

1-1-1988

Wide-Sense Nonblocking Networks

Paul Feldman

Massachusetts Institute of Technology

Joel Friedman

Princeton University

Nicholas Pippenger

Harvey Mudd College

Recommended Citation

Paul Feldman, Joel Friedman, and Nicholas Pippenger. "Wide-Sense Nonblocking Networks", *Society for Industrial and Applied Mathematics Journal of Discrete Mathematics*, 1, 158 (1988).

This Article is brought to you for free and open access by the HMC Faculty Scholarship at Scholarship @ Claremont. It has been accepted for inclusion in All HMC Faculty Publications and Research by an authorized administrator of Scholarship @ Claremont. For more information, please contact scholarship@cuc.claremont.edu.

WIDE-SENSE NONBLOCKING NETWORKS*

PAUL FELDMAN†, JOEL FRIEDMAN‡, AND NICHOLAS PIPPENGER§

Abstract. A new method for constructing wide-sense nonblocking networks is presented. Application of this method yields (among other things) wide-sense nonblocking generalized connectors with n inputs and outputs and size $O(n \log n)$, and with depth k and size $O(n^{1+1/k}(\log n)^{1-1/k})$.

Key words. nonblocking network, connection network, concentrator, generalizer

AMS(MOS) subject classifications. 94C15, 68E10, 05C35

1. Introduction. In this introductory section, we shall give an informal account of our results. More precise definitions will be given in later sections.

Informally, a “network” is an acyclic directed graph that allows connections to be made between “inputs” and “outputs.” A “nonblocking” network is one that allows connections to be created and destroyed dynamically, with each request for creation being satisfied without disturbing other previously established connections.

Given a network, we may consider a game between two players, the “blocker” (representing the collective activity of the users interconnected by the network) and the “router” (representing the activity of the controller of the network). The positions in this game are the states of the network, and the initial position is the state in which no connections are established. The players move alternately, starting with the blocker. The blocker moves by making a request to create or destroy a connection, and the router responds by satisfying this request (if possible). The goal of the blocker is to make a request that cannot be satisfied, and the goal of the router is to avoid this by satisfying every request. Roughly speaking, a “wide-sense” nonblocking network is one in which the router can indefinitely avoid losing. (In addition to wide-sense nonblocking networks, there are “strictly” nonblocking networks and “rearrangeable” networks. We do not need these notions to formulate our results, but we shall occasionally need them to discuss other results in the literature.)

The main contribution of this paper is a new principle, according to which certain “expanding graphs” are shown to be “generalized concentrators” (a special case of which are “concentrators”). For these networks, inputs request connection to one or more outputs, without specifying to which output they are to be connected. This principle will be presented in § 2.

In § 3, we apply the results of § 2 to construct “generalized connectors” (a special case of which are “connectors”). For these networks, inputs request connection to one or more specified outputs. Our results in this section, taken together with results of Friedman [F], show that wide-sense nonblocking connectors can have smaller size than strictly nonblocking connectors with the same number of inputs and outputs and the same depth. This answers a question that has been open since the distinction between “wide-sense” and “strictly” nonblocking networks was introduced by Beneš in 1962 [B1] (see also [B2], [B3]).

In § 4, we apply the results of § 2 to construct “bi-generalized connectors.” This new type of network provides greater interconnection facility than most other types

* Received by the editors June 17, 1987; accepted for publication (in revised form) November 6, 1987.

† Massachusetts Institute of Technology, Cambridge, Massachusetts 02139.

‡ Princeton University, Princeton, New Jersey 08544.

§ IBM Almaden Research Center, San Jose, California 95120-6099.

considered in the literature, and it serves as a natural setting to display the power of the new principle.

The results of §§ 3 and 4 concern networks with “limited depth,” in which the length of the routes from inputs to outputs is constrained. In § 5 we apply the results of § 2 to construct networks of “unlimited depth,” in which only the “size” (corresponding to the number of switches in the network) is to be minimized. Our results in this section, taken together with results of Shannon [S], show that wide-sense nonblocking generalized connectors need be larger than wide-sense nonblocking connectors by only a constant factor. The corresponding result for “rearrangeable” rather than “wide-sense nonblocking” networks has been known since the work of Ofman [O] in 1965. The corresponding result for “strictly nonblocking” rather than “wide-sense nonblocking” networks does not hold (see Bassalygo and Pinsker [BP2]).

A preliminary account of some of these results appears in our earlier paper [FFP].

2. Generalized concentrators. In this section we shall define networks precisely and present the principle on which most of the results of this paper are based. We shall base our formulation of a “network” on a directed graph with “inputs” and “outputs.” There is an alternative formulation based on an undirected graph with “terminals.” Our choice for this paper is technically simpler, though our results hold (with obvious modifications) for either formulation.

A *network* is an acyclic directed graph with a set of distinguished vertices called *inputs* and a disjoint set of distinguished vertices called *outputs*. A vertex that is neither an input nor an output will be called a *link*. We shall use the prefix “ (n, m) -” to the term “network” (and later to other nouns) to indicate that there are n inputs and m outputs; we shall use the prefix “ n -” in these cases in place of “ (n, n) -”

A *route* in a network is a directed path from an input to an output. The *size* of a network is the number of edges in it. The *depth* of a network is the maximum number of edges in a route in it.

Two routes are *compatible* if their intersection is an initial segment (possibly empty) of both of them. A *state* of a network is a set of pairwise compatible routes. (A state may be regarded as a forest of vertex-disjoint trees, with their roots at the inputs and their leaves at the outputs. This alternative viewpoint is taken in many papers on networks with “generalized” properties.) The states of a network are partially ordered by inclusion, and we shall use the terms *above* and *below* to refer to this partial order. A vertex or edge is *busy* in a state if it appears in some route of that state; otherwise, it is *idle*.

A *concentration request* in a network is an input of the network. A concentration request is *satisfied* by a route if the route originates at the input of the request. A *generalized concentration assignment* is a multiset of concentration requests. A generalized concentration assignment is *realized* by a state if each request is satisfied in the state by a number of routes equal to its multiplicity.

A generalized concentration assignment is (a, r) -*limited* if each request has multiplicity at most r and the sum of the multiplicities is at most a . A state is (a, r) -*limited* if the generalized concentration assignment it realizes is (a, r) -*limited*. A concentration request is (a, r) -*limited* in a state if the generalized concentration assignment obtained by adjoining the request to the generalized concentration assignment realized by the state is (a, r) -*limited*.

A *wide-sense nonblocking (a, r) -limited generalized concentrator* is a network for which there exists a set of distinguished states, called the *safe* states, with the following properties: (1) the empty state is safe; (2) any state below a safe state is itself safe; and (3) given any safe state and any (a, r) -limited concentration request in that state, there

exists a safe state above the given state containing another route satisfying the given request. (In our previous paper [FFP], “generalized concentrators” were called “generalizers.” This change has been adopted to increase the uniformity of our terminology.)

A *wide-sense nonblocking a -limited concentrator* is a wide-sense nonblocking $(a, 1)$ -limited generalized concentrator. When discussing concentrators, we shall use the adjective “ a -limited” in place of “ $(a, 1)$ -limited.”

The key to the results of this paper is the following proposition.

PROPOSITION 1. *If F is a network of depth one in which every set of x inputs, where $1 \leq x \leq 2a$, has at least $(r + s)x$ neighbors, then F is a wide-sense nonblocking (as, r) -limited generalized concentrator.*

Proof. Let $F = (V, W, E)$ be a network of depth one, with inputs V , outputs W , and edges E . Suppose that if $X \subseteq V$, with $|X| = x$ and $1 \leq x \leq 2a$, then $|E(X)| \geq (r + s)x$. (Here $E(X)$ denotes the set of neighbors of X , and $|Y|$ denotes the cardinality of Y .) We shall show that F is a wide-sense nonblocking (as, r) -limited generalized concentrator.

