Constructions of Given-Depth and Optimal Multirate Rearrangeably Nonblocking Distributors

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

***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 ATM switching systems.**

Rearrangeable multirate multicast switching networks are customarily called distribution networks, or distributors for short. It has been known for more than 15 years that distributors with cross-point complexity $O(n \log^2 n)$ can be constructed, where n is the number of inputs. The problem of constructing optimal distributors remains open thus far.

In this paper, we give a general method for constructing given-depth rearrangeable multirate distributors. One of the rewards of our construction method is a distributor with cross-point complexity $O(n \log n)$, which we then show to be optimal. We thus settle the aforementioned open problem.

I. Introduction

Multi-rate switching networks are switching networks that support varying bandwidth connections. The theory of multirate switching networks, perhaps started with the papers by Niestegge [1] and Melen and Turner [2], 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 [3]–[5]. Roughly speaking, as opposed to space switching where each connection request can only be carried on an internal or external link of a switch, the multirate switches allow for connections with varying "rates" or bandwidths to be carried on a single link, as long as the total connection rates 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 [6], represented by the (still open) conjecture by Chung and Ross in 1991 [7] 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 Konig's theorem for edge coloring bipartite graphs. Later developments on this conjecture and related problems were reported in [8]–[13]. See also [14], [15] for several related lines of research.

In the multicast and broadcast cases, there have been notably few known results, though. The works presented in [14], [16]–[18] concern conditions for the Clos network to be multicast capable. The study presented in [19] (the journal version is [20]) was the only one that deals directly with more general constructions and complexities of multicast multirate switching networks. In their paper, using Pippenger's network [21], the authors constructed a rearrangeable multirate distributor with cross-point complexity $O(n \log^2 n)$. (*Distributor*, also called *generalized connector*, is a standard name referring to multicast switching networks.)

The problem of constructing optimal multirate multicast switching networks remains open thus far. In this paper, we give a general method for constructing rearrangeable multirate distributors. One of the rewards of the method is a multicast distributor with cross-point complexity $O(n \log n)$. We then show that this is optimal, thus settling the aforementioned open problem.

The rest of the paper is organized as follows. Section II presents basic definitions and several fundamental compositions of networks. Section III gives the definition and construction of a special version of multirate concentrators, which is crucial for the later constructions of multirate distributors. Section IV contains the main results, including a general distributor construction given the network depth. The construction gives rise to a multirate n-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 [22], 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 [23], [24]. 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 each input and output* of

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: (yw43, hungngo, nguyen9)@cse.buffalo.edu

the network $\mathcal{N}$. Input nodes have capacity (normalized to) 1. The factor $1/\beta$ is often referred to as the *speed advantage* of the system. This *internal speedup* is a common technique for designing broadband switches [2], [25], [26].

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

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

As we are only concerned with distribution networks in this paper, the term "request" should be implicitly understood as "distribution request" henceforth, unless it is explicitly specified otherwise. 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 β, and (b) total request weight to any output does not exceed β, in other words

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

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

A *distribution route* (or just *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 (or rate) of R. A *state* of $\mathcal{N}$ 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* 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 are now ready to define the central notions of nonblockingness in the multirate environment. In defining different notions of distributors, we drop the "multirate" qualifier to avoid being too wordy. Distribution networks in this papers are implicitly understood as multirate distribution networks, unless explicitly stated otherwise.

A *rearrangeable (RNB) n-distributor* (or simply *n-distributor*) 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. The main theme of research in switching network has been to address the tradeoff between networks' depths and sizes. This paper is no exception.

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

B. Classical networks

In constructing multirate distributors, we will also need the notions of classical distributors, concentrators, superconcentrators.

The classical n-distributor is defined similar to the multirate distributor, except for the fact that $\beta = 1$ and all request weights are 1. Thus, in a classical network state the distribution routes are vertex disjoint trees. Also, since all request weights are 1, there is no need to include a weight to describe a request.

Given integers $n \geq m > 0$, 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. 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 X and outputs Y such that for any $S \subseteq X$ and $T \subseteq Y$ with $|S| = |T| = k$, there exist a set of k 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-superconcentrator and an n-superconcentrator of depth k, respectively.

