Rearrangeable and Nonblocking $[w, f]$-Distributors

Hung Q. Ngo, *Member, IEEE,* Yang Wang, *Member, IEEE,* and Dazhen Pan, *Member, IEEE*

***Abstract*— We formulate a graph model called $[w, f]$-*distributors* which is useful in analyzing the structures and comparing the quantitative complexities and qualitative features of optical multicast cross-connects. Using the formulation we show that two strictly nonblocking multicast optical cross-connects under two different models are equivalent topologically, even though one request model is much less restrictive than the other. We then investigate the tradeoff between the depth and the complexity of an optical multicast cross-connect using the model. Upper and lower complexity bounds are proved. In the process, we also give a generic recursive construction that can be used to construct optimal and near-optimal distributors. The recursive construction can also be used to construct cost-effective optical multicast cross-connects. Another important result that follows is the exact asymptotic behavior of the size of optimal $[w, f]$-connectors, the one-to-one communication version of $[w, f]$-distributors.**

I. Introduction

With the advances in the dense wavelength division multiplexing (DWDM) technology [1]–[3], the number of wavelengths per fiber has increased to up to hundreds or more, each of which can operate at 10Gbps (OC-192) or higher [4]–[6]. As a result, the raw transmission bandwidth has increased by more than four orders of magnitude over the last decade or so. However, the capacity of switches has only been up by a factor of ten, making them a bottleneck in the core optical network infrastructure [7]. How to design large capacity and large scale WDM switches cost-effectively thus becomes a challenging and important problem.

A very important component of the design process is to understand the qualitative features of a given cross-connect such as whether it is blocking or nonblocking. There are many existing designs which have different degrees of nonblockingness (strictly, widesense, rearrangeably) under various request models [8]–[15] and traffic patterns (unicast [9], [13], [15], [16], multicast [15]–[18]). A design can also be blocking as long as its blocking probability is below a certain threshold [11], [19]. Other desirable qualitative features include small cross-talk [12], or fault-tolerance [20]. Given an intuitively good design, one needs to know what qualitative features the design possesses. Most of the time, it is not trivial to answer such questions. We will see later that the graph model introduced in this paper helps, in several ways, answer this type of questions.

Given the large number of possible designs, it is also important to be able to compare different designs, and to know how close to be optimal a new design is in terms of cost-effectiveness. Because WDM switches may use a large variety of switching components such as semiconductor optical amplifiers (SOA) (often in conjunction with splitters) vs. array waveguide grating routers (AWGRs) (often in conjunction with tunable transceivers or wavelength converters) that affect the complexity of the switches, there is no single measure that can be used to represent the complexity of WDM switches. This is in contrast to electronic switches for which the number of cross-points has been commonly used to represent their complexities. Note that one cannot simply use the dollar cost as the ultimate measure either because it can be changed by the economic forces of supply and demand as well as technological advances. As shall be discussed, the proposed graph model can provide a good approximate measure on how "complex" a construction is.

To address the need for a generic model which can be used for both the qualitative and quantitative analyses of WDM switching networks, we have devised one such model for the one-to-one communication case and illustrated its usefulness in constructing new WDM optical cross-connect architectures [13].

A main contribution of this paper is to extend the graph model framework proposed in [13] to the one-to-many (multicast/broadcast) communication pattern. Our model also suggests interesting generalizations of classical switching network theory [21], [22]. Specifically, in this paper we make the following contributions:

(a) We formulate a graph model called $[w, f]$-*distributors* to analyze and compare WDM switches quantitatively and qualitatively. The formulations were done under three different request models, one of which is equivalent to the (multicast) circuit switching case; the other two models are more relaxed and are specific to optical switching. (Section IV.)

(b) We show that two strictly nonblocking WDM multicast cross-connects under two different models are equivalent topologically, even though one request model is much less restrictive than the other. (Section V.) This illustrates a point we made earlier about being able to analyze qualitative features of cross-connects using our framework.

(c) There is an inherent tradeoff between a WDM multicast switch's "depth" (which is proportional to signal attenuation, cross-talk) and its "size" (which approximates its complexity). We investigate this tradeoff in Sections V and VI.

(d) We give a generic recursive construction which can be used to construct WDM multicast switches of a given depth. The constructions also give optimal $[w, f]$-distributors and $[w, f]$-connectors (this question was left

Manuscript received – –, —-; revised – –, —-. The work of Hung Q. Ngo was supported in part by NSF CAREER Award CCF-0347565.

The authors are with the Computer Science and Engineering department, State University of New York at Buffalo, U.S.A. Emails: (hungngo, yw43, dpan)@cse.buffalo.edu

open in [13]). Moreover, all of the limited-depth constructions are near optimal (within a sub-linear factor).

The rest of the paper is organized as follows. Section II introduces intuitively basic concepts of WDM cross-connects (WXCs), the three request models, and nonblocking degrees. Section III motivates the graph models which will be rigorously defined in Section IV. Section V addresses several key complexity problems arising from the framework, including a proof of an equivalence relation and complexity lower bounds. Section VI explicitly constructs graphs with low complexity and give upper bounds. The ideas in this section can be used to construct actual WXCs of low complexity. Lastly, Section VII concludes the paper with a few remarks and discussions on future works.

II. Request models and nonblockingness

A general WDM cross-connect (WXC) consists of f input fibers each of which can carry a set $\Lambda = \{\lambda_1, \dots, \lambda_w\}$ of w wavelengths, and f' output fibers each of which can carry a set $\Lambda' = \{\lambda'_1, \dots, \lambda'_{w'}\}$ of w' wavelengths, where $fw = f'w'$. (See Figure 1.) This setting is often referred to as the

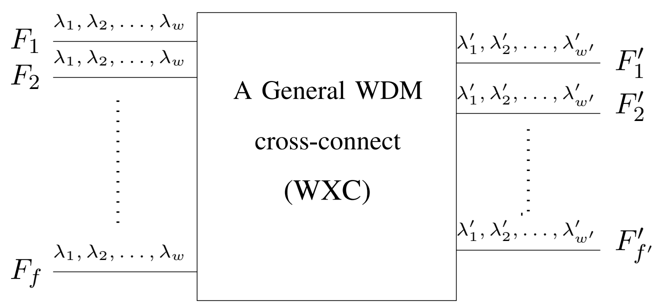

Fig. 1. Heterogeneous WDM Cross-Connect

heterogeneous case, which is needed to connect subnetworks from different manufacturers. Henceforth, let $n = fw = f'w'$, unless specified otherwise.

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. There are three main kinds of request models for multicast WDM switching networks. For lack of better names, we shall number them from 0 to 2, from the strictest to the more relaxed request models. The discussions in this section will be somewhat informal to convey the key ideas. The formal definitions of these models will be presented in the next section.

Request model 0: In this model a multicast request is of the form $(\lambda, F, \mathcal{P})$, where $\lambda \in \Lambda, F \in \mathcal{F}$, and $\mathcal{P} \subseteq \Lambda' \times \mathcal{F}'$. Basically, to satisfy this request one needs a multicast tree rooted at wavelength λ in fiber F whose leaves are the wavelengths λ' in the corresponding output fibers F' where $(\lambda', F') \in \mathcal{P}$.

Request model 1: This is the same as request model 0, except that no output fiber $F' \in \mathcal{F}'$ can appear more than once in $\mathcal{P}$. The restriction was made since in practical networks it is often not necessary to have a multicast connection going to the same output fiber on two different wavelengths [17], [23].

Request model 2: A multicast request is of the form $(\lambda, F, \mathcal{S})$, where $\lambda \in \Lambda, F \in \mathcal{F}$, and $\mathcal{S} \subseteq \mathcal{F}'$. Basically, in this case the request does not specify particular output wavelengths to be routed to. The request only needs to be routed to the given set $\mathcal{S}$ of output fibers. A multicast tree satisfying this request must have a leaf representing any one wavelength from each fiber in $\mathcal{S}$.

For each type of request models, three levels of nonblockingness are defined: rearrangeably nonblocking (RNB), wide-sense nonblocking (WSNB), and strictly nonblocking (SNB). Roughly speaking, an RNB switching network should be able to route a set of compatible requests given in advance. In the WSNB case, requests are nonblocking provided that they are routed according to some algorithm. In the SNB case, a new request compatible with the current network state can always be routed. These concepts are naturally carried over from classical circuit switching theory and thus should be familiar to readers who have been exposed to the theory. The reader is referred to [21], [22], [24], [25] for background materials on classical switching theory.

One might have expected that an optical switching network under model 2 would be less complex than that under model 1, and even further less complex than under model 0. Because, nonblocking under model 0 implies nonblocking under 1, which in turn implies nonblocking under 2. What is interesting and somewhat surprising is that this is not always the case, as we shall see later.

III. Motivations for a generic graph model

There have been a lot of studies on the constructions and characteristics of many different types of nonblocking multicast WXCs [8]–[12], [14]–[18], [26]–[32]. (The citations are by no means comprehensive, and we have already restricted ourselves to the multicast case. The reader is referred to [9], [13], [18] for more references on the unicast case.) The constructions from these references made use of various different types of optical components, such as arrayed waveguide grating routers (AWGR), limited-range wavelength converters (LWC), full-range wavelength converters (FWC), SOAs, optical add-drop multiplexers (OADM), wavelength selective cross-connects (WSC), wavelength interchangers (WI), directional couplers (DC), etc. It is clear that the task of comparing different designs is not easy. They make use of different optical switching components which oftentimes are tradeoffs. For instance, designs using SOAs and LWCs often have lower wavelength conversion cost than those that use AWGRs and LWCs; however, AWGRs are preferred over SOAs since AWGRs consume virtually no power.

In order to have a unified view of this wide landscape of design possibilities, we need a model that can capture the topological essence of most (if not all) the designs. As we have mentioned in the introduction, the model can be used for many different purposes, not just to compare designs topologically.

We will first briefly describe a graph model proposed by Ngo [13], which was used to study one-to-one communication is WDM switching networks. The main theme of this paper is to extend this model to the one-to-many communication case and investigate complexity and construction problems formulated from the model.

For any WDM switch design, we apply the following procedure to construct a directed acyclic graph (DAG) from the design: (a) replace each fiber by a set of vertices $\Lambda \cup \Lambda'$, which represents all possible wavelengths that can be carried on the fiber; (b) the edges of the DAG are defined according to the capability of switching components in the design. The edges connect wavelengths (i.e. vertices) on the inputs of each switching component to the wavelengths on the outputs in accordance with the functionality of the switching component.

We shall be somewhat brief on this construction. However, the reader will undoubtedly see the basic idea. As an example, Figure 2 shows how to turn an AWGR, an FWC, and a MUX into edges. Figure 3 shows a complete construction of the DAG from the design on the left.

