SIAM J. Discrete Math.
Vol. 20, No. 2, pp. 287–290
© 2006 Society for Industrial and Applied Mathematics

Downloaded 04/26/14 to 140.113.38.11. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

ON MULTICAST REARRANGEABLE 3-STAGE CLOS NETWORKS WITHOUT FIRST-STAGE FAN-OUT*

HONG-BIN CHEN† AND FRANK K. HWANG†

Abstract. For the multicast rearrangeable 3-stage Clos networks where input crossbars do not have fan-out capability, Kirkpatrick, Klawe, and Pippenger gave a sufficient condition and also a necessary condition which differs from the sufficient condition by a factor of 2. In this paper, we first tighten their conditions. Then we propose a new necessary condition based on the affine plane such that the necessary condition matches the sufficient condition for an infinite class of 3-stage Clos networks.

Key words. rearrange, Clos networks, multicast, affine plane

AMS subject classifications. 94C15, 05B05

DOI. 10.1137/05062336X

1. Introduction.

Consider a 3-stage Clos network $C(n_1, r_1, m, n_2, r_2)$, where the input stage consists of r_1 $n_1 \times m$ crossbars, the middle stage m $r_1 \times r_2$ crossbars, the output stage r_2 $m \times n_2$ crossbars, and where there exists one link between every pair of crossbars between two adjacent stages (see Figure 1).

The inlets of the input crossbars are the inputs of the network, and the outlets of the output crossbars are the outputs of the network. In the multicast traffic network, an input can appear in a request more than once. If the appearance is restricted to at most f times, the traffic is called an f-cast traffic. If there is no restriction, then it is called the *broadcast* traffic. A network is rearrangeable if any set of disjoint pairs of inputs and outputs can be simultaneously connected. If the calls come sequentially, rearrangeability means we can disconnect all existing connections and reroute them together with the new call simultaneously.

A crossbar is said to have the fan-out capability if the crossbar itself can route multicast traffic without blocking, i.e., any inlet can be connected to any number of idle outlets regardless of other connections. If the crossbars in a given stage perform only point-to-point connections, then we say the stage has no fan-out capability. Four models have been studied [3] on 3-stage Clos networks:

Model 0. no restriction on fan-out capability,
Model 1. input stage has no fan-out capability,
Model 2. middle stage has no fan-out capability,
Model 3. output stage has no fan-out capability.

Masson and Jordan [7] proved that $C(n_1, r_1, m, n_2, r_2)$ under model 2 is multicast rearrangeable if and only if $m \geq \max\{\min\{n_1 f, N_2\}, \min\{n_2, N1\}\}$. However, necessary and sufficient conditions are not known under the other models. Under model 1, Kirkpatrick, Klawe, and Pippenger [6] gave a sufficient condition for $C(n_1, r_1, m, n_2, r_2)$ to be multicast rearrangeable, and also a necessary condition which differs from the

*Received by the editors January 27, 2005; accepted for publication (in revised form) October 12, 2005; published electronically March 24, 2006.
http://www.siam.org/journals/sidma/20-2/62336.html

†Department of Applied Mathematics, National Chiao Tung University, Hsinchu 300, Taiwan (andan.am92g@nctu.edu.tw, fhwang@math.nctu.edu.tw). The work of the second author was partially supported by Republic of China, National Science Council grant NSC 92-2115-M-009-014.

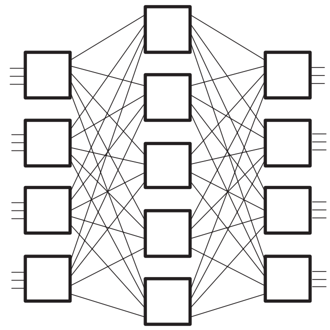

FIG. 1. $C(3,4,5,3,4)$.

sufficient condition by a factor of 2. In this paper, we tighten their conditions, then propose a new necessary condition which matches the sufficient condition for an infinite class of networks.

2. Main results. Kirkpatrick, Klawe, and Pippenger proved the following theorem.

THEOREM 2.1. (i) $C(n_1, r_1, m, n_2, r_2)$ *is broadcast rearrangeable under model* 1 *if* $m \geq n_1 + (n_2(n_2-1)r_2)^{1/2}$, (ii) $C(n_1, r_1, m, n_2, r_2)$ *is broadcast rearrangeable under model* 1 *only if* $m \geq \frac{n_1+(n_2(n_2-1)r_2)^{1/2}}{2}$.

In the following theorem we modify Theorem 2.1(i) by considering integrality of the number of crossbars. We also improve Theorem 2.1(ii).

