

# Constructions of Given-Depth and Optimal Multirate Rearrangeably Nonblocking Distributors

Yang Wang, Hung Q. Ngo, and Thanh-Nhan Nguyen  
Computer Science and Engineering,

201 Bell Hall,

State University of New York at Buffalo,  
Amherst, NY 14260, USA.

{yw43, hungngo, nguyen9}@cse.buffalo.edu

**Abstract**—The theory of multirate switching networks, started in the late 80s, has been very practically useful. In particular, it has served as the theoretical foundation for the development of most Asynchronous Transfer Mode (ATM) switching systems.

Rearrangeable multirate multicast switching networks are customarily called rearrangeable multirate distributors. It has been known for more than 18 years that rearrangeable multirate distributors with cross-point complexity  $O(n \log^2 n)$  can be constructed, where  $n$  is the number of inputs (and outputs) of the switching network. The problem of constructing optimal distributors remains open thus far.

In this paper, we give a general construction of rearrangeable multirate distributors with given depths. One of the rewards of our construction is a rearrangeable multicast distributor with cross-point complexity  $O(n \log n)$ . We shall also show that this cross-point complexity is optimal, thus settling the aforementioned open problem.

One of the key ingredients of our new construction is the notion of *multirate concentrators*. The second ingredient is a multirate version of the Pippenger network, which is a rearrangeable multirate distributor recursively constructed based on multirate concentrators. We shall show how to construct given-depth multirate concentrators and given-depth multirate Pippenger networks with small sizes. When the depth is chosen to optimize the size, we obtain the aforementioned  $O(n \log n)$  cross-point complexity.

## I. INTRODUCTION

Multirate switching networks are switching networks that support varying bandwidth connections. The theory of multirate switching networks, perhaps started with the papers by Niestegge [28] and Melen and Turner [20], has proved to be very useful in practice. For example, this theory has served as the theoretical foundation for the development of most Asynchronous Transfer Mode (ATM) switching systems from major ATM equipment manufacturer [5], [34], [35]. Roughly speaking, as opposed to space switching where each internal or external link of a switch can only carry one connection request at a time, the multirate switches allow for many connections with varying *rates* (or *bandwidths*) to be carried on a single link, as long as the total connection rate does not exceed the link's capacity.

In the unicast case, one particularly fruitful line of research on multirate switching networks has been on the multirate rearrangeability of the Clos network [4], represented by the (still open) conjecture by Chung and Ross in 1991 [3] which states

that the Clos network  $C(n, m, r)$  is multirate rearrangeably nonblocking when the number  $m$  of middle-stage switches is at least  $2n - 1$ . This conjecture is interesting because it points towards a possible generalization of the König's theorem for edge coloring bipartite graphs. Later developments on this conjecture and related problems were reported in [6], [8], [11], [19], [23], [26]. See also [9], [18] for several related lines of research.

In the multicast and broadcast cases, there have been notably few known results, though. The works presented in [15], [16], [18], [39] concern conditions for the Clos network to be multicast capable. The study presented in [21] (the journal version is [22]) was the only one that deals directly with more general constructions and complexities of multicast multirate switching networks. In their paper, using the Pippenger's network [30], the authors constructed a rearrangeable multirate distributor with cross-point complexity  $O(n \log^2 n)$ , where  $n$  is the number of inputs (and outputs) of the switching network. (*Distributor*, also called *distribution network* or *generalized connector*, is a standard name referring to multicast switching networks in the switching network literature [12], [13].)

The problem of constructing optimal-size rearrangeable multirate distributors remains open thus far. In this paper, we give a general construction of rearrangeable multirate distributors with given depths. (The notions of depths and rearrangeability shall be defined rigorously in a later section.) One of the rewards of our construction is a rearrangeable multicast distributor with cross-point complexity  $O(n \log n)$ . We shall also show that this size is optimal, thus settling the aforementioned open problem.

One of the key ingredients of our new construction is the notion of *multirate concentrators*. The second ingredient is a multirate version of the Pippenger network, which is a rearrangeable multirate distributor recursively constructed based on multirate concentrators. We shall show how to construct given-depth multirate concentrators and given-depth multirate Pippenger networks with small sizes. When the depth is chosen to optimize the size, we obtain the aforementioned  $O(n \log n)$  cross-point complexity.

The rest of the paper is organized as follows. Section II presents basic definitions and several fundamental compositions of networks. Section III defines a particular type of multirate concentrators, which is crucial in constructing

multirate distributors. Section IV contains the main results, including a general given-depth construction. The construction gives rise to a rearrangeable multirate distributor of size  $O(n \lg n)$  which is then shown to be optimal. Lastly, Section V concludes the paper with a few remarks and discussions on future works.

## II. PRELIMINARIES

### A. Multirate networks

In the rest of the paper, let  $[m] = \{1, \dots, m\}$  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 [38], for instance).

An  $(n_1, n_2)$ -network is a directed acyclic graph (DAG)  $\mathcal{N} = (V, E; X, Y)$ , where  $V$  is the set of vertices,  $E$  is the set of edges,  $X$  is a set of  $n_1$  nodes called *inputs*, and  $Y$  – disjoint from  $X$  – is a set of  $n_2$  nodes called *outputs*. The vertices in  $V - X \cup Y$  are *internal* vertices. The in-degrees of the inputs and the out-degrees of the outputs are zero. The *size* of a network is its number of edges. The size of a network is the equivalence of the *cross-point complexity* of a switch. The DAG model is standard for studying the complexity of switching networks [24], [25]. 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.

In the multirate environment, a constant  $\beta \leq 1$  is often used to represent the *capacity of the inputs and outputs* of  $\mathcal{N}$ . All internal nodes have capacity (normalized to) 1. The ratio  $1/\beta$  is often referred to as the *speed advantage* of the system. This *internal speedup* is a common technique for designing broadband switches [12], [13], [20].

