Rearrangeable Circuit-Switching Networks

by
Nicholas Pippenger

Technical Report 92-14
June 1992

Department of Computer Science
University of British Columbia
Rm 333 - 6356 Agricultural Road
Vancouver, B.C.
CANADA V6T 1Z2

Rearrangeable Circuit-Switching Networks

Nicholas Pippenger*

Department of Computer Science
The University of British Columbia
Vancouver, British Columbia V6T 1Z2
CANADA

Abstract: We present simple proofs of the basic results concerning the complexity of rearrangeable connectors and superconcentrators. We also define several types of networks whose connectivity properties interpolate between these extremes, and show that their complexities also interpolate.

* This research was partially supported by an NSERC Operating Grant and an ASI Fellowship Award.

1. Introduction

The word "network" is used in many senses. Samuel Johnson's definition ("Anything reticulated or decussated, at equal distances, with interstices between the intersections ...") is famous as an example of his style, but is little help to us. Other, more specialized definitions occur in operations research, circuit theory, logical design, and artificial intellegence. Our present interest, however, is with networks as mediators of communication. For this purpose it is satisfactory to model them as a species of directed graph, whence the following definition.

An *n-network* is a directed acyclic graph in which there are n distinguished vertices $x_1, \ldots, x_n$ with in-degree 0 called *inputs* and n distinguished vertices $y_1, \ldots, y_n$ with out-degree 0 called *outputs*. Vertices that are neither inputs nor outputs are called *links*.

The results of this paper all fit within the following framework. We define a "connectivity property" of networks, and then seek the smallest n-network (that is, the n-network with the smallest possible number of edges) having the given connectivity property. The minimum possible number of edges is of course a function of n (and perhaps other parameters involved in the connectivity property), and we shall be more interested in the order of growth of this function (to within constant factors) than in its numerical value for any particular value of n.

The terms "rearrangeable" and "circuit-switching" in the title refer to aspects of the connectivity properties we shall study, which all take the following form. We delineate a set of "tasks" for a network, where each task constitutes a set of "requests" to join certain inputs to certain outputs. We then ask that for each such task there be a set of vertex-disjoint paths that satisfy all of the requests in that task.

The term "circuit-switching" refers to the requirement that the paths fulfilling a given task must be disjoint, so that different requests are fulfilled by different parts of the network. (This may be contrasted with "packet-switching", where different requests may be fulfilled by the same part of the network at different times.) The term "rearrangeable" refers to the fact that the set of paths may depend on the entire task, so that the change of a single request in a task might occasion the rearrangement of all the paths. (This may be contrasted with "non-blocking" operation, where it must be possible to satisfy new requests without disturbing the paths that satisfy old ones.)

2. Connectors

A *connection n-assignment* is a sequence $v = (v_1, \ldots, v_n)$ such that (1) $v_i \in \{0, \ldots, n\}$ for all $1 \le i \le n$, and (2) for $1 \le u \le n$, $v_i = u$ for at most one $1 \le i \le n$. Say that two connection n-assignments v and w are *similar* if $\{v_1, \ldots, v_n\} = \{w_1, \ldots, w_n\}$.

Let v and w be connection n-assignments, and let N be an n-network. We shall say that N *realizes* the pair (v, w) if there exists a set P of vertex-disjoint paths from inputs to outputs in N such that (1) x_i is the origin of a path in P if and only if $v_i > 0$, (2) y_j is the destination of a path in P if and only if $w_j > 0$, and (3) if there is a path from x_i to y_j in P, then $v_i = w_j$. We may think of the non-zero elements of an assignment as "requests" of n different types, with zero elements indicating the absence of a request. A network then realizes a pair of assignments if it can "match" each input request with an output request of the same type.

Throughout this paper we shall assume that n, and later certain other parameters, are integral powers of 2, so that a proof "by induction on n" will typically proceed from n to $2n$, rather than from n to $n+1$. This fits well with our interest in orders of growth to within constant factors, and it will significantly simplify some of our proofs. Let $c(n)$ denote the minimum possible number of edges in an n-connector.

Theorem 2.1: We have $c(n) \le 4n \log_2 n$.

