavelength. Since WDM switching networks with no wavelength conversion cannot provide full WDM connectivity [2, 3], wavelength conversion should be included in the design of all-optical WDM switching networks. Thus, the overall hardware cost of a WDM optical switching network includes both the cost of switching elements and the cost of wavelength conversion. Some previous works [3, 7] have aimed at minimizing the number of switching cross-points and the number of wavelength converters. However, since wavelength converters are still expensive, how to further reduce the cost of wavelength conversion is a critical issue in designing WDM switching networks.

Thus far, researchers have considered two approaches for further reducing the cost of wavelength conversion. One approach is to adopt limited range wavelength converters (LWCs) [4, 8–12] instead of full range wavelength converters (FWCs). Clearly, LWCs are less expensive but less powerful than FWCs. A challenging task is to design WDM switching networks with full connection capability by using the less powerful LWCs. The recent designs of WDM switching networks in [4, 12] have achieved this design goal. One of the major differences between them is that the switching elements used in [4] are arrayed waveguide grating routers (AWGR); while the ones in [12] adopted SOA-based or MEMS-based switching elements. Since an AWGR switch is wavelength sensitive, the SOA-based or MEMS-based WDM switching networks have lower wavelength conversion cost than the AWGR-based ones. On the other hand, the advatange of AWGR-based designs is that they consume very little power. Also notice that both designs in [4, 12] used wavelength converters dedicated to each wavelength in the switching network.

Another approach to reducing the cost of wavelength conversion is to share wavelength converters among the fiber links of the switching network instead of dedicated to each link [8, 11, 13]. There are two architectures proposed for switching networks to share converters [8]. In share-per-node structure, all converters are collected in a single converter bank, and shared among all input/output wavelength pairs of the switching network. In share-per-fiber-link structure, each output fiber link is provided with a dedicated converter bank so that the wavelength converters in the converter bank are shared among input/output wavelength pairs related to this output fiber link. However, so far there has been no work on how many wavelength converters are necessary and sufficient for WDM switching networks with full unicast and multicast connection capabilities.

In this paper, by combining the two aforementioned approaches, we consider WDM switching networks with limited range wavelength converters shared among network input/output wavelengths. We will study the problem of determining the minimum number of LWCs needed to construct unicast and multicast switching networks under various nonblocking conditions and traffic patterns. As the results, we will give several optimal and near-optimal WDM switching network constructions with respect to the number of limited range wavelength converters.

The rest of the paper is organized as follows. Section II formulates our problem, and defines concepts and terminologies used throughout the paper. Sections III and IV address the problem for unicast switches. Section V addresses the problem for multicast switches. Section VI presents constructions of WDM optical switching networks that make use of the minimum number of limited-range wavelength converters as determined in previous sections. Lastly, Section VII concludes the paper and discusses some future research directions coming out of this work.

II. PRELIMINARIES

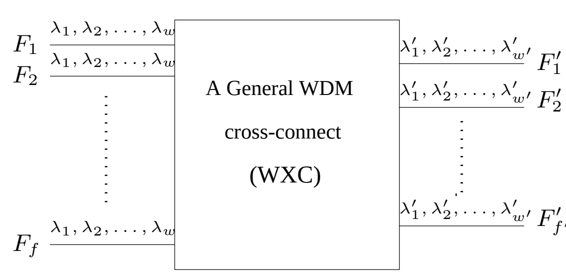

Fig. 1. Heterogeneous WDM Cross-Connect

A general WDM cross-connect (WXC) consists of f input fibers with w wavelengths on each fiber, and f' output fibers with w' wavelengths on each fiber, where $fw = f'w'$ (see Figure 1). The set of input wavelengths need not have any relation with the output counterpart. This kind of WXCs are referred to as the "heterogeneous WXCs," which are needed to optically switch data from different manufacturers [14].

In this paper, we consider the homogeneous situation where each WXC[1] has f input fibers and f output fibers, each of which can carry a set $\Lambda = \{\lambda_1, \dots, \lambda_w\}$ of w wavelengths. (Note that most of the ideas here apply to the heterogeneous case too. Let $\mathcal{F} = \{F_1, \dots, F_f\}$ and $\mathcal{F}' = \{F'_1, \dots, F'_f\}$ denote the set of input and output fibers, respectively.

In the (λ, F, F')-*request model* [4], 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.

The (λ, F, F')-request 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 other model is particularly useful for asynchronous switching of optical bursts using JET (just enough time) and void filling [4,15], as well as certain circuit-switching applications requiring specific QoS.

In this paper, we are concerned with the $(\lambda, F, \lambda', F')$-model, leaving the (λ, F, F')-model for a future work. We next briefly define the concepts of strictly nonblocking (SNB), wide-sense nonblocking (WSNB) and rearrangeably nonblocking (RNB) for both request models.

Consider a WXC with a few unicast connections already established. Under the $(\lambda, F, \lambda', F')$-model, a new unicast request $(\lambda, F, \lambda', F')$ is *valid* if λ is free in F and λ' is free in F'. A *request frame* is a set of requests such that no two requests are from the same input wavelength in the same input fiber, and no two requests are to the same output wavelength in the same output fiber. 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.

All the concepts defined above have their multicast counterparts. A *multicast request* is of the form $(\lambda, F, \mathcal{P})$ where $\lambda \in \Lambda$, $F \in \mathcal{F}$, and $\mathcal{P} \subseteq \Lambda \times \mathcal{F}'$ such that no $F' \in \mathcal{F}'$ appears more than once in $\mathcal{P}$. (That is, if (a, b) and (c, d) are different pairs in $\mathcal{P}$, then $b \neq d$.) This restriction was made since in practical networks it is not necessary to have a multicast connection going to the same output fiber on two different wavelengths [7,16]. A *multicast request frame* is a set of multicast requests such that no two different requests are from the same input wavelength in the same input fiber, and that each wavelength on an output fiber is requested at most once. (We shall be more rigorous in later sections.) The concepts of SNB, RNB, WSNB for the multicast case are defined similarly.

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

For $1 \leq d \leq w-1$, let LWC(d) denote a limited wavelength converter of degree d, namely a LWC(d) can convert λ_i to any λ_j where $|i-j| \leq d$. Note that this conversion cannot "loop around," i.e. λ_1 cannot be converted to λ_w when $d \leq w-2$. When $d = w-1$, we get a full wavelength converter.

We will address two questions: (a) under a certain request model, and certain traffic pattern (unicast, multicast), how many LWC(d) are necessary?; and (b) how do we construct an WXC with as small a number of LWC(d) as possible (sufficiency)?.

Throughout the paper, for any positive integer n, let $[n] := \{1, \dots, n\}$, and let S_n denote the set of all permutations on $[n]$. The graph theoretic terminologies and notations we adopt in this paper are standard (see, e.g., [17]).

The following lemmas are going to be used throughout the paper. We omit the proofs due to space limitation and their simplicities.

Lemma II.1. *For any $a_1, a_2, a_3 \in [w]$, we have*

$$\left\lceil \frac{|a_1 - a_3|}{d} \right\rceil \leq \left\lceil \frac{|a_1 - a_2|}{d} \right\rceil + \left\lceil \frac{|a_2 - a_3|}{d} \right\rceil$$

Lemma II.2. *For any $a_1, a_2, a_3, a_4 \in [w]$ where both a_1 and a_2 are strictly less than a_3 and a_4, we have*

$$\left\lceil \frac{|a_1 - a_2|}{d} \right\rceil + \left\lceil \frac{|a_3 - a_4|}{d} \right\rceil \leq \left\lceil \frac{|a_1 - a_4|}{d} \right\rceil + \left\lceil \frac{|a_2 - a_3|}{d} \right\rceil$$

Lemma II.3. *For any $a_1, a_2, a_3, a_4 \in [w]$ where $a_1 \leq a_2 \leq a_3 \leq a_4$, we have*

$$\left\lceil \frac{|a_1 - a_2|}{d} \right\rceil + \left\lceil \frac{|a_3 - a_4|}{d} \right\rceil \leq \left\lceil \frac{|a_1 - a_3|}{d} \right\rceil + \left\lceil \frac{|a_2 - a_4|}{d} \right\rceil$$

III. Unicast WDM cross-connects

Consider any set $\mathcal{D}$ of requests under the $(\lambda, F, \lambda', F')$-request model. Each request D in $\mathcal{D}$ is of the form $(\lambda_i, F_j, \lambda_{i'}, F'_{j'})$, where $1 \leq i, i' \leq w$, and $1 \leq j, j' \leq f$. A set $\mathcal{D}$ of requests is called a *request frame* if

$$\begin{aligned} |\{(i,j) : (\lambda_i, F_j, \lambda_{i'}, F'_{j'}) \in \mathcal{D}\}| &\leq 1, \ \forall i', j', \\ |\{(i',j') : (\lambda_i, F_j, \lambda_{i'}, F'_{j'}) \in \mathcal{D}\}| &\leq 1, \ \forall i, j. \end{aligned}$$

If equality holds for both of the above inequalities, we call the request frame a *full request frame*. In other words, $\mathcal{D}$ is a request frame if there is at most one request in $\mathcal{D}$ from a wavelength in a particular input fiber, and there is at most one request to a wavelength in a particular output fiber. In a full request frame, replace "at most one" by "exactly one" in the previous sentence.

