

1-1-1985

Some Graph-Colouring Theorems with Applications to Generalized Connection Networks

David G. Kirkpatrick
University of British Columbia

Maria M. Klawe
Harvey Mudd College

Nicholas Pippenger
Harvey Mudd College

Recommended Citation

Kirkpatrick, David G., Maria Klawe, and Nicholas Pippenger. "Some Graph-Colouring Theorems with Applications to Generalized Connection Networks." *SIAM Journal on Algebraic and Discrete Methods* 6, no. 4 (October 1985): 576-582.

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.

SOME GRAPH-COLOURING THEOREMS WITH APPLICATIONS TO GENERALIZED CONNECTION NETWORKS*

DAVID G. KIRKPATRICK‡, MARIA KLAWE† AND NICHOLAS PIPPENGER†

Abstract. With the aid of a new graph-colouring theorem, we give a simple explicit construction for generalized n -connectors with $2k - 1$ stages and $O(n^{1+1/k}(\log n)^{(k-1)/2})$ edges. This is asymptotically the best explicit construction known for generalized connectors.

1. Introduction. Our goals in this paper are twofold. Our first goal is to give a new construction for generalized connectors. Under certain circumstances this new construction is superior to all other known constructions. Our second goal is to use graph-colouring theorems to systematically derive old and new results on networks.

For the purposes of this introduction, we shall give some informal definitions. We shall give more precise and more general definitions in the next section. An n -network is an acyclic directed graph with n distinguished vertices called *inputs* and n other distinguished vertices called *outputs*. We shall be concerned with the minimum possible size (number of edges) and *depth* (number of edges in the longest path from an input to an output) that n -networks with various connectivity properties can possess. An n -connector is an n -network such that, for any one-to-one correspondence between certain inputs and distinct outputs, there exist vertex-disjoint paths joining each chosen input to the corresponding output. A *generalized n -connector* is an n -network such that, for any one-to-many correspondence between certain inputs and disjoint sets of outputs, there exist vertex-disjoint trees joining each chosen input to the corresponding set of outputs. An n -crossbar is an n -network with depth 1 and size n^2 , with an edge joining each input to each output. For both of the problems considered here, a crossbar provides a solution with small depth and large size. Our goal is to find alternate solutions with larger but limited depth and smaller size.

Let $f(n)$ denote the minimum possible size of an n -connector. It has long been known (see Beneš [2]) that $f(n) = O(n \log n)$ and $f(n) = \Omega(n \log n)$. (For the sharpest known estimates, see Pippenger [12] for the upper bound and Pippenger [13] for the lower bound.)

Let $g(n)$ denote the minimum possible size of a generalized n -connector. It was shown by Ofman [10] that $g(n) = O(n \log n)$ and by Pippenger [11] that $g(n) = f(n) + O(n)$. The first of these results is proved by an extension of the explicit construction used to show that $f(n) = O(n \log n)$. (For the best explicit construction known, see Dolev et al. [4].) The second result was established by a probabilistic construction. It is now possible to replace this by an explicit construction (see Gabber and Galil [5]), but in any case the constants involved are so large as to render the result completely impractical. (For the best probabilistic construction known, see Bassalygo [1].)

Let $f_k(n)$ denote the minimum possible size of an n -connector with depth at most k . It was shown by Pippenger and Yao [14] that

$$f_k(n) = O(n^{1+1/k}(\log n)^{1/k})$$

and that

$$f_k(n) = \Omega(n^{1+1/k}).$$

* Received by the editors May 14, 1984.

† IBM Research Laboratory, San Jose, California 95193.

‡ University of British Columbia, Vancouver, British Columbia, V6T 1W5, Canada.

(Throughout this paper, the constants implicit in the notation $O(\cdot \cdot \cdot)$ are independent of n , but may depend on k .) The upper bound here is proved by a probabilistic construction; the best explicit construction known gives

$$f_{2k-1}(n) = O(n^{1+1/k})$$

(see Pippenger [12]).

Let $g_k(n)$ denote the minimum possible size of a generalized n -connector with depth at most k . Dolev et al. [4] showed that

$$g_k(n) = O((n \log n)^{1+1/k})$$

by a probabilistic construction; they also showed that

$$g_{3k-2}(n) = O(n^{1+1/k})$$

by an explicit construction. Masson and Jordan [8] and Nassimi and Sahni [9] showed that

$$g_3(n) = O(n^{5/3})$$

by two quite different explicit constructions. Attempts to extend these constructions to depths greater than 3 do not give results competitive with the construction for depth $3k-2$ mentioned above.