Proof: We shall describe a recursive construction for an n-connector $C(n)$. For the basis of the recursion, we take $C(2)$ to be a complete bipartite graph with 2 inputs and 2 outputs. For the recursive step, we describe how to construct a $(2n)$-connector from 2 disjoint n-connectors and $8n$ additional edges. We will then have $c(2) \le 4$ together with the recurrence

$$c(2n) \le 2c(n) + 8n,$$

from which it follows by induction that $c(n) \le 4n \log_2 n$.

The recursive construction for $C(2n)$ is as follows. We take four disjoint sets each containing $2n$ vertices: the inputs $\{1, \ldots, 2n\}$, the links $\{\widehat{1}, \ldots, \widehat{2n}\}$ and $\{\widetilde{1}, \ldots, \widetilde{2n}\}$, and the outputs $\{\overline{1}, \ldots, \overline{2n}\}$. For $r \in \{1, \ldots, n\}$, we install a complete bipartite graph (called an *input node*) from $\{r, \ldots, n+r\}$ to $\{\widehat{r}, \ldots, \widehat{n+r}\}$, and a complete bipartite graph (called an *output node*) from $\{\widetilde{r}, \ldots, \widetilde{n+r}\}$ to $\{\overline{r}, \ldots, \overline{n+r}\}$. We also install a copy of $C(n)$ (called the *lower subnetwork*) from $\{\widehat{1}, \ldots, \widehat{n}\}$ to $\{\widetilde{1}, \ldots, \widetilde{n}\}$, and a copy of $C(n)$ (called the *upper subnetwork*) from $\{\widehat{n+1}, \ldots, \widehat{2n}\}$ to $\{\widetilde{n+1}, \ldots, \widetilde{2n}\}$.

To see that the resulting network is a $(2n)$-connector, we consider a pair (v, w) of connection $(2n)$-assignments. We may assume without loss of generality that $\{v_1, \ldots, v_{2n}\} =$

$\{w_1, \ldots, w_{2n}\} = \{1, \ldots, 2n\}$, since we can always add additional requests, then delete the resulting paths from the realization. We shall construct a bipartite graph from the n input nodes to the n output nodes by adding an *arc* from input node p to output node q for each request to connect an input of p (that is, p or $n+p$) to an output of q (that is, $\overline{q}$ or $\overline{n+q}$). Every node in the resulting bipartite graph has degree 2; thus the arcs may be assigned 2 colors (say, red and blue) in such a way that every node meets one arc of each color. The problem of constructing $2n$ vertex-disjoint paths in $C(2n)$ can now be reduced to that of constructing n vertex-disjoint paths in each copy of $C(n)$ by routing the requests corresponding to red arcs through the lower subnetwork, and those corresponding to blue arcs through the upper subnetwork. $\triangle$

The ideas underlying this proof of Theorem 2.1 were known to D. Slepian in 1952; they did not appear in print, however, until 1962, when the history of the problem was traced by Beneš [B]. By then they had been rediscovered at least twice, and they have continued to be rediscovered from time to time over the subsequent decades.

Theorem 2.2: For $n \geq 4$, we have $c(n) \geq (n/4)\log_2 n$.

Proof: In a network with c edges, there can be at most 2^c distinct sets of vertex-disjoint paths. Since there must be such a set for each of the $n!$ one-to-one correspondences between the n inputs an n outputs, we must have $2^c \geq n!$. Assuming that $n \geq 4$ is an integral power of 2, and noting that the $n/2$ largest factors in $n!$ are each at least $n/2$, we obtain $c \geq \log_2(n!) \geq (n/2)\log_2(n/2) \geq (n/4)\log_2 n$. $\triangle$

The idea behind this proof is due to Shannon [S].

3. Superconcentrators

A *superconcentration n-assignment* is a sequence $v = (v_1, \ldots, v_n)$ such that $v_i \in \{0,1\}$ for all $1 \leq i \leq n$. Say that two superconcentration n-assignments are *similar* if $\sum_{1\leq i\leq n} v_1 = \sum_{1\leq j\leq n} w_j$.

Let v and w be superconcentration n-assignments, and let N be an n-network. We shall say that N *realizes* the pair (v, w) if there exists a set P of vertex-disjoint paths from inputs to outputs in N such that (1) x_i is the origin of a path in P if and only if $v_i > 0$, and (2) y_j is the destination of a path in P if and only if $w_j > 0$. We may think of the non-zero elements of an assignment as "requests" of a single type, with zero elements indicating the absence of a request. A network then realizes a pair of assignments if it can "match" each input request with an output request.