Given an  $n$ -network  $\mathcal{N} = (V, E; X, Y)$ , a *distribution request* (or *multicast request*) is a triple

$$D = (x, S, w) \in X \times 2^Y \times [b, B].$$

The request is from input  $x$  to a subset  $S$  of outputs. The *weight* or *rate*  $w$  of the request satisfies  $b \leq w \leq B$  for some given lower- and upper-bounds  $0 \leq b < B \leq \beta \leq 1$ .

A *distribution assignment* is a set  $\mathcal{D}$  of requests satisfying the following conditions: (a) total weight of requests coming from any particular input does not exceed  $\beta$ , and (b) total request weight to any output does not exceed  $\beta$ ; to put it rigorously,

$$\begin{aligned} \sum_{(x, S, w) \in \mathcal{D}} w &\leq \beta, \quad \forall x \in X \\ \sum_{\substack{(x, S, w) \in \mathcal{D} \\ y \in S}} w &\leq \beta, \quad \forall y \in Y. \end{aligned}$$

A request  $D$  is *compatible* with a distribution assignment  $\mathcal{D}$  if and only if  $\mathcal{D} \cup \{D\}$  is also a distribution assignment.

A *distribution route*  $R$  for a request  $D = (x, S, w)$  is a (directed) tree rooted at  $x$  whose leaves are precisely the nodes in  $S$ . We also say that  $R$  *realizes*  $D$ , and call  $w$  the weight of  $R$ . A *state* of the  $n$ -network  $\mathcal{N}$  is a set  $\mathcal{R}$  of distribution routes,

where the total weight of routes containing any node does not exceed the capacity of that node. Each state of  $\mathcal{N}$  realizes a unique distribution assignment, one route per request. A distribution assignment  $\mathcal{D}$  is *realizable* if and only if 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 are now ready to define the fundamental notions of nonblockingness in the multirate environment. In defining different notions of distributors, we drop the “multirate” qualification to avoid being wordy. Distribution networks in this paper are implicitly understood as multirate distribution networks, unless specified otherwise.

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 if and only if 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.

Given the parameters  $\beta, b$ , and  $B$  as described above, let  $\text{mrd}_{\beta[b, B]}(n)$ ,  $\text{mwd}_{\beta[b, B]}(n)$ , and  $\text{msd}_{\beta[b, B]}(n)$  denote the minimum size of a multirate RNB, WSNB, and SNB  $n$ -distributor, respectively. In the given-depth case, let  $\text{mrd}_{\beta[b, B]}(n, k)$ ,  $\text{mwd}_{\beta[b, B]}(n, k)$ , and  $\text{msd}_{\beta[b, B]}(n, k)$  denote the minimum size of an RNB, WSNB, and SNB  $n$ -distributor with depth  $k$ , respectively.

In the special case when  $b = 0, B = \beta = 1$ , i.e. the case when there is no internal speedup and no request rate restriction, we will drop the subscripts  $\beta[0, B]$  and use  $\text{mrd}(\cdot)$ ,  $\text{mwd}(\cdot)$ ,  $\text{msd}(\cdot)$  to denote the corresponding functions.

### B. Classical networks

In constructing multirate distributors, we will also need the notions of (classical) concentrators and super-concentrators (see [12], [13], [25] for more details on these networks).

For positive integers  $n \geq m$ , recall that an  $(n, m)$ -concentrator is an  $(n, m)$ -network such that for any subset  $S$  of  $m$  inputs there exists a set of  $m$  vertex disjoint paths connecting  $S$  to the outputs. (It does not matter which particular outputs these disjoint paths connect the inputs in  $S$  to.) Let  $c(n, m)$  and  $c(n, m, k)$  denote the minimum sizes of an

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

Depth $k$	Size $s(n, k)$
2	$\Theta\left(\frac{n \log^2 n}{\log \log n}\right)$ [33]
3	$\Theta(n \log \log n)$ [1]
$2d, 2d+1, d \geq 2$	$\Theta(n \lambda(d, n))$ [7], [32]
In particular, for $k = 4, 5$	$\Theta(n \log^* n)$ [7], [32]
$\Theta(\alpha(n))$	$\Theta(n)$ [7]

$(n, m)$ -concentrator and an  $(n, m)$ -concentrator of depth  $k$ , respectively.

An  $(n, m)$ -*superconcentrator* is an  $(n, m)$ -network with inputs  $X$  and outputs  $Y$  such that for any  $S \subseteq X$  and  $T \subseteq Y$  with  $|S| = |T|$ , there are  $|S|$  paths which are vertex disjoint and which connect vertices in  $S$  to vertices in  $T$ . Thus, each vertex in  $S$  and  $T$  is an end point of exactly one of these paths. An  $n$ -*superconcentrator* is an  $(n, n)$ -superconcentrator. Let  $s(n)$  and  $s(n, k)$  denote the minimum sizes of an  $n$ -superconcentrator and an  $n$ -superconcentrator of depth  $k$ , respectively.

When  $n \geq m$ , an  $(n, m)$ -superconcentrator can be obtained from an  $n$ -superconcentrator by simply removing arbitrarily  $(n - m)$  outputs from the  $n$ -superconcentrator. Moreover, an  $(n, m)$ -superconcentrator is obviously an  $(n, m)$ -concentrator. (The converse does not necessarily hold.) Hence,

$$\begin{aligned} c(n, m) &\leq s(n), \\ c(n, m, k) &\leq s(n, k). \end{aligned}$$

Note that the concentrators and superconcentrators described above operate in the *space domain*, namely no two paths can share a vertex. Thus, the concentrators and superconcentrators defined above will also be referred to as space-domain concentrators and superconcentrators, in cases where we want to emphasize the domain they operate in.

When the network depth is not limited, it has been known for more than three decades that there are concentrators and superconcentrators of linear size [36], [37]. The constructions were based on *expanders*, whose applications in mathematics and computer science are numerous [10].