In this paper we shall show that

$$g_3(n) = O(n^{3/2}(\log n)^{1/2})$$

by an explicit construction. We extend this to

$$g_{2k-1}(n) = O(n^{1+1/k}(\log n)^{(k-1)/2}),$$

which differs merely by logarithmic factors from the corresponding bound for $f_{2k-1}(n)$.

It seems unlikely that the construction of the present paper will be useful in practice; to compare it with the competing constructions of Masson and Jordan [8], Nassimi and Sahni [9] and Dolev et al. [4], we observe that $n^{3/2}(\log_2 n)^{1/2} = n^{5/3}$ for n about 10^3 , $n^{4/3} \log n = n^{3/2}$ for n about 6×10^8 and $n^{5/4}(\log n)^{3/2} = n^{4/3}$ for n about 6×10^{37} . Nor is this result asymptotically the best; probabilistic constructions give sharper upper bounds for all fixed depths. We can, however, say that it is asymptotically the best explicit construction known, when the depth is fixed and the number of inputs and outputs is large.

Edge-colouring in bipartite graphs provides a vivid and convenient language for discussing connectors and their control algorithms. This relationship was observed by Lev, Pippenger and Valiant [7] (who used it to describe parallel control algorithms) and later by Hwang [6] (who did not, however, give any new results). We extend this method in the present paper by using hyperedge-colouring of bipartite hypergraphs to discuss generalized connectors. For the application of yet another combinatorial colouring problem to networks, see Dolev et al. [4].

2. Networks and graph-colouring problems. A (p, q) -network is an acyclic digraph with p distinguished vertices called *inputs* and q other distinguished vertices called *outputs*. Vertices that are neither inputs nor outputs will be called *links*.

A *request* is a pair comprising an input and an output. An *assignment* is a set of requests, no two of which have an input or output in common. A *generalized assignment* is a set of requests, no two of which have an output in common.

A *route* is a directed path from an input to an output. A *state* is a set of routes, no two of which have a vertex in common. A *generalized state* is a set of routes, any two of which have at most an initial segment of their vertices in common.

An assignment or generalized assignment is *realized* by a state or generalized state, respectively, if, for each request in the assignment, there is a route in the state from the input of the request to the output of the request.

A (p, q) -*connector* is a (p, q) -network for which every assignment is realized by a state. Let $f_k(p, q)$ denote the minimum possible size of a (p, q) -connector with depth at most k . A *generalized* (p, q) -connector is a (p, q) -network for which every generalized assignment is realized by a generalized state. Let $g_k(p, q)$ denote the minimum possible size of a generalized (p, q) -connector with depth at most k . A (p, q) -*crossbar* is a (p, q) -network with depth 1 and size pq , with an edge joining each input to each output.

We shall be particularly concerned with the three-stage construction shown in Fig. 1. The first stage comprises r (a, c) -subnetworks, the second stage comprises c (r, s) -subnetworks and the third stage comprises s (c, b) -subnetworks. An output of each subnetwork in the first stage is identified with an input of each subnetwork in the second stage to form a link and an output of each subnetwork in the second stage is identified with an input of each subnetwork in the third stage to form a link. The resultant is a (p, q) -network, where $p = ar$ and $q = bs$.