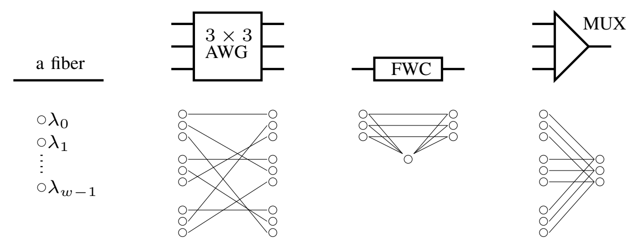

Fig. 2. Turning optical components into parts of a graph. A fiber is replaced by a set of vertices representing the wavelengths it can carry. Other components define edges connecting input wavelengths to output wavelengths. For the AWGR, MUX, and FWC, we illustrate with $w = 3$. Edges are directed from left to right.

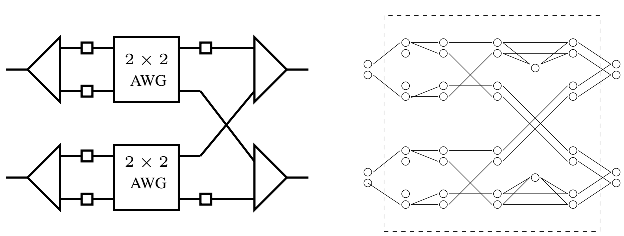

Fig. 3. A WDM switch design and its corresponding DAG.

It is easy to see that **a set of compatible routes from input wavelengths to output wavelengths correspond to a set of vertex disjoint paths from the inputs to the outputs of the DAG.**

There are two main parameters of the DAG, which capture the notion of "switch complexity" discussed earlier. The number of edges of the DAG, called the *size* of the DAG, is roughly proportional to the total cost of various components in the design. For example, an FWC corresponds to $3w$ edges while an WI corresponds to w^2 edges; a $w \times w$ AWGR corresponds to w^2 edges, while a $w \times w$ WDM crossbar corresponds to w^4 edges; etc. As WIs and WDM crossbars are more expensive than FWCs and AWGRs, this model makes sense. Other components follow the same trend.

The reader might have noticed that different components contribute different "weights" to the total cost, hence summing up the number of edges may not give the "right" cost. To answer this doubt, we make three points. Firstly, as argued earlier one cannot hope to have a perfect model which fits all needs, and part of the notion of cost is a business matter. Our first aim is at a more theoretical level. Secondly, this is the first step toward a good cost model. One certainly can envision weighted graphs as the next step. Thirdly, we surely can and should still use more traditional cost functions such as the direct counts of the number of each components and compare them individually.

The second measure on the DAG is its *depth*, i.e. the length of a longest path from any input to any output. as signals passing through different components of a design, they lose some power. The depth of the DAG hence reflects power loss, or in some cases even the signal delay. Again, different components impose different power loss factors. Hence, other information need to be taken into account to estimate power loss. However, it is clear that network depth is an important measure.

Last but not least, this DAG model provides a nice bridge between classical switching theory and WDM switching theory. As we shall see in later sections, this model helps us tremendously in answering qualitative questions about a particular construction. For example, if an wf-input wf-output DAG must have size $\Omega(f^2w^2)$ to be strictly nonblocking, then we know for certain that a construction of cost $o(f^2w^2)$ (reflected by the DAG's size) cannot be strictly nonblocking (for sufficiently large values of fw.)

IV. Distribution networks

In this section, we shall give more rigorous descriptions of the DAG models motivated from the last section. Our graphs will capture different degrees of nonblockingness and the tradeoff between size and depth of a network.

In the rest of the paper, define $[m] = \{1, \dots, m\}$ and $\mathbb{Z}_m = \{0, \dots, m-1\}$ for any positive integer m. For any finite set X, let 2^X denote the power set of X. For any positive integer k, we use $\binom{X}{k}$ to denote the set of all k-subsets of X. Graph theoretic terminologies we use here are fairly standard. See [33], for instance.

A. Request model 0

An (n_1, n_2)-*network* is a directed acyclic graph (DAG)

$$\mathcal{N} = (V, E; A, B),$$

where V is the set of vertices, E is the set of edges, A is a set of n_1 nodes called *inputs*, and B, disjoint from A, is a set of n_2 nodes called *outputs*. The vertices in $V - A \cup B$ are *internal* vertices. The in-degrees of the inputs and the out-degrees of the outputs are 0. The *size* of a network is its number of edges. The *depth* of a network is the maximum length of a path from an input to an output. For short, we call an (n, n)-network an n-network.

An n-network is meant to represent the DAG from last section under the request model 0, because in this request model each multicast request is to a specific set of wavelengths on the outputs. In this sense, each pair (λ', F'), where $\lambda' \in \Lambda'$ and $F' \in \mathcal{F}'$, is represented by an output in the n-network. (Recall $n = wf = w'f'$.) Later, we shall define $[w, f]$-networks which represent the DAG under the other two request models.

Given an n-network $\mathcal{N} = (A, B; V, E)$, a pair $D = (a, S) \in A \times 2^B$ is called a *distribution request* (or a *multicast request*) for $\mathcal{N}$ under model 0. When the request model under consideration is not ambiguous, we will drop the phrase "under request model i" for the sake of brevity. As we are only concerned with distribution networks in this paper, the term "request" should be implicitly understood as "distribution request" henceforth.

A *distribution assignment* is a set $\mathcal{D}$ of requests where no two requests share an input nor an output. A request $D = (a, S)$ is *compatible* with a distribution assignment $\mathcal{D}$ iff $\mathcal{D} \cup \{D\}$ is also a distribution assignment. A *distribution route* R for a request $D = (a, S)$ is a (directed) tree rooted at a whose leaves are precisely the nodes in S. We also say R *realizes* D. A *state* of $\mathcal{N}$ is a set $\mathcal{R}$ of vertex disjoint distribution routes. Each state of $\mathcal{N}$ realizes a unique distribution assignment, one route per request. A distribution assignment $\mathcal{D}$ is *realizable* iff there is a network state realizing it. A request is *compatible* with a state if it is compatible with the distribution assignment realized by the state.

A *rearrangeable (RNB) n-distributor* (or just n-distributor for short) is an n-network in which any distribution assignment is realizable.

A *strictly nonblocking (SNB) n-distributor* is an n-network $\mathcal{N}$ in which, given any network state $\mathcal{R}$ realizing a distribution assignment $\mathcal{D}$ and a new request D compatible with $\mathcal{D}$, there exists a route R such that $\mathcal{R} \cup \{R\}$ is a network state realizing $\mathcal{D} \cup \{D\}$.

As requests come and go, a strategy to pick new routes for new requests is called a *routing algorithm*. An n-network is called a *widesense nonblocking (WSNB) n-distributor* with respect to a routing algorithm $\mathbf{A}$ if $\mathbf{A}$ can always pick a new route for a new request compatible with the current network state. We can also replace $\mathbf{A}$ by a class of algorithms $\mathcal{A}$. In general, an n-network $\mathcal{N}$ is WSNB iff it is WSNB with respect to *some* algorithm.

We will consider two classes of functions on each network type: (a) the minimum size of a network, and (b) the minimum size of a network with a given depth. One of the key problems addressed in this paper is the tradeoff between networks' depths and their sizes.

Let $rd(n)$, $wd(n)$, and $sd(n)$ denote the minimum size of an RNB, WSNB, and SNB n-distributor, respectively. Let $rd(n, k)$, $wd(n, k)$, and $sd(n, k)$ denote the minimum size of an RNB, WSNB, and SNB n-distributor with depth k, respectively.

These classes of functions are well studied in the context of circuit switching networks (see, e.g., the surveys [22], [25]). The point we are making is that studying WDM switches under the request model 0 is in a sense the same as studying classical switching networks. A lot of results can be readily re-used. Table I summarizes the best bounds and constructions known to date.

Remark IV.1. In the literature, distributors are also called *generalized connectors*.

B. Request model 1

In this request model, each pair (λ, F) where $\lambda \in \Lambda$ is a wavelength in input fiber $F \in \mathcal{F}$ can still be thought of as an "input" to our graphs as in the previous request model. However, on the output side we do have to indicate the number f of fibers and the number of wavelengths w on each fiber.

A *$[w, f]$-network* is a wf-network (i.e. an n-network) $\mathcal{N} = (A, B; V, E)$ in which the set B of outputs is further partitioned into f subsets $B_1, \dots, B_f$ of size w each. Each set B_i represents an output fiber in the WDM switch. Members of B_i represent the wavelengths in the fiber.

Given a $[w, f]$-network $\mathcal{N}$, a pair

$$D = (a, S) \in A \times 2^B, \ S \neq \emptyset, \ \ |S \cap B_i| \leq 1, \ \forall i \in [f]$$

is called a (distribution) *request* for $\mathcal{N}$. The requirement that S has at most one element in each B_i represents the fact that the request goes to at most one specific wavelength in each output fiber. The size of S is called the *fanout* of the request.

The other concepts such as *distribution assignment*, *compatibility* between a new request and an assignment *distribution route*, and *network states* are defined in a similar fashion as in request model 0. Also, similar to the RNB, WSNB, and SNB $[w, f]$-distributors, we define RNB, WSNB, and SNB *$[w, f]_1$-distributors*.

Let $rd_1(w, f)$, $wd_1(w, f)$, and $sd_1(w, f)$ denote the minimum size of an RNB, WSNB, and SNB n-distributor, respectively. Let $rd_1(w, f, k)$, $wd_1(w, f, k)$, and $sd_1(w, f, k)$ denote the minimum size of an RNB, WSNB, and SNB $[w, f]_1$-distributor with depth k, respectively.

C. Request model 2

Given a $[w, f]$-network $\mathcal{N}$, a pair

$$D = (a, T) \in A \times 2^{[f]}, \ T \neq \emptyset$$

is called a *request* for $\mathcal{N}$. The set T specifies the subset of output fibers that the request needs to be routed to. In this request model, we are not concerned about the specific output wavelengths for the requests. Given a request $D = (a, T)$, define $a(D) = a$ and $T(D) = T$.

The concept of a distribution assignment under this model is not as straightforward as in the other two models. Let $\mathcal{D}$ be a set of requests; then, it is called a *distribution assignment* iff the following two conditions hold: (i) for any two different requests D_1 and D_2 in $\mathcal{D}$, $a(D_1) \neq a(D_2)$, and (ii) for any $i \in [f]$, the number of requests D in $\mathcal{D}$ whose $T(D)$ contains i is at most w. Rigorously, condition (ii) says that

$$|\{D \ : \ D \in \mathcal{D}, \ i \in T(D)\}| \leq w, \ \forall i \in [f].$$

The condition ensures that we are not "over-requesting" the output fiber B_i, because each of them can carry at most w signals on w different wavelengths.