THEOREM 2.2. (i) $C(n_1, r_1, m, n_2, r_2)$ *is broadcast rearrangeable under model* 1 *if* $m \geq \lceil (\frac{n_2 r_2}{n_2-1})^{1/2} \rceil (n_2-1) + (n_1-n_2+1)^+$, *where* $x^+ = \max\{x, 0\}$, (ii) $C(n_1, r_1, m, n_2, r_2)$ *is broadcast rearrangeable under model* 1 *only if* $m \geq \max\{n_1, \lfloor n_2/2 \rfloor \lfloor (2r_2)^{1/2} \rfloor\}$.

Proof. (i) The set L of requests each asking to connect to at least $\lceil (\frac{n_2 r_2}{n_2-1})^{1/2} \rceil$ outputs has size of at most

$$\frac{n_2 r_2}{\lceil (\frac{n_2 r_2}{n_2-1})^{1/2} \rceil} \leq \frac{n_2 r_2}{(\frac{n_2 r_2}{n_2-1})^{1/2}} = [n_2 r_2 (n_2-1)]^{1/2} \leq \left\lceil \left(\frac{n_2 r_2}{n_2-1} \right)^{1/2} \right\rceil (n_2-1).$$

Route each of these requests through a distinct middle crossbar. A request g other than these can ask for connections to a set O_g of at most $\lceil (\frac{n_2 r_2}{n_2-1})^{1/2} \rceil - 1$ output crossbars. Such a request has to be routed through a middle crossbar not taken by any of the at most $(\lceil (\frac{n_2 r_2}{n_2-1})^{1/2} \rceil - 1)(n_2-1)$ outputs on crossbars in O_g, nor by the $n_1 - 1$ inputs on the same input crossbar as g. Therefore

$$\left(\left\lceil \left(\frac{n_2 r_2}{n_2-1} \right)^{1/2} \right\rceil - 1 \right) (n_2-1) + (n_1-1) + 1 = \left\lceil \left(\frac{n_2 r_2}{n_2-1} \right)^{1/2} \right\rceil (n_2-1) + n_1 - n_2 + 1$$

middle crossbars are sufficient to route g. However, if $n_1 - n_2 + 1 < 0$, the number of middle crossbars still cannot be less than the number required to route L.

(ii) Construct a complete graph K_v with $v = \lfloor (2r_2)^{1/2} \rfloor$ vertices and $e = \binom{v}{2} \leq r_2$ edges. Label each edge by a distinct output crossbar. Take $c = \lfloor n_2/2 \rfloor$ copies of K_v, keeping the edge-labels intact, and label the vc vertices by the set $\{1, 2, \ldots, vc\}$. Identify each vertex u as a request $(O_{u1}, \ldots, O_{u(v-1)})$, where $O_{u1}, \ldots, O_{u(v-1)}$ are the labels of the $v-1$ edges incident to u. Note that each output crossbar appears in $2c \leq n_2$ requests.

Downloaded 04/26/14 to 140.113.38.11. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

Downloaded 04/26/14 to 140.113.38.11. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

Since every pair of requests intersect in at least one output crossbar, each of the vc requests must be routed through a distinct middle crossbar. ∎

Our construction in (ii) improves over that of [6] by increasing the number of edges in K_v. The corresponding graph in [6] contains only $r_2^{1/2}$ vertices, hence roughly $r_2/2$ edges.

COROLLARY 2.3. *Part* (ii) *is valid for f-cast traffic with $f \geq \lfloor (2r_2)^{1/2} \rfloor - 1$.*

Next we give a stronger necessary condition which is based on the observation that requests from the same input crossbar cannot use the same middle crossbar, even if the requests do not intersect (in output crossbars). Then the request graph we need to construct is no longer a complete graph, but a graph whose vertices can be partitioned into r_1 subsets such that an edge exists between every pair of vertices from different subsets. Such a graph corresponds to a resolvable block design.

A block design $B(v, b, r, k, 1)$ is a collection of b k-subsets (called blocks) of a v-set S, $k < v$, such that each pair of elements of S appears together in exactly one block and each element of S appears in exactly r blocks. $B(v, b, r, k, 1)$ is resolvable if the blocks can be partitioned into r orbits such that each element appears once in each orbit. For example, the following 12 blocks form a resolvable $B(9, 12, 4, 3, 1)$ which can be grouped into 4 orbits, each of 3 blocks, so that the blocks in each orbit together contain each element exactly once:

$$(\{1,2,3\},\{4,8,9\},\{5,6,7\}),$$
$$(\{1,5,8\},\{3,4,6\},\{2,7,9\}),$$
$$(\{1,4,7\},\{2,6,8\},\{3,5,9\}),$$
$$(\{1,6,9\},\{2,4,5\},\{3,7,8\}).$$

A block design $B(n^2, n^2 + n, n + 1, n, 1)$ is called an affine plane of order n. It is well known that every affine plane is resolvable and that an affine plane of order q exists whenever q is a prime power; see [1].