Let $s(n)$ denote the minimum possible number of edges in an n-superconcentrator.

Theorem 3.1: We have $s(n) \leq 90n$.

For the proof of this theorem, we shall need an auxiliary notion. Consider a bipartite n-network N, that is, an n-network in which there are no links, so that every edge is directed from an input to an output. For set X of inputs, we define the set $N(X)$ to be the set of outputs y such that, for some input $x \in X$, there is an edge from x to y. We shall say that N is an *n-expander* if, for every X with $\#(X) \leq n/3$, we have $\#(N(X)) \geq 2\#(X)+1$. (Many different definitions of "expander" appear in the literature; the one given here caters to the proof of Theorem 3.1, but is similar in spirit to all the others.) Let $e(n)$ denote the minimum possible number of edges in an n-expander.

Proposition 3.2: We have $e(n) \leq 15n$.

Proof: We take N to have inputs $A = \{1, \ldots, n\}$ and outputs $B = \{\overline{1}, \ldots, \overline{n}\}$. If π is a permutation of the set $\{1, \ldots, n\}$, the set $\{(i, \overline{\pi(i)}) : 1 \leq i \leq n\}$ will be called the *graph* of π. We let the edges E of N be the union of the 15 graphs of 15 independent uniformly distributed random permutations of $\{1, \ldots, n\}$. The bipartite n-network N clearly has at most $15n$ edges. We shall show that it is an n-expander with probability strictly greater than 0.

If N is not an n-expander, then there exists k in the range $1 \leq k \leq n/3$, $X \subseteq A$ with $\#(X) = k$ and $Y \subseteq B$ with $\#(Y) = 2k$ such that each of the $15k$ edges that meets a vertex in X also meets a vertex in Y. There are $\binom{n}{k}$ ways to choose X with $\#(X) = k$ and $\binom{n}{2k}$ ways to choose Y with $\#(Y) = 2k$, and for each such pair of choices, the probability that each edge that meets X also meets Y is $\left(\binom{2k}{k} \big/ \binom{n}{k}\right)^{15}$. Thus the probability that N is not an n-expander is at most

$$S = \sum_{1 \leq k \leq n/3} \binom{n}{k} \binom{n}{2k} \left(\binom{2k}{k} \Big/ \binom{n}{k} \right)^{15},$$

and it will suffice to show that the sum S is strictly less than 1.

Since $\binom{n}{2k} \leq \binom{n}{k}^2$, we have

$$S \leq \sum_{1 \leq k \leq n/3} \binom{2k}{k}^3 \left(\binom{2k}{k} \Big/ \binom{n}{k} \right)^{12},$$

and since $\binom{2k}{k} \leq 2^{2k}$, this becomes

$$S \leq \sum_{1 \leq k \leq n/3} 2^{6k} \left(\binom{2k}{k} \Big/ \binom{n}{k} \right)^{12}.$$

Furthermore, we have $\binom{2k}{k} \,/\, \binom{n}{k} = 2k(2k-1)\cdots(k+1)/n(n-1)\cdots(n-k+1) \le (2k/n)^k$, and thus

$$S \le \sum_{1\le k\le n/3} 2^{6k}\,(2k/n)^{12k}.$$

Since $2k/n \le 2/3$, we have

$$S \le \sum_{1\le k\le n/3} 2^{6k}\,(2/3)^{12k},$$

and since $2^{18}/3^{12} < 1/2$, we have

$$S < \sum_{1\le k\le n/3} (1/2) < 1,$$

which completes the proof. $\triangle$

Proof of Theorem 3.1: We shall describe a recursive construction for an n-superconcentrator $S(n)$. For the basis of the recursion, we take $S(n)$ to be a complete bipartite graph if $n \ge 16$. For the recursive step, we describe how to construct a $(2n)$-superconcentrator from an n-superconcentrator, a $(2n)$-expander, and four copies of an n-expander. We will then have $s(n) \le 16n$ for $n \le 16$, together with the recurrence

$$\begin{aligned} s(2n) &\le s(n) + e(2n) + 4e(n) \\ &\le s(n) + 90n, \end{aligned}$$

from which it follows by induction that $s(n) \le 90n$.