For the given-depth case, the asymptotic behaviors of all the  $s(n, k)$  were only completely devised recently. Table I 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

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

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

$$\begin{aligned} \lambda(d, n) &:= \log^{\overbrace{* \dots *}^{d-1}} n := \\ &\min\{l \geq 0 \mid \underbrace{\log^{\overbrace{* \dots *}^{d-2}} \dots \log^{\overbrace{* \dots *}^{d-2}} n}_l \leq 1\} \end{aligned}$$

The reader is referred to [7] for the definition of  $\alpha(n)$  (which is actually called  $\beta(n)$  in their paper, but we change its name to avoid confusion with our speedup parameter  $\beta$ ).

### C. Basic compositions of networks

Let  $\mathcal{N}_1$  and  $\mathcal{N}_2$  be any two  $(n, m)$ -networks. We use  $\mathcal{N}_1 \square \mathcal{N}_2$  to denote an  $(n, m)$ -network  $\mathcal{N}$  obtained by identifying the inputs of  $\mathcal{N}_1$  and  $\mathcal{N}_2$  in any one-to-one manner, and identifying the outputs of  $\mathcal{N}_1$  and  $\mathcal{N}_2$  in any one-to-one manner. See Figure 1(a) for an illustration. We refer to  $\mathcal{N}_1 \square \mathcal{N}_2$  as the *stacking* of  $\mathcal{N}_1$  and  $\mathcal{N}_2$ . When stacking  $k$  copies of a network  $\mathcal{N}$ , denote the result by  $\square^k \mathcal{N}$ .

Given any  $k$   $(n, m)$ -networks  $\mathcal{N}_1, \dots, \mathcal{N}_k$ , let  $\vdash (\mathcal{N}_1, \dots, \mathcal{N}_k)$ , called a *folding* of  $\mathcal{N}_1, \dots, \mathcal{N}_k$ , denote the  $(n, mk)$ -network obtained by identifying the inputs of  $\mathcal{N}_1, \dots, \mathcal{N}_k$  in any one-to-one fashion. In effect, we “paste” together the inputs of  $\mathcal{N}_1, \dots, \mathcal{N}_k$ . See Figure 1(b) for an illustration. When the  $\mathcal{N}_i$  are identical copies of the same  $(n, m)$ -network  $\mathcal{N}$ , we use  $\vdash^k \mathcal{N}$  to denote the result instead of writing  $\vdash (\mathcal{N}, \dots, \mathcal{N})$ .

Given an  $(n, m)$ -network  $\mathcal{M}$  and a  $(m, l)$ -network  $\mathcal{N}$ , a *concatenation* of  $\mathcal{M}$  and  $\mathcal{N}$ , denoted by  $\mathcal{M} \circ \mathcal{N}$ , is the  $(n, l)$ -network obtained by identifying the outputs of  $\mathcal{M}$  and the inputs of  $\mathcal{N}$  in any one-to-one fashion. (See Figure 1(c).)

## III. MULTIRATE CONCENTRATORS

There are several obvious ways to generalize the notion of space-domain concentrators to multirate concentrators. To avoid cumbersome notations, we will define here only a particular type of multirate concentrators which is used in the next section to construct multirate distributors.

Given positive integers  $n \geq m > 0$ . Consider an  $(n, m)$ -network  $\mathcal{C} = (V, E; X, Y)$ . A *multirate concentration request* is a pair  $(x, w)$ , where  $x$  is an input and  $0 < w \leq 1$  is the weight of the request. A path from  $x$  to some output is called a *route* realizing this request. A set of routes are *compatible* if the total weight of routes through any vertex is at most 1. (In other words, in the context of multirate concentrators we are defining, all nodes – including the inputs and outputs – have capacity 1.)

A *multirate concentration assignment* is a set of multirate concentration requests such that each input generates requests with total weight at most 1, and that the total weight of all requests is at most  $m/2$ . (The total weight limit of  $m/2$  is technically needed for the proof of Lemma III.1 to work. Thus, the limitation is due to technicality and it may seem a little unnatural at first.)

The network  $\mathcal{C}$  is called a *multirate  $(n, m)$ -concentrator* if and only if, given any multirate concentration assignment  $\mathcal{D}$  there exists a set of compatible routes realizing requests in the assignment.

**Lemma III.1.** *Let  $\mathcal{C}$  be any space-domain  $(n, m)$ -concentrator and  $\mathcal{S}$  be any space-domain  $(n, m)$ -superconcentrator. Then, the stacking  $\mathcal{C}(n, m) = \mathcal{C} \square \mathcal{S}$  is a multirate  $(n, m)$ -concentrator.*




Fig. 1. Illustrations of basic network compositions




Fig. 2. Construction of an  $(n, m)$ -multirate concentrator  $\mathcal{C}(n, m)$ .

*Proof.* The reader is referred to Figure 2 for an illustration of  $\mathcal{C}(n, m)$ . To prove this lemma, we will use a routing algorithm adapted from the CAP algorithm proposed in [20].

Let  $\mathcal{D}$  be any multirate concentration assignment. Note that the inputs of the multirate concentration requests in  $\mathcal{D}$  are not necessarily distinct. As long as there are still two requests  $(x, w_1)$  and  $(x, w_2)$  coming from the same input  $x$ , replace them by a new request  $(x, w_1 + w_2)$ . The new set of multirate concentration requests is still a (valid) concentration assignment. Moreover, a route realizing  $(x, w_1 + w_2)$  can be “decomposed” back into two routes with weights  $w_1$  and  $w_2$  to

satisfy the original requests  $(x, w_1)$  and  $(x, w_2)$ . Consequently, without loss of generality we can assume that the inputs of these multirate concentration requests are distinct.

Partition  $\mathcal{D}$  into two subsets  $\mathcal{D}_l$  (large requests) and  $\mathcal{D}_s$  (small requests), where  $\mathcal{D}_l$  consists of all requests with weights strictly greater than  $1/2$ , and  $\mathcal{D}_s$  consists of the rest of the requests. Let  $x = |\mathcal{D}_l|$  be the number of large requests and  $y = |\mathcal{D}_s|$  be the number of small requests. Respectively, let  $W_l$  and  $W_s$  be the total weights of requests in  $\mathcal{D}_l$  and  $\mathcal{D}_s$ .