For any request $D \in \mathcal{D}$, let $i(D)$ and $i'(D)$ denote the indices of the input and output wavelengths of D, and $j(D)$ and $j'(D)$ denote the indices of the input and output fibers of D. In other words, D can be written as $D = \left(\lambda_{i(D)}, F_{j(D)}, \lambda_{i'(D)}, F'_{j'(D)}\right)$.

In order to convert λ_i to $\lambda_{i'}$, we need at least $\left\lceil \frac{|i-i'|}{d} \right\rceil$ LWC(d). This observation leads to the following theorem.

Theorem III.1. *A strictly, wide-sense, or rearrangeably nonblocking WXC under the* $(\lambda, F, \lambda', F')$*-request model needs at least*

$$m_1(w, f, d) := \max_{\mathcal{D}} \sum_{D \in \mathcal{D}} \left\lceil \frac{|i(D) - i'(D)|}{d} \right\rceil \quad (1)$$

LWC(d)*, where the* max *goes over all request frames* $\mathcal{D}$*. Moreover,* $m_1(w, f, d)$ *is also the sufficient number of LWC*(d) *to construct a rearrangeably or a wide-sense non-blocking WXC.*

Proof. The fact that $m_1(w, f, d)$ is necessary is obvious from the observation above. Sufficiency is shown with the constructions shown in Section VI. □

Theorem III.1 gave a characterization of $m_1(w, f, d)$, the necessary and sufficient number of LWC(d) for a rearrangeably and wide-sense non-blocking WXC. However, maximizing a function over all possible request frames (an exponential number) does not give us a good handle on $m_1(w, f, d)$. The following lemma characterizes $m_1(w, f, d)$ in a much better fashion.

Lemma III.2.

$$m_1(w, f, d) = f \max_{\pi \in S_w} \sum_{i=1}^{w} \left\lceil \frac{|i - \pi(i)|}{d} \right\rceil \quad (2)$$

Proof. For any request frame $\mathcal{D}$, define its "cost" by

$$c(\mathcal{D}) = \sum_{D \in \mathcal{D}} \left\lceil \frac{|i(D) - i'(D)|}{d} \right\rceil.$$

Then, $m_1(w, f, d) = \max_{\mathcal{D}} c(\mathcal{D})$ by Theorem III.1. Consider the request frame

$$\mathcal{D}_\sigma = \{(\lambda_i, F_j, \lambda_{\sigma(i)}, F'_j) \mid i \in [w], j \in [f]\},$$

$$\text{where } \sigma = \arg\max_{\pi \in S_w} \left\{ \sum_{i=1}^{w} \left\lceil \frac{|i - \pi(i)|}{d} \right\rceil \right\}.$$

Then, it is straightforward that

$$m_1(w, f, d) \geq c(\mathcal{D}_\sigma) = f \max_{\pi \in S_w} \sum_{i=1}^{w} \left\lceil \frac{|i - \pi(i)|}{d} \right\rceil.$$

Next, we show that $m_1(w, f, d)$ is at most the right hand side. Let $\mathcal{D}$ be a request frame which maximizes $c(\mathcal{D})$ so that $m_1(w, f, d) = c(\mathcal{D})$. If $\mathcal{D}$ is not a full request frame, we can add more requests into $\mathcal{D}$ to make it full without reducing the cost $c(\mathcal{D})$. Consequently, without loss of generality we can assume that $\mathcal{D}$ is a full request frame.

To this end, construct a bipartite multi-graph $G = (X \cup Y; E)$, where X and Y are the color classes, $X = \{x_1, \dots, x_w\}$, $Y = \{y_1, \dots, y_w\}$, and there is a copy of an edge $(x_s, y_t) \in E$ for every request in $\mathcal{D}$ of the form $(\lambda_s, F_j, \lambda_t, F'_{j'})$, for any j, j'. This way, every edge of G represents one request in $\mathcal{D}$.

Because $\mathcal{D}$ is a full request frame, G is f-regular. By König's theorem [18], G is f-edge-colorable. Let $[f]$ be the set of colors. Let M_j be the set of edges with color j, then M_j is a perfect matching of G. Each perfect matching of G corresponds naturally to a permutation $\pi \in S_w$, where $\pi(s) = t$ for each edge (x_s, y_t) in the perfect matching. Let π_j be the permutation corresponding to M_j. We have

$$\begin{aligned} m_1(w, f, d) &= \sum_{D \in \mathcal{D}} \left\lceil \frac{|i(D) - i'(D)|}{d} \right\rceil \\ &= \sum_{(x_s, y_t) \in E} \left\lceil \frac{|s - t|}{d} \right\rceil = \sum_{j=1}^{f} \sum_{(x_s, y_t) \in M_j} \left\lceil \frac{|s - t|}{d} \right\rceil \\ &= \sum_{j=1}^{f} \sum_{i=1}^{w} \left\lceil \frac{|i - \pi_j(i)|}{d} \right\rceil \leq f \max_{\pi \in S_w} \sum_{i=1}^{w} \left\lceil \frac{|i - \pi(i)|}{d} \right\rceil. \end{aligned}$$

□

IV. An explicit formula for $m_1(w, f, d)$

From the result of Lemma III.2, it is easy to see that the function $m_1(w, f, d)$ can be computed in polynomial time with a standard maximum weighted matching algorithm (such as the algorithm in [19]). However, an explicit formula for $m_1(\cdot)$ would obviously be much more desirable.

For convenience, define a function $g : S_w \to \mathbb{R}$ by

$$g(\pi) = \sum_{i=1}^{w} \left\lceil \frac{|i - \pi(i)|}{d} \right\rceil. \quad (3)$$

We first consider the case when w is an even number, say $w/2 = k \in \mathbb{Z}$. Lemma III.2 basically says that

$$m_1(w, f, d) \geq f \cdot g(\pi), \quad \text{for any } \pi \in S_w,$$

and that the bound is best possible. Our objective is thus to find a $\pi \in S_w$ that maximizes $g(\pi)$. For any $\pi \in S_w$ and any subset $X \subseteq [w]$, let $\pi(X) = \{\pi(i) \mid i \in X\}$.

The following lemma restricts the search space for π that maximizes $g(\pi)$.