A request D is *compatible* with a distribution assignment $\mathcal{D}$ iff $\mathcal{D} \cup \{D\}$ is also a distribution assignment. A *distribution route* R for a request $D = (a, T)$ is a (directed) tree rooted at a which has exactly one leaf node in each $B_i, i \in T$. We also say R *realizes* D. A *state* of $\mathcal{N}$ is a set $\mathcal{R}$ of vertex disjoint distribution routes. Each state of $\mathcal{N}$ realizes a unique

TABLE I

A SUMMARY OF BEST KNOWN RESULTS ON VARIOUS COMPLEXITY MEASURES FOR DISTRIBUTORS. THE COLUMN "GAP?" INDICATES IF THERE IS STILL AN ASYMPTOTIC GAP BETWEEN THE LOWER AND THE UPPER BOUNDS OF A FUNCTION. THE ANSWER "NO" DOES NOT MEAN THAT THE PROBLEM FOR THAT NETWORK TYPE IS SOLVED. IN ALL CASES IT REMAINS A GREAT CHALLENGE TO CONSTRUCT OPTIMAL NETWORKS. ENTRIES MARKED WITH "?" ARE NOT KNOWN EXCEPT TRIVIAL BOUNDS.

Function	Lower bound	Upper bound	Size of explicit constructions	Gap?
$rd(n)$	$\Omega(n \lg n)$, as [34] showed $rd(n) = c(n) + O(n)$	$O(n \lg n)$	$\Theta(n \lg n)$ [35]	NO
$rd(n, k)$	$\Omega(n^{1+\frac{1}{k}})$ (since $rd(n, k) \geq rc(n, k)$	$O((n \log n)^{1+\frac{1}{k}})$ [36]	$O(n^{1+\frac{1}{j}})$ ($k = 3j - 2$, [36]) $O(n^{\frac{5}{3}})$ ($k = 3$, [37], [38]) $O(n^{\frac{3}{2}}(\log n)^{\frac{1}{2}})$ ($k = 3$, [39]) $O(n^{1+\frac{1}{j}}(\log n)^{\frac{j-1}{2}})$ ($k = 2j - 1$, [39], also [40])	YES
$wd(n)$	?	?	?	YES
$wd(n, k)$	$\Omega(n^{1+\frac{1}{k}})$ (since $wd(n, k) \geq wc(n, k)$)	$O(n^{1+\frac{1}{k}}(\log n)^{1-\frac{1}{k}})$ [41]	$O(n^{5/3})$ (for $k = 2$, [37], [41]) $O(n^{1+2/j})$ (for $k = j^2 - 3j + 3$, [38], [42]) $O(n^{1+\frac{1}{k}+o(1)})$ [43], see also [44]–[50]	YES
$sd(n)$	$\Omega(n^2)$ [51]	$\Theta(n^2)$	Trivial	NO
$sd(n, k)$	?	?	?	YES

distribution assignment, one route per request. A distribution assignment $\mathcal{D}$ is *realizable* iff there is a network state realizing it. A request is *compatible* with a state if it is compatible with the distribution assignment realized by the state.

We can now define RNB, WSNB, and SNB $[w, f]_2$-*distributors* in a similar fashion as the previous two request models. The functions $rd_2(\cdot), wd_2(\cdot)$, and $sd_2(\cdot)$ are also similarly defined.

V. LOWER BOUNDS AND AN EQUIVALENCE RELATION

A. Straightforward observations

The following observations follow straightforwardly from the definitions.

Proposition V.1. *Let $n = wf$, and k be any positive integer. We have*

$$rd(n) \leq wd(n) \leq sd(n),$$

$$rd(n, k) \leq wd(n, k) \leq sd(n, k),$$

$$rd_1(w, f) \leq wd_1(w, f) \leq sd_1(w, f),$$

$$rd_1(w, f, k) \leq wd_1(w, f, k) \leq sd_1(w, f, k),$$

$$rd_1(w, f) \leq wd_1(w, f) \leq sd_1(w, f),$$

$$rd_1(w, f, k) \leq wd_1(w, f, k) \leq sd_1(w, f, k).$$

Proposition V.2. *Let $n = wf$. We have*

$$rd(n) \geq rd_1(w, f) \geq rd_2(w, f),$$

$$wd(n) \geq wd_1(w, f) \geq wd_2(w, f),$$

$$sd(n) \geq sd_1(w, f) \geq sd_2(w, f).$$

Also, for any depth k,

$$rd(n, k) \geq rd_1(w, f, k) \geq rd_2(w, f, k),$$

$$wd(n, k) \geq wd_1(w, f, k) \geq wd_2(w, f, k),$$

$$sd(n, k) \geq sd_1(w, f, k) \geq sd_2(w, f, k).$$

Beside these simple relations, the sizes of depth-1 distributors are also simple to derive. We omit the proof of the following proposition. Note also that the rest of the paper only consider the SNB and RNB cases, leaving the more elusive WSNB case for future works.

Proposition V.3. *Let $n = wf$, then*

$$rd(n, 1) = sd(n, 1) = rd_1(w, f, 1) = sd_1(w, f, 1) = sd_2(w, f, 1) = (wf)^2. \quad (1)$$

For $rd_2(w, f, 1)$, the situation is a little different. For integers $n \geq m$, an (n, m)-*concentrator* is an (n, m)-network where $n \geq m$, such that for any subset S of m inputs there exists a set of m vertex disjoint paths connecting S to the outputs. It is easy to show that the optimal size of a depth-1 (n, m)-concentrator is exactly $m(n - m + 1)$ [13]. Consider an RNB $[w, f]_2$-distributor $\mathcal{N} = (V, E; A, B)$ of depth 1. The restriction of the graph to V and B_i for any $i \in [f]$ is obviously a (wf, w)-concentrator. So, the minimum size of a depth-1 RNB $[w, f]_2$-distributor is at least $f \cdot w(wf - w + 1) = wf(wf - w + 1)$. Conversely, we can identify, in a one-to-one manner, the inputs of f separate (wf, w)-concentrators of depth 1 and size $w(wf - w + 1)$ to obtain a $[w, f]_2$-distributor. The outputs of each concentrator form a separate output band B_i. We have just shown the following proposition.

Proposition V.4. *We have*

$$rd_2(w, f) = wf(wf - w + 1). \quad (2)$$

The above propositions and known results in Table I already gave us some preliminary bounds for the functions we are studying. The rest of this paper is devoted to proving bounds beyond these simple bounds, and to constructing some optimal and near-optimal $[w, f]$-networks.

B. Strictly nonblocking distributors

The following theorem essentially shows that being SNB in the more relaxed request model 2 gives us no advantage than being in model 1 as far as network cost is concerned.

Theorem V.5. *Let w, f be positive integers where $f \geq 2$. Then, a $[w, f]$-network is an SNB $[w, f]_1$-distributor if and only if it is an SNB $[w, f]_2$-distributor.*

Proof. It is obvious that an SNB $[w, f]_1$-distributor is also an SNB $[w, f]_2$-distributor. We now show the converse.

Let $\mathcal{N} = (V, E; A, B)$ be an SNB $[w, f]_2$-distributor. (Recall that the partition $B = B_1 \cup \cdots \cup B_f$ is implicit.) Let $\mathcal{R}$ be a state of $\mathcal{N}$ under model 1, namely $\mathcal{R}$ is a set of vertex disjoint trees whose roots are inputs and whose leaves are outputs of $\mathcal{N}$ with at most one leaf in each "output band" B_i. Let $D = (a, S)$ be a request compatible with $\mathcal{R}$ under model 1, namely under $\mathcal{R}$ the input a is free, and each output $s \in S$ is free and belongs to a distinct output band B_i. For each $s \in S$, let $B_{i(s)}$ denote the output band of which s is a member. By the definition of request model 1, $B_{i(s)} \neq B_{i(s')}$ for any two members $s \neq s'$ of S.

We need to show that there is a tree R rooted at a with leaves S, and R is vertex disjoint from the trees in $\mathcal{R}$. The main idea of the proof is to show that there is a state $\mathcal{S}$ of $\mathcal{N}$ satisfying the following conditions:

(i) $\mathcal{R} \subseteq \mathcal{S}$,
(ii) a is free in $\mathcal{S}$, and
(iii) each s in S is the only free output in $B_{i(s)}$.

Suppose such a state $\mathcal{S}$ can be constructed. Consider the request (a, T), where $T = \{i(s) \mid s \in S\}$. This request is compatible with $\mathcal{S}$ under model 2. Since $\mathcal{N}$ is an SNB $[w, f]_2$-distributor, there is a tree R realizing (a, T) compatible with $\mathcal{S}$. This is the tree we are looking for, as the leaves of the tree have to be precisely those in S, and any route R compatible with $\mathcal{S}$ is also compatible with $\mathcal{R}$.

To show the existence of such a state $\mathcal{S}$, let us consider two cases as follows.

Case 1: *there is some route in $\mathcal{R}$ with more than one leaf.* Let X and Y be the number of free inputs and outputs in $\mathcal{R}$, respectively. Then, $a \in X$ and $|X| > |Y|$, because the total numbers of inputs and outputs are the same ($= n$). Let k be an integer such that B_k has some free output in $\mathcal{R}$. Let x be a member of $X - \{a\}$. The request (x, k) is compatible with $\mathcal{R}$; hence, there is a route R_1 from x to some output in B_k for which $\mathcal{R} \cup \{R_1\}$ is a state. Repeat this process $|Y|$ times, we will have a state $\mathcal{R}' = \mathcal{R} \cup \{R_1, \dots, R_{|Y|}\}$ in which there is no more free outputs, yet a is still free. Now, remove from $\mathcal{R}'$ all routes whose endpoints are those in S, we get the desired state $\mathcal{S}$.

Case 2: *all routes in $\mathcal{R}$ are one-to-one routes.* This is a much trickier case, as $|X| = |Y|$ and a has to be involved in the "filling up" process. As in case 1, we make requests of the form (x, k), $x \in X$. The vertex a is somewhat special, we make sure that a request $(a, i(\bar{s}))$ is created first, for some $\bar{s} \in S$. The rest of the (x, k) requests are arbitrary as before. For each $x \in X$, let R_x denote the route corresponding for the request (x, k). As in case 1, let $\mathcal{R}' = \mathcal{R} \cup \{R_x \mid x \in X\}$ be the final state.

If R_a ends at $\bar{s}$, then we are lucky. Remove from $\mathcal{R}'$ the routes R_a and all the R_x which end at some s in S, we get the desired state $\mathcal{S}$.

If we are not lucky, R_a ends at some t in $B_{i(\bar{s})}$, where $t \neq \bar{s}$. Let $\bar{a} \in X$ be an input such that $R_{\bar{a}}$ ends at $\bar{s}$. Let b be any input whose corresponding route in $\mathcal{R}'$ is $(b, v_1, \dots, v_p, u)$, where $u \in B_j$ for some $j \neq i(\bar{s})$. (Since $f \geq 2$, that there is some $j \neq i(\bar{s})$.)

Let $\mathcal{S}' = \mathcal{R}' - \{R_a, R_{\bar{a}}, R_b\}$, which is a network state.

Claim: there is an $(\bar{a}, t)$-route compatible with $\mathcal{S}'$.