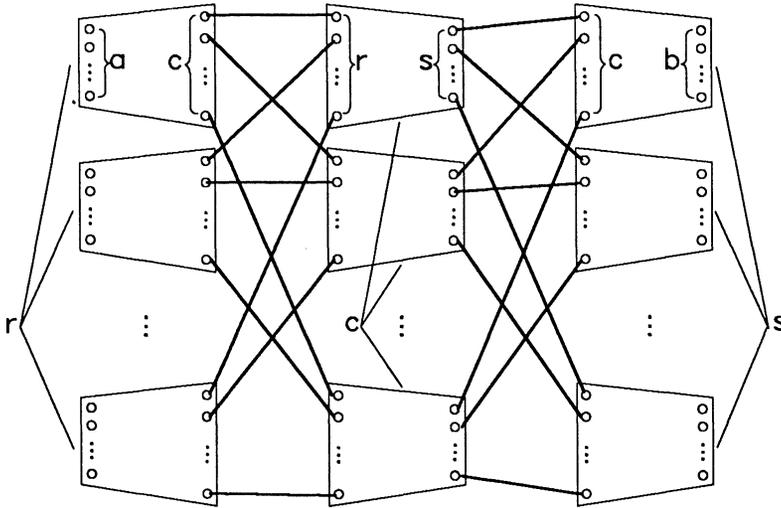


FIG. 1. The interconnection of subnetworks in three stages to form a network. The thick lines joining subnetworks represent identifications of inputs and outputs to form links, not edges.

We shall be concerned with conditions under which, if various subnetworks are connectors or generalized connectors, the resultant is a connector or generalized connector. We shall systematically obtain such conditions by reducing them to graph-colouring problems.

An (r, s, a, b) -*graph* is a bipartite graph (R, S, E) with vertices $R = \{1, \dots, r\}$ and $S = \{1, \dots, s\}$ and (possibly multiple) edges E such that at most a edges are incident with each vertex in R and at most b edges are incident with each vertex in S .

A *c-colouring* of a bipartite graph is an assignment of the colours $C = \{1, \dots, c\}$ to the edges of the graph such that the edges incident with any vertex are assigned distinct colours.

An (r, s, a, b) -hypergraph is a bipartite hypergraph (R, S, E) with vertices $R = \{1, \dots, r\}$ and $S = \{1, \dots, s\}$ and (possibly multiple) hyperedges E such that each hyperedge is incident with exactly one vertex in R , at most a hyperedges are incident with each vertex in R and at most b hyperedges are incident with each vertex in S .

A c -colouring of a bipartite hypergraph is an assignment of the colours $C = \{1, \dots, c\}$ to the hyperedges of the graph such that the hyperedges incident with any vertex are assigned distinct colours.

A c -hypercolouring of a bipartite hypergraph (R, S, E) is an assignment of subsets of the colours $C = \{1, \dots, c\}$ to the hyperedges of the graph such that the hyperedges incident with any vertex in R are assigned disjoint sets of colours and the hyperedges incident with any vertex in S are assigned sets of colours possessing a system of distinct representatives.

The reductions of network conditions to graph-colouring problems are very trite. Rather than present four such reductions in detail, we shall describe the paradigm once. Given an assignment or generalized assignment, we construct a graph or hypergraph by associating a vertex in R with each subnetwork in the first stage, a vertex in S with each subnetwork in the third stage and an edge or hyperedge in E with each request or maximal set of requests having an input in common. We then apply a graph-colouring theorem to obtain a colouring or hypercolouring of this graph or hypergraph. We associated a colour in C with each subnetwork in the second stage. From the colouring or hypercolouring, we construct assignments or generalized assignments for the subnetworks. By hypothesis, there are states or generalized states realizing these assignments or generalized assignments. We then patch together these states or generalized states to obtain a state or generalized state for the resultant. In each case, the graph-colouring theorem is a transparent paraphrase of the network condition.

3. Some graph-colouring theorems. The following proposition is perhaps the oldest result concerning connecting networks (see Beneš [2, Thm. 3.1]).

PROPOSITION 1. *If all the subnetworks are connectors, then the resultant is a connector if*

$$c \geq \max \{a, b\}.$$

This proposition reduces to the following well-known graph-colouring theorem (see Berge [3, Chap. 12, Thm. 2]).

THEOREM 2. *Every (r, s, a, b) -graph has a c -colouring if*

$$c \geq \max \{a, b\}.$$

Taking all subnetworks in Proposition 1 to be crossbars and choosing $a = b = r = s = \lceil n^{1/2} \rceil$ yields n -connectors with depth 3 and size $O(n^{3/2})$. Taking the subnetworks in the first and third stages to be crossbars and constructing the subnetworks in the second stage recursively yields n -connectors with depth $2k - 1$ and size $O(n^{1+1/k})$. This is asymptotically the best explicit construction known for connectors with limited depth.

The following proposition is due to Masson and Jordan [8].