Lemma IV.1. *Suppose $w/2 = k \in \mathbb{Z}$. Let $A = [k]$ and $B = [w] - [k]$. The function $g(\pi)$ is maximized at some permutation $\pi \in S_w$ where $\pi(A) = B$ and $\pi(B) = A$.*

Proof. Consider a permutation $\pi \in S_w$ that maximizes $g(\pi)$. Obviously, if $\pi(A) = B$, then $\pi(B) = A$, and vice versa. Suppose there is some $a \in A$ such that $\pi(a) \in A$, then there is some $b \in B$ such that $\pi(b) \in B$. Note that a and $\pi(a)$ are both strictly less than b and $\pi(b)$. Hence, by Lemma II.2,

$$\left\lceil \frac{|a-\pi(a)|}{d} \right\rceil + \left\lceil \frac{|b-\pi(b)|}{d} \right\rceil \le \left\lceil \frac{|a-\pi(b)|}{d} \right\rceil + \left\lceil \frac{|b-\pi(a)|}{d} \right\rceil.$$

Consider the permutation π' defined by $\pi'(i) = \pi(i)$ for all $i \in [w] - \{a, b\}$, $\pi'(a) = \pi(b)$, and $\pi'(b) = \pi(a)$. Then, $g(\pi') \ge g(\pi)$. Basically, π' is obtained from π by exchanging the images of a and b under π, keeping the rest the same. We can keep exchanging images of these kinds of pairs a, b to finally get a permutation π with $\pi(A) = B$ (and thus $\pi(B) = A$) while not reducing the cost $g(\pi)$. $\square$

To this end, we can restrict our attention to permutations $\pi \in S_w$ consisting of two parts: $\pi|_A : A \to B$ and $\pi|_B : B \to A$, where $\pi|_X$ is the mapping obtained by restricting π to the subset X of $[w]$. (Note that $\pi|_A$ and $\pi|_B$ are one-to-one correspondences.) Let $S_{A,B}$ denote the set of all one-to-one correspondences between A and B. Then, due to symmetry,

$$\max_{\pi \in S_w} g(\pi) = 2 \max_{\sigma \in S_{A,B}} \sum_{a \in A} \left\lceil \frac{\sigma(a) - a}{d} \right\rceil. \qquad (4)$$

Consequently, to find an optimal π we can just find an optimal $\sigma \in S_{A,B}$ with respect to the right hand side of (4), then let $\pi|_A = \sigma$ and $\pi|_B = \sigma^{-1}$.

For any number $i \in [w]$, let

$$q_i = \left\lfloor \frac{i}{d} \right\rfloor, \quad r_i = i \bmod d = i - d \left\lfloor \frac{i}{d} \right\rfloor. \qquad (5)$$

Hence, $0 \le r_i \le d-1, \forall i \in [w]$. Moreover, for any $a \in A$ and $b \in B$,

$$\left\lceil \frac{b-a}{d} \right\rceil = \begin{cases} q_b - q_a & \text{if } r_b \le r_a \\ q_b - q_a + 1 & \text{if } r_b > r_a \end{cases}$$

Consequently, for any $\sigma \in S_{A,B}$, we have

$$\begin{aligned} & \sum_{a \in A} \left\lceil \frac{\sigma(a) - a}{d} \right\rceil \\ = \; & \sum_{a \in A} (q_{\sigma(a)} - q_a) + \sum_{a \in A} \left\lceil \frac{r_{\sigma(a)} - r_a}{d} \right\rceil \\ = \; & \left(\sum_{b \in B} q_b - \sum_{a \in A} q_a \right) + |\{a \in A : r_{\sigma(a)} - r_a > 0\}|. \end{aligned}$$

The first term is independent of σ. Thus, we can just concentrate on finding a $\sigma \in S_{A,B}$ which maximizes

$$t(\sigma) = |\{a \in A : r_{\sigma(a)} - r_a > 0\}|. \qquad (6)$$

Then, recalling relations (2) and (4), we conclude with a simpler characterization of $m_1(w, f, d)$ in the following lemma.

Lemma IV.2. *When $w/2 \in \mathbb{Z}^+$, we have*

$$m_1(w, f, d) = 2f \left(\sum_{b \in B} q_b - \sum_{a \in A} q_a \right) + 2f \max_{\sigma \in S_{A,B}} t(\sigma) \qquad (7)$$

We now obtain a simple consequence of this lemma.

Theorem IV.3. *$m_1(w, f, 1) = \frac{1}{2} f w^2$ when w is even.*

Proof. When $d = 1$, $r_i = 0, \forall i \in [w]$, hence $t(\sigma) = 0, \forall \sigma \in S_{A,B}$. Moreover, $q_i = i, \forall i \in [w]$. Consequently, (7) gives

$$m_1(w, f, 1) = 2f \left(\sum_{b \in B} b - \sum_{a \in A} a \right) = \frac{1}{2} f w^2.$$

$\square$

When $d \ge 2$, the situation is not as simple. Let R_A be the multiset of all $r_a, a \in A$, and R_B be the multiset of all $r_b, b \in B$. For any integer r, let $\mu_A(r)$ and $\mu_B(r)$ denote the multiplicities of r in R_A and R_B, respectively. For example, for $w = 8, d = 3$, we have

$$\begin{aligned} A &= \{1,2,3,4\} & B &= \{5,6,7,8\} \\ R_A &= \{1,2,0,1\} & R_B &= \{2,0,1,2\} \\ \mu_A(0) = 1, \mu_A(1) &= 2 & \mu_B(0) &= 1, \mu_B(1) = 1, \\ \mu_A(2) &= 1 & \mu_B(2) &= 2. \end{aligned}$$

First, we make a simple yet important observation in the following lemma. This idea is the key to finding an explicit formula for $m_1(w, f, d)$!

Lemma IV.4. *Suppose w is even and $d \geq 2$. Then, for any $\sigma \in S_{A,B}$, we have*

$$t(\sigma) \leq \frac{w}{2} - \mu_B(0) - \max\{0, \mu_B(1) - \mu_A(0)\}. \quad (8)$$

Proof. By definition, $t(\sigma)$ is the number of pairs $(r_{\sigma(a)}, r_a)$ for which $r_{\sigma(a)} > r_a$, where a goes over all elements of A. There are exactly $w/2$ of these pairs. The numbers $r_{\sigma(a)}$ are picked one by one from the multiset R_B. Thus, for any σ there are $\mu_B(0)$ of the $r_{\sigma(a)}$ that are equal to 0. When $r_{\sigma(a)} = 0$, $r_{\sigma(a)} \leq r_a$ and thus the pair will not be counted toward $t(\sigma)$. This explains the term $-\mu_B(0)$ in the right hand side of (8).

Similarly, there are $\mu_B(1)$ of the $r_{\sigma(a)}$ that are equal to 1. When $r_{\sigma(a)} = 1$, $r_{\sigma(a)} \leq r_a$ unless $r_a = 0$, and there are only $\mu_A(0)$ of the r_a which are 0. Thus, the best we can do is to pair up as many 1's in R_B with 0's in R_A as possible. If $\mu_B(1) - \mu_A(0) \leq 0$, then we have enough 0's in R_A to pair up with 1's in R_B. When $\mu_B(1) - \mu_A(0) > 0$, there must be at least $\mu_B(1) - \mu_A(0)$ 1's left in R_B, and the corresponding pairs will not be counted toward $t(\sigma)$. This explains the term $-\max\{0, \mu_B(1) - \mu_A(0)\}$. □

What is amazing is that inequality (8) is the best possible, as the following lemma shows. The lemma (and its proof) basically gives an explicit formula for $m_1(w, f, d)$.

Lemma IV.5. *For even w and $d \geq 2$, we have*

$$\max_{\sigma \in S_{A,B}} t(\sigma) = \frac{w}{2} - \mu_B(0) - \max\{0, \mu_B(1) - \mu_A(0)\}. \quad (9)$$

Proof. With the result of Lemma IV.4 in mind, we only need to find a σ such that

$$t(\sigma) = \frac{w}{2} - \mu_B(0) - \max\{0, \mu_B(1) - \mu_A(0)\}.$$

In fact, we only need to find a way to pair up elements of R_A and R_B in a one-to-one fashion, so that the number of pairs $(r_a, r_{\sigma(a)})$ with $r_a < r_{\sigma(a)}$ is equal to the right hand side of the above relation.

Recall the example $w = 8, d = 3$ above. The upper bound gives $t(\sigma) \leq 4 - 1 - \max\{0, 1 - 1\} = 3$. The following pairing achieves this bound:

$$\left[\begin{array}{c|cccc} r_a & 1 & 2 & 0 & 1 \\ \hline r_{\sigma(a)} & 2 & 0 & 1 & 2. \end{array}\right]$$

Implicitly, we get (at least) a σ that achieves the bound: $\sigma(i) = i + 4, 1 \leq i \leq 4$. The general case is not as simple, but the above example gives the main ingredients.