The first step of the proof is the definition of the set of safe states; for this we shall need some auxiliary definitions. Let X be a set of inputs. We shall define the *assets* of X in S , $A_S(X)$, to be the number of idle neighbors of X in S . Let v be an input in V . We shall define the *liabilities* of v in S , $B_S(v)$, to be r minus the number of busy edges directed out of v in S . We shall define the *liabilities* of X in S , $B_S(X)$, to be the sum of $B_S(v)$ over inputs v in X . We define the *balance* of X in S , $C_S(X)$, to be $A_S(X) - B_S(X)$. We shall say that X is *solvent* in S if $C_S(X) \geq 0$, that X is *critical* in S if $C_S(X) = 0$, and that X is *bankrupt* in S if $C_S(X) < 0$. Finally, we shall say that S is *safe* if, for every $X \subseteq V$ with $|X| = x$ and $1 \leq x \leq 2a$, X is solvent in S .

Let O denote the empty state. We shall show that O is safe. Let X be a set of inputs with $|X| = x$ and $1 \leq x \leq 2a$. By the hypothesis on F , $|E(X)| \geq (r + s)x \geq rx$. Since all vertices are idle in O , $A_O(X) \geq rx$. Since $B_O(X) = rx$, X is solvent, and thus O is safe.

Let S be a state, let (v, w) be a route not in S , and suppose that $T = S \cup \{(v, w)\}$ is a safe state. We shall show that S is safe. We observe that (1) $A_T(X) = A_S(X) - 1$ if $w \in E(X)$, otherwise $A_T(X) = A_S(X)$; and (2) $B_T(X) = B_S(X) + 1$ if $v \in X$, otherwise $B_T(X) = B_S(X)$. Since $v \in X$ implies $w \in E(X)$, $C_S(X) \geq C_T(X) \geq 0$, and thus S is safe.

Finally, let S be a safe state and let v be an (as, r) -limited concentration request in S . We shall show that for some idle neighbor w of v in S , the state $S \cup \{(v, w)\}$ is safe. Suppose, to obtain a contradiction, that for every idle neighbor w of v , the state $S \cup \{(v, w)\}$ is unsafe. Then for every idle neighbor w of v , there is a set of inputs X_w with $|X_w| \leq 2a$ such that X_w is bankrupt in $S \cup \{(v, w)\}$. Of course, since S is safe, for every w , X_w is solvent in S . From these conditions, and from observations (1) and (2) above, it is easy to see that for every idle neighbor w of v we have (i) w belongs to $E(X_w)$, but v does not belong to X_w , and (ii) X_w is critical in S . We shall need the following lemmas concerning critical sets.

LEMMA 1.1. *If X is critical in S , then $|X| < a$.*

Proof. Suppose, to obtain a contradiction, that X with $|X| = x$ and $a \leq x \leq 2a$ is unsafe in S . By the hypothesis on F , $|E(X)| \geq (r + s)x$. Since the request v is (as, r) -limited in S , fewer than sa outputs are busy in S , so $A_S(X) > (r + s)x - sa$. Since $B_S(X) \leq B_O(X) = rx$ and $x \geq a$, we have $C_S(X) > 0$, contradicting the assumption that X is critical in S . \square

LEMMA 1.2. *For any state S , $C_S(X)$ is submodular in X ; that is, $C_S(X \cup Y) + C_S(X \cap Y) \leq C_S(X) + C_S(Y)$.*

Proof. Since $B_S(X)$ is defined by a sum over the inputs in X , it is *modular* in X ; that is, $B_S(X \cup Y) + B_S(X \cap Y) \leq B_S(X) + B_S(Y)$. Thus it remains to show that $A_S(X)$

is submodular in X . This follows from the relations $E(X \cup Y) = E(X) \cup E(Y)$ and $E(X \cap Y) \subseteq E(X) \cap E(Y)$, and the identity $|X \cup Y| + |X \cap Y| = |X| + |Y|$. \square

LEMMA 1.3. *If S is a safe state, $|X|, |Y| \leq 2a$, and X and Y are critical in S , then $|X \cup Y| \leq 2a$ and $X \cup Y$ is critical in S .*

Proof. Since X and Y are critical for S , Lemma 1.1 implies $|X|, |Y| < a$, and thus $|X \cup Y| \leq 2a$. Thus it remains to show that $X \cup Y$ is critical in S . Since X and Y are critical, $C_S(X) + C_S(Y) = 0$. By Lemma 1.2, $C_S(X \cup Y) + C_S(X \cap Y) \leq 0$. Since S is safe and $|X \cap Y| < a$, $X \cap Y$ is solvent. Thus $X \cup Y$ is critical. \square

Returning to the proof of Proposition 1, let the set of inputs X be the union of X_w over all idle neighbors w of v . Since each X_w is critical in S , Lemma 1.3 implies that X is critical in S , and Lemma 1.1 implies $|X| < a$. Now consider the set $X \cup \{v\}$. Since S is safe and $|X \cup \{v\}| \leq a$, $X \cup \{v\}$ is solvent in S . On the other hand, (i) and (ii) above imply that every idle neighbor of v belongs to $E(X)$, but v does not belong to X . Thus $A_S(X \cup \{v\}) = A_S(X)$ and, since the request v is (a, r) -limited in S ,

$$B_S(X \cup \{v\}) = B_S(X) + 1.$$

Thus, since X is critical in S , $X \cup \{v\}$ is bankrupt in S . This contradiction completes the proof of Proposition 1. \square

The argument used in the proof of Proposition 1 has also found application to some combinatorial problems concerning fault-tolerant networks (see Friedman and Pippenger [FP]) and processor-interconnection networks (see Peleg and Upfal [PU]). This argument suffers the defect that it provides no efficient algorithm for finding safe routes. Indeed, determining the solvency of a state is a minor variant of a co-NP-complete problem (see Blum et al. [BKVPY]). We shall present some results (Theorems 4 and 6) that avoid this defect; it is avoided in another way in the work of Peleg and Upfal [PU].

A wide-sense nonblocking (a, r) -limited generalized concentrator remains one if any number of inputs (together with the edges incident with them) are deleted and any number of outputs (with no incident edges) are added.

If F is an (n, m) -network and G is an (m, l) -network, let $F \circ G$ denote a network obtained by identifying the outputs of F with the inputs of G in any one-to-one fashion. If F is an n -network, define $F^{\circ k}$ by letting $F^{\circ 1} = F$ and $F^{\circ k} = F \circ F^{\circ(k-1)}$ for $k \geq 2$.

PROPOSITION 2. *If F is a wide-sense nonblocking (a, r) -limited generalized m -concentrator and G is a wide-sense nonblocking (a, s) -limited generalized m -concentrator, then $F \circ G$ is a wide-sense nonblocking (a, rs) -limited generalized m -concentrator.*

Proof. Let the safe states of $F \circ G$ be those in which the induced states of F and G are safe. Clearly, the empty state is safe and any state below a safe state is itself safe. Let S be a safe state and let v be an (a, rs) -limited concentration request in S . Let $S(F)$ be the state of F induced by S and let $S(G)$ be the state of G induced by S .

Since v is (a, rs) -limited, S contains at most $a - 1$ routes, at most $rs - 1$ of which are from v . Thus either (1) there are routes in $S(F)$ from v to at most $r - 1$ outputs of F , or (2) there is a route in $S(F)$ from v to some output w of F and there are routes in $S(G)$ from w to at most $s - 1$ outputs of G .

In case (1), the concentration request v is (a, r) -limited in $S(F)$, so there is a safe state $T(F)$ above $S(F)$ containing a route from v to some output w of F that is idle in $S(F)$. The concentration request w is then (a, s) -limited in $S(G)$, so there is a safe state $T(G)$ above $S(G)$ containing a route from w to some output of G that is idle in $S(G)$. Then the state obtained by adjoining to S the concatenation of the route from v to w in $T(F)$ with the route from w in $T(G)$ is a safe state above S containing a route from v to an output of $F \circ G$ that is idle in S .