PROPOSITION 3. *If the subnetworks in the first and third stages are generalized connectors and those in the second stage are connectors, then the resultant is a generalized connector if*

$$c \geq \max \{as, b\}.$$

This proposition reduces to Theorem 2 in the same manner as Proposition 1. Taking all subnetworks to be crossbars and choosing $a = s = \lceil n^{1/3} \rceil$ and $b = r = \lceil n^{2/3} \rceil$ yields generalized n -connectors with depth 3 and size $O(n^{5/3})$.

PROPOSITION 4. *If the subnetworks in the second and third stages are generalized connectors and those in the first stage are connectors, then the resultant is a generalized connector if*

$$c \geq a + (b(b-1)s)^{1/2}.$$

This proposition reduces to the following theorem.

THEOREM 5. *Every (r, s, a, b) -hypergraph has a c -colouring if*

$$c \geq a + (b(b-1)s)^{1/2}.$$

Proof. Let G be an (r, s, a, b) -hypergraph. Let G^* be the hypergraph containing those hyperedges of G that are incident with more than $(bs/(b-1))^{1/2}$ vertices in S . There are at most $(b(b-1)s)^{1/2}$ hyperedges in G^* . Let each hyperedge in G^* be assigned a distinct colour. It will suffice to show that this c -colouring of G^* can be extended to a c -colouring of G .

We shall prove this by induction on the number of hyperedges that are in G but not in G^* . If there are no such hyperedges, then we are done. If there is at least one such hyperedge, let G' be a hypergraph obtained by deleting one such hyperedge H from G . By inductive hypothesis, the c -colouring of G^* can be extended to a c -colouring of G' . It will suffice to show that it can be extended to a c -colouring of G .

Let τ be a vertex in R or S . Let us say that a colour is *good* at τ if it is not assigned to a hyperedge incident with τ in G' . Let H be incident with the vertex ρ in R . All but at most $a-1$ colours are good at ρ . The hyperedge H is incident with at most $(bs/(b-1))^{1/2}$ vertices in S . For each such vertex σ , all but at most $b-1$ colours are good at σ . Since $c > (a-1) + (b(b-1)s)^{1/2}$, at least one colour is good at every vertex with which H is incident. Assigning this colour to H yields a c -colouring of G . \square

Taking all subnetworks in Proposition 4 to be crossbars and choosing $a = s = \lceil n^{2/3} \rceil$ and $b = r = \lceil n^{1/3} \rceil$ again yields generalized n -connectors with depth 3 and size $O(n^{5/3})$.

Remark. It is not hard to see that Theorem 5 is, to within a constant factor, the best possible. Clearly at least a colours may be necessary. We shall show that at least $\frac{1}{2}(b(b-1)s)^{1/2}$ colours may be necessary. It will follow that at least $\max\{a, \frac{1}{2}(b(b-1)s)^{1/2}\} \geq \frac{1}{2}(a + (b(b-1)s)^{1/2})$ colours may be necessary.

To show that at least $\frac{1}{2}(b(b-1)s)^{1/2}$ colours may be necessary, it will suffice to construct a hypergraph (V, E) with at most s vertices in V and at least $\frac{1}{2}(b(b-1)s)^{1/2}$ hyperedges in E such that there are at most b hyperedges incident with any vertex and every pair of hyperedges is incident with a common vertex. It will in fact suffice to construct a hypergraph (V, H) with at most s vertices in V and at least $s^{1/2}$ hyperedges in H such that there are at most 2 hyperedges incident with any vertex and every pair of hyperedges is incident with a common vertex, for then we may take E to contain $\lfloor b/2 \rfloor \geq \frac{1}{2}(b(b-1))^{1/2}$ copies of each hyperedge in H .

Let $t = \lceil s^{1/2} \rceil$. Then $t(t-1)/2 \leq s$. Let K be a complete graph with t vertices and $t(t-1)/2$ edges. Construct (V, H) by associating a vertex in V with each edge in K and a hyperedge in H with each vertex in K (a vertex in V is incident with a hyperedge in H if and only if the associated edge is incident with the associated vertex in K). \square

PROPOSITION 6. *If all subnetworks are generalized connectors, then the resultant is a generalized connector if*

$$c \geq (a-1)\lfloor \log_2(2s) \rfloor + 2b.$$

This proposition reduces to the following theorem.