Let $k = w/2$. As usual, write $k = q_k d + r_k$. (Recall that we assumed $d \geq 2$.) Then,

$$R_A = \{1, 2, \dots, d-1, 0, \cdots, 1, \dots, d-1, 0, 1, \dots, r_k\},$$

where there are q_k groups of $\{1, \dots, d-1, 0\}$ and the last group is $\{1, \dots, r_k\}$. Thus,

$$\mu_A(r) = \begin{cases} q_k + 1 & 1 \leq r \leq r_k \\ q_k & r_k < r \leq d-1, \text{ or } r = 0. \end{cases} \quad (10)$$

One nice thing is to have $\mu_A(0) = q_k$, regardless of r_k. The numbers $\mu_B(r)$ are not so nice, broken up into a few cases. In general, R_B looks like

$$\{r_k+1, \dots, d-1, 0, \dots, d-1, \cdots, 0, ..., d-1, 0, \dots, 2r_k\},$$

in which there are $q_k - 1$ groups of $\{0, 1, \dots, d-1\}$ in the middle, a group $\{r_k + 1, \dots, d-1\}$ in the beginning, and the last group is $\{0, \dots, 2r_k\}$.

When $r_k = d - 1$, the first group is empty. When $r_k = 0$, the last group does not contain a 1. When $2r_k = d$, the last group contains two 0's and a 1. When $2r_k > d$, the last group contains two 0's and two 1's. We formally consider these cases as follows.

Case 1: $r_k = d - 1$.

When $d = 2$ we get $\mu_B(0) = q_k + 1$, and $\mu_B(1) = q_k$. The upper bound is $t(\sigma) \leq k - q_k - 1$, which can be achieved with the following σ:

$$\sigma(i) = i + k, \;\; 1 \leq i \leq k.$$

When $d > 2$ we get $\mu_B(0) = \mu_B(1) = q_k + 1$. The upper bound is $t(\sigma) \leq k - q_k - 2$, which can be achieved with the following σ:

$$\sigma(i) = \begin{cases} i + k + 2 & 1 \leq i \leq k - 2 \\ i + 2 & k - 1 \leq i \leq k. \end{cases}$$

Case 2: $2r_k < d$. We have $\mu_B(0) = \mu_B(1) = q_k$. The upper bound is $t(\sigma) \leq k - q_k$, which is certainly obtainable with the following σ:

$$\sigma(i) = \begin{cases} k + 1 + i & 1 \leq i \leq k - 1 \\ k + 1 & i = k \end{cases}, \text{when } r_k = 0$$

$$\sigma(i) = \begin{cases} k + 1 - r_k + i & r_k \leq i \leq k \\ 2k + 1 - r_k + i & 1 \leq i \leq r_k - 1 \end{cases}, \text{when } r_k > 0$$

Case 3: $2r_k = d$. We have $\mu_B(0) = q_k + 1$, $\mu_B(1) = q_k$. The upper bound is $t(\sigma) \leq k - q_k - 1$, which is achievable with the following σ:

$$\sigma(i) = \begin{cases} k + 1 - r_k + i & r_k \leq i \leq k \\ 2k + 1 - r_k + i & 1 \leq i \leq r_k - 1. \end{cases}$$

Case 4: $2r_k > d$. We have $\mu_B(0) = \mu_B(1) = q_k + 1$. The upper bound is $t(\sigma) \le k - q_k - 2$, which is achievable with the following σ:

$$\sigma(i) = \begin{cases} k + 1 - r_k + i & r_k \le i \le k \\ 2k + d + 1 - 2r_k + i & 1 \le i \le 2r_k - d - 1 \\ 2k + d + 2 - 3r_k + i & 2r_k - d \le i \le r_k - 1. \end{cases}$$

□

The previous lemma gives explicit formulas for the expression $\max_{\sigma \in S_{A,B}} t(\sigma)$, which is either $k - q_k$, $k - q_k - 1$, or $k - q_k - 2$, depending on the relationships between d and k. Combining Lemma IV.5 with Lemma IV.2, noting the fact that $q_k = \left\lfloor \frac{w}{2d} \right\rfloor$ and $r_k = \frac{w}{2} - d \left\lfloor \frac{w}{2d} \right\rfloor$, we obtain the following theorem for the even-w case.

Theorem IV.6. *Suppose $\frac{w}{2} \in \mathbb{Z}$, and $2 \le d \le w - 1$. Let*

$$\begin{aligned} r &= \frac{w}{2} - d \left\lfloor \frac{w}{2d} \right\rfloor \\ \bar{m}_1 &= 2f \sum_{i=1}^{w/2} \left(\left\lfloor \frac{\frac{w}{2} + i}{d} \right\rfloor - \left\lfloor \frac{i}{d} \right\rfloor \right) + 2f \left(\frac{w}{2} - \left\lfloor \frac{w}{2d} \right\rfloor \right). \end{aligned}$$

Then, we have

(i) $m_1(w, f, d) = \bar{m}_1$ *when* $d > 2r$.
(ii) $m_1(w, f, d) = \bar{m}_1 - 2f$ *when* $d = 2r$.
(iii) $m_1(w, f, d) = \bar{m}_1 - 4f$ *when* $d - 1 = r > 1$ *or* $d < 2r$.

To this end, let us consider the case when w is odd. Suppose $w = 2k + 1$, where k is a positive integer. (The case when $w = 1$ is trivial, as no wavelength converter is needed.) Let

$$A = [k], B = [2k{+}1] - [k], A' = [k{+}1], B' = [2k{+}1] - [k{+}1].$$

Note that $[w] = [2k + 1] = A \cup B = A' \cup B'$. We follow roughly the same path as the even-w case.

Lemma IV.7. *Suppose $(w - 1)/2 = k \in \mathbb{Z}^+$. The function $g(\pi)$ is maximized at some permutation $\pi \in S_w$ where $\pi(A) = B'$ and $\pi(B) = A'$.*

Proof. Consider any permutation π that maximizes $g(\pi)$. Suppose $\pi(k + 1) = \bar{a}, \pi(\bar{b}) = k + 1$ for some $\bar{a}, \bar{b} \in [w]$.

Case 1: $\bar{a} = \bar{b} = k + 1$. If there is some $a \in A$ with $\pi(a) \in A'$, then there must be some $b \in B - \{k + 1\}$ with $\pi(b) \in B'$. Note that both a and $\pi(a)$ are strictly less than b and $\pi(b)$. Similar to the argument in Lemma IV.1, we can switch the images of a and b to get a new π while not reducing $g(\pi)$. Keep doing so until there is no more $a \in A$ with $\pi(a) \in A'$, we obtain a permutation π with the desired property.

Case 2: $\bar{a} < k + 1$ and $\bar{b} > k + 1$. If there is some $a \in A$ with $\pi(a) \in A'$, then, since π is a permutation, we have $\pi(a) \in A' - \{k{+}1, \bar{a}\}$. The set $A' - \{k{+}1, \bar{a}, a\}$ has $k - 2$ elements, hence it cannot contain all the π-images of elements of $B - \{k + 1, \bar{b}\}$ which has $k - 1$ elements. Hence, there must be some $b \in B - \{k{+}1, \bar{b}\}$ with $\pi(b) \in B'$. We again can switch the images of a and b.

Case 3: $\bar{a} > k + 1$ and $\bar{b} < k + 1$. Noting that $g(\pi) = g(\pi^{-1})$, we can replace π by π^{-1} and go back to case 2.

Case 4: $\bar{a} < k + 1$ and $\bar{b} < k + 1$. In this case, the set $A' - \{k{+}1, \bar{a}\}$ of size $k - 1$ cannot hold all the images of the set $B - \{k + 1\}$ of size k. Hence, there must be some $b > k + 1$ with $\pi(b) > k + 1$. Since $\bar{b}$ and $k + 1$ are both strictly smaller than b and $\pi(b)$, we can switch the images of $\bar{b}$ and b while not reducing $g(\pi)$. We then go back to case 2.