For $n \geq m$, an *(n, m)-superconcentrator* is a network obtained by removing any $(n - m)$ outputs from an n-superconcentrator. Obviously, an (n, m)-superconcentrator is an (n, m)-concentrator. Hence,

$$c(n, m) \leq s(n), \tag{1}$$
$$c(n, m, k) \leq s(n, k). \tag{2}$$

Note that the concentrators and superconcentrators described above operate in the *space domain* or the *circuit switching environment*, namely no two paths can share a vertex. It has been known for more than 3 decades that there are concentrators and superconcentrators of linear sizes [27], [28]. The constructions were based on a class of graphs called *expanders*, whose applications in mathematics and computer science are numerous [29].

For the fixed depth case, the asymptotic behaviors of all the $s(n, k)$ were only completely characterized recently. Table I summarizes the results. The function $\lambda(d, n)$ is the inverse of functions in the Ackerman hierarchy: they are increasing

TABLE I

MINIMUM SIZE OF n-SUPERCONCENTRATORS WITH DEPTH k

Depth k	Size $s(n,k)$
2	$\Theta\left(\frac{n\log^2 n}{\log\log n}\right)$ [30]
3	$\Theta\left(n\log\log n\right)$ [31]
$2d, 2d+1,\ d\geq 2$	$\Theta\left(n\lambda(d,n)\right)$ [32], [33]
In particular, for $k=4,5$	$\Theta\left(n\log^* n\right)$ [32], [33]
$\Theta(\alpha(n))$	$\Theta(n)$ [33]

extremely slowly. The reader is referred to [33] for the definitions of $\lambda(d,n)$ and $\alpha(n)$ (which is actually called $\beta(n)$ in their paper, but we change its name to avoid confusion with our speedup parameter β).

C. Basic compositions of networks

Let $\mathcal{N}_1$ and $\mathcal{N}_2$ be any two (n,m)-networks. We use $\mathcal{N}_1 \bowtie \mathcal{N}_2$ 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. We refer to $\mathcal{N}_1 \bowtie \mathcal{N}_2$ as the *stacking* of $\mathcal{N}_1$ and $\mathcal{N}_2$. When stacking k copies of a network $\mathcal{N}$, denote the result by $\bowtie^k \mathcal{N}$.

Given any k (n,m)-networks $\mathcal{N}_1,\dots,\mathcal{N}_k$, let $\vdash (\mathcal{N}_1,\dots,\mathcal{N}_k)$ 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 $\mathcal{N}_1,\dots,\mathcal{N}_k$.) 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 $\mathcal{N}$, 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.

III. MULTIRATE CONCENTRATORS

There are several obvious ways to generalize the notion of classical concentrators to multirate concentrators. To avoid cumbersome notations, we will define here only a particular type of multirate concentrators which are used in later sections to construct good multirate distributors.

Given integers $n \geq m > 0$. Consider an (n,m)-network $\mathcal{C} = (V,E;X,Y)$. A *concentration request* is a pair (x,w), where x is an input and $w \leq 1$ 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 containing any vertex is at most 1. A *concentration assignment* is a set of 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 network $\mathcal{C}$ is called an (n,m)-*multirate concentrator* if and only if, for each concentration assignment $\mathcal{D}$ there exists a set of compatible routes realizing requests in the assignment.

Lemma 1. *Let $\mathcal{C}$ be any (n,m)-concentrator and $\mathcal{S}$ be any (n,m)-superconcentrator. Then, $\mathcal{C}(n,m) = \mathcal{C} \bowtie \mathcal{S}$ is an (n,m)-multirate concentrator.*

Proof. The reader is referred to Figure 1 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 [2].