THEOREM 7. *Every (r, s, a, b) -hypergraph has a c -hypercolouring if*

$$c \geq (a-1)\lfloor \log_2(2s) \rfloor + 2b.$$

Proof. Let G be an (r, s, a, b) -hypergraph. We shall consider two cases. If $\lfloor \log_2(2s) \rfloor \geq 2b$, then $c \geq (a - 1) \lfloor \log_2(2s) \rfloor + 2b \geq ab$. We can assign disjoint sets of b colours to each of the a or fewer hyperedges incident with each vertex in R . A vertex in S is incident with at most b hyperedges, so the sets of colours assigned to these hyperedges will possess a system of distinct representatives.

Suppose, then, that $\lfloor \log_2(2s) \rfloor \leq 2b$. We shall prove that G has a c -hypercolouring in which $\lfloor \log_2(2s) \rfloor$ colours are assigned to each hyperedge. The proof is by induction on the number of hyperedges in G . If G has no hyperedges, then we are done. If G has at least one edge, let G' be a hypergraph obtained by deleting one hyperedge H incident with a vertex ρ in R . By inductive hypothesis, G' has a c -hypercolouring in which $\lfloor \log_2(2s) \rfloor$ colours are assigned to each hyperedge. It will suffice to show that this c -hypercolouring can be extended to a c -hypercolouring of G by assigning a set of $\lfloor \log_2(2s) \rfloor$ colours to H .

Let us say that a colour is *good* at ρ if it does not belong to the union of the sets of colours assigned to the hyperedges incident with ρ in G' . All but at most $(a - 1) \lfloor \log_2(2s) \rfloor$ colours are good at ρ , so at least $2b$ colours are good at ρ . Let σ be a vertex in S . Let us say that a colour is *good* at σ if the sets of colours assigned to the hyperedges incident with σ possess a system of distinct representatives not containing that colour. For each σ in S , all but at most b colours are good at σ . Let K be a set of colours that are good at ρ . Assigning the set K to the hyperedge H will yield a c -hypercolouring of G if, for every σ in S , some colour in K is good at σ . There are at least $(2b) \cdots (2b - k + 1)/k!$ ways of choosing $k \leq 2b$ colours that are good at ρ . For each σ in S , at most $(b) \cdots (b - k + 1)/k!$ of these do not contain a colour that is good at σ . Thus at most $s(b) \cdots (b - k + 1)/k!$ of them do not yield a c -hypercolouring of G . If $k = \lfloor \log_2(2s) \rfloor \leq 2b$, then

$$s < 2^k \leq (2b) \cdots (2b - k + 1)/(b) \cdots (b - k + 1),$$

and there exists a set K of k colours that yields a c -hypercolouring of G . \square

Taking all subnetworks in Proposition 6 to be crossbars and choosing $a = s = \lceil (n/\log_2 n)^{1/2} \rceil$ and $b = r = \lceil (n \log_2 n)^{1/2} \rceil$ yields generalized n -connectors with depth 3 and size $O(n^{3/2}(\log n)^{1/2})$. More generally, we can prove by induction on k that $g_{2k-1}(n) = O(n^{1+1/k}(\log n)^{(k-1)/2})$. If $k = 1$, and n -crossbar provides the basis $g_1(n) = O(n^2)$. If $k \geq 2$, we apply Proposition 6 with $a = \lceil n^{1/k}(\log n)^{(k-3)/2} \rceil$, $b = \lceil n^{1/k}(\log n)^{(k-1)/2} \rceil$, $r = \lceil n^{(k-1)/k}/(\log n)^{(k-3)/2} \rceil$ and $s = \lceil n^{(k-1)/k}/(\log n)^{(k-1)/2} \rceil$. We have $r \geq s$, and $g_{2k-3}(r, s) \leq \lceil r/s \rceil g_{2k-3}(s)$, as can be seen by identifying the outputs of $\lceil r/s \rceil$ disjoint generalized s -connectors. Thus, by inductive hypothesis, we have $g_{2k-3}(r, s) = O(rs^{1/(k-1)}(\log s)^{(k-2)/2})$. Taking the subnetworks in the first and third stages to be crossbars and constructing the subnetworks in the second stage recursively completes the induction. This is asymptotically the best construction known for generalized connectors with limited depth. It matches, to within logarithmic factors, the best explicit construction known for connectors.