In case (2), the concentration request w is (a, s) -limited in $S(G)$, so there is a safe state $T(G)$ above $S(G)$ containing a route from w to an output of G that is idle in $S(G)$. Then the state obtained by adjoining to S the concatenation of the route from v to w in $S(F)$ with the route from w in $T(G)$ is a safe state above S containing a route from v to an output of $F \circ G$ that is idle in S . \square

3. Generalized connectors. In this section we shall use Proposition 1 to construct wide-sense nonblocking generalized connectors. Our first goal is to define these networks.

A *connection request* in a network is an ordered pair comprising an input and an output. A connection request is *satisfied* by a route if the route originates at the input of the request and terminates at the output of the request. A *generalized connection assignment* is a set of connection requests, no two of which have an output in common. A generalized connection assignment is *realized* by a state if each request of the generalized connection assignment is satisfied by some route of the state.

A generalized connection assignment is (a, r) -limited if it contains at most a requests, of which at most r have an input in common. Observe that a state is (a, r) -limited if the generalized connection assignment that it realizes is (a, r) -limited. A connection request is (a, r) -limited in a state if (1) the output of the request is idle in the state and (2) the generalized connection assignment obtained by adjoining the connection request to the generalized connection assignment realized by the state is (a, r) -limited.

A *wide-sense nonblocking (a, r) -limited generalized connector* is a network for which there exists a set of distinguished states, called the *safe states*, with the following properties: (1) the empty state is safe; (2) any state below a safe state is itself safe; and (3) given any safe state and any (a, r) -limited connection request in that state, there exists a safe state above the given state containing a route satisfying the given request.

A *wide-sense nonblocking a -limited connector* is a wide-sense nonblocking $(a, 1)$ -limited generalized connector. When discussing (un)generalized connectors, we shall use the term “connection assignment” in place of “generalized connection assignment” and the adjective “ a -limited” in place of “ $(a, 1)$ -limited.” We shall also use the term “connector” in place of “ ∞ -limited connector.”

Our main result in this section is the following.

THEOREM 3. *For every fixed $k \geq 1$ and $n \rightarrow \infty$, a wide-sense nonblocking generalized n -connector exists with depth k and size $O(n^{1+1/k}(\log n)^{1-1/k})$.*

The bound of Theorem 3 is almost the best possible; Pippenger and Yao [PY] have shown that size $\Omega(n^{1+1/k})$ is necessary (even for networks satisfying a much weaker property, that of being “rearrangeable n -shifters” with depth k).

The previously best upper bound for wide-sense nonblocking generalized connectors was obtained from a construction due to Pippenger [P1] and Nassimi and Sahni [NS]. It yields wide-sense nonblocking generalized connectors with depth $j^2 - 3j + 3$ and size $O(n^{1+2/j})$. The special case of depth three and size $O(n^{5/3})$ was given (using a different method) by Masson and Jordan [MJ].

The construction of Pippenger [P1] and Nassimi and Sahni [NS] requires wide-sense nonblocking (un)generalized connectors as components. The previously best upper bound for these networks was obtained from constructions due to Clos [Cl], Cantor [Ca], and Pippenger [P2]. It yields networks of depth $2j - 1$ and size $O(n^{1+1/j})$. (These networks satisfy a stronger property, that of being “strictly nonblocking n -connectors” with depth $2j - 1$.)

The best upper bound for wide-sense nonblocking connectors is now given by Theorem 3. The best upper bound for “rearrangeable generalized connectors” is also now given by Theorem 3 (the previously best bound was $O(n \log n)^{1+1/k}$, due to Dolev et

al. [DDPW]). Indeed, the bound of Theorem 3 is not much larger than the currently best bound of $O(n^{1+1/k}(\log n)^{1/k})$ for “rearrangeable connectors,” due to Pippenger and Yao [PY].

The lower bound of $\Omega(n^2)$ (for any depth) due to Bassalygo and Pinsker [BP2] shows that Theorem 3 cannot be extended to “strictly nonblocking generalized connectors.” The lower bound of $\Omega(n^{1+1/(k-1)})$ (for depth k) due to Friedman [F] shows that it cannot even be extended to “strictly nonblocking connectors.”

We need the following lemma for the proof of Theorem 3.

LEMMA 3.1. *For every n and a , there exists a network of depth one with n inputs and $8a$ outputs in which (1) every input has $d = \lceil \log_2(2n^3) \rceil$ neighbors, and (2) every set of x inputs with $1 \leq x \leq 2a$ has at least $2x$ neighbors.*

Proof. Consider a random graph in which each input has a randomly chosen set of d outputs as neighbors, with all sets of d outputs being equally likely and with the random choices for distinct inputs being independent. Clearly (1) is satisfied; it remains to show that the probability that (2) is satisfied is strictly positive. The probability that (2) fails is at most

$$\sum_{1 \leq x \leq 2a} \binom{n}{x} \binom{8a}{2x} \left(\binom{2x}{d} / \binom{8a}{d} \right)^x.$$

Using the estimates

$$\binom{n}{x} \binom{8a}{2x} \leq n^{3x}$$

and

$$\binom{2x}{d} / \binom{8a}{d} \leq \left(\frac{2x}{8a} \right)^d \leq \left(\frac{1}{2} \right)^d,$$

we see that this probability is at most $\sum_{1 \leq x \leq 2a} (1/2)^x < 1$. \square

A wide-sense nonblocking (a, r) -limited generalized connector remains one if any number of inputs and outputs (together with the edges incident with them) are deleted.

The following propositions and their subsequent use are the counterparts for wide-sense nonblocking generalized connectors of a method used by Pippenger [P1] for “strictly nonblocking generalized connectors” and by Nassimi and Sahni [NS] for “rearrangeable generalized connectors.”

PROPOSITION 3.2. *If F is a wide-sense nonblocking a -limited (n, m) -concentrator and G is a wide-sense nonblocking (a, r) -limited generalized (m, l) -connector, then $F \circ G$ is a wide-sense nonblocking (a, r) -limited generalized (n, l) -connector.*

Proof. Let the safe states of $F \circ G$ be those in which the induced states of F and G are safe. Clearly, the empty state is safe and any state below a safe state is itself safe. Let S be a safe state and let (v, w) be an (a, r) -limited connection request in S . Let $S(F)$ be the state of F induced by S and let $S(G)$ be the state of G induced by S . Either (1) v is idle in S , or (2) S contains a route originating at v .

In case (1), the concentration request v is a -limited in $S(F)$, so there is a safe state $T(F)$ above $S(F)$ containing a route from v to an output u of F that is idle in $S(F)$. The connection request (u, w) is (a, r) -limited in $S(G)$, so there is a safe state $T(G)$ above $S(G)$ containing a route from u to w . Then the state obtained by adjoining to S the concatenation of the route from v to u in $T(F)$ with the route from u to w in $T(G)$ is a safe state above S containing a route from v to w .

In case (2), the route originating at v passes through some link u that is an input of G . The connection request (u, w) is (a, r) -limited in $S(G)$, so there is a safe state $T(G)$ above $S(G)$ containing a route from u to w . Then the state obtained by adjoining to S the concatenation of the route from v to u in $S(F)$ with the route from u to w in $T(G)$ is a safe state above S containing a route that from v to w . \square

If F is an (n, m) -network and G is an (n, l) -network, let $F \vdash G$ denote an $(n, m + l)$ -network obtained by identifying the inputs of F with the inputs of G in any one-to-one fashion. If F is any network, define F^{t+1} by letting $F^{t+1} = F$ and $F^{t+2} = F \vdash F^{t+1}$ for $t \geq 2$.

PROPOSITION 3.3. *If F is a wide-sense nonblocking (a, r) -limited generalized (n, m) -connector and G is a wide-sense nonblocking (a, r) -limited generalized (n, l) -connector, then $F \vdash G$ is a wide-sense nonblocking (a, r) -limited generalized $(n, m + l)$ -connector.*

Proof. Let the safe states of $F \vdash G$ be those that induce safe states of F and G . Clearly, the empty state is safe and any state below a safe state is itself safe. Let S be a safe state and let the connection request (v, w) be (a, r) -limited in S . Without loss of generality, we may suppose that w is an output of F .