Case 5: $\bar{a} > k{+}1$ and $\bar{b} > k{+}1$. This case is symmetric to case 4. □

Let $S_{A,B'}$ (respectively $S_{A',B}$) denote the set of all one-to-one correspondences between A and B' (respectively B and A'). Lemma IV.7 implies

$$\max_{\pi \in S_w} g(\pi) = \max_{\sigma \in S_{A,B'}} \sum_{a \in A} \left\lceil \frac{\sigma(a) - a}{d} \right\rceil + \max_{\tau \in S_{A',B}} \sum_{a \in A'} \left\lceil \frac{\sigma(a) - a}{d} \right\rceil \tag{11}$$

If we can find an optimal σ and an optimal τ in the above expression, an optimal π with respect to $g(\pi)$ can be constructed by setting $\pi|_A = \sigma$ and $\pi|_B = \tau^{-1}$. We thus can get an analog of Lemma IV.2 for the case when w is odd. We omit the proof as it is similar to that of Lemma IV.2.

Lemma IV.8. *When $(w - 1)/2 \in \mathbb{Z}^+$, we have*

$$m_1(w, f, d) = 2f \left(\sum_{b \in B'} q_b - \sum_{a \in A} q_a \right) + f \max_{\sigma \in S_{A,B'}} t_1(\sigma) + f \max_{\tau \in S_{A',B}} t_2(\tau), \tag{12}$$

where $t_1(\sigma) = |\{a \in A : r_{\sigma(a)} - r_a > 0\}|$, and $t_2(\tau) = |\{a \in A' : r_{\tau(a)} - r_a > 0\}|$.

The following result is immediate from the lemma.

Theorem IV.9. $m_1(w, f, 1) = \frac{1}{2} f (w - 1)^2$ *when w is odd.*

We can now assume $d \ge 2$. The followings are analogs of Lemma IV.5 for the functions t_1 and t_2.

Lemma IV.10. *When w is odd and $d \ge 2$,*

$$\max_{\sigma \in S_{A,B'}} t_1(\sigma) = \frac{w}{2} - \mu_{B'}(0) - \max\{0, \mu_{B'}(1) - \mu_A(0)\}. \tag{13}$$

Proof. The fact that $t_1(\sigma) \leq \frac{w}{2} - \mu_{B'}(0) - \max\{0, \mu_{B'}(1) - \mu_A(0)\}$ for all $\sigma \in S_{A,B'}$ can be shown similar to that of Lemma IV.5. Thus, we only need to specify a σ that achieves equality. Write $k = q_k d + r_k$ as usual, then

$$\mu_A(r) = \begin{cases} q_k + 1 & 1 \leq r \leq r_k \\ q_k & r_k < r \leq d-1, \text{ or } r = 0. \end{cases} \tag{14}$$

The key is that $\mu_A(0) = q_k$, regardless of r_k. In general, $R_{B'}$ looks as follows.

$$R_{B'} = \{r_k+2, \dots, d-1, 0, \dots, d-1, \cdots, 0, \dots, 2r_k+1\},$$

in which there are $q_k - 1$ groups of $\{0, \dots, d-1\}$ in the middle, a group $\{r_k + 2, \dots, d-1\}$ in the beginning, and the last group is $\{0, \dots, 2r_k+1\}$. Similar to Lemma IV.5, we break up the cases as follows.

Case 1: $r_k = d-1$. In this case, we have $\mu_{B'}(0) = q_k$, and $\mu_{B'}(1) = q_k + 1$. The upper bound is $t_1(\sigma) \leq k - q_k - 1$, which can be achieved with the following σ:

$$\sigma(i) = \begin{cases} i + k + 2 & 1 \leq i \leq k-1 \\ k + 2 & i = k. \end{cases}$$

Case 2: $r_k = d - 2$.

If $d = 2$, then $\mu_{B'}(0) = \mu_{B'}(1) = q_k$. The upper bound is $t_1(\sigma) \leq k - q_k$, which can be achieved with the following σ:

$$\sigma(i) = i + k + 1, \;\; 1 \leq i \leq k.$$

If $d = 3$, then $\mu_{B'}(0) = q_k + 1$, and $\mu_{B'}(1) = q_k$. The upper bound is $t_1(\sigma) \leq k - q_k - 1$, which can be achieved with the following σ:

$$\sigma(i) = \begin{cases} i + k + 3 & 1 \leq i \leq k-2 \\ k + 1 & i = k-1 \\ k + 2 & i = k. \end{cases}$$

If $d > 3$, then $\mu_{B'}(0) = \mu_{B'}(1) = q_k$. The upper bound is $t_1(\sigma) \leq k - q_k - 2$, which can be achieved with the same σ as the case $d = 3$.

Case 3: $r_k \leq d - 3$ and $2r_k + 1 < d$. In this case, we have $\mu_{B'}(0) = \mu_{B'}(1) = q_k$. The upper bound is $t_1(\sigma) \leq k - q_k$, which can be achieved with the following σ:

$$\sigma(i) = \begin{cases} i + 2k + 1 - r_k & 1 \leq i \leq r_k \\ i + k + 1 - r_k & r_k + 1 \leq i \leq k. \end{cases}$$

Case 4: $r_k \leq d - 3$ and $2r_k + 1 = d$. In this case, we have $\mu_{B'}(0) = q_k + 1$ and $\mu_{B'}(1) = q_k$. The upper bound is $t_1(\sigma) \leq k - q_k - 1$, which can be achieved with the same σ as that in case 3.

Case 5: $r_k \leq d - 3$ and $2r_k + 1 > d$. In this case, we have $\mu_{B'}(0) = \mu_{B'}(1) = q_k + 1$. The upper bound is $t_1(\sigma) \leq k - q_k - 2$, which can be achieved with the following σ:

$$\sigma(i) = \begin{cases} i + k + d - r_k + 2 & 1 \leq i \leq k - d + r_k - 1 \\ i + d - r_k + 2 & k - d + r_k \leq i \leq k. \end{cases}$$

□

The following lemma can be shown similarly. We omit the proof due to space limitation.

Lemma IV.11. *When w is odd and $d \geq 2$,*

$$\max_{\tau \in S_{A',B}} t_2(\tau) = \frac{w}{2} - \mu_B(0) - \max\{0, \mu_B(1) - \mu_{A'}(0)\}. \tag{15}$$

Lemmas IV.10, IV.11, and IV.8 finally characterize $m_1(w, f, d)$ for the case when w is odd.

Theorem IV.12. *Suppose $\frac{w-1}{2} \in \mathbb{Z}$, and $2 \leq d \leq w - 1$.*

$$\begin{aligned} r &= \frac{w-1}{2} - d \left\lfloor \frac{w-1}{2d} \right\rfloor \\ \bar{m}_1 &= 2f \sum_{i=1}^{\frac{w-1}{2}} \left(\left\lfloor \frac{\frac{w-1}{2} + i}{d} \right\rfloor - \left\lfloor \frac{i}{d} \right\rfloor \right) + \\ &\quad 2f \left(w - 1 - 2 \left\lfloor \frac{w-1}{2d} \right\rfloor \right). \end{aligned}$$

Then, we have

(i) $m_1(w, f, d) = \bar{m}_1$ *when* $0 < r \leq d-3$ *and* $2r+1 < d$.

(ii) $m_1(w, f, d) = \bar{m}_1 - 2f$ *when* $r = 0$.

(iii) $m_1(w, f, d) = \bar{m}_1 - 4f$ *when either* $r = d - 1$ *or* $r = d - 2 = 1$ *or* $d = 2r + 1 \leq 2d - 5$.

(iv) $m_1(w, f, d) = \bar{m}_1 - 8f$ *when either* $r = d - 2 > 1$ *or* $d < 2r + 1 \leq 2d - 5$.

Remark IV.13. It is possible to put the results of both Theorems IV.6 and IV.12 in closed form (i.e. without the summations). However, the given presentations of the theorems are easier to follow. For instance, for the even w case in Theorem IV.6, if $2r < d$, then $\bar{m}_1$ can be calculated to be $\frac{3}{2}dq^2 - \frac{1}{2}dq + 3rq + q$, where $q = \lfloor w/(2d) \rfloor$. What we do want to notice is that in both cases, $m_1(w, f, d) = \Theta(fw^2/d)$!