The recursive construction for $S(n)$ is as follows. We take four disjoint sets of vertices: the inputs $\{1,\ldots,2n\}$, the links $\{\widehat{1},\ldots,\widehat{n}\}$ and $\{\widetilde{1},\ldots,\widetilde{n}\}$, and the outputs $\{\overline{1},\ldots,\overline{2n}\}$. We install a $(2n)$-expander from $\{1,\ldots,2n\}$ to $\{\overline{1},\ldots,\overline{2n}\}$, and four copies of an n-expander: from $\{1,\ldots,n\}$ to $\{\widehat{1},\ldots,\widehat{n}\}$, from $\{n+1,\ldots,2n\}$ to $\{\widehat{1},\ldots,\widehat{n}\}$, from $\{\overline{1},\ldots,\overline{n}\}$ to $\{\widetilde{1},\ldots,\widetilde{n}\}$, and from $\{\overline{n+1},\ldots,\overline{2n}\}$ to $\{\widetilde{1},\ldots,\widetilde{n}\}$. Finally, we install an n-superconcentrator from $\{\widehat{1},\ldots,\widehat{n}\}$ to $\{\widetilde{1},\ldots,\widetilde{n}\}$.

To see that the resulting network is a $(2n)$-superconcentrator, we consider a pair (v,w) of similar superconcentration $(2n)$-assignments. To find a set of vertex-disjoint paths satisfying these requests, we shall first find as many direct paths as possible through the $(2n)$-expander. There will remain a pair (v',w') of similar superconcentration $(2n)$-assignments corresponding to requests that were not satisfied in this way. We shall show below that at most $2\lfloor n/3\rfloor+1$ such requests remain. We shall then find as large a matching as possible between inputs r with $v'_r = 1$ and distinct links in $\{\widehat{1},\ldots,\widehat{n}\}$, through the two n-expanders joining these sets. We shall show below that this matching accomodates all the remaining requests. Similarly, we shall find a matching between outputs $\overline{r}$ such that $w'_{\overline{r}} = 1$ and distinct links in $\{\widetilde{1}, \ldots, \widetilde{n}\}$. The remainders of the paths satisfying the requests of v', w') are then furnished by the n-superconcentrator.

To see that at most $2\lfloor n/3 \rfloor + 1$ requests remain in (v', w'), we consider any maximal pair (v', w') of similar superconcentration $(2n)$- assignments with the property that no edge of the $(2n)$-expander joins a requesting input of v' to a requesting output of w'. Suppose that there are more than $2\lfloor n/3 \rfloor + 1$ requesting inputs in v'. Since $\lfloor 2n/3 \rfloor \leq 2\lfloor n/3 \rfloor + 1$, we can find a set X of comprising exactly $\lfloor 2n/3 \rfloor$ requesting inputs in v'. Since $\lfloor 2n/3 \rfloor \leq 2n/3$, these inputs are joined by edges of the $(2n)$-expander to at least $2\lfloor 2n/3 \rfloor + 1$ of the $2n$ outputs. This leaves at most $2n - (2\lfloor 2n/3 \rfloor + 1) \leq 2\lfloor n/3 \rfloor + 1$ outputs that include the requesting outputs in w'.

To see that all of the requesting inputs in v' can be matched to distinct links in $\{\widehat{1}, \ldots, \widehat{n}\}$, we use the marriage theorem, whereby it suffices to show that, for every $k \leq 2\lfloor n/3 \rfloor + 1$, every set of k requesting inputs in v' is joined by edges in the two n-expanders (the *upper* and *lower* n-expanders) to at least k distinct links. Of k such requesting inputs, we may assume without loss of generality that at least $\lfloor k/2 \rfloor$ belong to the upper n-expander. Since $\lfloor k/2 \rfloor \leq \lfloor n/3 \rfloor \leq n/3$, these $\lfloor k/2 \rfloor$ inputs are joined by edges in the upper n-expander to at least $2\lfloor k/2 \rfloor + 1 \geq k$ distinct links. $\triangle$

The idea behind the proof of Proposition 3.2 is due to Pinsker [P1]. The result of Theorem 3.1 is due to Valiant [V]; the outline of the proof we have given is due to Pippenger [P2].