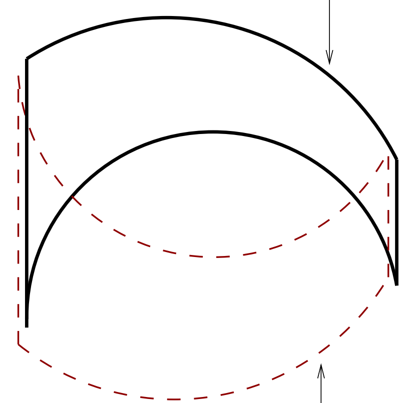

Fig. 1. Construction of an (n,m)-multirate concentrator $\mathcal{C}(n,m)$.

Let $\mathcal{D}$ be any concentration assignment. Note that the inputs of these requests are not necessarily different from one another. 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 requests is still a valid concentration assignment. Moreover, a valid route for (x,w_1+w_2) can be "decomposed" back into two routes with weights w_1 and w_2 to satisfy the requests (x,w_1) and (x,w_2). Consequently, we can assume that the inputs of these requests are distinct.

Partition $\mathcal{D}$ into two subsets $\mathcal{D}_1$ and $\mathcal{D}_2$, where $\mathcal{D}_1$ consists of all requests with weights $> 1/2$. Let $x=|\mathcal{D}_1|$, $y=|\mathcal{D}_2|$. For $i=1,2$, let W_i be the total weight of requests in $\mathcal{D}_i$. Then, because $\mathcal{D}$ is a concentration assignment,

$$m/2 \geq W_1+W_2 > x/2+W_2. \tag{3}$$

The set of requests in $\mathcal{D}_1$ can be routed through the concentrator $\mathcal{C}$ so that no two routes share a vertex. Thus, the vertex capacity constraint is satisfied.

Next, we route the requests in $\mathcal{D}_2$ through the superconcentrator $\mathcal{S}$ to the other $m-x$ outputs that are unused after routing $\mathcal{D}_1$. These requests are routed using the CAP algorithm.

Let $s=m-x$. Partition the y requests of $\mathcal{D}_2$ into $t=\lceil y/s\rceil$ groups of size s each, with possibly the last group having less than s members. Assume the weights for these requests are $w_1 \geq w_2 \geq \cdots \geq w_y$. The partition is such that 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 any group of requests, in the (n,m)-superconcentrator $\mathcal{S}$ there are s vertex disjoint paths joining the inputs of the requests in the group to some s outputs. We will use these paths as routes realizing the requests in the group. This ensures that 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 might carry is at most

$$\begin{aligned} & w_1 + w_{s+1} + \cdots + w_{(t-1)s+1} \\ \leq \quad & \frac{1}{2} + \frac{w_1 + \cdots + w_s}{s} + \cdots + \frac{w_{(t-2)s+1} + \cdots + w_{(t-1)s}}{s} \\ \leq \quad & \frac{1}{2} + \frac{W_2}{s} \leq \frac{1}{2} + \frac{m/2 - x/2}{m - x} = 1. \end{aligned}$$

The last inequality follows from (3). □

Corollary 2. *An (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 classical n-superconcentrator yields an (n, m)-superconcentrator, which is also an (n, m)-concentrator. Thus, in fact our (n, m)-multirate concentrator of depth k is of size at most $2s(n, k)$. □

IV. Rearrangeable multirate distributors

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

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

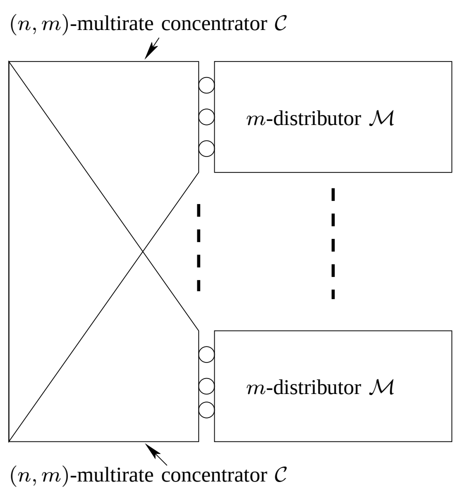

Fig. 2. Recursive construction of distributors with capacity-1 inputs.

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. The following construction is the multirate version of Pippenger's network [21].