Because  $\mathcal{D}$  is a multirate concentration assignment, we have

$$m/2 \geq W_l + W_s \geq x/2 + W_s. \quad (1)$$

The set  $\mathcal{D}_l$  of large requests can be routed through the space-domain  $(n, m)$ -concentrator  $\mathcal{C}$  so that no two routes share a vertex. Thus, the vertex capacity constraint is satisfied. Every output which is not involved in these  $x$  paths are said to be free. There are  $m - x$  free outputs.

Next, we route the small requests in  $\mathcal{D}_s$  through the space-domain superconcentrator  $\mathcal{S}$ . The small requests will be routed to the  $m - x$  free outputs.

First, divide the  $y$  small requests in  $\mathcal{D}_s$  into  $t = \lceil y/s \rceil$  groups of size  $s = m - x$  each, with possibly the last group having size  $< s$ . The division is done as follows. Without loss of generality, suppose the weights for these requests are  $w_1 \geq w_2 \geq \dots \geq w_y$ . Then, the first group consists of  $s$  largest weights  $w_1, \dots, w_s$ ; the second group consists of the next  $s$  largest weights  $w_{s+1}, \dots, w_{2s}$ ; and so forth.

Because  $s \leq m$ , for every group of requests there are  $s$  vertex disjoint paths in the space-domain  $(n, m)$ -superconcentrator  $\mathcal{S}$  which join the inputs of the requests in the group to the  $s$  free outputs. These paths are routes realizing the corresponding requests in the group. Consequently, no two routes for requests in the same group share any vertex.

To this end, we need to show that no vertex of  $\mathcal{S}$  carries routes with total weight exceeding 1. In the worst case, a vertex carries one request from each group. Thus, the maximum

weight a vertex is forced to carry is at most

$$\begin{aligned}
& w_1 + w_{s+1} + \dots + w_{(t-1)s+1} \\
& \leq \frac{1}{2} + \frac{w_1 + \dots + w_s}{s} + \dots + \frac{w_{(t-2)s+1} + \dots + w_{(t-1)s}}{s} \\
& \leq \frac{1}{2} + \frac{W_s}{s} \\
& \leq \frac{1}{2} + \frac{m/2 - x/2}{m - x} \\
& = 1.
\end{aligned}$$

The last inequality follows from (1).  $\square$

**Corollary III.2.** A  $(n, m)$ -multirate concentrator of depth  $k$  can be constructed with the same asymptotic complexity as  $s(n, k)$  shown in Table I.

*Proof.* This follows directly from the fact that, removing any  $n - m$  outputs from a space-domain  $n$ -superconcentrator yields an  $(n, m)$ -superconcentrator, which is also an  $(n, m)$ -concentrator.  $\square$

#### IV. REARRANGEABLE MULTIRATE DISTRIBUTORS

We will recursively make use of the multirate concentrator  $\mathcal{C}(n, m)$  constructed in the previous section to construct rearrangeable multirate distributors.

##### A. Distributors for the case $B \leq \beta \leq 1/2$

In this subsection, we show how to construct multirate distributors under the condition  $B \leq \beta \leq 1/2$ . In fact, we will construct a slightly stronger distributor, where the capacity of input nodes are allowed to be 1. Obviously, any distributor with capacity-1 inputs is also a distributor with capacity- $\beta$  inputs. The outputs' capacities remain equal to  $\beta$ , which is at most  $1/2$ .




Fig. 3. Recursive construction of multirate distributors with capacity-1 inputs.

In the following lemma, we ignore the issue of divisibility for the sake of clarity. It is simple but tedious to deal directly with divisibility. Basically, if  $n$  is not divisible by  $m$ , we can always find  $l \in \mathbb{N}$  such that  $lm < n < (l + 1)m$ . Then we can construct an  $n'$ -distributor where  $n' = (l + 1)m$ . Removing

arbitrarily  $n' - n$  inputs and  $n' - n$  outputs from an  $n'$ -distributor yields an  $n$ -distributor of the same asymptotic size, because  $n = \Theta(n')$ .

The following construction is the multirate version of Pippenger network [30].

**Lemma IV.1.** Let  $m$  be a factor of  $n$ . Let  $\mathcal{C}(n, m)$  be the multirate concentrator constructed in the previous section. Let  $\mathcal{M}$  be any multirate  $m$ -distributor with capacity-1 inputs. Then, the network

$$\mathcal{N} = \vdash^{n/m} (\mathcal{C}(n, m) \circ \mathcal{M})$$

is a multirate  $n$ -distributor with capacity-1 inputs. Note that, we only consider the case when  $b \leq B \leq \beta \leq 1/2$ .

*Proof.* The reader is referred to Figure 3 for an illustration of  $\mathcal{N}$ . For  $1 \leq i \leq n/m$ , let  $Y_i$  denote the set of outputs of the  $i$ th copy of the multirate  $m$ -distributor  $\mathcal{M}$ .

Consider a distribution assignment  $\mathcal{D}$ . Partition  $\mathcal{D}$  into  $n/m$  subsets  $\mathcal{D} = \mathcal{D}_1 \cup \mathcal{D}_2 \cup \dots \cup \mathcal{D}_{n/m}$  as follows. For each request  $D = (x, T, w) \in \mathcal{D}$  and  $i \in \{1, \dots, n/m\}$ , let

$$T_i = T \cap Y_i.$$

Then, if  $T_i \neq \emptyset$  we add the multirate distribution request  $(x, T_i, w)$  into  $\mathcal{D}_i$ . Note that, if we can find routes realizing all of the requests in  $\mathcal{D}_1, \dots, \mathcal{D}_{n/m}$ , then the natural union of those routes will realize  $\mathcal{D}$ . For example, to realize the request  $D$  above, simply take the union of the routes realizing  $(x, T_1, w), \dots, (x, T_{n/m}, w)$ .

The idea is to use the first concentrator and distributor to realize  $\mathcal{D}_1$ , the second concentrator and distributor for  $\mathcal{D}_2$ , and so on. Since the construction is symmetric, we only need to show how to construct routes realizing  $\mathcal{D}_1$ .

