A bound on the number of middle-stage crossbars in *f*-cast rearrangeable Clos networks*

Marek Kubale, Paweł Obszarski
Faculty of Electronics, Telecommunications and Informatics
Gdańsk University of Technology
Narutoicza 11/12, 80-233 Gdańsk
POLAND
kubale@eti.pg.gda.pl
pawel.obszarski@eti.pg.gda.pl

Abstract: - In 2006 Chen and Hwang gave a necessary and sufficient condition under which a three-stage Clos network is rearrangeable for broadcast connections. Assuming that only crossbars of the first stage have no fan-out property, we give similar conditions for *f*-cast Clos networks, where *f* is an arbitrary but fixed invariant of the network. Such assumptions are valid for some practical switching systems, e.g. high-speed crossconnects. We also recognize the complexity status for a related routing problem. In our considerations we introduce the hypergraph edge coloring model, which is a suitable mathematical idealization for the three-stage Clos networks.

Key-Words: - Clos network, crossbar, *f*-cast traffic, hypergraph coloring, rearrangeable network.

1 Introduction

Rearrangeable switching networks are of prac-tical interest since they can be operated as non-blocking networks. Rearrangeability expresses a limiting effectiveness of crosspoint usage in switching systems.

Three-stage interconnection networks of this type are constructed as follows:

- The first (input) stage consists of r_1 crossbars (switches) each with n_1 inlets and m outlets. We say that such a crossbar is of size n_1 x m.
- The second (middle) stage consists of m crossbars each of size r_1 x r_2.
- The third (output) stage consists of r_2 crossbars each of size m x n_2.
- There exists exactly one link between each middle crossbar and each input and output crossbar.

A three-stage Clos network C(n_1,r_1,m,n_2,r_2) is shown in Fig. 1. The inlets of the first stage (input) crossbars are the inputs of the network, and the outlets of the third stage (output) crossbars are the outputs of the network. In the multicast traffic network an input can appear in a request more than once. If this appearance is restricted to at most *f* times, the traffic is called an *f-cast traffic* or simply1:*f traffic*. A crossbar is said to have the 1:*f* property if the crossbar itself can route *f*-cast traffic without blocking, i.e., any idle inlet can be connected to any set of *f* or fewer idle outlets regardless of other connections.

In the following our basic notation and terminology follows that of Hwang [3].

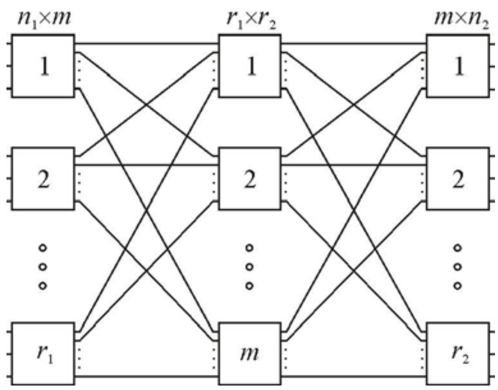

Fig. 1. C(n_1,r_1,m,n_2,r_2)

In this paper we consider Clos networks without input stage fan-out, i.e., operating under the so-called model 1 [2]. The first attempt to estimate the number of middle switches (often denoted by m) in such rearrangeable networks has been done by Jajszczyk [4]. However, he underestimated the value of m which is necessary in the networks. The first theoretical result concerning broadcast networks operating under model 1 is due to Kirkpatrick, Klawe, and Pippenger [5]. They gave a sufficient condition and also a necessary condition which differs from the sufficient condition by a factor of 2. Chen and Hwang [2] tightened their conditions.

*This paper has been partially supported by Narodowe Centrum Nauki under contract DEC-2011/02/A/ST6/00201

More precisely, they formulated a necessary and sufficient condition for m which is valid for all rearrangeable broadcast networks of Clos type. Their formula holds tight also for f-cast Clos networks $C(n_1,r_1,m,n_2,r_2)$, provided that $f \geq \lfloor (2r_2)^{1/2} \rfloor - 1$. Clearly, broadcast networks could be used as f-cast networks, however, the crosspoint usage in this case would be wasteful.

In what follows we give a new formula on the number m in $C(n_1,r_1,m,n_2,r_2)$ to be f-cast re-arrangeable. For this reason we introduce a mathematical model of a state of $C(n_1,r_1,m,n_2,r_2)$, namely a hypergraphs coloring, which allows to estimate the value of m on one hand, and to control the multicast Clos networks, on the other. We show that this formula is tight in the worst case. Finally, we conclude with some remarks on the complexity of control algorithms for rearranging f-cast Clos networks in a tabular form.