THEOREM 2.4. *For every prime power q, and every M in the range $1 \leq M \leq r$, there exist $n_1 \geq q, r_1 \geq M, n_2 \geq M$, and $r_2 \geq q^2$ such that $C(n_1, r_1, m, n_2, r_2)$ is broadcast rearrangeable only if $m \geq q(q + 1)$.*

Proof. For a prime power q we know that there exists an affine plane of order q, i.e., a resolvable block design $B(q^2, q^2 + q, q + 1, q, 1)$. Identify the elements as the output crossbars, the orbits as the input crossbars, and the blocks as inputs, while the elements in a block i represent the output crossbars input i requests to connect. Note that each output crossbar appears in $M \leq n_2$ requests. Hence the given set of requests are legitimate.

By our construction, two requests from different input crossbars (orbits) intersect in one output crossbar and hence must be routed through different middle crossbars. Requests from the same input crossbar do not intersect in any output crossbar, but still have to be routed through different middle crossbars since they share the input crossbar. Therefore the total number of middle crossbars required is at least the number of requests constructed above, which is $q(q + 1)$. ∎

Introduction of the parameter M is just to broaden the applicability of Theorem 2.4, i.e., n_2 does not have to be equal r, but can be less.

Now we show that the necessary condition of Theorem 2.4 matches the sufficient condition in Theorem 2.2(i). Theorem 2.4 shows the necessary condition is $m \geq q(q+1)$. From Theorem 2.2(i), setting $n_1 = \frac{v}{k} = q, n_2 = M = q+1$, and $r_2 = v = q^2$, the sufficient condition is

$$m \geq \left\lceil \left(\frac{Mq^2}{q}\right)^{1/2} \right\rceil q + (q - (q+1) + 1)^+ = \lceil ((q+1)q)^{1/2} \rceil q = q(q+1),$$

same as the necessary condition.

COROLLARY 2.5. *Theorem* 2.4 *holds for f-cast traffic with $f \geq q$.*

3. Conclusions. The current necessary condition for broadcast rearrangeable 3-stage Clos networks differs from the sufficient condition by a factor of 2. We tightened these conditions such that they match for an infinite class of 3-stage Clos networks. This shows that our tightened conditions cannot be further improved for general parameters.

While the main results obtained for multicast rearrangeable 3-stage Clos networks are for broadcast networks so far, our arguments for necessary conditions are valid also for f-cast networks for some specific f as shown in the corollaries, thus starting the study of f-cast rearrangeable 3-stage Clos networks, which has been a vacuum so far.

The model-1 model can be interpreted in two ways. One is that the input crossbars do not have the fan-out capability, and thus perhaps can be obtained with a cheaper cost. The other is that they do have the fan-out capability, but our routing algorithm chooses not to use it. This type of routing algorithm has been used in the mixed-requirement model where all point-to-point requests meet the strictly nonblocking requirement and all f-cast requests for $f \geq 2$ meet the rearrangeable requirement [2, 4, 5].

REFERENCES

[1] I. ANDERSON, *Combinatorial designs: Construction methods*, Ellis Horwood Limited, New York, 1990.

[2] D.-Z. DU AND H. Q. NGO, *An extension of DHH-Erdös conjecture on cycle-plus-triangle graphs*, Taiwanese J. Math., 6 (2002), pp. 261–267.

[3] F. K. HWANG, *The mathematical theory of nonblocking switching networks*, 2nd ed., World Scientific, Hackensack, NJ, 2004.

[4] F. K. HWANG, S. C. LIAW, AND L. D. TONG, *Strictly nonblocking* 3*-stage Clos networks with some rearrangeable multicast capability*, IEEE Trans. Commun., 51 (2003), pp. 1765–1767.

[5] F. K. HWANG AND C. H. LIN, *Broadcasting in a* 3*-stage point-to-point nonblocking networks*, Inter. J. Rel. Qual. Safety Engin., 2 (1995), pp. 299–307.

[6] D. G. KIRKPATRICK, M. KLAWE, AND N. PIPPENGER, *Some graph-coloring theorems with applications to generalized connection networks*, SIAM J. Algebraic Discrete Methods, 6 (1985), pp. 576–582.

[7] G. M. MASSON AND B. W. JORDAN, JR., *Generalized multi-stage connection networks*, Networks, 2 (1972), pp. 191–209.

Downloaded 04/26/14 to 140.113.38.11. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php