4. Subconnectors

An *n-assignment* is a sequence $v = (v_1, \ldots, v_n)$ such that $v_i \in \{0, \ldots, n\}$ for all $1 \leq i \leq n$. Let v and w be n-assignments, and let N be an n-network. We shall say that N *realizes* the pair (v, w) if there exists a set P of vertex disjoint paths from inputs to outputs in N such that (1) x_i is the origin of a path in P if and only if $v_i > 0$, (2) y_j is the destination of a path in P if and only if $w_j > 0$, and (3) if there is a path from x_i to y_j in P, then $v_i = w_j$. We may think of the non-zero elements of an assignment as "requests" of n different types, with zero elements indicating the absence of a request. A network then realizes a pair of assignments if it can "match" each input request with an output request of the same type.

Let $J(u, v) = \#(\{i : u = v_i\})$ denote the number of requests of type u in v. Say that requests v and w are *similar* if $J(u, v) = J(u, w)$ for all $1 \leq u \leq n$. If a pair (v, w)

is realized by some network, then v and w must be similar. We shall be interested in networks for which the converse (or limited versions of it) holds.

An assignment v has *average* m if $v_1 + \cdots + v_n \le mn$. (We are departing from normal usage, which would describe such an assignment as having average at most m.) An *m-average n-subconnector* is an n-network that realizes every pair of similar n-assignments with average m.

We shall assume henceforth that m as well as n is an integral power of 2. Let $h(n, m)$ denote the minimum possible number of edges in an m-average n-subconnector.

Theorem 4.1: We have $h(n, m) \le 1456n \log_2(32m))$.

To prove Theorem 4.1 it will be convenient to have a somewhat less flexible network available as a building block.

An assignment v has *maximum* m if $v_i \le m$ for every $1 \le i \le n$. (We are again departing from normal usage, which would describe such an assignment as having maximum at most m.) An *m-maximum n-subconnector* is an n-network that realizes every pair of similar n-assignments with maximum m.

Let $g(n, m)$ denote the minimum possible number of edges in an m-maximum n-subconnector.

Theorem 4.2: We have $g(n, m) \le 728n \log_2(4m)$.

We shall prove Theorem 4.2 later. For now let us see how to use it to prove Theorem 4.1.

Proof of Theorem 4.1: We shall describe a recursive construction for an m-average n-subconnector $H(n, m)$. For the basis of the recursion, we take $H(n, n)$ to be the n-connector $C(n)$. For the recursive step, we describe how to construct an m-average $(2n)$-subconnector from a $(2m)$-maximum $(2n)$-subconnector, a $(2m)$-average n-subconnector, two copies of an n-superconcentrator, and $2n$ additional edges. We then have $h(n, n) \le c(n) \le 4n \log_2 n$ from Theorem 2.1 and

$$\begin{aligned} h(2n, m) &\le h(n, 2m) + g(2n, 2m) + 2s(n) + 2n \\ &\quad h(n, 2m) + 1456n \log_2(8m) + 182n \\ &\quad h(n, 2m) + 1456n \log_2(16m) \end{aligned}$$

from Theorems 4.2 and 3.1. It follows by induction that $h(n, m) \le 1456n \log_2(32m)$.

The recursive construction for $H(2n, m)$ is as follows. We take four disjoint sets of vertices: the inputs $\{1, \ldots, 2n\}$, the links $\{\widehat{1}, \ldots, \widehat{n}\}$ and $\{\widetilde{1}, \ldots, \widetilde{n}\}$, and the outputs $\{\overline{1}, \ldots, \overline{2n}\}$. We install a $(2m)$-maximum $(2n)$-subconnector from $\{1, \ldots, 2n\}$ to

$\{\overline{1},\ldots,\overline{2n}\}$, n edges from $\{1,\ldots,n\}$ to $\{\widehat{1},\ldots,\widehat{n}\}$, n edges from $\{\overline{1},\ldots,\overline{n}\}$ to $\{\widetilde{1},\ldots,\widetilde{n}\}$, and two copies of an n-superconcentrator: from $\{n+1,\ldots,2n\}$ to $\{\widehat{1},\ldots,\widehat{n}\}$, and from $\{\overline{n+1},\ldots,\overline{2n}\}$ to $\{\widetilde{1},\ldots,\widetilde{n}\}$. Finally, we install a $(2m)$-average n-subconnector from $\{\widehat{1},\ldots,\widehat{n}\}$ to $\{\widetilde{1},\ldots,\widetilde{n}\}$.