Firstly, note that the total weight of requests from  $\mathcal{D}_1$  is at most  $m/2$ , because there are at most  $m$  outputs involved in these requests, each with capacity  $\beta \leq 1/2$ . Thus, there are compatible routes in  $\mathcal{C}(n, m)$  joining each input  $x$  of a request  $(x, T_1, w)$  in  $\mathcal{D}_1$  to an output  $f(x)$  of  $\mathcal{C}(n, m)$ . Note also that, for two different inputs  $x$  and  $x'$ ,  $f(x)$  and  $f(x')$  might be the same.

Secondly, construct a distribution assignment  $\mathcal{D}'_1$  for the corresponding (i.e. first) multirate  $m$ -distributor as follows. For each output  $c$  of  $\mathcal{C}(n, m)$ , define

$$\begin{aligned}
w(c) &= \sum_{\substack{(x, T, w) \in \mathcal{D}_1 \\ f(x) = c}} w \\
T(c) &= \bigcup_{\substack{(x, T, w) \in \mathcal{D}_1 \\ f(x) = c}} T.
\end{aligned}$$

Then, define

$$\mathcal{D}'_1 = \{(c, T(c), w(c)) \mid w(c) > 0\}.$$

By definition of compatibility, the total weight of compatible routes to any output  $c$  of  $\mathcal{C}(n, m)$  is at most 1. Consequently,  $\mathcal{D}'_1$  is a valid distribution assignment, which can be realized by some network state  $\mathcal{R}'_1$  of  $\mathcal{D}'_1$ .

Finally, each request  $(x, T_1, w)$  in  $\mathcal{D}_1$  can be realized by concatenating the route in  $\mathcal{C}(n, m)$  from  $x$  to  $f(x)$  with the corresponding branch of the route in  $\mathcal{R}'_1$  from  $f(x)$  to  $T_1$ .  $\square$

The following theorem can be made slightly better with more careful calculus. We state a somewhat weaker version for the sake of clarity.

**Theorem IV.2.** For  $\beta \leq 1/2$ , we can construct  $n$ -distributors of

- (a) depth  $k = 3$  and size  $O\left(n^{3/2} \frac{\log n}{\sqrt{\log \log n}}\right)$ .
- (b) depth  $k = 4$  and size  $O\left(n^{3/2} \sqrt{\log \log n}\right)$ .
- (c) depth  $k = 5$  and size  $O\left(n^{4/3} \frac{\log^{4/3} n}{(\log \log n)^{2/3}}\right)$ .
- (d) depth  $k = 6$  and size  $O\left(n^{4/3} (\log n)^{2/3}\right)$ .
- (e) any depth  $k \geq 3$  and size  $O\left(n^{1+1/j} \frac{(\log n)^{1+1/j}}{(\log \log n)^{1-1/j}}\right)$ , where  $j = \lceil k/2 \rceil$ .
- (f) size  $O(n \log n)$ .

*Proof.* The reader is referred to Table I and Corollary III.2 when examining the following reasoning.

- (a) Let  $m = \sqrt{n} \frac{\log n}{\sqrt{\log \log n}}$ . Choose  $\mathcal{C}(n, m)$  of depth 2 and size  $O\left(n \frac{\log^2 n}{\log \log n}\right)$ . Choose  $\mathcal{M}$  to be the complete  $m \times m$  bipartite graph.
- (b) Let  $m = \sqrt{n \log \log n}$ . Choose  $\mathcal{C}(n, m)$  of depth 3 and size  $O(n \log \log n)$ . Choose  $\mathcal{M}$  to be the complete  $m \times m$  bipartite graph.
- (c) Let  $m = n^{2/3} \frac{(\log n)^{2/3}}{(\log \log n)^{1/3}}$ . Choose  $\mathcal{C}(n, m)$  of depth 2, and  $\mathcal{M}$  the depth-3  $m$ -distributor constructed in part (a).
- (d) Let  $m = n^{2/3} \frac{\log \log n}{(\log n)^{2/3}}$ . Choose  $\mathcal{C}(n, m)$  of depth 3, and  $\mathcal{M}$  the depth-3  $m$ -distributor constructed in part (a).
- (e) We induct on  $k$ . For  $2 \leq k \leq 6$ , the previous cases serve as the bases for our induction hypothesis. When,  $k = 2j$  with  $j \geq 4$ , choose  $m = n^{1-1/j} \frac{(\log \log n)^{(2j-1)/(j+1)}}{\log n}$ ,  $\mathcal{C}(n, m)$  of depth 3 and size  $O(n \log \log n)$ , and  $\mathcal{M}$  to be the depth- $(k-3)$   $m$ -distributor inductively constructed. The case when  $k = 2j - 1$  is similar.
- (f) In this case, we choose  $m = n/2$ ,  $\mathcal{C}(n, m)$  to be the linear size multirate concentrator (with depth  $\alpha(n)$  as in Table I). The network  $\mathcal{M}$  is recursively constructed this way. Suppose the  $\mathcal{C}(n, m)$  are of size  $cn$  for some constant  $c$ . The total size is then

$$2 \cdot cn + 4 \cdot c \frac{n}{2} + \dots + 2^{\log n} c \frac{n}{2^{\log n - 1}} = O(n \log n).$$

$\square$

### B. Distributors for the general case

**Lemma IV.3.** Let  $S$  be a set of  $k$  positive real numbers  $\{w_1, \dots, w_k\}$ , where  $w_i \leq 1/2$ ,  $\forall i \in [k]$ , and  $\sum_{i=1}^k w_i \leq 1$ . Then,  $S$  can be partitioned into at most 4 subsets, each of whose sums is at most  $1/2$ .

*Proof.* Let  $S_i = \{w_i\}$  for each  $i \in [k]$ . We will gradually merge these  $S_i$  until there are only at most four sets left with the desired property. For each set  $X$ , let  $w(X)$  denote the sum of elements in  $X$ . Call  $X$  a type- $j$  set if  $1/2^{j+1} < w(X) \leq 1/2^j$ .