Remark. It is possible to improve Theorem 7 as regards constant factors but we do not know whether it is, to within a constant factor, best possible. As we saw in the proof of the Theorem, ab colours are always sufficient. We shall show that $(a - 1)b + 1$ colours may be necessary, if r and s are large enough.

Let $\binom{x}{y}$ denote the binomial coefficient $x(x - 1) \cdots (x - y + 1)/y!$. Given a and b , construct the hypergraph H as follows. There are $r = b \lfloor (a - 1)b, b - 1 \rfloor$ vertices in R , each incident with a hyperedges, for a total of ar hyperedges. For each set of b hyperedges, there is a vertex in S incident with just those b hyperedges, for a total of

$s = [ar, b]$ vertices in S . Suppose that H can be hypercoloured with $c = (a - 1)b$ colours. We shall derive a contradiction.

Because of the vertices in R , each of the c colours can be assigned to at most r hyperedges. It follows that at most $cr/b = (a - 1)r$ hyperedges are assigned b or more colours, and thus that at least $ar - (a - 1)r = r$ hyperedges are each assigned at most $b - 1$ colours. There are $[c, b - 1]$ ways of choosing $b - 1$ colours; thus some set of $r/[c, b - 1] = b$ hyperedges are assigned colours from some set of $b - 1$ colours. These b hyperedges are incident with some vertex in S , but their colours cannot possess a system of distinct representatives. This contradiction shows that at least $(a - 1)b + 1$ colours are necessary. \square

The foregoing example excludes the possibility that $c = O(\max\{a, b\})$ is sufficient (as is the case for edge-colouring bipartite graphs), but it does not exclude the possibility that $c = O(\max\{r, s, a, b\})$ is sufficient. If this were true, then taking all subnetworks to be crossbars and choosing $a = b = r = s = \lceil n^{1/2} \rceil$ would yield generalized n -connectors with depth 3 and size $O(n^{3/2})$. Taking the subnetworks in the first and third stages to be crossbars and constructing the subnetworks in the second stage recursively would yield generalized n -connectors with depth $2k - 1$ and size $O(n^{1+1/k})$. This would match, to within a constant factor, the best explicit construction known for connectors.

REFERENCES

- [1] L. A. BASSALYGO, *Asymptotically optimal switching circuits*, Prob. Info. Trans., 17 (1981), pp. 206-211.
- [2] V. E. BENEŠ, *Mathematical Theory of Connecting Networks and Telephone Traffic*, Academic Press, New York, 1965.
- [3] C. BERGE, *Graphs and Hypergraphs*, North-Holland, Amsterdam, 1973.
- [4] D. DOLEV, C. DWORK, N. PIPPENGER AND A. WIGDERSON, *Superconcentrators, generalizers and generalized connectors with limited depth*, 15th ACM Symposium on Theory of Computing, 1983, pp. 42-51.
- [5] O. GABBER AND Z. GALIL, *Explicit constructions of linear-sized superconcentrators*, J. Comput. Sys. Sci., 22 (1981), pp. 407-420.
- [6] F. K. HWANG, *Control algorithms for rearrangeable close networks*, IEEE Trans. Comm., 31 (1983), pp. 952-954.
- [7] G. LEV, N. PIPPENGER AND L. G. VALIANT, *A fast parallel algorithm for routing in permutation networks*, IEEE Trans. Comput., 30 (1981), pp. 93-100.
- [8] G. M. MASSON AND B. W. JORDAN, JR., *Generalized multi-stage connection networks*, Networks, 2 (1972), pp. 191-209.
- [9] 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.
- [10] Yu. P. OFMAN, *A universal automaton*, Trans. Moscow Math. Soc., 14 (1965), pp. 200-215.
- [11] N. PIPPENGER, *Generalized connectors*, SIAM J. Comput., 7 (1978), pp. 510-514.
- [12] ———, *On rearrangeable and non-blocking switching networks*, J. Comput. System Sci., 17 (1978), pp. 145-162.
- [13] ———, *A new lower bound for the number of switches in rearrangeable networks*, this Journal, 1 (1980), pp. 164-167.
- [14] N. PIPPENGER AND A. C.-C. YAO, *Rearrangeable networks with limited depth*, this Journal, 3 (1982), pp. 411-417.