To see that the resulting network is a m-average $(2n)$-subconnector, consider a pair (v,w) of similar m-average $(2n)$-assignments. We observe that at most one-half of the requests (that is, at most n) can have types that exceed twice the average type (that is, m). These requests form a pair of similar $(2m)$-average n-assignments, for which paths will be found through the $(2m)$-average n-subconnector; the others form a pair of similar $(2m)$-maximum $(2n)$-assignments, for which paths will be found through the $(2m)$-maximum n-subconnector. It remains to show how the requests with types exceeding m can be joined to distinct links in $\{\widehat{1},\ldots,\widehat{n}\}$ and $\{\widetilde{1},\ldots,\widetilde{n}\}$. For this, it suffices first to join requests at inputs in $\{1,\ldots,n\}$ and outputs in $\{\overline{1},\ldots,\overline{n}\}$ through the $2n$ additional edges, then to join the remaining requests through the two n-superconcentrators. △

To prove Theorem 4.2 it will be convenient to have a still less flexible network available as a building block.

We shall say that an assignment v is *smooth* if $J(u,v)$ is an integral power of 2 whenever $u>0$ and $J(u,v)>0$. An *m-maximum n-infraconnector* is an n-network that realizes every pair of similar smooth n-assignments with maximum m.

Let $f(n,m)$ denote the minimum possible number of edges in an m-maximum n-infraconnector.

Theorem 4.3: We have $f(n,m)\le 364n\log_2(2m)$.

We shall prove Theorem 4.3 later. For now let us see how to use it to prove Theorem 4.2.

Proof of Theorem 4.2: We shall describe a recursive construction for an m-maximum n-subconnector $G(n,m)$. For the basis of the recursion, we take $G(n,n)$ to be the n-connector $C(n)$. For the recursive step, we describe how to construct an m-maximum $(2n)$-subconnector from an m-maximum $(2n)$-infraconnector, an m-maximum n-subconnector, two copies of an n-superconcentrator, and $2n$ additional edges. We then have $g(n,n)\le c(n)\le 4n\log_2 n$ from Theorem 2.1 and

$$g(2n,m)\le g(n,m)+f(2n,m)+2s(n)+2n$$
$$g(n,m))+728n\log_2(2m)+182n$$
$$g(n,m))+728n\log_2(4m)$$

from Theorems 4.2 and 3.1. It follows by induction that $g(n,m) \le 728n\log_2(4m)$.

The recursive construction for $G(2n,m)$ is as follows. We take four disjoint sets of vertices: the inputs $\{1,\ldots,2n\}$, the links $\{\widehat{1},\ldots,\widehat{n}\}$ and $\{\widetilde{1},\ldots,\widetilde{n}\}$, and the outputs $\{\overline{1},\ldots,\overline{2n}\}$. We install an m-maximum $(2n)$-infraconnector from $\{1,\ldots,2n\}$ to $\{\overline{1},\ldots,\overline{2n}\}$, n edges from $\{1,\ldots,n\}$ to $\{\widehat{1},\ldots,\widehat{n}\}$, n edges from $\{\overline{1},\ldots,\overline{n}\}$ to $\{\widetilde{1},\ldots,\widetilde{n}\}$, and two copies of an n-superconcentrator: from $\{n+1,\ldots,2n\}$ to $\{\widehat{1},\ldots,\widehat{n}\}$, and from $\{\overline{n+1},\ldots,\overline{2n}\}$ to $\{\widetilde{1},\ldots,\widetilde{n}\}$. Finally, we install an m-maximum n-subconnector from $\{\widehat{1},\ldots,\widehat{n}\}$ to $\{\widetilde{1},\ldots,\widetilde{n}\}$.