Now, consider the sets  $S_i, i \in [k]$ . For any  $j > 1$ , as long as there are two sets  $S_i, S_{i'}$  of type  $j$ , merge  $S_i$  and  $S_{i'}$ . The merge yields a type- $(j-1)$  set. When it is no longer possible to merge, we have at most 3 sets of type-1, and at most 1 set of type- $j$  for each  $j > 1$ . To this end, merge all sets of type- $j$  for all  $j > 1$ . Because  $1/4 + 1/8 + \dots = 1/2$ , the result of this last merge has sum at most  $1/2$ .  $\square$

**Lemma IV.4.** Let  $\mathcal{M}$  be any classical  $n$ -distributor. Let  $\mathcal{N}$  be the distributor constructed in Lemma IV.1. Then, the stacking  $\mathcal{M} \square (\square^4 \mathcal{N})$  of  $\mathcal{M}$  and 4 copies of  $\mathcal{N}$  is a multirate  $n$ -distributor for any  $0 \leq b < B \leq \beta \leq 1$ .

*Proof.* Consider any distribution assignment  $\mathcal{D}$ . For each output vertex  $y$ , consider the set of weights of requests involving this vertex. Partition this weight set into at most 5 classes. Class 0 consists of (at most) one weight which is  $> 1/2$ . Partition all the weights  $\leq 1/2$  into 4 sets using Lemma IV.3, then label the sets classes 1 to 4.

For each request  $(x, T, w) \in \mathcal{D}$ , partition  $T$  into at most 5 classes  $T_0, T_1, \dots, T_4$ , where  $y \in T_i$  iff the weight  $w$  belongs to class  $i$  of output vertex  $y$ . In effect, we decompose the request  $(x, T, w)$  into 5 separate requests  $(x, T_i, w)$ .

The idea is to route the set of all  $(x, T_0, w)$  using the classical  $n$ -distributor. The routes in the classical distributor are vertex disjoint, hence they will certainly satisfy the vertex capacity constraint. Moreover, each output has at most one request with weight  $> 1/2$ , implying that the set of requests  $(x, T_0, w)$  is valid for the distributor.

Then, route all requests  $(x, T_i, w)$  using the  $i$ th copy of  $\mathcal{N}$ . Note that the requests that a copy of  $\mathcal{N}$  is responsible for were chosen so that each output has total requested weight at most  $1/2$ . Hence,  $\mathcal{N}$  can handle them easily by Lemma IV.1.  $\square$

Figure 4 illustrates the construction used in IV.4. Note that our construction works regardless of the values of  $\beta, B$ , and  $b$ . If  $\beta \leq 1/2$  then we do not need the classical distributor in the stacking. However, asymptotically this fact does not reduce the size of the multirate distributor.




Fig. 4. Construction of multirate distributor

Theorem IV.2 and Lemma IV.4 give the key result of this paper.

**Theorem IV.5.** For any  $b \leq B \leq \beta$ , and for any  $k \geq 3$ , we can construct a depth- $k$  multirate  $n$ -distributor of size  $O\left(n^{1+1/j} \frac{(\log n)^{1+1/j}}{(\log \log n)^{1-1/j}}\right)$ , where  $j = \lceil k/2 \rceil$ . This means

$$\text{mrd}_{\beta[b,B]}(n, k) = O\left(n^{1+1/j} \frac{(\log n)^{1+1/j}}{(\log \log n)^{1-1/j}}\right), \quad j = \lceil k/2 \rceil. \quad (2)$$

Furthermore, we can also construct a multirate  $n$ -distributor of size  $O(n \log n)$ . Thus,

$$\text{mrd}_{\beta[b,B]}(n) = O(n \log n). \quad (3)$$

*Proof.* Consider first the fixed-depth case. For the 4 copies of  $\mathcal{N}$  in Lemma IV.4, we use part (e) of Theorem IV.4. For classical  $n$ -distributor of depth  $k$ , we can use the constructions in [17] (see also [14]) with size  $O(n^{1+\frac{1}{j}} (\log n)^{\frac{j-1}{2}})$ , which is asymptotically slightly smaller than our depth- $k$  distributors from part (e) of Theorem IV.4.

If there is no restriction in the network depth, we use part (f) of Theorem IV.4 for  $\mathcal{N}$ . The classical  $n$ -distributor of size  $O(n \log n)$  has been constructed in [29].  $\square$

### C. On the optimality of our distributor

For classical distributors, it has been known for a long time that every  $n$ -distributor must have size  $\Omega(n \log n)$  [31]. Is it possible that, due to the internal speedup factor of  $1/\beta$ , one can construct multirate  $n$ -distributors with size asymptotically better than  $O(n \log n)$ ? For example, when  $1/\beta$  is extremely large (compared to  $n$ ) it is easy to see that **one** internal node is sufficient because this node's capacity can handle all requests.

In the following theorem, we show that, when  $B$  is a constant, we cannot do better than  $O(n \log n)$ , implying that our result in Theorem IV.5 is optimal!

**Theorem IV.6.** Suppose  $0 = b < B = \beta \leq 1$ , where  $B$  is a constant. Given any multirate  $n$ -distributor of size  $g(n)$ , we can construct a classical  $n$ -distributor of size  $O(g(n))$  with the same depth. Thus, any asymptotic lowerbound for classical distributors is also an asymptotic lowerbound for multirate distributors, whether or not the depth is specified.

In particular, a multirate  $n$ -distributor must have size  $\Omega(n \log n)$ ; that is,

$$\text{mrd}_{\beta[b,B]} = \Omega(n \log n).$$

*Proof.* Let  $c = \lfloor 1/B \rfloor$ , which is a constant. Let  $\mathcal{N}$  be any multirate  $n$ -distributor of size  $g(n)$ . We will construct a classical  $n$ -distributor  $\mathcal{M}$  of size  $c^2 g(n) = O(g(n))$  and the same depth. By the aforementioned result,  $c^2 g(n) = \Omega(n \log n)$ ; thus,  $g(n) = \Omega(n \log n)$ , completing the proof.