Let $S(F)$ be the state of F induced by S . The connection request (v, w) is (a, r) -limited in $S(F)$, so there is a safe state $T(F)$ above $S(F)$ containing a route from v to w . Then the state obtained by adjoining this route to S is a safe state above S containing a route from v to w . \square

Proof of Theorem 3. We proceed by induction on k . For $k = 1$, the conclusion is trivial. For $k \geq 2$, take $a = \lceil n^{1-1/k}(\log n)^{1/k} \rceil$ in Lemma 3.1 and $r = s = 1$ in Proposition 1 to obtain a wide-sense nonblocking concentrator F with n inputs, $8a$ outputs, depth one, and size $O(n \log n)$. By inductive hypothesis, there exists a wide-sense nonblocking generalized $(8a)$ -connector G with depth $k - 1$ and size

$$O(a^{1+1/(k-1)}(\log a)^{1-1/(k-1)}) = O(n \log n).$$

By Proposition 3.2, $H = F \circ G$ is a wide-sense nonblocking generalized $(n, 8a)$ -connector with depth k and size $O(n \log n)$. Let $t = \lceil n/8a \rceil = O(n^{1/k}/(\log n)^{1/k})$. By Proposition 3.3, H^{t+1} is a wide-sense nonblocking generalized $(n, 8at)$ -connector with depth k and size $O(n^{1+1/k}(\log n)^{1-1/k})$. Deleting all but n outputs of this network completes the proof. \square

The networks of Theorem 3 are obtained by the nonconstructive Lemma 3.1; moreover, since they rely on Proposition 1, no efficient algorithm is known for finding safe routes. The following remaining results of this section mitigate these defects in some special cases.

THEOREM 4. *Wide-sense nonblocking generalized n -connectors can be explicitly constructed with depth two and size $O(n^{5/3})$, and with depth three and size $O(n^{11/7})$.*

For the proof of Theorem 4 we shall need the following lemma.

LEMMA 4.1. *For every prime power q , a (q^3, q^2) -network with depth one can be explicitly constructed in which every input has q neighbors and every pair of distinct inputs has at most two common neighbors.*

Proof. Let $\text{GF}(q)$ be the field with q elements. Let the inputs be the elements of $\text{GF}(q)^3$ and let the outputs be the elements of $\text{GF}(q)^2$. Let the input (a, b, c) have the output (x, y) as a neighbor if and only if $ax^2 + bx + c = y$ (where all arithmetic is in $\text{GF}(q)$).

Every input has q neighbors, one for each possible value of x . If (x, y) is a common neighbor of (a, b, c) and (a', b', c') , then $ax^2 + bx + c = y$ and $a'x^2 + b'x + c' = y$. Subtracting these equations, we have $(a - a')x^2 + (b - b')x + (c - c') = 0$. If (a, b, c) is distinct from (a', b', c') , the coefficients of this quadratic equation do not all vanish,

and thus it can be satisfied by at most two values of x . Thus this pair of inputs can have at most two common neighbors, one for each of these values of x . \square

Let $F(q)$ denote the network constructed in the proof of Lemma 4.1.

PROPOSITION 4.2. *The network $F(5^\lambda)$ is a wide-sense nonblocking $(2 \cdot 5^{2\lambda-2})$ -limited concentrator.*

Proof. Let the safe states be those in which every input has at most $3 \cdot 5^{\lambda-1}$ neighbors that terminate routes originating at other inputs. Clearly, the empty state is safe and any state below a safe state is itself safe. Consider a safe state S in which some concentration request v cannot be satisfied. We shall show that at least $2 \cdot 5^{2\lambda-2}$ outputs are busy in S .

Since v is idle in S , v has at most $3 \cdot 5^{\lambda-1}$ busy neighbors in S . Thus it has at least $2 \cdot 5^{\lambda-1}$ idle neighbors in S . Say that an output w is *critical* if w is a neighbor of v , w is idle in S , and $S \cup \{(v, w)\}$ is not safe. Then there are at least $2 \cdot 5^{\lambda-1}$ critical outputs.

Say that an input u is *critical* if u is distinct from v and u has $3 \cdot 5^{\lambda-1}$ neighbors that terminate routes that originate at inputs other than u . Each critical output w must be the neighbor of some critical input, else $S \cup \{(v, w)\}$ would be safe. Each critical input u can have at most two critical outputs as neighbors, else the distinct inputs v and u would have more than two common neighbors. Thus there must be at least $5^{\lambda-1}$ critical inputs.

Let X be a set of $5^{\lambda-1}$ critical inputs. Each critical input in X has at least $3 \cdot 5^{\lambda-1}$ busy neighbors. Furthermore, each of the $\binom{5^{\lambda-1}}{2}$ pairs of distinct critical inputs in X has at most two busy neighbors in common. Thus there are at least

$$5^{\lambda-1}(3 \cdot 5^{\lambda-1}) - \binom{5^{\lambda-1}}{2} \cdot 2 \geq 2 \cdot 5^{2\lambda-2}$$

busy outputs in S . \square

Proof of Theorem 4. To prove the first assertion, we may assume that $n = 5^{3\lambda}$ for some integer $\lambda \geq 1$. Let $q = 5^\lambda$ and $a = 2 \cdot 5^{2\lambda-2}$. Let G be a wide-sense nonblocking generalized (q^2, a) -connector with depth one and size $q^2a = O(n^{4/3})$ (in which each input is joined by an edge to each output). By Proposition 3.2, $H = F(q) \circ G$ is a wide-sense nonblocking generalized (n, a) -connector with depth two and size $O(n^{4/3})$. Let $t = \lceil n/a \rceil = O(n^{1/3})$. By Proposition 3.3, H^{t^t} is a wide-sense nonblocking generalized (n, at) -connector with depth 2 and size $O(n^{5/3})$. Deleting all but n outputs completes the proof of the first assertion.

To prove the second assertion, we may assume that $n = 5^{7\mu}$ for some integer $\mu \geq 1$. Let $q = 5^{3\mu}$ and $a = 2 \cdot 5^{6\mu-2}$. By deleting all but n inputs from $F(q)$, we obtain a wide-sense nonblocking a -limited (n, q^2) -concentrator F' with depth one and size $O(n^{10/7})$. Let G' be a wide-sense nonblocking generalized (q^2, a) -connector with depth two and size $O(n^{10/7})$ constructed by taking $\lambda = 2\mu$ in the proof of the first assertion and deleting all but a outputs. By Proposition 3.2, $H = F' \circ G'$ is a wide-sense nonblocking generalized (n, a) -connector with depth three and size $O(n^{10/7})$. Let $t = \lceil n/a \rceil = O(n^{1/7})$. By Proposition 3.3, H^{t^t} is a wide-sense nonblocking generalized (n, at) -connector with depth three and size $O(n^{11/7})$. Deleting all but n outputs completes the proof of the second assertion. \square

We could extend Theorem 4 to yield size $O(n^{23/25})$ for depth four, and so forth. As $k \rightarrow \infty$, the sequence of exponents approaches, but remains greater than, $3/2$. In the following section, however, we shall obtain size $O(n^{3/2})$ for depth four, so this extension is of little interest.

4. Bi-generalized connectors. Our efforts to replace the probabilistic argument of the preceding section by an explicit construction lead naturally to consideration of the

best currently known explicit constructions for expanding graphs, due to Lubotzky, Phillips, and Sarnak [LPS] and Margulis [M2]. The networks produced with the aid of these constructions have greater size than those produced in the preceding section; they possess, however, a stronger property than that of being generalized connectors. We shall call them “bi-generalized connectors.” Our next goal is to define these networks.