V. Multicast WDM cross-connects

Consider any set $\mathcal{D}$ of (multicast) requests. Each request D in $\mathcal{D}$ is of the form $(\lambda_i, F_j, \mathcal{P})$, where $i \in [w]$, and $j \in f$, and $\mathcal{P} \subseteq \Lambda \times \mathcal{F}'$ such that $|\{i' : (\lambda_{i'}, F'_{j'}) \in \mathcal{P}\}| \leq 1, \forall j' \in [f]$. As usual, we use $i(D), j(D)$, and $\mathcal{P}(D)$ to denote i, j, and $\mathcal{P}$. Let $I'(D)$ denote the multiset of requested output wavelengths, namely $I'(D) := \{i' \mid (\lambda_{i'}, F'_{j'}) \in \mathcal{P}(D), \text{ for some } j'\}$. The size of $I'(D)$ $(= |\mathcal{P}|)$ is often called the *fanout* of the request D. Multicast requests with fanouts equal to 1 are nothing but unicast requests.

A set $\mathcal{D}$ of requests is called a *multicast request frame* if, for all $i \in [w]$ and $j \in [f]$,

$$|\{D \in \mathcal{D} : i(D) = i \text{ and } j(D) = j\}| \leq 1,$$

and, for all $i' \in [w]$ and $j' \in [f]$,

$$|\{D \in \mathcal{D} : (\lambda_{i'}, F'_{j'}) \in \mathcal{P}(D)\}| \leq 1.$$

If equality holds for the later inequality for all i', j', we call the request frame a *full request frame*. In other words, $\mathcal{D}$ is a request frame if there is at most one request in $\mathcal{D}$ from a wavelength in a particular input fiber, and there is at most one request to a wavelength in a particular output fiber. In a full request frame, each of the output wavelengths is involved in some request.

Recall that, if D is a unicast request, then we need at least $\lceil \frac{i(D) - i'(D)}{d} \rceil$ LWC(d). What is the corresponding number if D were are multicast request? Consider a multicast request D where $I'(D) = \{i'_1, \dots, i'_k\}$, and $i(D) = i$. Without loss of generality, assume

$$i'_1 \leq \cdots \leq i'_s \leq \mathbf{i} < i'_{s+1} \leq \cdots \leq i'_k,$$

where $0 \leq s \leq k$. If $s = 0$, then all $i' \in I'(D)$ are strictly greater than i. If $s = k$, then all $i' \in I'(D)$ are at most i.

In order to satisfy this request, a set of paths must be established between λ_i on input fiber F_j and all the $\lambda_{i'_t}, 1 \leq t \leq k$, on the corresponding output fibers. The union of these paths typically form the *routing tree* for this request. Construct the *wavelength conversion tree* T corresponding to the routing tree in the following manner. The tree T has nodes labeled with wavelengths, and is rooted at λ_i. Each node of T represents one wavelength on the routing tree. There is an edge connecting λ_x to λ_y in T if λ_x was converted to λ_y on some path of the routing tree. The number of edges of T is thus the number of wavelength converters used to satisfy the request. The tree T has the property that if (λ_x, λ_y) is an edge, then $|x - y| \leq d$, and that λ_i and all $\lambda_{i'_t}$, $t = 1, \dots, k$, appear as nodes of T. We say that the tree T *realizes* the request D. (In a sense, this tree is a Steiner tree spanning all wavelengths in $I'(D) \cup \{\lambda_i\}$.)

To find the necessary number of LWC(d) needed for the request D, we would like to determine the minimum number of edges (i.e. LWC(d)) of a wavelength conversion tree for D. Since each wavelength only needs to appear in the tree once (we can just identify occurrences of the same wavelength), multiplicities in the multiset $I'(D) \cup \{i(D)\}$ do not matter. We can assume without loss of generality that $i'_1 < \cdots < i'_s < i < i'_{s+1} < \cdots < i'_k$.

Another interesting observation is that we can turn T into a tree in which each vertex is of degree at most 2 (and thus is a path), while maintaining the fact that it is a wavelength conversion tree. Suppose T has a vertex labeled λ_x with degree at least 3. Let $\lambda_{x_1}, \lambda_{x_2}, \lambda_{x_3}$ be some three neighbors of λ_x in T. By the pigeonhole principle, two of the x_1, x_2, x_3 must be larger than x or smaller than x. Without loss of generality, assume $x_1 < x_2 < x$. Then, $|x - x_1| \leq d$ and $|x - x_2| \leq d$ imply $|x_1 - x_2| \leq d$. We thus can connect λ_{x_1} and λ_{x_2}, and remove the edge $(\lambda_{x_1}, \lambda_x)$ from the tree. The resulting tree still spans $I'(D)$ with all the desired properties, and with the same number of edges. In case $x_1 > x_2 > x$, we apply the same procedure.

Repeated applications of this procedure will produce a tree with maximum degree 2 if we can show that we do not run into an infinite loop. For each edge $e = (\lambda_x, \lambda_y)$ of T with $x < y$, define the *crossing weight* $w(e)$ be the number of nodes λ_z in T for which $x < z < y$. The procedure described above reduces the total crossing weight by one each time. Hence, repeated applications of the procedure shall terminate.

To this end, we can assume that T is a path. Let $i_1, \dots, i_{k+1}$ be the occurrences of elements of $I'(D) \cup \{i(D)\}$ as we visit T from one end to another. It is easy to see that the number of edges of T is at least

$$\sum_{t=1}^{k} \left\lceil \frac{|i_t - i_{t+1}|}{d} \right\rceil \geq \left\lceil \frac{|i'_1 - i'_2|}{d} \right\rceil + \left\lceil \frac{|i'_2 - i'_3|}{d} \right\rceil + \cdots + \left\lceil \frac{|i'_s - \mathbf{i}|}{d} \right\rceil + \left\lceil \frac{|\mathbf{i} - i'_{s+1}|}{d} \right\rceil + \left\lceil \frac{|i'_{s+1} - i'_{s+2}|}{d} \right\rceil + \cdots + \left\lceil \frac{|i'_{k-1} - i'_k|}{d} \right\rceil.$$

We can now conclude with an important lemma.

Lemma V.1. *Let D be a multicast request where $I'(D) = \{i'_1, \dots, i'_k\}$, $i = i(D)$, such that*

$$i'_1 \leq \cdots \leq i'_s \leq \mathbf{i} < i'_{s+1} \leq \cdots \leq i'_k.$$

Then, the number of LWC(d) necessary and sufficient to satisfy D is

$$\bar{c}(D) := \left\lceil \frac{|i'_1 - i'_2|}{d} \right\rceil + \left\lceil \frac{|i'_2 - i'_3|}{d} \right\rceil + \cdots + \left\lceil \frac{|i'_s - \mathbf{i}|}{d} \right\rceil + \left\lceil \frac{|\mathbf{i} - i'_{s+1}|}{d} \right\rceil + \left\lceil \frac{|i'_{s+1} - i'_{s+2}|}{d} \right\rceil + \cdots + \left\lceil \frac{|i'_{k-1} - i'_k|}{d} \right\rceil.$$

Proof. Necessity was shown above, noticing that identical elements of $I'(D)$ contributes zero to the cost $\bar{c}(D)$. The basic idea to show sufficiency is that we can convert λ_i to λ_s, λ_s to λ_{s-1}, until we get to λ_1, using exactly

$$\left\lceil \frac{|i'_1 - i'_2|}{d} \right\rceil + \left\lceil \frac{|i'_2 - i'_3|}{d} \right\rceil + \cdots + \left\lceil \frac{|i'_s - \mathbf{i}|}{d} \right\rceil$$

LWC(d). Similarly, we get the other half of the sum $\bar{c}(D)$. The signal on λ_i only needs to be split once. The point shall be clearer when we present the construction in Section VI. □

To satisfy all requests in a multicast request frame $\mathcal{D}$, the number of LWC(d) needed is at least $\bar{c}(\mathcal{D}) = \sum_{D\in\mathcal{D}} \bar{c}(D)$. When $\mathcal{D}$ is a unicast request frame, $\bar{c}(\mathcal{D}) = c(\mathcal{D})$. Hence, $\max_{\mathcal{D}} \bar{c}(\mathcal{D}) \geq m_1(w, f, d)$. What is interesting is that we can also show $\max_{\mathcal{D}} \bar{c}(\mathcal{D}) \leq m_1(w, f, d)$, which - along with our construction in the next Section - gives an analog of Theorem III.1 for the multicast case.