The network  $\mathcal{M}$  is constructed as follows. Replace each internal vertex  $v$  of  $\mathcal{N}$  by  $c$  copies  $v_1, \dots, v_c$ . For each edge  $(u, v)$  of  $\mathcal{N}$ , do the following:

- if  $u$  is an input and  $v$  is an output of  $\mathcal{N}$ , add the edge  $(u, v)$  to  $\mathcal{M}$ ;
- if  $u$  is an input and  $v$  is an internal vertex of  $\mathcal{N}$ , create  $c$  new edges  $(u, v_1), \dots, (u, v_c)$  in  $\mathcal{M}$ ;

- if  $u$  is an internal vertex and  $v$  is an output of  $\mathcal{N}$ , create  $c$  new edges  $(u_1, v), \dots, (u_c, v)$  in  $\mathcal{M}$ ;
- and lastly, if both  $u$  and  $v$  are internal vertices, then create  $c^2$  new edges  $(u_i, v_j)$  in  $\mathcal{M}$ , for all  $i, j \in [c]$ .

We need to show that  $\mathcal{M}$  is indeed a classical  $n$ -distributor. Consider any distribution assignment  $\mathcal{D}$  in the space domain. This assignment consists of requests of the form  $(x, T)$ , where  $T$  is a subset of the outputs. Each output can only be requested at most once. Now, create a distribution assignment  $\mathcal{D}'$  for  $\mathcal{N}$  as follows. For each request  $(x, T)$  in  $\mathcal{D}$ , create a request  $(x, T, B)$  and add to  $\mathcal{D}'$ . Let  $\mathcal{R}'$  be a network state of  $\mathcal{N}$  realizing  $\mathcal{D}'$ . Obviously each internal node of  $\mathcal{N}$  belongs to at most  $c$  routes in  $\mathcal{R}'$ . Thus, from the routes in  $\mathcal{R}'$  we can construct a set of routes realizing  $\mathcal{D}$  for  $\mathcal{M}$  easily because each vertex  $v$  of  $\mathcal{N}$  has  $c$  copies in  $\mathcal{M}$ .  $\square$

We show an example of the transformation from a multirate distributor to a classical distributor for  $c = 2$  in Figure 5.




Fig. 5. An example of the transformation

## V. DISCUSSIONS

Just like in the classical case, there are still small gaps between the upper and lower bounds of depth- $k$  distributors. These are still open problems. The reader is referred to [25] for more details. With more careful computation, the results of Theorem IV.5 for given depths can be made better. Another open problem is the asymptotic sizes of multirate distributors when  $\beta$  is not a constant.

On a more practical side, it is not difficult to recursively construct good concentrators and superconcentrators from  $2 \times 2$ -switching elements. However, these constructions are based on expanders [10]. Even though several explicitly constructed expanders are known, converting them to an arrangement of switching elements is very tedious and unlikely to yield a design suitable for physical layout. Some steps toward realizing concentrators using sparse crossbars have been made (e.g., [40]). However, the problem of constructing concentrators and superconcentrators which not only have linear crosspoint complexity but also are suitable for physical layout is an important open problem in switching network theory.



Perhaps the most important contribution in this paper is the particular notion of multirate concentrators defined and constructed in Section III. Without this new idea, the proofs would not “go through.”

Our result is yet again an illustration of the power of Pippenger’s network in constructing multicast switching networks. The basic structure of Pippenger’s network is shown in Figure 3. This topological structure was used in [21] for the  $O(n \log^2 n)$ -cross-point complexity construction of multirate distributors (the best result prior to our work), in [31] for constructing optimal space-domain distributors, and in [27] for constructing optimal wavelength-domain distributors. If we ignore the new notion of multirate concentrators, structurally the route selection algorithm on a Pippenger network for a distribution assignment is similar in the space-domain [31], wavelength-domain [27], and multirate domain (this paper). This is a by-product of the fact that all three constructions are based on the same topological idea. Some technical details are different from one domain to the next though.

The topological structure of Pippenger’s network also has manifestations in contexts other than multicast switching network, where the concentrator/distributor pairs might be replaced with some other pairs. For example, the (inverse) Banyan network and its topological equivalences (Delta, butterfly, etc.) have this basic structure [13]. The multi-butterfly network defined and constructed in [2] is a very close cousin of the Pippenger’s network.

#### ACKNOWLEDGEMENTS

We would like to thank an anonymous referee for detailed comments which have helped improve our presentation. Our work was supported in part by NSF CAREER Award CCF-0347565.