Lemma 3. *Let m be a factor of n. Let $\mathcal{C}$ be an (n, m)-multirate concentrator. Let $\mathcal{M}$ be any multirate m-distributor with capacity-1 inputs. Then, the network $\mathcal{N} = \vdash^{n/m} (\mathcal{C} \circ \mathcal{M})$ is an n-distributor with capacity-1 inputs. Note that, we only consider the case when $B \leq \beta \leq 1/2$.*

Proof. The reader is referred to Figure 2 for an illustration of $\mathcal{N}$. Consider a distribution assignment $\mathcal{D}$. Partition $\mathcal{D}$ into n/m subsets $\mathcal{D} = \mathcal{D}_1 \cup \mathcal{D}_2 \cup \cdots \cup \mathcal{D}_{n/m}$ as follows. For each request $D = (x, T, w) \in \mathcal{D}$ and $i \in \{1, \ldots, n/m\}$, let

$$T_i = T \cap \{(i-1)m + 1, (i-1)m + 2, \ldots, im\}$$

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

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 show how to construct routes realizing $\mathcal{D}_1$.

Firstly, notice 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 $\mathcal{C}$ joining each input x of a request (x, T_1, w) in $\mathcal{D}_1$ to an output $f(x)$ of $\mathcal{C}$. For two inputs x and x', $f(x)$ and $f(x')$ might be the same, though.

Secondly, construct a distribution assignment $\mathcal{D}'_1$ for the corresponding m-distributor as follows. For each request (x, T_1, w) in $\mathcal{D}_1$, add $(f(x), T_1, w)$ to $\mathcal{D}'_1$. By definition of compatibility, the total weight of compatible routes to any output of $\mathcal{C}$ is at most 1. Consequently, $\mathcal{D}'_1$ is a valid distribution assignment, which can be realized by some network state $\mathcal{R}'_1$.

Finally, each request (x, T_1, w) in $\mathcal{D}_1$ can be realized by the concatenation of the route in $\mathcal{C}$ and the corresponding route in $\mathcal{R}'_1$. □

We can now construct multirate distributors for the case $B \leq \beta \leq 1/2$. The following theorem can be made slightly better with more careful calculus. We state a somewhat weaker version for the sake of clarity.

Theorem 4. *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 2 when examining the following reasoning. We will use the construction of Lemma 3.

(a) Let $m = \sqrt{n} \frac{\log n}{\sqrt{\log \log n}}$. Choose $\mathcal{C}$ of depth 2 and size $O\left(n \frac{\log^2 n}{\log \log n}\right)$. Chose $\mathcal{M}$ to be the complete $m \times m$ bipartite graph.
(b) Let $m = \sqrt{n \log \log n}$. Choose $\mathcal{C}$ of depth 3 and size $O(n \log \log n)$. Chose $\mathcal{M}$ 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 $\mathcal{C}$ 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 $\mathcal{C}$ 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 $m = n^{1-1/j}\frac{(\log\log n)^{(2j-1)/(j+1)}}{\log n}$, $\mathcal{C}$ 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 part, we choose $m = n/2$, $\mathcal{C}$ 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 $\mathcal{C}$ are of size cn for some constant c. The total size is thus

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

□

B. Distributors for the general case

We first need a technical lemma.

Lemma 5. *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 four 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 results in 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 $1/4 + 1/8 + \cdots = 1/2$, the resulting set of this last merge has sum at most $1/2$. □

Lemma 6. *Let $\mathcal{M}$ be any classical n-distributor. Let $\mathcal{N}$ be a distributor as in Lemma 3. Then, the stacking of $\mathcal{M}$ and 4 copies of $\mathcal{N}$ is a multirate n-distributor for any parameters $0 \leq b < B \leq \beta \leq 1$.*

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 5, 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 $\mathcal{M}$. 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 ith 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 3. □

Note that our construction works regardless of the values of β, 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.

Theorem 4 and Lemma 6 give the key result of this paper.