To prove the claim, consider the state $\mathcal{S}' \cup \{R_a\}$. The request $(b, i(\bar{s}))$ is compatible with the state. Moreover, $\bar{s}$ is the only free output in $B_{i(\bar{s})}$ in the state $\mathcal{S}'$. Thus, there is a $(b, \bar{s})$-route $R_{b\bar{s}}$ such that $\mathcal{S}' \cup \{R_a, R_{b\bar{s}}\}$ is a state. Now, in the state $\mathcal{S}' \cup \{R_{b\bar{s}}\}$ the output t is the only free output in $B_{i(\bar{s})}$. Hence, the compatible request $(\bar{a}, i(\bar{s}))$ has to be routed to t. Hence, there is an $(\bar{a}, t)$-route compatible with $\mathcal{S}'$ as claimed.

To this end, we further consider two cases as follows.

(2a) among all $(\bar{a}, t)$-routes which are compatible with $\mathcal{S}'$, there is a route $R_{\bar{a}t}$ which is vertex disjoint from R_b, i.e. $R_{\bar{a}t}$ does not contain any vertex $v_1, \dots, v_p$. In this case, $\mathcal{S}' \cup \{R_b, R_{\bar{a}t}\}$ is a state in which a and $\bar{s}$ are the only free vertices. A request $(a, i(\bar{s}))$ then brings us back to the "lucky" situation considered earlier.

(2b) every $(\bar{a}, t)$-route compatible with $\mathcal{S}'$ intersects R_b at some point. Let $R_{\bar{a}t}$ be such an $(\bar{a}, t)$-route whose last intersection vertex on $(v_1, \dots, v_p)$ has the largest index, say v_q, where $1 \leq q \leq p$. Then, $R_{\bar{a}t}$ is composed of two parts: the part from $\bar{a}$ to v_q, and the part from v_q to t.

Now, let R_{bt} be the (b, t)-path consisting of the part $(b, v_1, \dots, v_q)$ concatenated with the (v_q, t)-part of $R_{\bar{a}t}$. Then, certainly $\mathcal{S}' \cup \{R_{bt}\}$ is a state in which the request $(a, i(\bar{s}))$ is valid, and $\bar{s}$ is the only free output in $B_{i(\bar{s})}$. Hence, there is an $(a, \bar{s})$-route $R_{a\bar{s}}$ which is compatible with $\mathcal{S}' \cup \{R_{bt}\}$.

If the route $R_{a\bar{s}}$ is vertex disjoint from R_b, then the request $(\bar{a}, i(\bar{s}))$ under the state $\mathcal{S}' \cup \{R_b, R_{a\bar{s}}\}$ can only be satisfied by routing $\bar{a}$ to t. The resulting state brings us back to the "lucky" situation.

For the contrary, suppose $R_{a\bar{s}}$ intersects R_b at some vertex. Then, due to the fact that $R_{a\bar{s}}$ is vertex disjoint from R_{bt}, the vertices in the intersection must all come after v_q. Let $v_{q'}$, $q' > q$, be a vertex in the intersection of $R_{a\bar{s}}$ and R_b. Let $R_{b\bar{s}}$ be the route obtained by concatenating the route $(b, v_1, \dots, v_{q'})$ and the $(v_{q'}, \bar{s})$-part of $R_{a\bar{s}}$. In the state $\mathcal{S}' \cup \{R_{b\bar{s}}\}$, the request $(\bar{a}, i(\bar{s}))$ is valid. A route realizing this request must intersect R_b (since we are in case 2b) at a point after $v_{q'}$ (since we are in the state $\mathcal{S}' \cup \{R_{b\bar{s}}\}$), contradicting the maximality of q. □

Corollary V.6. *Given positive integers w, f, and k, we have $sd_1(w, f) = sd_2(w, f)$ and $sd_1(w, f, k) = sd_2(w, f, k)$.*

The reader might be wondering if the same "equivalence" theorem holds for model 0 also. It seems that SNB is too strong a condition which makes all models equivalent. This turns out to be false, as shown below.

Theorem V.7. *Let $n = wf$. There is a $[w, f]_1$-distributor* *which is not an n-distributor. This implies that there is a $[w, f]_2$-distributor which is not an n-distributor.*

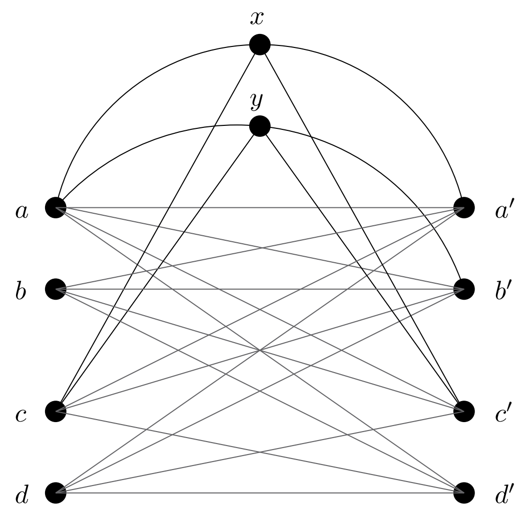

Fig. 4. An example showing that sd and sd_1/sd_2 are not equivalent.

Proof. Since every $[w, f]_1$-distributor is also a $[w, f]_2$-distributor, the second assertion is trivial. To show the first assertion, consider the $[w, f]$-networks shown in Figure 4. The network has $w = f = 2$, $B_1 = \{a', b'\}$, and $B_2 = \{c', d'\}$. It is almost a complete bipartite graph, except for the missing link (c, c'), and the links that are incident to x and y.

It is not difficult to see that the network is a $[w, f]_1$-distributor. The only request that may not be realizable is a request from c to c' (and possibly another output in B_1). For this request to be valid in a network state under model 1, c' must be free. Moreover, either x is free or y is free, implying that the request to c' can be routed. If both are not, then there is a route from a to both a' and b', contradicting the fact that we're under model 1.

On the other hand, under model 0 it is possible to have a network state in which the tree from a goes to x, y and then a', b'. In this case the request from c to c' cannot be realized. □

The above theorem does not show that $sd(n) > sd_1(w, f)$ nor $sd(n, k) > sd_1(w, f, k)$, however. It is still possible that $sd(n) = sd_1(w, f)$ or $sd(n, k) = sd_1(w, f, k)$ for some (or all) k. We leave the question open.

C. Rearrangeable distributors

Let $A = \{a_0, \dots, a_{n-1}\}$ and $B = \{b_0, \dots, b_{n-1}\}$. An *$n$-shifter* is an n-network $G = (V, E; A, B)$ such that for each $k \in \{0, \dots, n-1\}$, there are n vertex disjoint paths joining a_i to $b_{(i+k) \bmod n}$, for $i = 0, \dots, n$.

Let $n = wf$. For each $q \in \mathbb{Z}_f$ define a function

$$\phi_q(i) = ((i + q) \pmod{f}) + 1,\ i \in \mathbb{Z}_n.$$

Basically, for each fixed value of q, ϕ_q is a function assigning each number i in $\mathbb{Z}_n$ to a number in $[f]$. In fact, it is easy to verify that, for any $k \in [f]$ and any $q \in \mathbb{Z}_f$, we have

$$|\{i \in \mathbb{Z}_n \ : \ \phi_q(i) = k\}| = w.$$

In words, the number of i which ϕ_q maps to k is exactly w. Again, let $A = \{a_0, \dots, a_{n-1}\}$ and $B = \{b_0, \dots, b_{n-1}\}$. A *$[w, f]$-shifter* is a $[w, f]$-network $G = (V, E; A, B)$ such that, for any $q \in \mathbb{Z}_f$, there are n vertex disjoint paths joining a_i to some vertex in $B_{\phi_q(i)}$.

Note that a $[1, n]$-shifter is an n-shifter. Pippenger and Yao [52] showed that an n-shifter of depth k must have at least $kn^{1+1/k}$ edges. We shall use their idea to show that $[w, f]$-shifters must have size at least $kwf^{1+1/k}$.

Let $T_k(f)$ be a directed rooted tree with f leaves and depth at most k where all edges are directed to the direction of the leaves. Let $P_1, \dots, P_f$ be the f paths from the root to the leaves of $T_k(f)$. Define

$$\Delta(T_k(f)) := \sum_{j=1}^{f} \sum_{v \in P_j} \text{out-degree}(v). \tag{3}$$

The following lemma is from [52].

Lemma V.8 (Pippenger-Yao)**.** $\Delta(T_k(f)) \geq kf^{1+\frac{1}{k}}$.

Now, we are ready to prove the $[w, f]$-version of Pippenger-Yao's theorem.

Lemma V.9. *Every $[w, f]$-shifter of depth k must have size at least $kwf^{1+1/k}$.*

Proof. Because $\mathcal{N}$ is a $[w, f]_2$-distributor, for each $q \in \mathbb{Z}_f$ there are n vertex disjoint paths P_{iq}, $i \in \mathbb{Z}_n$, such that P_{iq} joins a_i to some vertex in $B_{\phi_q(i)}$.

Fix an i, assemble all f paths P_{iq} into a tree T_i (keeping only the initial common segments of the paths), then T_i is a tree with f leaves and depth at most k. To this end, for $1 \leq i \leq n$, $1 \leq q \leq f$, and $e \in E$, let

$$\mu(i, q, e) := \begin{cases} 1 & \text{if } e \text{ is an arc emitted from a node on } P_{iq} \\ 0 & \text{otherwise.} \end{cases}$$

For each vertex $v \in V$, let out-degree$_{T_i}(v)$ denote the out-degree of v in T_i. It is easy to see the following

$$\sum_{e \in E} \mu(i, q, e) \geq \sum_{v \in P_{iq}} \text{out-degree}_{T_i}(v). \tag{4}$$

Basically, the left hand side counts also some arcs not in T_i (but starts on P_{iq}).

Summing (4) over $i = 0, \dots, n-1$ and $q = 0, \dots, f-1$, we get

$$\begin{aligned} \sum_{i=0}^{n-1} \sum_{q=0}^{f-1} \sum_{e \in E} \mu(i, q, e) \quad &\geq \quad \sum_{i=0}^{n-1} \sum_{q=0}^{f-1} \sum_{v \in P_{iq}} \text{out-degree}_{T_i}(v) \\ &= \quad \sum_{i=0}^{n-1} \Delta(T_i) \\ &\geq \quad nkf^{1+1/k}. \end{aligned} \tag{5}$$

The last inequality comes from Lemma (V.8).

On the other hand, since the paths P_{iq} for a fixed q are vertex disjoint, we have

$$\sum_{i=0}^{n-1} \mu(i, q, e) \leq 1.$$

Consequently,

$$\sum_{i=0}^{n-1}\sum_{q=0}^{f-1}\sum_{e\in E}\mu(i,q,e)=\sum_{q=0}^{f-1}\sum_{e\in E}\sum_{i=0}^{n-1}\mu(i,q,e)\le f|E|. \quad (6)$$

Together, (5) and (6) lead to $|E|\ge kwf^{1+1/k}$ as desired. □

Because every $[w,f]_1$-distributor is a $[1,wf]$-shifter and every $[w,f]_2$-distributor is a $[w,f]$-shifter, the following theorem follows straightforwardly.