#### REFERENCES

- [1] N. ALON AND P. PUDLÁK, *Superconcentrators of depths 2 and 3; odd levels help (rarely)*, J. Comput. System Sci., 48 (1994), pp. 194–202.
- [2] S. ARORA, F. T. LEIGHTON, AND B. M. MAGGS, *On-line algorithms for path selection in a nonblocking network*, SIAM J. Comput., 25 (1996), pp. 600–625.
- [3] S.-P. CHUNG AND K. W. ROSS, *On nonblocking multirate interconnection networks*, SIAM J. Comput., 20 (1991), pp. 726–736.
- [4] C. CLOS, *A study of non-blocking switching networks*, Bell System Tech. J., 32 (1953), pp. 406–424.
- [5] P. COPPO, M. D’AMBROSIO, AND R. MELEN, *Optimal cost/performance design of atm switches.*, IEEE/ACM Trans. Netw., 1 (1993), pp. 566–575.
- [6] J. R. CORREA AND M. X. GOEMANS, *An approximate König’s theorem for edge-coloring weighted bipartite graphs*, in Proceedings of the 36th Annual ACM Symposium on Theory of Computing, New York, 2004, ACM, pp. 398–406 (electronic).
- [7] D. DOLEV, C. DWORK, N. PIPPENGER, AND A. 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.
- [8] D. Z. DU, B. GAO, F. K. HWANG, AND J. H. KIM, *On multirate rearrangeable Clos networks*, SIAM J. Comput., 28 (1999), pp. 464–471 (electronic).
- [9] J. A. FINGERHUT, S. SURI, AND J. S. TURNER, *Designing least-cost nonblocking broadband networks*, J. Algorithms, 24 (1997), pp. 287–309.
- [10] S. HOORY, N. LINIAL, AND A. WIGDERSON, *Expander graphs and their applications*, Bull. Amer. Math. Soc. (N.S.), 43 (2006), pp. 439–561 (electronic).
- [11] X.-D. HU, X.-H. JIA, D.-Z. DU, AND F. K. HWANG, *Monotone routing in multirate rearrangeable Clos networks*, J. Parallel Distrib. Comput., 61 (2001), pp. 1382–1388.
- [12] J. H. HUI, *Switching and traffic theory for integrated broadband networks*, Kluwer Academic Publishers, Boston/Dordrecht/London, 1990.
- [13] F. K. HWANG, *The mathematical theory of nonblocking switching networks*, World Scientific Publishing Co. Inc., River Edge, NJ, 1998.
- [14] R. M. KARP, *Three-stage generalized connectors*, SIAM J. Discrete Math., 5 (1992), pp. 259–272.
- [15] D. S. KIM AND D.-Z. DU, *Multirate multicast switching networks*, Theoret. Comput. Sci., 261 (2001), pp. 241–251. Computing and combinatorics (Taipei, 1998).
- [16] D. S. KIM AND D.-Z. DU, *Multirate multicast switching networks.*, Theor. Comput. Sci., 261 (2001), pp. 241–251.
- [17] D. G. KIRKPATRICK, M. KLAWE, AND N. PIPPENGER, *Some graph-colouring theorems with applications to generalized connection networks*, SIAM J. Algebraic Discrete Methods, 6 (1985), pp. 576–582.
- [18] S. C. LIEW, M.-H. NG, AND C. W. CHAN, *Blocking and nonblocking multirate Clos switching networks.*, IEEE/ACM Trans. Netw., 6 (1998), pp. 307–318.
- [19] G.-H. LIN, D.-Z. DU, X.-D. HU, AND G. XUE, *On rearrangeability of multirate Clos networks*, SIAM J. Comput., 28 (1999), pp. 1225–1231 (electronic).
- [20] R. MELEN AND J. S. TURNER, *Nonblocking multirate networks*, SIAM J. Comput., 18 (1989), pp. 301–313.
- [21] R. MELEN AND J. S. TURNER, *Nonblocking multirate distribution networks.*, in INFOCOM, 1990, pp. 1234–1241.
- [22] R. MELEN AND J. S. TURNER, *Nonblocking multirate distribution networks*, IEEE Transactions on Communications, 41 (1993), pp. 362–369.
- [23] H. Q. NGO, *A new routing algorithm for multirate rearrangeable Clos networks*, Theoret. Comput. Sci., 290 (2003), pp. 2157–2167.
- [24] ———, *Multiwavelength distribution networks*, in Proceedings of the 2004 Workshop on High Performance Switching and Routing (HPSR), Phoenix, Arizona, U.S.A., April 2004, IEEE, pp. 186–190.
- [25] H. Q. NGO AND D.-Z. DU, *Notes on the complexity of switching networks*, in Advances in Switching Networks, D.-Z. Du and H. Q. Ngo, eds., vol. 5 of Network Theory and Applications, Kluwer Academic Publishers, 2001, pp. 307–367.
- [26] H. Q. NGO AND V. H. VU, *Multirate rearrangeable Clos networks and a generalized bipartite graph edge coloring problem*, SIAM Journal on Computing, 32 (2003), pp. 1040–1049.
- [27] H. Q. NGO, Y. WANG, AND D. PAN, *Rearrangeable and nonblocking  $[w, f]$ -distributors*, IEEE/ACM Transactions on Networking, (2008). Accepted for publication.
- [28] G. NIESTEGGE, *Nonblocking multirate switching networks*, in Proceedings of the 5th ITC Seminar on Traffic Engineering for ISDN Design and Planning, 1987, pp. 449–458.
- [29] J. P. OFMAN, *A universal automaton*, Trudy Moskov. Mat. Obšč., 14 (1965), pp. 186–199.
- [30] N. PIPPENGER, *The complexity of switching networks*, PhD thesis, Department of Electrical Engineering, Massachusetts Institute of Technology, Cambridge, Massachusetts, 1973.
- [31] ———, *Generalized connectors*, SIAM J. Comput., 7 (1978), pp. 510–514.
- [32] P. PUDLÁK, *Communication in bounded depth circuits*, Combinatorica, 14 (1994), pp. 203–216.
- [33] J. RADHAKRISHNAN AND A. TA-SHMA, *Bounds for dispersers, extractors, and depth-two superconcentrators*, SIAM J. Discrete Math., 13 (2000), pp. 2–24 (electronic).
- [34] J. TURNER AND N. YAMANAKA, *Architectural choices in large scale ATM switches*, IEICE Trans. Commun., E81-B (1998), pp. 120–137.
- [35] J. S. TURNER AND R. MELEN, *Multirate Clos networks*, Communications Magazine, IEEE, 41 (2003), pp. 38–44.
- [36] L. G. VALIANT, *On non-linear lower bounds in computational complexity*, in Seventh Annual ACM Symposium on Theory of Computing (Albuquerque, N. M., 1975), Assoc. Comput. Mach., New York, 1975, pp. 45–53.
- [37] ———, *Graph-theoretic properties in computational complexity*, J. Comput. System Sci., 13 (1976), pp. 278–285. Working papers presented at the ACM-SIGACT Symposium on the Theory of Computing (Albuquerque, N. M., 1975).
- [38] D. B. WEST, *Introduction to graph theory*, Prentice Hall Inc., Upper Saddle River, NJ, 1996.
- [39] Y. YANG, *An analysis model on nonblocking multirate broadcast networks.*, in International Conference on Supercomputing, 1994, pp. 256–263.



- [40] A. YAVUZ ORUC AND H. HUANG, *Crosspoint complexity of sparse crossbar concentrators*, IEEE Trans. Inform. Theory, 42 (1996), pp. 1466–1471.