Two routes are *bi-compatible* if their intersection is an initial segment or a final segment (possibly empty) of both of them. A *bi-state* is a set of pairwise bi-compatible routes. (A bi-state may be regarded as two forests with their leaves identified, one with its roots at the inputs and the other with its roots at the outputs.)

A *bi-generalized connection assignment* is a set of connection requests. A bi-generalized connection assignment is *realized* by a bi-state if each request of the assignment is satisfied by some route of the bi-state.

A bi-generalized connection assignment is *a-limited* if it contains at most a connection requests. A bi-state is *a-limited* if the bi-generalized connection assignment it realizes is *a-limited*. A connection request is *a-limited* in a bi-state if the bi-generalized connection assignment obtained by adjoining the connection request to the bi-generalized assignment realized by the bi-state is *a-limited*.

A *wide-sense nonblocking a-limited bi-generalized connector* is a network for which there exists a set of distinguished bi-states, called the *safe* bi-states, with the following properties: (1) the empty bi-state is safe; (2) any bi-state below a safe bi-state is itself safe; and (3) given any safe bi-state and any *a-limited* connection request in that bi-state, there exists a safe bi-state above the given bi-state and containing a route satisfying the given request.

We shall use the term “bi-generalized connector” in place of “bi-generalized ∞ -limited connector.”

Our main result in this section is the following.

THEOREM 5. *For every fixed $k \geq 2$ and $n \rightarrow \infty$, a wide-sense nonblocking bi-generalized n -connector can be explicitly constructed with depth k and size $O(n^{1+2/k})$.*

For the proof of Theorem 5 we shall need the following lemma.

LEMMA 5.1. *Let p and q be distinct primes congruent to one modulo four. Then a $(q + 1)$ -network of depth one can be explicitly constructed in which every input has $p + 1$ outputs as neighbors, every output is a neighbor of $p + 1$ inputs and every set of x inputs with $1 \leq x \leq 12(q + 1)/(p + 1)$ has at least $x(p + 1)/16$ neighbors.*

Proof. By the results of Lubotzky, Phillips, and Sarnak [LPS], an undirected graph G can be explicitly constructed in which there are $q + 1$ vertices, $p + 1$ edges incident with each vertex and in which the adjacency matrix $M(G)$ of G has $p + 1$ as a simple eigenvalue and has all other eigenvalues at most $2p^{1/2}$ in absolute value. Let the network F be the double cover of G ; that is, let the inputs and outputs of F each correspond to the vertices of G and let the input v have the output w as a neighbor if and only if the vertices of G corresponding to v and w are adjacent. Clearly F is a $(q + 1)$ -network of depth one in which every input has $p + 1$ neighbors and every output is a neighbor of $p + 1$ inputs. Let $M(F)$ denote the matrix whose rows are indexed by inputs of F , whose columns are indexed by outputs of F and whose (v, w) entry is the number of edges from v to w . Then $M(F)^T M(F) = M(G)^2$ has $(p + 1)^2$ as a simple eigenvalue and all other eigenvalues at most $4p$ in absolute value. Let $\epsilon = 4p/(p + 1)^2$ and let $f(\xi) = \xi/(\xi + \epsilon(1 - \xi))$. By the result of Tanner [T], if X is a set of inputs and $F(X)$ is the set of neighbors of X , then $|F(X)|/(q + 1) \geq f(|X|/(q + 1))$. Let $\delta = 4/(p + 1)$. Then $\epsilon \leq \delta$ and $f(3\delta) \geq 3/4$. Since f is concave, if $|X| \leq 3\delta(q + 1)$, then $|F(X)| \geq (q + 1)f(|X|/(q + 1)) \geq (|X|/3\delta)f(3\delta) \geq |X|/4\delta$. \square

A wide-sense nonblocking bi-generalized a -limited connector remains one if any number of inputs and outputs (together with the edges incident with them) are deleted.

A t -approximate a -limited m -juncator is an m -network in which, given any $(a - 1)$ -limited bi-state S , any set V containing t inputs and any set W containing t outputs, there exists a route ρ from some $v \in V$ to some $w \in W$ such that $S \cup \{\rho\}$ is a bi-state.

A t -approximate a -limited juncator remains one if any number of inputs and outputs (together with the edges incident with them) are deleted.

If F is an (n, m) -network, let F^* denote the (m, n) -network obtained by reversing the edges and exchanging the inputs and outputs of F . Clearly, $F^{**} = F$.

If S is a bi-state of F , let S^* denote the bi-state of F^* obtained by reversing the routes of S . Clearly, $S^{**} = S$.

Two routes are *anti-compatible* if their intersection is a final segment (possibly empty) of both of them. An *anti-state* is a set of pairwise anti-compatible routes. Clearly, if S is a state, then S^* is an anti-state.

PROPOSITION 5.2. *Let F be a wide-sense nonblocking $(a - 1 + t, t)$ -limited generalized (n, m) -concentrator and let G be an t -approximate a -limited m -juncator. Then $F \circ G \circ F^*$ is a wide-sense nonblocking a -limited bi-generalized n -connector.*

Proof. Let S be a bi-state of $F \circ G \circ F^*$. Say that S is safe if (1) S induces a safe state $S(F)$ of F and (2) S induces an anti-state $S(F^*)$ of F^* and $S(F^*)^*$ is a safe state of F . Clearly, the empty bi-state is safe and any bi-state below a safe bi-state is itself safe. Let S be a safe bi-state of $F \circ G \circ F^*$ and let (v, w) be an a -limited connection request in S . Since (v, w) is a -limited, S contains at most $a - 1$ routes.

The concentration request v is (a, t) -limited in $S(F)$, so there is a safe state $T(F)$ above $S(F)$ in which v has routes to all the outputs in some set V of t outputs of F . Similarly, there is an anti-state $T(F^*)$ above $S(F^*)$ in which all the inputs in some set W of t inputs of F^* have routes to w and such that $T(F^*)^*$ is a safe state for F .

The bi-state $S(G)$ of G induced by S contains at most $a - 1$ routes. Since G is a t -approximate a -limited m -juncator, there is a bi-state $T(G)$ above $S(G)$ containing a route from some $v' \in V$ to some $w' \in W$. Let the route ρ be the concatenation of the route from v to v' in $T(F)$, the route from v' to w' in $T(G)$, and the route from w' to w in $T(F^*)$. Then $S \cup \{\rho\}$ is a safe bi-state above S containing a route from v to w . \square

PROPOSITION 5.3. *Let F be an m -network with depth one in which every set of t inputs has at least $\lceil (m + a)/2 \rceil$ neighbors. Then $F \circ F^*$ is a t -approximate a -limited m -juncator.*

Proof. Let S be an $(a - 1)$ -limited bi-state of $F \circ F^*$. If V is a set of t inputs, then V has as neighbors all the outputs of F in some set V' of $\lceil (m + a)/2 \rceil$ outputs. Similarly, if W is a set of t outputs, then W has as neighbors all the inputs of F^* in some set W' of $\lceil (m + a)/2 \rceil$ inputs. Then $V' \cap W'$ contains at least $\lceil (m + a)/2 \rceil + \lceil (m + a)/2 \rceil - m \geq a$ links between F and F^* . Since S is $(a - 1)$ -limited, at most $a - 1$ links can be busy. Thus there is some idle link u in $V' \cap W'$. The link u has as neighbors some input v in V and some output w in W . Let ρ be the route from v through u to w in $F \circ F^*$. Then $S \cup \{\rho\}$ is a bi-state above S containing a route from $v \in V$ to $w \in W$. \square

COROLLARY 5.4. *Let q be a prime congruent to one modulo four with $32n \leq q + 1 \leq 128n$. For every $t \leq n$, a t -approximate n -limited $(q + 1)$ -juncator can be explicitly constructed with depth 2 and size $O(n^2/t)$.*

Proof. Let p be the smallest prime congruent to one modulo four such that $p + 1 \geq 12(q + 1)/t$. Then $p + 1 = O(n/t)$. Let F be the network constructed in the proof of Lemma 5.1. Then F has depth one and size $O(n^2/t)$. Furthermore, every set of t inputs has at least $3(q + 1)/4$ neighbors. It follows from Proposition 5.3 that F is a t -approximate n -limited $(q + 1)$ -juncator. \square

Proof of Theorem 5 for k even. For $k = 2$, the claim is trivial, so suppose $k \geq 4$.