Theorem V.10. *For any integer $k\ge 2$, a depth-k $[w,f]_1$-distributor must have size at least $k(wf)^{1+1/k}$. In other words,*

$$rd_1(w,f,k)\ge k(wf)^{1+1/k}. \quad (7)$$

Similarly, a depth-k $[w,f]_2$-distributor must have size at least $kwf^{1+1/k}$, i.e.

$$rd_2(w,f,k)\ge kwf^{1+1/k}. \quad (8)$$

Corollary V.11. *For $k\ge 2$,*

$$rd_1(w,f)\ge ewf\ln(wf), \quad (9)$$

and

$$rd_2(w,f)\ge ewf\ln f, \quad (10)$$

where e is the base of the natural log.

Proof. For any integer m, the function $g(x)=xm^{1+1/x}$, with $x\ge 1$, is minimized at $x=\ln m$. □

What is quite amazing is that both the lower bounds in Corollary V.11 are asymptotically optimal, as we will see in the next section.

VI. Constructions and Upper Bounds

Asymptotically, the lower bounds for $rd_1(w,f)$ and $rd_1(w,f,k)$ are the same as those of $rd(n)$ and $rd(n,k)$ in Table I. The gaps between the upper and lower bounds of $rd(n)$ and $rd(n,k)$ are quite small: within poly-log for $rd(n,k)$, and exactly the same for $rd(n)$. Since every n-distributor is also a $[w,f]_1$-distributor, we can just use the upper bounds and constructions of n-distributors for $[w,f]_1$-distributors. All upper bounds in Table I can be used for $rd_1(w,f)$ and $rd_1(w,f,k)$. In fact, Corollary V.11 gives the exact asymptotic for $rd_1(w,f)$:

Theorem VI.1. *We have*

$$rd_1(w,f)=\Theta(wf\log(wf)). \quad (11)$$

This is attainable with the classic construction shown in Figure 5 using an n-connector in the second part [34]. The reader is referred to [25], [34] for the definitions of connectors and generalizers, and constructions of generalizers of linear size. Note also that generalizers are also referred to as *copy networks*, *replicators*, or *generalized concentrators* [22], [53].

The situation with $[w,f]_2$-distributors are quite different, however. The lower bound $kwf^{1+1/k}$ of $rd_2(w,f,k)$ is missing a factor $w^{1/k}$ as compared to the corresponding lower bound for $rd(n,k)$. Moreover, in practice there are hundreds

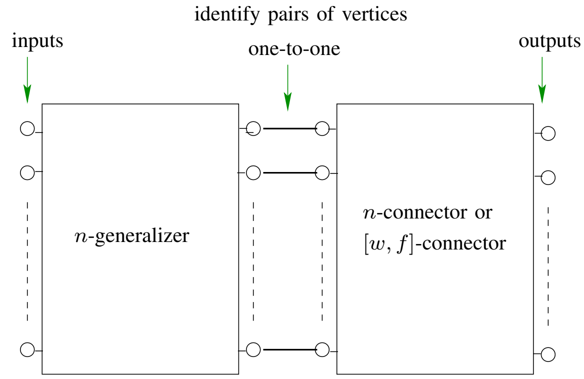

Fig. 5. Constructing distributors from generalizers and connectors

of wavelengths per fiber while there are only tens of fibers per cross-connect, which means $w\gg f$. Consequently, we need to construct $[w,f]_2$-distributors which are quite different than the existing constructions of n-distributors in order to get near the target lower bounds.

The so-called $[w,f]$-connectors were introduced in [13] to capture the one-to-one communication case of request model 2. Given a $[w,f]$-network $\mathcal{N}$, a pair $D=(a,i)\in A\times[f]$ is called a (connection) *request* for $\mathcal{N}$. The number i is called the *output fiber number* of D. A set $\mathcal{D}$ of requests is called a *request frame* iff no two requests share an input, and for any $i\in[f]$, we have

$$|\{a\mid(a,i)\in\mathcal{D}\}|\le w.$$

Here, a request frame is the one-to-one version of the distribution assignment. A request $D=(a,k)$ is *compatible* with a request frame $\mathcal{D}$ iff $\mathcal{D}\cup\{D\}$ is also a request frame. A *route* R for a request $D=(a,k)$ is a path from a to some vertex b in B_k. We also say R *realizes* D. A *state* of $\mathcal{N}$ is a set $\mathcal{R}$ of vertex disjoint routes. Each state of $\mathcal{N}$ realizes a request frame. A request frame $\mathcal{D}$ is *realizable* iff there is a network state realizing it.

A *rearrangeable* $[w,f]$*-connector* is a $[w,f]$-network in which the request frame

$$\mathcal{D}=\{(a,\sigma(a))\mid a\in A\}$$

is realizable for any mapping $\sigma:A\to[f]$ such that

$$|\{a\mid\sigma(a)=k\}|=w,\ \forall k\in[f].$$

It is easy to see that the classic construction also works for request model 2, i.e. we can replace the n-connector in Figure 5 by a $[w,f]$-connector to get a $[w,f]_2$-distributor. Formally, we have

Theorem VI.2. *Let $n=wf$. Concatenating an n-generalizer and a $[w,f]$-connector gives us a $[w,f]_2$-distributor.*

Let $\overline{rc}(w,f)$ be the minimum size of a $[w,f]$-connector. The construction and Corollary V.11 gives

$$ewf\ln f\le rd_2(w,f)\le\Theta(wf)+\overline{rc}(w,f). \quad (12)$$

So, if $\overline{rc}(w,f)=O(wf\log(wf))$, then $rd_2(w,f)=\Theta(wf\log(wf))$. On the other hand, [13] has shown

TABLE II
KNOWN RESULTS ON $s(n, k)$

Depth k	Size $s(n, k)$
2	$\Theta\left(\frac{n \log^2 n}{\log\log n}\right)$ [56]
3	$\Theta(n \log\log n)$ [57]
$2d, 2d+1, d \geq 2$	$\Theta(n\lambda(d, n))$ [36], [58]
In particular, for $k = 4, 5$	$\Theta(n \log^* n)$ [36], [58]
$\Theta(\beta(n))$	$\Theta(n)$ [36]

that $\overline{rc}(w, f) = \Omega(wf \log(wf)$; hence, if $\overline{rd}(w, f) = O(wf \log(wf))$ then both $\overline{rc}(w, f)$ and $rd_2(w, f)$ are $\Theta(wf \log(wf)$, which is a very strong result and turns out to hold true as we shall see later. We will take a different path to prove this important result.

Using the classic construction we cannot give very good bounds for the limited depth case. To get a linear-size generalizer, we already have to use at least a logarithmic depth. We could, in principle, use the limited depth constructions of $[w, f]$-connectors given in [13] for the second component of the construction; and, construct generalizer with given depth and size larger than linear. However, it is quite messy to get this idea to work.

In what follows, we will give a new recursive construction for $[w, f]_2$-distributors, which can be used to prove strong upper bounds for $rd_2(w, f, k)$ and to give asymptotically optimal $rd_2(w, f)$ as well as $\overline{rc}(w, f)$. For preparation, we will need the classic constructs of superconcentrators and concentrators.

Recall an (n, m)-*concentrator* is an (n, m)-network where $n \geq m$, such that for any subset S of m inputs there exists a set of m vertex disjoint paths connecting S to the outputs. Let $c(n, m)$ and $c(n, m, k)$ denote the minimum sizes of an (n, m)-concentrator and an (n, m)-concentrator of depth k, respectively. An n-*superconcentrator* is an n-network with inputs A and outputs B such that for any $S \subseteq A$ and $T \subseteq B$ with $|S| = |T| = c$, there exist a set of c vertex disjoint paths connecting vertices in S to vertices in T. Let $s(n)$ and $s(n, k)$ denote the minimum sizes of an (n, m)-superconcentrator and an (n, m)-superconcentrator of depth k, respectively. Removing $n - m$ outputs from an n-superconcentrator gives an (n, m)-concentrator. Hence,

$$c(n, m) \leq s(n), \tag{13}$$
$$c(n, m, k) \leq s(n, k). \tag{14}$$

It has been known for more than 3 decades that there are concentrators and superconcentrators of linear size [54], [55]. The constructions were based on *expanders*, whose applications in various areas of mathematics and computer science are numerous. For the fixed depth case, the asymptotic behaviors of all the $s(n, k)$ were only completely devised recently. Table II summarizes the results. The function $\lambda(d, n)$ is the inverse of functions in the Ackerman hierarchy: they are increasing extremely slowly. They can be defined as follows. Let

$$\log^* n := \min\{l \geq 0 \mid \underbrace{\log \ldots \log n}_{l} \leq 1\}$$

where the logarithms are to base 2. By induction on k, define

$$\lambda(d, n) := \log^{\overbrace{* \cdots *}^{d-1}} n := \min\{l \geq 0 \mid \underbrace{\log^{\overbrace{* \cdots *}^{d-2}} \ldots \log^{\overbrace{* \cdots *}^{d-2}}}_{l} n \leq 1\}$$

We shall use limited-depth concentrators to recursively construct $[w, f]_2$-distributors. For any k (n, m)-networks $\mathcal{N}_1, \ldots, \mathcal{N}_k$, let $\vdash (\mathcal{N}_1, \ldots, \mathcal{N}_k)$ denote the (n, mk)-network obtained by identifying the inputs of $\mathcal{N}_1, \ldots, \mathcal{N}_k$ in any one-to-one fashion. If the $\mathcal{N}_i$ are identical copies of the same (n, m)-network $\mathcal{N}$, then we use $\vdash^k \mathcal{N}$ to denote the result instead of writing $\vdash (\mathcal{N}, \ldots, \mathcal{N})$. Given an (n, m)-newtork $\mathcal{M}$ and a (m, l)-network $\mathcal{N}$, let $\mathcal{M} \circ \mathcal{N}$ be the network obtained by identifying the outputs of $\mathcal{M}$ and the inputs of $\mathcal{N}$ in any one-to-one fashion.

Now, let $n = wf$ and x be an integer divisible by f. Let $\mathcal{C}$ be a $(wf, wf/x)$-network and $\mathcal{M}$ be a $[w, f/x]$-network. Then,

$$\mathcal{N} = \vdash^x (\mathcal{C} \circ \mathcal{M})$$

is a $[w, f]$-network where the output bands of $\mathcal{N}$ are the union of the ouput bands of the x copies of $\mathcal{M}$. (See Figure 6, where $\mathcal{M}$ are distributors.)

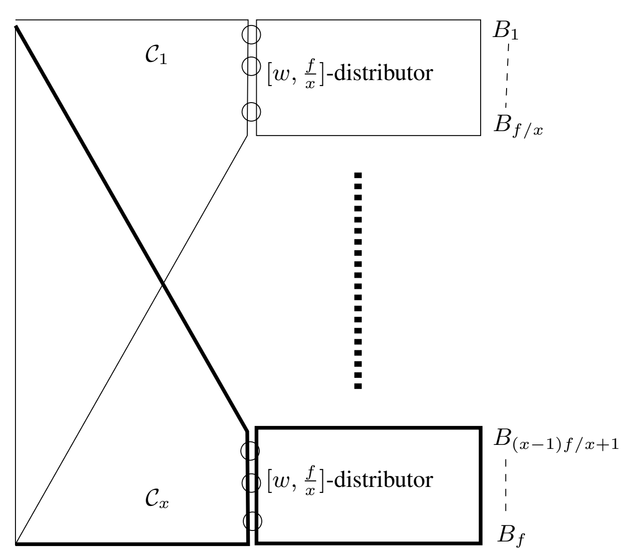

Fig. 6. Recursive construction of $[w, f]_2$-distributors.