To see that the resulting network is a m-maximum $(2n)$-subconnector, we observe that a pair (v,w) of similar m-maximum $(2n)$-assignments can be split into a pair of similar smooth m-maximum $(2n)$-assignments together with a pair (v',w') of similar m-maximum $(2n)$-assignments in which there are at most n requests. (The largest integral power of 2 not exceeding l is at least $l/2$, which leaves at most $l/2$ requests for (v',w').) The smooth m-maximum $(2n)$-assignments can be routed through the $(2n)$-infraconnector, while the remaining requests (v',w') can be routed through the additional edges, the n-superconcentrators, and the n-subconnector, as in the proof of Theorem 4.1. △

For the proof of Theorem 4.3 we shall need the following intuitively obvious lemma.

Lemma 4.4: Let $2^\kappa \ge 2^{\lambda_1} \ge \cdots \ge 2^{\lambda_t} \ge 1$ be integral powers of 2 in non-increasing order with $2^{\lambda_1} + \cdots + 2^{\lambda_t} > 2^\kappa$. Let s be the smallest index such that $2^{\lambda_1} + \cdots + 2^{\lambda_s} \ge 2^\kappa$. Then $2^{\lambda_1} + \cdots + 2^{\lambda_s} = 2^\kappa$.

Proof: If $s = 1$ the lemma is obvious, so suppose that $s \ge 1$. Then $2^{\lambda_1} + \cdots + 2^{\lambda_{s-1}} < 2^\kappa$. Let $\Delta = 2^\kappa - (2^{\lambda_1} + \cdots + 2^{\lambda_{s-1}})$. Since $2^\kappa \ge 2^{\lambda_1} \ge \cdots \ge 2^{\lambda_t} \ge 1$, 2^{λ^s} divides each of $2^\kappa, 2^{\lambda_1} + \cdots + 2^{\lambda_{s-1}}$, and there fore also divides Δ. Since $\Delta \ge 1$, this implies $s^{\lambda_s} \le \Delta$, which yields the conclusion of the lemma. △

Proof of Theorem 4.3: We shall describe a recursive construction for an m-maximum n-infraconnector $F(n,m)$. For the bases of the recursion, we take $F(n,n)$ to be the n-connector $C(n)$, and take $F(n,1)$ to be the n-superconcentrator $S(n)$. For the recursive step, we describe how to construct a $(2m)$-maximum $(2n)$-infraconnector from a $(2m)$-maximum n-infraconnector, an m-maximum n-infraconnector, four sets of n additional edges, and four copies of an n-superconcentrator. We then have $f(n,n) \le c(n) \le 4n\log_2 n$ from Theorem 2.1, $f(n,1) \le s(n) \le 90n$ from Theorem 3.1, and

$$\begin{aligned} f(2n,2m) &\le f(n,2m) + f(n,m) + 4s(n) + 4n \\ &\le f(n,2m) + f(n,m) + 364n \end{aligned}$$

from Theorem 3.1. It follows by induction that $f(n,m) \leq 364n\log_2(2m)$.

The recursive construction for $F(2n, 2m)$ is as follows. We take four disjoint sets each containing $2n$ vertices: the inputs $\{1, \ldots, 2n\}$, the links $\{\widehat{1}, \ldots, \widehat{2n}\}$ and $\{\widetilde{1}, \ldots, \widetilde{2n}\}$, and the outputs $\{\overline{1}, \ldots, \overline{2n}\}$. We install four sets of n additional edges: from $\{1, \ldots, n\}$ to $\{\widehat{1}, \ldots, \widehat{n}\}$, from $\{n+1, \ldots, 2n\}$ to $\{\widehat{n+1}, \ldots, \widehat{2n}\}$, from $\{\widetilde{1}, \ldots, \widetilde{n}\}$ to $\{\overline{1}, \ldots, \overline{n}\}$, and from $\{\widetilde{n+1}, \ldots, \widetilde{2n}\}$ to $\{\overline{n+1}, \ldots, \overline{2n}\}$. We install four n-superconcentrators: from $\{1, \ldots, n\}$ to $\{\widehat{n+1}, \ldots, \widehat{2n}\}$, from $\{n+1, \ldots, 2n\}$ to $\{\widehat{1}, \ldots, \widehat{n}\}$, from $\{\widetilde{1}, \ldots, \widetilde{n}\}$ to $\{\overline{n+1}, \ldots, \overline{2n}\}$, and from $\{\widetilde{n+1}, \ldots, \widetilde{2n}\}$ to $\{\overline{1}, \ldots, \overline{n}\}$. Finally, we install a $(2m)$-maximum n-infraconnector from $\{\widehat{1}, \ldots, \widehat{n}\}$ to $\{\widetilde{1}, \ldots, \widetilde{n}\}$, and an m-maximum n-infraconnector from $\{\widehat{n+1}, \ldots, \widehat{2n}\}$ to $\{\widetilde{n+1}, \ldots, \widetilde{2n}\}$.