Let q be a prime congruent to one modulo four such that $32n \leq q + 1 \leq 64n$. Let p be a prime congruent to one modulo four such that $32n^{2/k} \leq p + 1 \leq 64n^{2/k}$. Let F be the network constructed in the proof of Lemma 5.1.

Let $r = \lceil (p + 1)/32 \rceil$ and $a = \lfloor 6(q + 1)/(p + 1) \rfloor$. Then $ar \geq 2n$. By Lemma 5.1, every set of x inputs with $1 \leq x \leq 2a$ has at least $2rx$ neighbors. Then by Proposition 1 with $s = r$, F is a wide-sense nonblocking $(2n, r)$ -limited generalized $(q + 1)$ -concentrator.

Let $j = (k - 2)/2$ and $t = \lceil n^{1-2/k} \rceil$. Since $r \geq n^{2/k}$, $r^j \geq t$. By Proposition 2, $H = F^{\circ j}$ is a wide-sense nonblocking $(2n, t)$ -limited generalized $(q + 1)$ -concentrator.

By Corollary 5.4, a t -approximate n -limited $(q + 1)$ -network G can be explicitly constructed with depth two and size $O(n^2/t) = O(n^{1+2/k})$. By Proposition 5.2, $H \circ G \circ H^*$ is a wide-sense nonblocking n -limited bi-generalized $(q + 1)$ -connector with depth k and size $O(n^{1+2/k})$. Deleting all but n inputs and n outputs from this network completes the proof. \square

COROLLARY 5.5. *Let q be a prime congruent to one modulo four with $32n \leq q + 1 \leq 128n$. For every $t \leq n$, a t -approximate n -limited $(q + 1)$ -juncator can be explicitly constructed with depth one and size $O(n^3/t^2)$.*

Proof. In the network constructed in Corollary 5.4, replace each route of length two by a route of length one, then delete all links and edges incident with them. The resulting network is clearly a t -approximate n -limited $(q + 1)$ -juncator with depth one. Since in the original network each of the $(q + 1)$ links has $O(n/t)$ inputs and $O(n/t)$ outputs as neighbors, the resulting network has size $O(n^3/t^2)$. \square

Proof of Theorem 5 for k odd. For $k = 1$, the claim is trivial, so suppose $k \geq 3$.

Let q be a prime congruent to one modulo four such that $32n \leq q + 1 \leq 64n$. Let p be a prime congruent to one modulo four such that $32n^{2/k} \leq p + 1 \leq 64n^{2/k}$. Let F be the network constructed in the proof of Lemma 5.1.

Let $r = \lceil (p + 1)/32 \rceil$ and $a = \lfloor 6(q + 1)/(p + 1) \rfloor$. Then $ar \geq 2n$. By Lemma 5.1, every set of x inputs with $1 \leq x \leq 2a$ has at least $2rx$ neighbors. Then by Proposition 1 with $s = r$, F is a wide-sense nonblocking $(2n, r)$ -limited generalized $(q + 1)$ -concentrator.

Let $j = (k - 1)/2$ and $t = \lceil n^{1-1/k} \rceil$. Since $r \geq n^{2/k}$, $r^j \geq t$. By Proposition 2, $H = F^{\circ j}$ is a wide-sense nonblocking $(2n, t)$ -limited generalized $(q + 1)$ -concentrator.

By Corollary 5.4, a t -approximate n -limited $(q + 1)$ -network G can be explicitly constructed with depth one and size $O(n^3/t^2) = O(n^{1+2/k})$. By Proposition 5.2, $H \circ G \circ H^*$ is a wide-sense nonblocking n -limited bi-generalized $(q + 1)$ -connector with depth k and size $O(n^{1+2/k})$. Deleting all but n inputs and n outputs from this network completes the proof. \square

Since the networks constructed in Theorem 5 rely on Proposition 1, no efficient algorithm is known for finding safe routes. The remaining results of this section mitigate this defect in some special cases.

THEOREM 6. *Wide-sense nonblocking bi-generalized n -connectors can be explicitly constructed with depth three and size $O(n^{5/3})$, and with depth four and size $O(n^{3/2})$.*

For the proof of Theorem 6 we shall need the following lemma.

LEMMA 6.1. *For every prime power q , a (q^3) -network of depth one can be explicitly constructed in which every input has q^2 neighbors and any two distinct inputs have at most q common neighbors.*

Proof. Let $\text{GF}(q)$ be the field with q elements. Let the outputs correspond to the q^3 points in the three-dimensional affine space over $\text{GF}(q)$. Let the inputs correspond to any q^3 planes in this space. (There are $q^3 + q^2 + q$ planes in the space.) Let the input v have the output w as a neighbor if the plane corresponding to v contains the point corresponding to w . Clearly, each plane contains q^2 points and any two distinct planes have at most the q points of a line in common. \square

Let $F_3(q)$ denote the network constructed in the proof of Lemma 6.1.

PROPOSITION 6.2. *The network $F_3(2^\lambda)$ is a wide-sense nonblocking $(2^{3\lambda-5}, 2^{2\lambda-1})$ -limited generalized concentrator.*

Proof. Let the safe states be those in which every input has at most $2^{2\lambda-2}$ neighbors that terminate routes originating at other inputs. Clearly the empty state is safe, and any state below a safe state is itself safe. Let S be a safe state, let v be a $(2^{3\lambda-5}, 2^{2\lambda-1})$ -limited concentration request in S , and suppose that there is no safe state above S containing a route from v to an output that is idle in S . We shall derive a contradiction.

Since v is $(2^{3\lambda-5}, 2^{2\lambda-1})$ -limited in S , there are routes from v to at most $2^{\lambda-1}$ neighbors of v . Since S is safe, at most $2^{2\lambda-2}$ other neighbors of v are busy. Say that an output w is *critical* if w is a neighbor of v , w is idle in S and $S \cup \{(v, w)\}$ is not safe. Then there are at least $2^{2\lambda} - 2^{2\lambda-1} - 2^{2\lambda-2} = 2^{2\lambda-2}$ critical outputs.

Say that an input is *critical* if u is distinct from v and $2^{2\lambda-2}$ neighbors of u terminate routes that originate at inputs other than u . Each critical output w must be the neighbor of some critical input, else $S \cup \{(v, w)\}$ would be safe. Each critical input u can have at most 2^λ critical outputs as a neighbor, else the distinct inputs v and u would have more than 2^λ common neighbors. Thus there must be at least $2^{\lambda-2}$ critical inputs.

Let X be a set of $2^{\lambda-2}$ critical inputs. Each critical input in X has at least $2^{2\lambda-2}$ busy neighbors. Furthermore, each of the $\binom{2^{\lambda-2}}{2}$ pairs of distinct critical inputs in X has at most 2^λ busy neighbors in common. Thus there are at least

$$2^{\lambda-2} \cdot 2^{2\lambda-2} - \binom{2^{\lambda-2}}{2} \cdot 2^\lambda \geq 2^{3\lambda-5}$$

busy outputs in S . This contradicts the assumption that v is $(2^{3\lambda-5}, 2^{2\lambda-1})$ -limited in S . \square

Proof of the first assertion of Theorem 6. We may assume that $n = 2^{3\lambda-6}$ for some integer $\lambda \geq 3$. Let $t = 2^{2\lambda-1} = O(n^{2/3})$. Let F be the network obtained by deleting all but n inputs from $F_3(2^\lambda)$. By Proposition 6.2, F is a wide-sense nonblocking $(2n, t)$ -limited generalized $(n, 64n)$ -concentrator with depth one and size $O(n^{5/3})$.