Theorem VI.3. *Let x be an integer divisible by f. If $\mathcal{C}$ is a $(wf, wf/x)$-concentrator and $\mathcal{M}$ is a $[w, f/x]_2$-distributor, then $\mathcal{N} = \vdash^x (\mathcal{C} \circ \mathcal{M})$ is a $[w, f]_2$-distributor. In the special case when $x = f$, we let $\mathcal{M}$ be an empty network, namely $\mathcal{N} = \vdash^x \mathcal{C}$.*

Proof. Consider a distribution assignment $\mathcal{D}$ (under request model 2). Partition $\mathcal{D}$ into x subsets $\mathcal{D} = \mathcal{D}_1 \cup \cdots \cup \mathcal{D}_x$ as follows. For each request $D = (a, T) \in \mathcal{D}$ and each $i \in [x]$, let

$$T_i = T \cap \{(i-1)f/x + 1, (i-1)f/x + 2, \ldots, if/x\}$$

Then, add (a, T_i) into $\mathcal{D}_i$, unless $T_i = \emptyset$. For any $i \in [x]$

$$\begin{aligned} |\mathcal{D}_i| &\leq \sum_{(a,S)\in\mathcal{D}_i} |S| \\ &= \sum_{(a,T)\in\mathcal{D}} |T \cap \{(i-1)f/x+1, \ldots, if/x\}| \\ &= \sum_{j=(i-1)f/x+1}^{if/x} |\{(a,T) \in \mathcal{D} \ : \ j \in T\}| \\ &\leq \sum_{j=(i-1)f/x+1}^{if/x} w \\ &= wf/x \end{aligned}$$

The second inequality follows because $\mathcal{D}$ is a distribution assignment under model 2. For each $i \in [x]$, let A_i be the set of inputs of the requests in $\mathcal{D}_i$, namely

$$A_i = \{a \mid (a, S) \in \mathcal{D}_i \text{ for some } S\}$$

Define $d_i = |A_i| = |\mathcal{D}_i|$. For each $i \in [x]$, let $P^i_1, \ldots, P^i_{d_i}$ be d_i vertex-disjoint paths on the ith concentrator C_i from the inputs A_i to some set A'_i of d_i outputs of the the concentrator (which are also inputs of the ith $[w, f/x]_2$-distributor). If a is an input in A_i, let a' denote the other end of the path in $\{P^i_1, \ldots, P^i_{d_i}\}$ that starts from a.

Create a distribution assignment $\mathcal{D}'_i$ for the $[w, f/x]$-distributor by replacing each request $(a, S) \in \mathcal{D}_i$ by (a', T_1). Now, let $|\mathcal{R}|'_i$ be a network state of the $[w, f/x]_2$-distributor realizing $\mathcal{D}'_i$. Clearly, the union $\mathcal{R}_i$ of $\mathcal{R}'_i$ and the paths $P^i_1, \ldots, P^i_{d_i}$ is a network state realizing all the requests in $\mathcal{D}_i$. Finally, $\mathcal{R} = \mathcal{R}_1 \cup \cdots \cup \mathcal{R}_x$ is a network state realizing the original distribution assignment $\mathcal{D}$. □

Probably, the most important Corollary of the theorem is the following, which gives us the optimal asymptotic estimates for both $\overline{rc}(w, f)$ and $rd_2(w, f)$.

Corollary VI.4. *We have*

$$\overline{rc}(w, f) = \Theta(wf \log f), \quad (15)$$

$$rd_2(w, f) = \Theta(wf \log f). \quad (16)$$

Both are attainable using the construction shown in the previous theorem.

Proof. We will show by induction on f that there is a constant C so that $rd_2(w, f) \leq Cwf \log f$. This claim will prove both of the relations. Without loss of generality, we can assume that f is a power of 2. If not, replace f by the least power of 2 greater than f, which is not more than twice f. In which case, we only increase the constant a little bit.

Now, repeatedly apply the recursive construction by setting $x = 2$, using linear size concentrators at each step. We know that (n, m)-concentrators of size Cn exist for some constant C (which is about 20 for the current best construction [25]). The total size of our recursive construction is then

$$2Cwf + 4Cwf/2 + \cdots + 2^{\log f} Cwf/2^{\log f - 1} = Cwf \log f.$$

□

We next look at the limited-depth case. It is worth emphasizing again that we could have used the upper bounds for $rd(n, k)$ as upper bounds for $rd_2(w, f, k)$. Indeed, many of the upper bounds for $rd(n, k)$ are also good for $rd_2(w, f, k)$ when w is relatively small compared to f. Here, being "good" means that the bounds are close to the corresponding lower bound. For instance, the upper bound $O(wf \log(wf))$ would be just as good as $O(wf \log f)$ when $w \leq f^c$ for some fixed constant c. In what follows, we will establish upper bounds for $rd_2(w, f, k)$ keeping in mind that $w \gg f$ as is the case in practice. Theoretically, we will choose the smaller of the two kinds of bounds.

Corollary VI.5. *We have*

$$rd_2(w, f, 2) = O\left(wf^2 \frac{\log^2(wf)}{\log\log(wf)}\right) \quad (17)$$

$$rd_2(w, f, 3) = O\left(wf^2 \log\log(wf)\right) \quad (18)$$

$$rd_2(w, f, 4) = O\left(wf^{3/2} \frac{\log^2(wf)}{\log\log(wf)}\right) \quad (19)$$

$$rd_2(w, f, 5) = O\left(wf^{3/2} \frac{\log^2(wf)}{\log\log(wf)}\right) \quad (20)$$

Proof. In the depth-2 and depth-3 case we apply our recursive construction with $x = f$ depth-2 or depth-3 (wf, w)-concentrators (and thus we do not need to use $[w, f/x]_2$-distributors). Upper bounds (17) and (18) follow immediately.

In the depth-4 case we apply our recursive construction with $x = \sqrt{f}$ depth-2 concentrators in the first stage. In the second stage we use depth-2 $[w, f/x]_2$-distributors just constructed to show relation (17). In total, the size of our depth-4 construction is

$$\sqrt{f} \cdot \Theta\left(wf \frac{\log^2(wf)}{\log\log(wf)}\right) + \sqrt{f} \cdot \Theta\left(w(\sqrt{f})^2 \frac{\log^2(w\sqrt{f})}{\log\log(w\sqrt{f})}\right) = O\left(wf^{3/2} \frac{\log^2(wf)}{\log\log(wf)}\right).$$

The equality follows from the fact that $\log^2 n / \log\log n$ is an increasing function. The reader might be suspicious of the case when $\sqrt{f}$ is not an integer. We can resolve this situation as follows. Note that, for any integer $f' \geq f$, we can obtain a $[w, f]_2$-distributor from a $[w, f']_2$-distributor by removing $(wf' - wf)$-inputs and $(f' - f)$ output bands. Now, let

$$f' = 4^{\lceil \log_4 f \rceil} \leq 4f$$

and apply the above construction to get a $[w, f']_2$-distributor. This is possible since $\sqrt{f'}$ is now an integer. The result follows since $O\left(wf'^{3/2} \frac{\log^2(wf')}{\log\log(wf')}\right)$ is also $O\left(wf^{3/2} \frac{\log^2(wf)}{\log\log(wf)}\right)$ because $f' \leq 4f$. We will not repeat this point further in later constructions.

The depth-5 case can be done similarly, using $\sqrt{f}$ depth-2 concentrators in the first stage and depth-3 $[w, \sqrt{f}]_2$-distributors in the second stage. □

The previous corollary can be used as the base for constructing larger depth distributors. The following corollary can be made better with finer case-by-case analysis. In the case of depth-6 distributors, for example, we could use $x = f^{1/3}$ depth-2 concentrators in the first stage and x depth-4 $[w, f^{2/3}]_2$-distributors in the second stage, resulting in an upper bound of

$$rd_2(w, f, 6) = O\left(wf^{1+1/3}\frac{\log^2(wf)}{\log\log(wf)}\right). \tag{21}$$

Or, we could use $x = \sqrt{f}$ depth-3 concentrators in the first stage and x depth-3 $[w, \sqrt{f}]_2$-distributors in the second stage, resulting in an upper bound of

$$rd_2(w, f, 6) = O\left(wf^{1+1/2}\log\log(wf)\right). \tag{22}$$

Because

$$\log\log(wf) < \frac{\log^2(wf)}{\log\log(wf)},$$

(21) is not necessarily better than (22). Instead of doing case-by-case analysis, we shall give an upper bound whose formula is "cleanest" and sufficient to convey the main idea.

Corollary VI.6. *For any integer $k \geq 4$,*

$$rd_2(w, f, k) = O\left(\lceil k/4\rceil wf^{1+1/\lfloor k/4\rfloor}\log^*(wf)\right) \tag{23}$$

Proof. We prove this by induction. When $k = 4, 5, 6, 7$ we can directly use $x = f$ concentrators of size $O(wf\log^* wf)$ in the first stage and empty distributors in the second stage. Let $j = \lfloor k/4\rfloor$. When $k \geq 8$, use $x = f^{1/j}$ depth-4 concentrators in the first stage, and x $[w, f^{1-1/j}]_2$-distributors in the second stage. Recursively use depth-4 concentrators for the second stage distributors until the depth is between 4 and 7, in which case we apply the base-case construction. □

Certainly we do not have to use depth-4 concentrators as a basis for induction. We could have used depth-2, 3, or even depth d for any fixed integer d. The effect of doing so would be to increase of reduce the power of f and the logarithmic term. Table III summarizes the bounds for rd_1 and rd_2 shown in this paper.