To see that the resulting network is a $(2m)$-maximum $(2n)$-infraconnector, it suffices to show that a pair of similar smooth $(2m)$-maximum $(2n)$-assignments can be split into a pair of similar smooth $(2m)$-maximum $(2n)$-assignments with at most n requests and a pair of similar smooth m-maximum $(2n)$-assignments with at most n requests. The first pair can then be routed through the $(2m)$-maximum n-infraconnector, and the second pair through the m-maximum n-infraconnector, in each case as in the proof of Theorem 4.2.

To exhibit the split in question, we simply consider the types in non-increasing order according to the number of requests, considering all requests of a given type together, and making the split at the first moment at which we have either considered at least m types or considered at least n requests. If we split at a moment when we have considered at least m types, then at the first such moment we have considered exactly m types, and at most n requests. If we split at a moment when we have considered at least n requests, then Lemma 4.4 guarantees that at the first such moment we will have considered exactly n requests, and at most m requests. In either case we obtain the desired split. $\triangle$

The problems dealt with in this section were proposed by C. J. Smyth.

5. Conclusion

We have discussed several types of rearrangeable circuit-switching networks, always with "network" in its sense of "acyclic directed graph". There are numerous variants of the problems discussed here in which additional restrictions are imposed on the graphs. One may, for example, restrict the lengths of paths from inputs to outputs, or require the graph to be planar. Much of the literature on the first type of restriction is surveyed in the review paper by Pippenger [P3].

A planar n-superconcentrator can be implemented as a simple n-by-n "grid" with $O(n^2)$ edges. Lipton and Tarjan [LT] showed, using their "planar separator theorem", that $\Omega(n^2)$ edges are necessary in a planar n-superconcentrator.

For a planar n-connector, Cutler and Shiloach [CS] gave a construction (again based on a grid) using $O(n^3)$ edges, and Klawe and Leighton [KL] showed that $\Omega(n^3)$ edges are necessary (their proof uses the planar separator theorem, and expanders as well).

The simplest conjecture that agrees with these results at the "endpoints" $m = 1$ and $m = n$ seems to be that, for a planar m-average (or m-maximum) n-subconnector, $\Theta(n^2 m)$ edges are necessary and sufficient. Is this true?

6. References

[B] V. E. Beneš, "Optimal Rearrangeable Multistage Connecting Networks", *Bell Sys. Tech. J.*, 43 (1964) 1641–1656.

[CS] M. Cutler and Y. Shiloach, "Permutation Layout", *Networks*, 8 (1978) 253–278.

[KL] M. Klawe and T. Leighton, "A Tight Lower Bound on the Size of Planar Permutation Networks", *SIGAL*, 1 (1990) 281–287.

[LT] R. J. Lipton and R. E. Tarjan, "A Separator Theorem for Planar Graphs", *SIAM J. Appl. Math.*, 36 (1979) 177–189.

[P1] M. S. Pinsker, "On the Complexity of a Concentrator", *Proc. Internat. Teletraffic Conf.*, 7 (1973) 318/1–4.

[P2] N. Pippenger, "Superconcentrators", *SIAM J. Computing*, 6 (1977) 298–304.

[P3] N. Pippenger, "Communication Networks", in J.van Leeuwen (Ed.), *Handbook of Theoretical Computer Science—Volume A: Algorithms and Complexity*, Elsevier, Amsterdam, 1990.

[S] C. E. Shannon, "Memory Requirements in a Telephone Exchange", *Bell Sys. Tech. J.*, 29 (1950) 343–349.

[V] L. G. Valiant, "Graph-Theoretic Properties in Computational Complexity", *J. Comp. and Sys. Sci.*, 13 (1976) 278–285.