Let q be a prime congruent to one modulo four with $64n \leq q + 1 \leq 128n$. By Corollary 5.5, a t -approximate n -limited $(q + 1)$ -juncter can be explicitly constructed with depth one and size $O(n^3/t^2) = O(n^{5/3})$. Let G be the network obtained by deleting all but $64n$ inputs and $64n$ outputs from this network. By Proposition 5.2, $F \circ G \circ F^*$ is a wide-sense nonblocking bi-generalized n -connector with depth three and size $O(n^{5/3})$, which completes the proof of the first assertion. \square

To prove the second assertion we shall need the following lemma.

LEMMA 6.3. *For every prime power q , a (q^2) -network of depth one can be explicitly constructed in which every input has q neighbors and any two distinct inputs have at most one common neighbor.*

Proof. Let $\text{GF}(q)$ be the field with q elements. Let the outputs correspond to the q^2 points in the affine plane over $\text{GF}(q)$. Let the inputs correspond to any q^2 lines in this plane. (There are $q^2 + q$ lines in the plane.) Let the input v have the output w as a neighbor if the line corresponding to v contains the point corresponding to w . Clearly, each line contains q points and any two distinct lines have at most one point in common. \square

Let $F_2(q)$ denote the network constructed in the proof of Lemma 6.3.

PROPOSITION 6.4. *The network $F_2(2^\lambda)$ is a wide-sense nonblocking $(2^{2\lambda-5}, 2^{\lambda-1})$ -limited generalized concentrator.*

Proof. Let the safe states be those in which every input has at most $2^{\lambda-2}$ neighbors that terminate routes originating at other inputs. Clearly the empty state is safe, and any state below a safe state is itself safe. Let S be a safe state, let v be a $(2^{2\lambda-5}, 2^{\lambda-1})$ -limited

concentration request in S , and suppose that there is no safe state above S containing a route from v to an output that is idle in S . We shall derive a contradiction.

Since v is $(2^{2\lambda-5}, 2^{\lambda-1})$ -limited in S , there are routes from v to at most $2^{\lambda-1}$ neighbors of v . Since S is safe, at most $2^{\lambda-2}$ other neighbors of v are busy. Say that an output w is *critical* if w is a neighbor of v , w is idle in S and $S \cup \{(v, w)\}$ is not safe. Then there are at least $2^\lambda - 2^{\lambda-1} - 2^{\lambda-2} = 2^{\lambda-2}$ critical outputs.

Say that an input u is *critical* if u is distinct from v and $2^{\lambda-2}$ neighbors of u terminate routes that originate at inputs other than u . Each critical output w must be the neighbor of some critical input, else $S \cup \{(v, w)\}$ would be safe. Each critical input u can have at most one critical output as a neighbor, else the distinct inputs v and u would have more than one common neighbor. Thus there must be at least $2^{\lambda-2}$ critical inputs.

Let X be a set of $2^{\lambda-2}$ critical inputs. Each critical input in X has at least $2^{\lambda-2}$ busy neighbors. Furthermore, each of the $\binom{2^{\lambda-2}}{2}$ pairs of distinct critical inputs in X has at most one busy neighbor in common. Thus there are at least

$$2^{\lambda-2} \cdot 2^{\lambda-2} - \binom{2^{\lambda-2}}{2} \cdot 1 \geq 2^{2\lambda-5}$$

busy outputs in S . This contradicts the assumption that v is $(2^{2\lambda-5}, 2^{\lambda-1})$ -limited in S . \square

Proof of the second assertion of Theorem 6. We may assume that $n = 2^{2\lambda-6}$ for some integer $\lambda \geq 4$. Let $t = 2^{\lambda-1} = O(n^{1/2})$. Let F be the network obtained by deleting all but n inputs from $F_2(2^\lambda)$. By Proposition 6.4, F is a wide-sense nonblocking $(2n, t)$ -limited generalized $(n, 64n)$ -concentrator with depth one and size $O(n^{3/2})$.

Let q be a prime congruent to one modulo four with $64n \leq q + 1 \leq 128n$. By Corollary 5.4, a t -approximate n -limited $(q + 1)$ -juncter can be explicitly constructed with depth two and size $O(n^2/t) = O(n^{3/2})$. Let G be the network obtained by deleting all but $64n$ inputs and $64n$ outputs from this network. By Proposition 5.2, $F \circ G \circ F^*$ is a wide-sense nonblocking bi-generalized n -connector with depth four and size $O(n^{3/2})$, which completes the proof of the second assertion. \square

5. Networks with unlimited depth. In previous sections we have used Proposition 1 to construct various networks with limited depth. In this section we shall use it in situations where depth is not constrained. We shall show that wide-sense nonblocking generalized n -connectors and bi-generalized n -connectors can be constructed with size $O(n \log n)$. That these results are best possible (to within constant factors) is shown by a result of Shannon [S] to the effect that even “rearrangeable n -connectors” require size $\Omega(n \log n)$. (Indeed, Pippenger and Valiant [PV] have shown that even networks satisfying a still weaker property, that of being “rearrangeable n -shifters,” must have size $\Omega(n \log n)$.)

The previously best upper bound for wide-sense nonblocking generalized connectors was $O(n(\log n)^2)$, due to Pippenger [P1]. This construction requires wide-sense nonblocking (ungeneralized) connectors as components. That size $O(n \log n)$ is sufficient for these networks (and indeed, even for networks that satisfy a stronger property, that of being “strictly nonblocking connectors”) is due to Bassalygo and Pinsker [BP1]. The proof of Bassalygo and Pinsker is nonconstructive (in the same sense as Lemma 3.1), but subsequent results by others ([M1], [GG], [JM], [LPS], [M2]) yield an explicit construction (see Pippenger [P3] for details).

That our results cannot be extended to “strictly nonblocking generalized connectors” is shown by the lower bound of $\Omega(n^2)$, due to Bassalygo and Pinsker [BP2].

The key to our results in this section is the following proposition.

PROPOSITION 7. *A wide-sense nonblocking generalized $(n, 18n)$ -concentrator can be explicitly constructed with depth $O(\log n)$ and size $O(n \log n)$.*

This proposition has the following immediate corollary.

COROLLARY 7.1. *A wide-sense nonblocking generalized n -connector can be explicitly constructed with depth $O(\log n)$ and size $O(n \log n)$.*

To prove Corollary 7.1 we shall need the following proposition.

PROPOSITION 7.2. *Let F be a wide-sense nonblocking (a, r) -limited generalized (n, m) -concentrator and let G be a wide-sense nonblocking (m, n) -connector. Then $F \cdot G$ is a wide-sense nonblocking (a, r) -limited generalized n -connector.*

Proof. Let the safe states of $F \cdot G$ be those that induce safe states of F and G . Clearly, the empty state is safe and any state below a safe state is itself safe. Let S be a safe state of $F \cdot G$ and let the connection request (v, w) be (a, r) -limited in S .

The concentration request v is (a, r) -limited in $S(F)$, so there exists a safe state $T(F)$ above $S(F)$ containing a route from v to an output u of F that is idle in $S(F)$. The connection request (u, w) is a -limited in $S(G)$, so there exists a safe state $T(G)$ above $S(G)$ containing a route from u to w .

The state obtained by adjoining to S the concatenation of the route from v to u in $T(F)$ with the route from u to w in $T(G)$ is a safe state above S containing a route from v to w . \square

Proof of Corollary 7.1. Let F be the network constructed in the proof of Proposition 7, and let G be an explicitly constructed nonblocking $(18n, n)$ -connector with depth $O(\log n)$ and size $O(n \log n)$ (see [P3] for a construction). By Proposition 7.2, $F \cdot G$ satisfies the claim. \square

We can also extend the foregoing corollary to bi-generalized connectors.