TABLE III
A SUMMARY OF BOUNDS SHOWN IN THIS PAPER FOR rd_1 AND rd_2

Function	Lower bound	Upper bound
$rd_1(w, f)$	$\Theta(wf\log wf)$	
$rd_1(w, f, k)$	$k(wf)^{1+\frac{1}{k}}$	$O((wf\log wf)^{1+\frac{1}{k}})$ [36]
$rd_2(w, f)$	$\Theta(wf\log f)$	
$rd_2(w, f, 2)$	$2wf^{1+1/2}$	$O\left(wf^2\frac{\log^2(wf)}{\log\log(wf)}\right)$
$rd_2(w, f, 3)$	$3wf^{1+1/3}$	$O\left(wf^2\log\log(wf)\right)$
$rd_2(w, f, 4)$	$4wf^{1+1/4}$	$O\left(wf^{1+1/2}\frac{\log^2(wf)}{\log\log(wf)}\right)$
$rd_2(w, f, 5)$	$5wf^{1+1/5}$	$O\left(wf^{1+1/2}\frac{\log^2(wf)}{\log\log(wf)}\right)$
$rd_2(w, f, 6)$	$6wf^{1+1/6}$	$O\left(wf^{1+1/3}\frac{\log^2(wf)}{\log\log(wf)}\right)$ $O\left(wf^{1+1/2}\log\log(wf)\right)$
$rd_2(w, f, 7)$	$7wf^{1+1/7}$	$O\left(wf^{1+1/3}\frac{\log^2(wf)}{\log\log(wf)}\right)$ $O\left(wf^{1+1/2}\log\log(wf)\right)$
$rd_2(w, f, k)$, $k \geq 8$	$kwf^{1+1/k}$	$O\left(wf^{1+1/\lfloor k/4\rfloor}\log^*(wf)\right)$ $O\left(wf^{1+1/\lfloor k/3\rfloor}\log\log(wf)\right)$ $O\left(wf^{1+1/\lfloor k/2\rfloor}\frac{\log^2(wf)}{\log\log(wf)}\right)$
$rd_2(w, f, k)$, $k \geq d$, d fixed	$kwf^{1+1/k}$	$O\left(wf^{1+1/\lfloor k/d\rfloor}\lambda(d, n)\right)$

VII. DISCUSSIONS

There are several ways to realize the theoretical recursive construction presented in the previous section into an actual multicast WXC. The basic component is obviously an (n, m)-concentrator of a specific depth. If we know the precise structure of a concentrator, it is easy to construct an "optical concentrator." For example, a $(4, 2)$-concentrator can be constructed as shown in Figure VII. The "size" complexity in our sense translates directly into the number of SOAs in such construction. One can easily compare our constructions with existing ones (see also the table at the end of [18]). There is a small difference in the number of stages as defined here (i.e. the "depth") and number of stages as defined in that paper (the number of columns of optical components). Table IV gives the precise comparisons of constructions given in this paper and known multicast OXC constructions.

All current constructions of good concentrators are based on the existence of expanders. Good expanding graphs are often shown to exist using the probabilistic method [59], [60]. Fortunately, a few good explicit constructions of expanders are known (the Margulis-Gabber-Galil construction [61]–[63], and the recent Zig-Zag product [64]). Unfortunately, finding vertex disjoint paths on distributors constructed using expanders are not at all simple. Consequently, one of the major future research direction is to give constructions of $[w, f]$-distributors (or $[w, f]$-connectors for that matter) whose routing algorithm is simple and effective, so that truly fast and cost-effective WXCs can be designed based on these constructions.

REFERENCES

[1] Rajiv Ramaswami and Kumar Sivarajan, *Optical Networks: A Practical Perspective (Second Edition)*, Morgan Kaufmann Publishers, San Francisco, CA, 2001.

[2] Thomas E. Stern and Kirshna Bala, *Multiwavelength Optical Networks: A Layered Approach*, Prentice Hall PTR, Upper Saddle River, NJ, 1999.

[3] B. Mukherjee, *Optical Communication Networks*, McGraw-Hill, New York, NY, 1997.

[4] Lucent Technologies Press Release, "Lucent Technologies unveils untra-high-capacity optical system; Time Warner Telecom first to announce it will deploy the system," 2001, http://www.lucent.com/press/0101/010117.nsa.html.

[5] 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.

[6] Lucent Technologies Website, "What is dense wave division multiplexing (DWDM)," 2002, http://www.bell-labs.com/technology/lightwave/dwdm.html.

[7] Amit Singhal and Raj Jain, "Terabit switching: a survey of techniques and current products," *Computer communications*, vol. 25, no. 8, pp. 547–556, 2002.

[8] G. Wilfong, B. Mikkelsen, C. Doerr, and M. Zirngibl, "WDM cross-connect architectures with reduced complexity," *Journal of Lightwave Technology*, vol. 17, no. 10, pp. 1732–1741, Oct 1999.

TABLE IV
THE COST COMPARISON BETWEEN DIFFERENT CONSTRUCTIONS. FOR AWGR, $x \otimes y$ MEANS x AWGRS WITH SIZE y.

	#SOAs	#WCs	#AWGRs	Depth	Cost of Graph
ICC-RNB-1 $f \geq w$	$O(f^2 w)$	$O(f^2 w)$	$f \otimes fw$	8	$O(f^2 w^2)$
ICC-RNB-1 $f < w$	$O(fw^2)$	$O(f^2 w)$	$f \otimes f^2$	8	$O(fw^3)$
ICC-RNB-2	$O(f^2 w)$	$O(f^{\frac{3}{2}} w)$	$O(\sqrt{f} \otimes \sqrt{f} w)$	8	$O(f^2 w^2)$
Ref [17]	$O(f^{\frac{3}{2}} w^2 \frac{\lg f}{\lg \lg f})$	$O(fw)$	N/A	12	$O(f^{\frac{3}{2}} w^3 \frac{\lg f}{\lg \lg f})$
$rd_2(w,f)$	$O(wf \log f)$	$O(fw)$	N/A	$O(\log f \log(wf))$	$\Theta(w^2 f + wf \log f)$
$rd_2(w,f,2)$	$O\left(wf^2 \frac{\log^2(wf)}{\log\log(wf)}\right)$	$O(fw)$	N/A	6	$O\left(w^2 f + wf^2 \frac{\log^2(wf)}{\log\log(wf)}\right)$
$rd_2(w,f,3)$	$O\left(wf^2 \log\log(wf)\right)$	$O(fw)$	N/A	7	$O\left(w^2 f + wf^2 \log\log(wf)\right)$
$rd_2(w,f,4)$	$O\left(wf^{1+1/2} \frac{\log^2(wf)}{\log\log(wf)}\right)$	$O(fw)$	N/A	8	$O\left(w^2 f + wf^{1+1/2} \frac{\log^2(wf)}{\log\log(wf)}\right)$
$rd_2(w,f,5)$	$O\left(wf^{1+1/2} \frac{\log^2(wf)}{\log\log(wf)}\right)$	$O(fw)$	N/A	9	$O\left(w^2 f + wf^{1+1/2} \frac{\log^2(wf)}{\log\log(wf)}\right)$
$rd_2(w,f,6)$	$O\left(wf^{1+1/3} \frac{\log^2(wf)}{\log\log(wf)}\right)$ $O\left(wf^{1+1/2} \log\log(wf)\right)$	$O(fw)$	N/A	10	$O\left(w^2 f + wf^{1+1/3} \frac{\log^2(wf)}{\log\log(wf)}\right)$ $O\left(w^2 f + wf^{1+1/2} \log\log(wf)\right)$
$rd_2(w,f,7)$	$O\left(wf^{1+1/3} \frac{\log^2(wf)}{\log\log(wf)}\right)$ $O\left(wf^{1+1/2} \log\log(wf)\right)$	$O(fw)$	N/A	11	$O\left(w^2 f + wf^{1+1/3} \frac{\log^2(wf)}{\log\log(wf)}\right)$ $O\left(w^2 f + wf^{1+1/2} \log\log(wf)\right)$
$rd_2(w,f,k)$	$O\left(wf^{1+1/\lfloor k/4 \rfloor} \log^*(wf)\right)$ $O\left(wf^{1+1/\lfloor k/3 \rfloor} \log\log(wf)\right)$ $O\left(wf^{1+1/\lfloor k/2 \rfloor} \frac{\log^2(wf)}{\log\log(wf)}\right)$	$O(fw)$	N/A	$k+4$	$O\left(w^2 f + wf^{1+1/\lfloor k/4 \rfloor} \log^*(wf)\right)$ $O\left(w^2 f + wf^{1+1/\lfloor k/3 \rfloor} \log\log(wf)\right)$ $O\left(w^2 f + wf^{1+1/\lfloor k/2 \rfloor} \frac{\log^2(wf)}{\log\log(wf)}\right)$

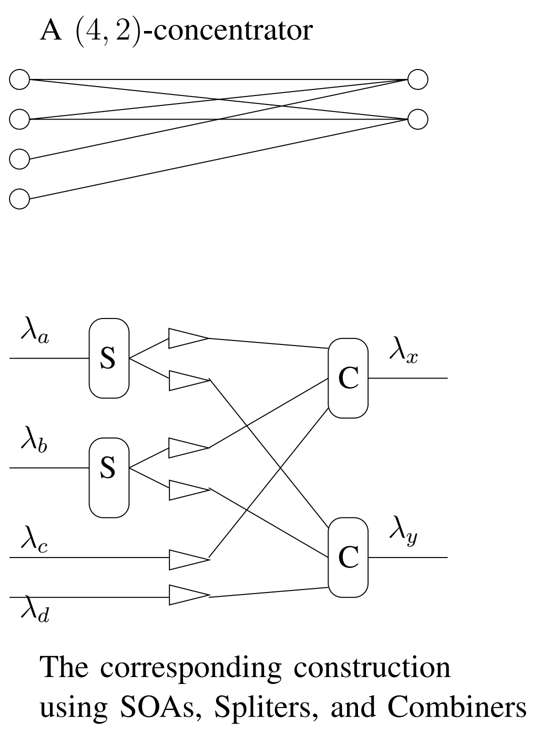

Fig. 7. Sample realization of a concentrator. S is a splitter. C is a combiner.

[9] Hung Q. Ngo, Dazhen Pan, and Chunming Qiao, "Constructions and analyses of nonblocking wdm switches based on arrayed waveguide grating and limited wavelength conversion," *IEEE/ACM Transactions on Networking*, vol. 14, no. 1, pp. 205–217, 2006.

[10] April Rasala and Gordon Wilfong, "Strictly non-blocking WDM cross-connects," in *Proceedings of the Eleventh Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'2000, San Francisco, CA)*, New York, 2000, pp. 606–615, ACM.

[11] Jeyashankher Ramamirtham and Jonathan S. Turner, "Design of wavelength converting switches for optical burst switching," in *Proceedings of the 21st Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*. 2002, vol. 2, pp. 1162–1171, IEEE.

[12] M.M. Vaez and C.-T. Lea, "Strictly nonblocking directional-coupler-based switching networks under crosstalk constraint," *IEEE Trans. Comm.*, vol. 48, no. 2, pp. 316–323, Feb 2000.

[13] Hung Q. Ngo, "WDM switching networks, rearrangeable and nonblocking $[w, f]$-connectors," *SIAM Journal on Computing*, vol. 35, no. 3, pp. 766–785, 2005-2006.

[14] G. Maier and A. Pattavina, "Design of photonic rearrangeable networks with zero first-order switching-element-crosstalk," *IEEE Trans. Comm.*, vol. 49, no. 7, pp. 1248–1279, Jul 2001.

[15] Hung Q. Ngo, Dazhen Pan, and Yuanyuan Yang, "Optical switching networks with minimum number of limited range wavelength converters," in *Proceedings of the 24rd Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, Miami, Florida, U.S.A., March 2005, vol. 2, pp. 1128–1138, IEEE.

[16] Hung Q. Ngo, Dazhen Pan, and Yuanyuan Yang, "Optical switching networks with minimum number of limited range wavelength converters," *IEEE/ACM Transactions on Networking*, vol. 15, no. 4, pp. 969–979, 2007.