Lemma V.2. *We have*

$$\max_{\mathcal{D}} \bar{c}(\mathcal{D}) = m_1(w, f, d),$$

where the $\max$ *is over all multicast request frames* $\mathcal{D}$.

Proof. As noted above, it is sufficient to show $\max_{\mathcal{D}} \bar{c}(\mathcal{D}) \leq m_1(w, f, d)$. Consider a multicast request frame $\mathcal{D}$ that maximizes $\bar{c}(\mathcal{D})$. If $\mathcal{D}$ is also a unicast request frame (i.e. all requests have fanouts one), then the inequality certainly holds. To show that it holds in general, we shall gradually turn $\mathcal{D}$ into a unicast request frame without reducing $\bar{c}(\mathcal{D})$. Consider a request $D \in \mathcal{D}$ with fanout at least 2. As usual, assume $I'(D) = \{i'_1, \dots, i'_k\}$, $i = i(D)$, and $i'_1 \leq \cdots \leq i'_s \leq \mathbf{i} < i'_{s+1} \leq \cdots \leq i'_k$. Since D has fanout at least two, there must be at least one free input wavelength λ_t on some input fiber $F_{\bar{j}}$, namely $(\lambda_t, F_{\bar{j}})$ is not part of any request in $\mathcal{D}$. Let

$$i' = \begin{cases} i'_s & \text{if } s \neq 0,\ s \neq k, \text{ and } t > i \\ i'_{s+1} & \text{if } s \neq 0,\ s \neq k, \text{ and } t \leq i \\ i'_k & \text{if } s = k \\ i'_1 & \text{if } s = 0. \end{cases}$$

Replace D by two requests D_1 and D_2 such that $i(D_1) = i, j(D_1) = j(D), I'(D_1) = I'(D) - \{i'\}$, and that $i(D_2) = t, j(D_2) = \bar{j}, I'(D_2) = \{i'\}$, matching the output fiber for all members of $I'(D)$. Then, it is straightforward from Lemmas II.2 and II.3 that $\bar{c}(D) \leq \bar{c}(D_1) + \bar{c}(D_2)$. Repeated applications of this replacement eventually yield a unicast request frame without reducing the $\bar{c}$ cost of $\mathcal{D}$. □

Theorem V.3. *A strictly, wide-sense, or rearrangeably nonblocking multicast WXC under the* $(\lambda, F, \lambda', F')$*-request model needs at least* $m_1(w, f, d)$ *Moreover,* $m_1(w, f, d)$ *is also the sufficient number of LWC*(d) *to construct a rearrangeably or a wide-sense non-blocking multicast WXC.*

Proof. Necessity is obvious. Sufficiency comes from the construction in Section VI and the previous lemma. □

VI. Constructions with minimum number of LWC(d)

In this section, we will describe a construction that makes use of the so-called *converter pool*, which works for both the unicast and the multicast case. Consider the construction shown in Figure 2. The construction has two main components: the *LWC pool* component, and the *switching fabrics* component. The detailed design of the LWC pool component is shown in Figure 3. The basic idea is that all input wavelengths have access to all m converters (LWC(d)) in the pool. The converters are interconnected with splitters and combiners, so that an optic signal can be split as many times as we want, and it can go through as many LWC(d) as needed. The main objective of the LWC pool component is to use the least number of LWC(d) to convert input wavelengths to the requested output wavelength (in the unicast case), or the requested output wavelengths (in the multicast case). After the wavelength conversions are done, all output wavelengths come at the inputs of the switching fabrics. We can then use an appropriate switching fabrics construction such as those in [3, 4, 7] to finish the design. Let

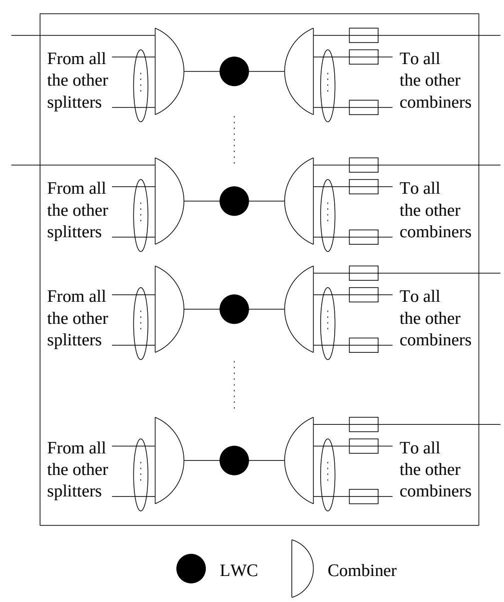

Fig. 3. A possible construction of the $LWC(d)$ pool.

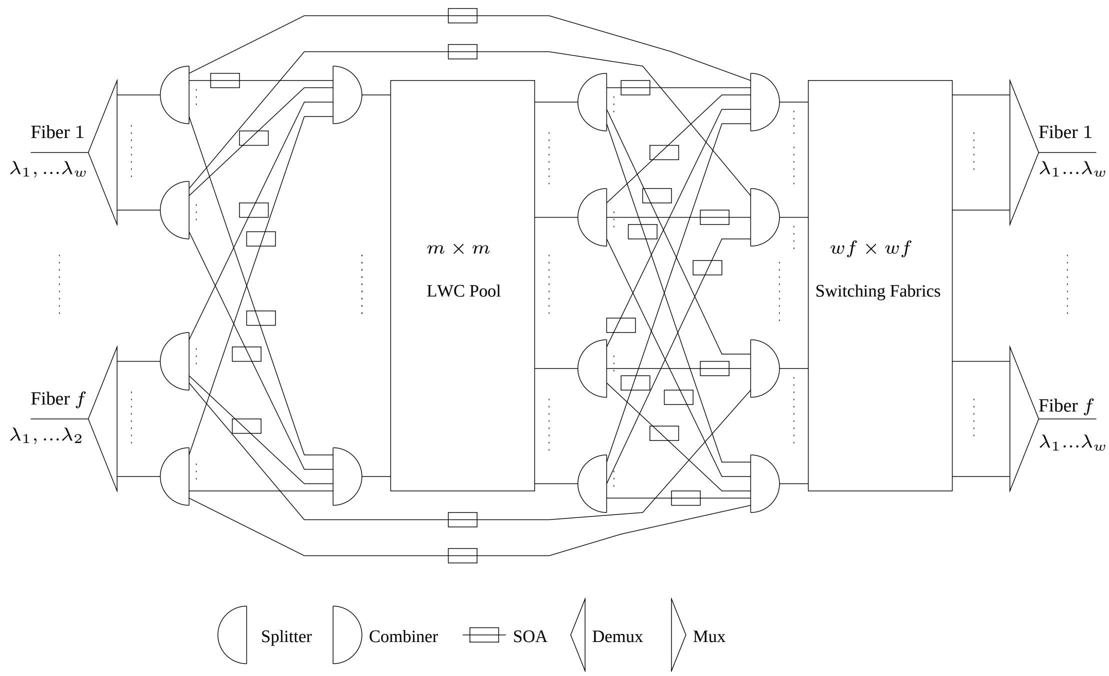

Fig. 2. Switching network with minimum number of LWC(d)s

us use $\mathcal{N}(m, w, f, d)$ to denote the construction described above. The switching fabrics does not have to convert wavelengths as all wavelengths arriving at its inputs are ready to come out to the right fiber, hence we only need it to be WSNB or RNB in the switching circuit sense [20, 21]. One can use a broadcast-and-select type of construction [7] to have an optical crossbar (instead of normal crossbars as in the circuit switching case). Instead of the number of cross-points, we have the same number of SOAs. This observation leads to the following results.

The following theorem completes the second part of Theorem III.1.