Theorem 7. *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*

$$\text{mrd}_{\beta[b,B]}(n, k) = O\left(n^{1+1/j}\frac{(\log n)^{1+1/j}}{(\log\log n)^{1-1/j}}\right). \quad (4)$$

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

$$\text{mrd}_{\beta[b,B]}(n) = O(n \log n). \quad (5)$$

Proof. Consider first the fixed-depth case. For the 4 copies of $\mathcal{N}$ in Lemma 6, we use part (e) of Theorem 4. For classical n-distributor of depth k, we can use the constructions in [34] (see also [35]) 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 4.

If there is no restriction in the network depth, we use part (f) of Theorem 4 for $\mathcal{N}$. The classical n-distributor $\mathcal{M}$ of size $O(n \log n)$ has was constructed in [36]. □

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 $\Omega(n \log n)$ [37]. 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 (this means the speedup factor $1/\beta$ is bounded), we cannot do better than $O(n \log n)$, implying that our result in Theorem 7 is optimal!

Theorem 8. *Suppose $0 = b < B \leq \beta \leq 1$, where B is a constant. Given any multirate n-distributor of size $t(n)$, we can construct a classical n-distributor of size $O(t(n))$ 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,

$$\text{mrd}_{\beta[b,B]} = \Omega(n \log n). \quad (6)$$

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

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 [23] for more details. With more careful computation, the results of Theorem 7 for given depths can be made better. Another open problem is the asymptotic sizes of multirate distributors when β is not a constant.

Last but not least, the wide-sense nonblocking case for multirate distributors is still wide open for further research.

References

[1] G. Niestegge, "Nonblocking multirate switching networks," in *Proceedings of the 5th ITC Seminar on Traffic Engineering for ISDN Design and Planning*, pp. 449–458, 1987.

[2] R. Melen and J. S. Turner, "Nonblocking multirate networks," *SIAM J. Comput.*, vol. 18, no. 2, pp. 301–313, 1989.

[3] J. S. Turner and R. Melen, "Multirate clos networks," *Communications Magazine, IEEE*, vol. 41, no. 10, pp. 38–44, 2003.

[4] P. Coppo, M. D'Ambrosio, and R. Melen, "Optimal cost/performance design of atm switches.," *IEEE/ACM Trans. Netw.*, vol. 1, no. 5, pp. 566–575, 1993.

[5] J. Turner and N. Yamanaka, "Architectural choices in large scale ATM switches," *IEICE Trans. Commun.*, vol. E81-B, no. 2, pp. 120–137, 1998.

[6] C. Clos, "A study of non-blocking switching networks," *Bell System Tech. J.*, vol. 32, pp. 406–424, 1953.

[7] S.-P. Chung and K. W. Ross, "On nonblocking multirate interconnection networks," *SIAM J. Comput.*, vol. 20, no. 4, pp. 726–736, 1991.

[8] D. Z. Du, B. Gao, F. K. Hwang, and J. H. Kim, "On multirate rearrangeable Clos networks," *SIAM J. Comput.*, vol. 28, no. 2, pp. 464–471 (electronic), 1999.

[9] G.-H. Lin, D.-Z. Du, X.-D. Hu, and G. Xue, "On rearrangeability of multirate Clos networks," *SIAM J. Comput.*, vol. 28, no. 4, pp. 1225–1231 (electronic), 1999.

[10] X.-D. Hu, X.-H. Jia, D.-Z. Du, and F. K. Hwang, "Monotone routing in multirate rearrangeable clos networks," *J. Parallel Distrib. Comput.*, vol. 61, no. 9, pp. 1382–1388, 2001.

[11] H. Q. Ngo, "A new routing algorithm for multirate rearrangeable Clos networks," *Theoret. Comput. Sci.*, vol. 290, no. 3, pp. 2157–2167, 2003.

[12] H. Q. Ngo and V. H. Vu, "Multirate rearrangeable Clos networks and a generalized bipartite graph edge coloring problem," *SIAM Journal on Computing*, vol. 32, no. 4, pp. 1040–1049, 2003.

[13] 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), pp. 398–406 (electronic), ACM, 2004.