[17] 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.

[18] Dazhen Pan, Vishal Anand, and Hung Q. Ngo, "Cost-effective constructions for nonblocking wdm multicast switching networks," in *Proceedings of the 2004 IEEE International Conference on Communications (ICC)*, Paris, France, June 2004, vol. 3, pp. 1801–1805, IEEE.

[19] Xiaohong Jiang, Hong Shen, Md. Mamun ur Rashid Khandker, and Susumu Horiguchi, "Blocking behaviors of crosstalk-free optical banyan networks on vertical stacking," *IEEE/ACM Transactions on Networking*, vol. 11, no. 6, pp. 982–993, 2003.

[20] Feng Cao and Ding-Zhu Du, "Fault-tolerant routing and multicasting in butterfly networks," in *SAC '99: Proceedings of the 1999 ACM symposium on Applied computing*. 1999, pp. 455–460, ACM Press.

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

[22] Nicholas Pippenger, "Communication networks," in *Handbook of theoretical computer science, Vol. A*, pp. 805–833. Elsevier, Amsterdam, 1990.

[23] 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.

[24] Frank K. Hwang, *The mathematical theory of nonblocking switching networks*, World Scientific Publishing Co. Inc., River Edge, NJ, 1998.
[25] 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., vol. 5 of *Network Theory and Applications*, pp. 307–367. Kluwer Academic Publishers, 2001.
[26] Hung Q. Ngo, "Multiwavelength distribution networks," in *Proceedings of the 2004 Workshop on High Performance Switching and Routing (HPSR)*, Phoenix, Arizona, U.S.A., April 2004, pp. 186–190, IEEE.
[27] Yao Li and Ting Wang, "Optical nonblocking multicast and broadcast interconnects using multidimensional multiplexing concepts," *J. Lightwave Technol.*, vol. 14, no. 2, pp. 129–138, Feb 1996.
[28] Chunling Zhou and Yuanyuan Yang, "Wide-sense nonblocking multicast in a class of regular optical WDM networks," *IEEE Trans. Comm.*, vol. 50, no. 1, pp. 126–134, Jan 2002.
[29] Xiangdong Qin and Yuanyuan Yang, "A cost-effective construction for WDM multicast switching networks," in *Proceedings of the 2002 IEEE International Conference on Communications (ICC 2002)*, New York, 1979, vol. 5, pp. 2902–2906, IEEE.
[30] R.K. Pankaj, "Wavelength requirements for multicasting in all-optical networks," *IEEE/ACM Trans. Networking*, vol. 7, no. 3, pp. 414–424, Jun 1999.
[31] Achille Pattavina and Gino L. Tesei, "Modelling the blocking behavior of multicast clos networks," in *Proceedings of the 22th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*. 2003, IEEE.
[32] Achille Pattavina and Gino L. Tesei, "Multicast nonblocking switching networks," *IEEE Trans. Comm.*, vol. 50, no. 8, pp. 1240–1243, Aug 2002.
[33] Douglas B. West, *Introduction to graph theory*, Prentice Hall Inc., Upper Saddle River, NJ, 1996.
[34] Nicholas Pippenger, "Generalized connectors," *SIAM J. Comput.*, vol. 7, no. 4, pp. 510–514, 1978.
[35] Ju. P. Ofman, "A universal automaton," *Trudy Moskov. Mat. Obšč.*, vol. 14, pp. 186–199, 1965.
[36] Danny Dolev, Cynthia Dwork, Nicholas Pippenger, and Avi Wigderson, "Superconcentrators, generalizers and generalized connectors with limited depth (preliminary version)," in *Proceedings of the Fifteenth Annual ACM Symposium on Theory of Computing*, Boston, Massachusetts, apr 1983, pp. 42–51.
[37] Gerald M. Masson and Jr. B. W. Jordan, "Generalized multi-stage connection networks," *Networks*, vol. 2, pp. 191–209, 1972.
[38] David Nassimi and Sartaj Sahni, "Parallel permutation and sorting algorithms and a new generalized connection network," *J. Assoc. Comput. Mach.*, vol. 29, no. 3, pp. 642–667, 1982.
[39] David G. Kirkpatrick, Maria Klawe, and Nicholas Pippenger, "Some graph-colouring theorems with applications to generalized connection networks," *SIAM J. Algebraic Discrete Methods*, vol. 6, no. 4, pp. 576–582, 1985.
[40] Richard M. Karp, "Three-stage generalized connectors," *SIAM J. Discrete Math.*, vol. 5, no. 2, pp. 259–272, 1992.
[41] Paul Feldman, Joel Friedman, and Nicholas Pippenger, "Wide-sense nonblocking networks," *SIAM J. Discrete Math.*, vol. 1, no. 2, pp. 158–173, 1988.
[42] Nicholas Pippenger, *The complexity of switching networks*, Ph.D. thesis, Department of Electrical Engineering, Massachusetts Institute of Technology, Cambridge, Massachusetts, 1973.
[43] Avi Wigderson and David Zuckerman, "Expanders that beat the eigenvalue bound: explicit construction and applications," *Combinatorica*, vol. 19, no. 1, pp. 125–138, 1999.
[44] Amnon Ta-Shma, "On extracting randomness from weak random sources (extended abstract)," in *Proceedings of the Twenty-eighth Annual ACM Symposium on the Theory of Computing (Philadelphia, PA, 1996)*, New York, 1996, pp. 276–285, ACM.
[45] Amnon Ta-Shma, "Almost optimal dispersers," in *STOC '98 (Dallas, TX)*, pp. 196–202. ACM, New York, 1999.
[46] Amnon Ta-Shma, David Zuckerman, and Shmuel Safra, "Extractors from Reed-Muller codes," in *42nd IEEE Symposium on Foundations of Computer Science (Las Vegas, NV, 2001)*, pp. 638–647. IEEE Computer Soc., Los Alamitos, CA, 2001.
[47] Noam Nisan and Amnon Ta-Shma, "Extracting randomness: a survey and new constructions," *J. Comput. System Sci.*, vol. 58, no. 1, part 2, pp. 148–173, 1999, Twenty-eighth Annual ACM Symposium on the Theory of Computing (Philadelphia, PA, 1996).
[48] Amnon Ta-Shma, "Almost optimal dispersers," *Combinatorica*, vol. 22, no. 1, pp. 123–145, 2002.
[49] Amnon Ta-Shma, "Storing information with extractors," *Inform. Process. Lett.*, vol. 83, no. 5, pp. 267–274, 2002.
[50] Aravind Srinivasan and David Zuckerman, "Computing with very weak random sources," *SIAM J. Comput.*, vol. 28, no. 4, pp. 1433–1459 (electronic), 1999.
[51] Leonid A. Bassalygo and Mark S. Pinsker, "Asymptotically optimal networks for generalized rearrangeable switching and generalized switching without rearrangements," *Problemy Peredači Informacii*, vol. 16, pp. 94–98, 1980.
[52] Nicholas Pippenger and Andrew C. C. Yao, "Rearrangeable networks with limited depth," *SIAM J. Algebraic Discrete Methods*, vol. 3, no. 4, pp. 411–417, 1982.
[53] J. H. Hui, *Switching and traffic theory for integrated broadband networks*, Kluwer Academic Publishers, Boston/Dordrecht/London, 1990.
[54] Leslie G. Valiant, "Graph-theoretic properties in computational complexity," *J. Comput. System Sci.*, vol. 13, no. 3, pp. 278–285, 1976, Working papers presented at the ACM-SIGACT Symposium on the Theory of Computing (Albuquerque, N. M., 1975).
[55] Leslie G. Valiant, "On non-linear lower bounds in computational complexity," in *Seventh Annual ACM Symposium on Theory of Computing (Albuquerque, N. M., 1975)*. 1975, pp. 45–53, Assoc. Comput. Mach., New York.
[56] Jaikumar Radhakrishnan and Amnon Ta-Shma, "Bounds for dispersers, extractors, and depth-two superconcentrators," *SIAM J. Discrete Math.*, vol. 13, no. 1, pp. 2–24 (electronic), 2000.
[57] Noga Alon and Pavel Pudlák, "Superconcentrators of depths 2 and 3; odd levels help (rarely)," *J. Comput. System Sci.*, vol. 48, no. 1, pp. 194–202, 1994.
[58] Pavel Pudlák, "Communication in bounded depth circuits," *Combinatorica*, vol. 14, no. 2, pp. 203–216, 1994.
[59] Joel Friedman, "A proof of Alon's second eigenvalue conjecture," in *Proceedings of the Thirty-Fifth Annual ACM Symposium on Theory of Computing*, New York, 2003, pp. 720–724 (electronic), ACM.
[60] Mark S. Pinsker, "On the complexity of a concentrator," in *Proceedings of the 7th International Teletraffic Conference, Stockholm, June 1973*, 1973, pp. 318/1–318/4.
[61] Ofer Gabber and Zvi Galil, "Explicit constructions of linear size superconcentrators," in *20th Annual Symposium on Foundations of Computer Science (San Juan, Puerto Rico, 1979)*, pp. 364–370. IEEE, New York, 1979.
[62] S. Jimbo and A. Maruoka, "Expanders obtained from affine transformations," *Combinatorica*, vol. 7, no. 4, pp. 343–355, 1987.
[63] G. A. Margulis, "Explicit group-theoretic constructions of combinatorial schemes and their applications in the construction of expanders and concentrators," *Problemy Peredachi Informatsii*, vol. 24, no. 1, pp. 51–60, 1988.
[64] Omer Reingold, Salil Vadhan, and Avi Wigderson, "Entropy waves, the zig-zag graph product, and new constant-degree expanders," *Ann. of Math. (2)*, vol. 155, no. 1, pp. 157–187, 2002.

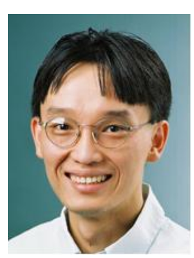

Hung Q. Ngo received his B.S. degree in Computer Engineering from Ho Chi Minh city Univ. of Tech., Viet Nam, in 1995, M.S. degree in Mathematics and Ph.D. degree in Computer Science from the Univ. of Minnesota at Twin Cities in 2001. He has been an assistant professor at SUNY Buffalo since. His research interests include networking, algorithms, and combinatorics. He has published about 50 papers in leading technical journals and conference proceedings. He has served in various NSF panels and technical program committees of professional conferences. His research has been supported by NSF and DARPA.

Yang Wang received his B.S and M.S degrees respectively in 2002 and 2005 from the Department of Mathematics, Zhejiang Univ., Hangzhou, China. He is working toward a Ph. D. degree in Computer Science and Engineering at the State University of New York at Buffalo. His research interests include the design and analysis of switching networks and optical switching networks.

Dazhen Pan Dazhen Pan received his B.S. degree in Computer Science in 1999 from Nankai University, China. He is currently working toward to Ph.D. degree in Computer Science and Engineering at State University of New York at Buffalo. His research is on complexity, constructions, and routing algorithms of optical switching networks.