Theorem VI.1. *Let $m_1 = m_1(w, f, d)$. Consider the construction $\mathcal{N}(m_1, w, f, d)$. If the switching fabrics is rearrangeably or wide-sense nonblocking in the circuit switching sense, then $\mathcal{N}(m_1, w, f, d)$ is rearrangeably or wide-sense nonblocking under the $(\lambda, F, \lambda', F')$-request model, respectively.*

Proof. We show the WSNB case. The RNB case is shown similarly. Our strategy to route a new request $D = (\lambda_i, F_j, \lambda_{i'}, F'_{j'})$, given some previously routed network state, is as follows. First, use $\left\lceil \frac{|i-i'|}{d} \right\rceil$ free LWC(d) to convert λ_i to $\lambda_{i'}$. Second, at the last LWC(d), set up the SOAs to route this signal to a free input of the switching fabrics. The fact that the switching fabrics is WSNB implies that we can also find a route from this free input to the corresponding requested output. We just have to make sure that we have the sufficient number of LWC(d) so that the routing does not get blocked. This is precisely why $m = m_1(w, f, d)$ was needed. □

We also have the corresponding theorem for the multicast case.

Theorem VI.2. *Let $m_1 = m_1(w, f, d)$. Consider the construction $\mathcal{N}(m_1, w, f, d)$. If the switching fabrics is multicast rearrangeably or wide-sense nonblocking in the circuit switching sense, then $\mathcal{N}(m_1, w, f, d)$ is multicast rearrangeably or wide-sense nonblocking under the $(\lambda, F, \lambda', F')$-request model, respectively.*

Proof. The proof is very much the same as that of the previous theorem. Consider request D with the usual $i'_1 \leq \cdots \leq i'_s \leq i < i'_{s+1} \leq \cdots \leq i'_k$. We can first split the signal on λ_i into two branches (unless $s = 0$ or $s = k$). The first branch converts λ_i to $\lambda_{i'_s}$, $\lambda_{i'_s}$ to $\lambda_{i'_{s-1}}$, and so on. The second branch does the symmetric operation. The total number of LWC(d) needed is precisely $\bar{c}(D)$. The number of LWC(d) is sufficient due to Lemma V.2. Every time we get to a λ_x, $x \in I'(D)$, the signal is split so that λ_x comes to a free input of the switching fabrics, while the other copy continues with the branch. If $s = 0$ or $s = k$, then there is only one conversion branch. The switching fabrics finishes the rest of the work. Note that $\mathcal{N}(m_1, w, f, d)$ is multicast capable, yet the switching fabrics only needs to be a unicast network. The converter pool already does the splitting for us. □

The last thing we would like to mention is that the total number of SOAs used in $\Theta(m^2)$ for the LWC pool, and

$\Theta(fw\lg(fw))$ for the best construction of the circuit-type switching fabrics. Our results earlier showed that $m^2 = m_1^2 = \Theta(f^2w^4/d^2)$. Hence, the SOA cost is dominated by the LWC pool.

VII. Concluding remarks and future works

We have completely characterized the minimum number of LWC(d) needed for wide-sense and rearrangeably nonblocking unicast and multicast optical cross-connects. The construction given in the previous section is not the only construction that minimizes the number of wavelength converters, and it may not be the simplest in terms of physical layout. We leave this question for future work.

Another interesting point that comes from this paper is that there seems to be an intrinsic trade-off between the number of LWC(d) and the number of SOAs used. Recent nonblocking constructions proposed in [4], for example, did not make use of any SOAs at all. Others [7,12] used fewer SOAs than the construction given here. The drawback is that previous WXCs used wider range and non-uniform limited wavelength converters. Having too many SOAs not only complicates physical layout of the network, but also consumes powers, leading to signal attenuation. Investigating this trade-off is another future research topic. For instance, how can we construct nonblocking WXCs which make use of a some more LWC(d) but much less SOAs?

The strictly nonblocking case is not yet completely characterized. It is easy to see, for example, that $fw\lceil\frac{w-1}{d}\rceil$ LWC(d) are sufficient for a unicast construction. What one can do is to use $\lceil\frac{w-1}{d}\rceil$ LWC(d) to simulate a full-range wavelength converter, and then the rest can be done with a strictly nonblocking circuit-type of switching fabrics. This number is about twice the necessary number $m_1(w, f, d)$. (We also would like to note that this upper bound can be reduced a little, but it is still more than $m_1(w, f, d)$.)

Last but not least, the (λ, F, F')-request model was not address in this paper. One can pose exactly the same types of questions for this request model.

References

[1] 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.

[2] B. Mukherjee, "WDM optical communication networks: progress and challenges," *IEEE J. Select. Areas Commun.*, vol. 18, no. 10, pp. 1810–1824, 2000.

[3] Y. Yang and J. Wang, "WDM optical switching networks using sparse crossbars," in *Proceedings of the 23th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM'2004, Hong Kong)*. 2004, IEEE.

[4] Hung Q. Ngo, Dazhen Pan, and Chunming Qiao, "Nonblocking WDM switches based on arrayed waveguide grating and limited wavelength conversion," in *Proceedings of the 23rd Conference of the IEEE Communications Society (INFOCOM'2004, HongKong)*. 2004, IEEE.

[5] D.J. Bishop, C.R. Giles, and G.P. Austin, "The Lucent LambdaRouter: MEMS technology of the future here today," *IEEE Communication Magazine*, pp. 75–79, Mar 2002.

[6] Zhigang Jing H. Jonathan Chao, Kung-Li Deng, "A petabit photonic switch (p3s)," in *Proceedings of IEEE Infocom'03*, San Francisco, USA., April 2003.

[7] Yuanyuan Yang, Jianchao Wang, and Chunming Qiao, "Nonblocking WDM multicast switching networks," *IEEE Trans. Para. and Dist. Sys.*, vol. 11, no. 12, pp. 1274–1287, Dec 2000.

[8] Kuo-Chun Lee and Victor O. K. Li, "A wavelength-convertible optical network," *J. Lightwave Technol.*, vol. 11, pp. 962–970, May/June 1993.

[9] J. Yates, J. Lacey, D. Everitt, and M. Summerfield, "Limited-range wavelength translation in all-optical networks," in *Proceedings of the 15th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM'96)*. 1996, vol. 3, pp. 954–961, IEEE.

[10] R. Ramaswami and G. Sasaki, "Multiwavelength optical networks with limited wavelength conversion," *IEEE/ACM Transactions on Networking (TON)*, vol. 6, pp. 744–754, 1998.

[11] B. Ramamurthy and B. Mukherjee, "Wavelength conversion in WDM networking," *IEEE J. Select. Areas Commun.*, vol. 16, pp. 1061–1073, Sep 1998.

[12] Y. Yang and J. Wang, "Designing WDM optical interconnects with full connectivity by using limited wavelength conversion," in *Proceedings of the 18th IEEE International Parallel and Distributed Processing Symposium (IPDPS'04)*. 2004, IEEE.

[13] J. Elmirghani and H. Mouftah, "All-optical wavelength conversion: Technologies and applications in DWDM networks," *IEEE Communication Magazine*, vol. 38, pp. 86–92, Mar 2000.

[14] 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.

[15] 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.

[16] Yang Wang and Yuanyuan Yang, "Multicasting in a class of multicast-capable WDM networks," *J. Lightwave Technol.*, vol. 13, no. 2, pp. 128–141, Feb 2002.

[17] Douglas B. West, *Introduction to graph theory*, Prentice Hall Inc., Upper Saddle River, NJ, 1996.

[18] D. König, "Über graphen und ihre anwendung auf determinantentheorie und mengenlehre," *Math. Ann.*, vol. 77, pp. 453–465, 1916.

[19] John E. Hopcroft and Richard M. Karp, "An $n^{5/2}$ algorithm for maximum matchings in bipartite graphs," *SIAM J. Comput.*, vol. 2, pp. 225–231, 1973.

[20] V. E. Beneš, *Mathematical theory of connecting networks and telephone traffic*, Academic Press, New York, 1965, Mathematics in Science and Engineering, Vol. 17.

[21] Hung Q. Ngo and Ding-Zhu Du, "Notes on the complexity of switching networks," in *Advances in Switching Networks*, D.-Z. Du and Hung Q. Ngo, Eds., pp. 307–367. Kluwer Academic Publishers, 2001.