2 Bipartite hypergraphs

Hypergraph H is a pair (V,E) in which V represents a set of vertices and E is a multiset of *hyperedges* (or simply *edges*), where an edge $e \in E$ is a nonempty subset of the vertex set. We say that edge e and vertex v are *incident* if $v \in e$. Two edges that have a vertex in common are said to be *adjacent*. The cardinality $|e|$ of edge e is called the *dimension* of e. The maximum dimension of an edge in hypergraph H is denoted by $\psi(H)$. A hypergraph is *d-uniform* if each edge is of cardinality d. The number of edges to which vertex v belongs is its *vertex degree* $\deg(v)$. By *hypergraph degree* $\Delta(H)$ we mean the maximum degree among all vertices in H.

An edge *k-coloring* of a hypergraph H is a function $c: E \to \{1,2,...,k\}$ such that for any two edges $e, f \in E$, $e \cap f \neq \emptyset$ we have $c(e) \neq c(f)$. In other words k-coloring is an assignment of k colors to the edges such that the edges that share a vertex get different colors. We say that a coloring c is *optimal* if it uses the minimum number of colors. Such a minimum number of colors is called the *chromatic index* $\chi'(H)$.

For a hypergraph H, the *line graph* $L(H)$ is a simple graph representing adjacency between hyperedges in H. More precisely, each hyperedge of H is assigned a vertex in $L(H)$ and two vertices in $L(H)$ are adjacent if and only if their corresponding hyperedges in H have a vertex in common. It is easy to notice that an edge-coloring of a hypergraph H is equivalent to vertex coloring of its line graph $L(H)$.

A hypergraph H is said to be *bipartite*, if there exists a subset of its vertices such that each edge shares exactly one vertex with this set. We will call this set the first (*input*) partition of H. The remaining vertices are called the second (*output*) partition. For the input partition we define the *input degree* $\Delta_i(H)$ that represents maximum degree of vertex in the first partition. The *output degree* $\Delta_o(H)$ is defined analogously.

3 Hypergraphs and Clos networks

Edge coloring problem for bipartite hypergraphs can be used to model Clos networks control. The crossbars from the input stage can be associated with vertices of the first partition, and the crossbars from the output stage one can associate with vertices of the second partition. Requests (calls) are represented by hyperedges. Note that each request involves one crossbar from the first stage and at most f crossbars from the third stage, hence each edge is spanned on exactly one vertex from the first partition and a number of vertices from the second partition. Moreover, for such a hypergraph we have $\psi(H) \leq 1+f$. If all the calls are of type $1{:}f$ then the corresponding hypergraph is $(1+f)$-uniform.

The edge coloring problem allows to operate Clos networks since middle-stage crossbars can be identified with colors. Hence edge coloring decides which middle crossbars should be assigned to the requests. Finally, as long as hypergraphs that need maximum number of colors are considered, the model allows to estimate the number of middle crossbars necessary and sufficient for a Clos network to be f-cast rearrangeable. For this reason we assume that each call in the network is of type $1{:}f$. This situation is shown in Figs. 2 and 3 for example, where $f = 3$.

4 Bounds on the number of middle stage crossbars

Proper configuration of the middle stage crossbars is essential in the process of designing of re-arrangeable Clos networks.

An upper bound on the chromatic index can be expressed as follows:

$$\chi'(H) \leq \Delta_i(H) + (\Delta_o(H) - 1)\,(\psi(H) - 1) \tag{1}$$

To justify this inequality it is enough to observe that the maximum number of possible neighbors of any hyperedge is $\Delta_i(H) - 1 + (\psi(H) - 1)(\Delta_o(H) - 1)$

ISBN: 978-1-61804-355-9

since an edge meets at most $\Delta_i(H) - 1$ neighbors on vertices from the first partition and $(\psi(H) - 1)(\Delta_o(H) - 1)$ neighbors on vertices from the second partition. Assuming that in the worst case all the neighbors have pairwise different colors there must be one more color for the considered edge.