[14] S. C. Liew, M.-H. Ng, and C. W. Chan, "Blocking and nonblocking multirate clos switching networks.," *IEEE/ACM Trans. Netw.*, vol. 6, no. 3, pp. 307–318, 1998.

[15] J. A. Fingerhut, S. Suri, and J. S. Turner, "Designing least-cost nonblocking broadband networks," *J. Algorithms*, vol. 24, no. 2, pp. 287–309, 1997.

[16] D. S. Kim and D.-Z. Du, "Multirate multicast switching networks," *Theoret. Comput. Sci.*, vol. 261, no. 2, pp. 241–251, 2001. Computing and combinatorics (Taipei, 1998).

[17] D. S. Kim and D.-Z. Du, "Multirate multicast switching networks.," *Theor. Comput. Sci.*, vol. 261, no. 2, pp. 241–251, 2001.

[18] Y. Yang, "An analysis model on nonblocking multirate broadcast networks.," in *International Conference on Supercomputing*, pp. 256–263, 1994.

[19] R. Melen and J. S. Turner, "Nonblocking multirate distribution networks.," in *INFOCOM*, pp. 1234–1241, 1990.

[20] R. Melen and J. S. Turner, "Nonblocking multirate distribution networks," *IEEE Transactions on Communications*, vol. 41, pp. 362–369, Feb. 1993.

[21] N. Pippenger, *The complexity of switching networks*. PhD thesis, Department of Electrical Engineering, Massachusetts Institute of Technology, Cambridge, Massachusetts, 1973.

[22] D. B. West, *Introduction to graph theory*. Upper Saddle River, NJ: Prentice Hall Inc., 1996.

[23] 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*, pp. 307–367, Kluwer Academic Publishers, 2001.

[24] H. Q. Ngo, "Multiwavelength distribution networks," in *Proceedings of the 2004 Workshop on High Performance Switching and Routing (HPSR)*, (Phoenix, Arizona, U.S.A.), pp. 186–190, IEEE, 2004.

[25] F. K. Hwang, *The mathematical theory of nonblocking switching networks*. River Edge, NJ: World Scientific Publishing Co. Inc., 1998.

[26] J. H. Hui, *Switching and traffic theory for integrated broadband networks*. Kluwer Academic Publishers, Boston/Dordrecht/London, 1990.

[27] L. 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).

[28] L. G. Valiant, "On non-linear lower bounds in computational complexity," in *Seventh Annual ACM Symposium on Theory of Computing (Albuquerque, N. M., 1975)*, pp. 45–53, Assoc. Comput. Mach., New York, 1975.

[29] S. Hoory, N. Linial, and A. Wigderson, "Expander graphs and their applications," *Bull. Amer. Math. Soc. (N.S.)*, vol. 43, no. 4, pp. 439–561 (electronic), 2006.

[30] J. Radhakrishnan and A. Ta-Shma, "Bounds for dispersers, extractors, and depth-two superconcentrators," *SIAM J. Discrete Math.*, vol. 13, no. 1, pp. 2–24 (electronic), 2000.

[31] N. Alon and P. Pudlák, "Superconcentrators of depths 2 and 3; odd levels help (rarely)," *J. Comput. System Sci.*, vol. 48, no. 1, pp. 194–202, 1994.

[32] P. Pudlák, "Communication in bounded depth circuits," *Combinatorica*, vol. 14, no. 2, pp. 203–216, 1994.

[33] 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), pp. 42–51, apr 1983.

[34] D. G. Kirkpatrick, M. Klawe, and N. Pippenger, "Some graph-colouring theorems with applications to generalized connection networks," *SIAM J. Algebraic Discrete Methods*, vol. 6, no. 4, pp. 576–582, 1985.

[35] R. M. Karp, "Three-stage generalized connectors," *SIAM J. Discrete Math.*, vol. 5, no. 2, pp. 259–272, 1992.

[36] J. P. Ofman, "A universal automaton," *Trudy Moskov. Mat. Obšč.*, vol. 14, pp. 186–199, 1965.

[37] N. Pippenger, "Generalized connectors," *SIAM J. Comput.*, vol. 7, no. 4, pp. 510–514, 1978.