COROLLARY 7.3. *A wide-sense nonblocking bi-generalized n -connector can be explicitly constructed with depth $O(\log n)$ and size $O(n \log n)$.*

To prove Corollary 7.3 we shall need the following proposition.

PROPOSITION 7.4. *Let F be a wide-sense nonblocking (a, a) -limited generalized (n, m) -concentrator and let G be a wide-sense nonblocking a -limited m -connector. Then $F \cdot G \cdot F^*$ is a wide-sense nonblocking bi-generalized n -connector.*

Proof. Let S be a bi-state of $F \cdot G \cdot F^*$. Say that S is safe if (1) S induces a safe state $S(F)$ of F ; (2) S induces a safe state $S(G)$ of G ; and (3) S induces an anti-state $S(F^*)$ of F^* and $S(F^*)^*$ is a safe state of F . Clearly, the empty bi-state is safe and any bi-state below a safe bi-state is itself safe. Let S be a safe bi-state of $F \cdot G \cdot F^*$ and let the connection request (v, w) be a -limited in S .

The concentration request v is a -limited in $S(F)$, so there exists a safe state $T(F)$ above $S(F)$ containing a route from v to an output v' of F that is idle in $S(F)$. Similarly, there is an anti-state $T(F^*)$ above $S(F^*)$ containing a route from an input w' of F^* that is idle in $S(F^*)$ to w and such that $T(F^*)^*$ is a safe state for F . The connection request (v', w') is a -limited in $S(G)$, so there exists a safe state $T(G)$ above $S(G)$ containing a route from v' to w' .

Let the route ρ be the concatenation of the route from v to v' in $T(F)$, the route from v' to w' in $T(G)$, and the route from w' to w in $T(F^*)$. Then $S \cup \{\rho\}$ is a safe bi-state above S containing a route from v to w . \square

Proof of Corollary 7.3. Let F be the network constructed in the proof of Proposition 7, and let G be an explicitly constructed nonblocking $(18n)$ -connector with depth $O(\log n)$ and size $O(n \log n)$ (see [P3] for a construction). By Proposition 7.4, $F \cdot G \cdot F^*$ satisfies the claim. \square

Proof of Proposition 7. Let $p = 53$ and let q be a prime congruent to one modulo four with $9n \leq q + 1 \leq 18n$. Let F be the network of depth one and size $O(n)$ constructed in the proof of Lemma 5.1. Then every set of x inputs with $1 \leq x \leq 2n$ has at least $3x$ neighbors in F . By Proposition 1 with $r = 2$ and $s = 1$, F is a wide-sense nonblocking $(n, 2)$ -limited generalized $(q + 1)$ -concentrator.

Let $k = \lceil \log_2 n \rceil$. By Proposition 2, F^{*k} is a wide-sense nonblocking (n, n) -limited generalized $(q + 1)$ -concentrator. Deleting all but n inputs and adding $18n - (q + 1)$ isolated outputs completes the proof. \square

The results of this paper that rely on Proposition 1 do not provide efficient algorithms for finding safe routes. In special cases (Theorems 4 and 6) we have been able to avoid this defect. It may be possible to extend the methods used in these special cases to obtain a version of Proposition 1 that allows an efficient algorithm for finding safe routes, but such an extension has eluded our efforts thus far.

REFERENCES

- [B1] V. E. BENEŠ, *Heuristic remarks and mathematical problems regarding the theory of switching systems*, Bell System Tech. J., 41 (1962), pp. 1201–1247.
- [B2] ———, *Semilattice characterization of nonblocking networks*, Bell System Tech. J., 52 (1973), pp. 697–706.
- [B3] ———, *Blocking states in connecting networks made of square switches arranged in stages*, Bell System Tech. J., 60 (1981), pp. 511–521.
- [BKPVY] M. BLUM, R. M. KARP, C. H. PAPIDIMITRIOU, O. VORNBERGER, AND M. YANNAKAKIS, *The complexity of superconcentrators*, Inform. Process. Lett., 13 (1981), pp. 164–167.
- [BP1] L. A. BASSALYGO AND M. S. PINSKER, *Complexity of an optimum nonblocking switching network without reconnections*, Problems Inform. Transmission, 9 (1974), pp. 64–66.
- [BP2] ———, *Asymptotically optimal networks for generalized rearrangeable switching and generalized switching without rearrangement*, Problemy Peredachi Informatsii, 16 (1980), pp. 94–98.
- [Ca] D. G. CANTOR, *On non-blocking switching networks*, Networks, 1 (1971), pp. 367–377.
- [Cl] C. CLOS, *A study of non-blocking switching networks*, Bell System Tech. J., 32 (1953), pp. 406–424.
- [DDPW] D. DOLEV, C. DWORK, N. PIPPENGER, AND A. WIDGERSON, *Superconcentrators, generalizers and generalized connectors with limited depth*, ACM Symposium on Theory of Computing, 15 (1983), pp. 42–51.
- [F] J. FRIEDMAN, *A lower bound for strictly non-blocking networks*, Combinatorica, to appear.
- [FFP] P. FELDMAN, J. FRIEDMAN, AND N. PIPPENGER, *Non-blocking networks*, ACM Symposium on Theory of Computing, 18 (1986), pp. 247–254.
- [FP] J. FRIEDMAN AND N. PIPPENGER, *Expanding graphs contain all small trees*, Combinatorica, 7 (1987), pp. 71–76.
- [GG] O. GABBER AND Z. GALIL, *Explicit construction of linear-sized superconcentrators*, J. Comput. System Sci., 22 (1981), pp. 407–420.
- [JM] S. JIMBO AND A. MARUOKA, *Expanders obtained from affine transformations*, ACM Symposium on Theory of Computing, 17 (1985), pp. 88–97.
- [LPS] A. LUBOTZKY, R. PHILLIPS, AND P. SARNAK, *Explicit expanders and the Ramanujan conjectures*, ACM Symposium on Theory of Computing, 18 (1986), pp. 240–246.
- [M1] G. A. MARGULIS, *Explicit constructions of concentrators*, Problems Inform. Transmission, 9 (1973), pp. 325–332.
- [M2] ———, *Arithmetic groups and graphs without short cycles*, Problemy Peredachi Informatsii, to appear.
- [MJ] G. M. MASSON AND B. W. JORDAN, JR., *Generalized multi-stage connection networks*, Networks, 2 (1972), pp. 191–209.
- [NS] D. NASSIMI AND S. SAHNI, *Parallel permutation and sorting algorithms and a new generalized connection network*, J. Assoc. Comput. Mach., 29 (1982), pp. 642–667.
- [O] YU. P. OFMAN, *A universal automaton*, Trans. Moscow Math. Soc., 14 (1965), pp. 200–215.
- [P1] N. PIPPENGER, *The complexity theory of switching networks*, Ph.D. thesis, Department of Electrical Engineering, Massachusetts Institute of Technology, Cambridge, MA, 1973.

- [P2] ———, *On rearrangeable and nonblocking switching networks*, J. Comput. System Sci., 17 (1978), pp. 145–162.
- [P3] ———, *Telephone switching networks*, AMS Proc. Symposium on Applied Mathematics, 26 (1982), pp. 101–133.
- [PU] D. PELEG AND E. UFFAL, *Constructing disjoint paths on expander graphs*, ACM Symposium on Theory of Computing, 19 (1987), pp. 264–273.
- [PV] N. PIPPENGER AND L. G. VALIANT, *Shifting graphs and their applications*, J. Assoc. Comput. Mach., 23 (1976), pp. 423–432.
- [PY] N. PIPPENGER AND A. C.-C. YAO, *Rearrangeable networks with limited depth*, SIAM J. Algebraic Discrete Methods, 3 (1982), pp. 411–417.
- [S] C. E. SHANNON, *Memory requirements in a telephone exchange*, Bell System Tech. J., 29 (1950), pp. 343–349.
- [T] R. M. TANNER, *Explicit construction of concentrators from generalized N -gons*, SIAM J. Algebraic Discrete Methods, 5 (1984), pp. 287–293.