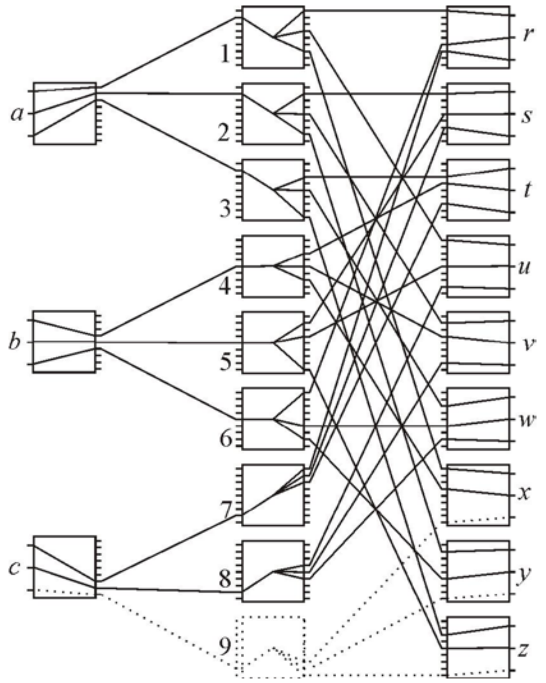

Fig. 2. Clos network C(3,3,8,3,9) with 8 calls. Call (c, x, y, z) is blocked, hence the 9-th middle crossbar is needed.

In relation with Clos networks the restriction on Δ_i may stand for the maximal number of calls from a single crossbar of the first stage. Similarly Δ_o restricts the number of calls leaving any crossbar of the third stage. Therefore, inequality (1) may be expressed in terms of Clos network parameters in the following manner:

$$m \leq n_1 + (n_2 - 1) f \qquad (2)$$

The upper bound in (1) is sharp. Fig. 3 presents a nontrivial hypergraph H which follows from the blocking state of Clos network of Fig. 2. Each of the nine edges is depicted as a quadruple of vertices joined by a unique line. We have here $\Delta_i(H) = 3$, $\Delta_o(H) = 3$, $\psi(H) = 4$, and $\chi'(H) = 9$. This hypergraph can be thought of as the first in an infinite series of hypergraphs for which bound (1) holds with equality.

Bound (1) is also tight for infinitely many values of $\psi(H)$. In fact, for each odd integer d there is a d-uniform bipartite hypergraph H of degree $\Delta(H) = 2$, for which bound (1) holds with equality. An example of such a hypergraph with $\psi(H) = 5$ is shown in Fig. 4. In terms of Clos network parameters, the above statement means that for each even value of f there is a rearrangeable Clos network

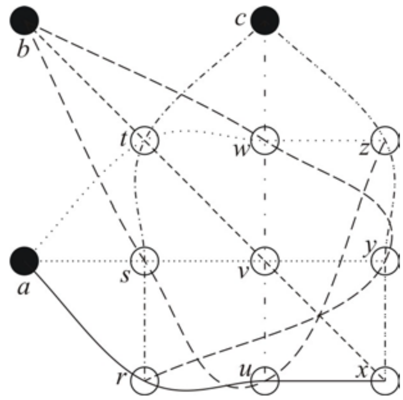

Fig. 3. Bipartite hypergraph that represents the calls and their incidences appearing in Clos network from Fig. 2.

C(n_1,r_1,m,n_2,r_2) such that bound (2) holds sharp. For example, the corresponding Clos network for hypergraph of Fig. 4 is C(2,3,6,2,12).

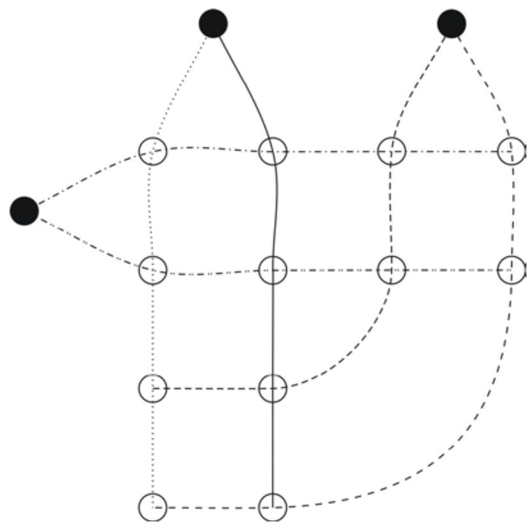

Fig. 4. Bipartite hypergraph with $\Delta_i(H) = \Delta_o(H) = 2$ and $\psi(H) = 5$ for which bound (1) is tight.

5 Complexity status for Clos networks control

In [6] we have investigated the complexity status for the operation of 2-cast Clos networks. Table 1 collects the results for various input and output degrees that may be directly transferred to Clos networks with $f = 2$.

Table 1. The complexity of edge coloring depending on the in- and out-degrees for $f = 2$.

	$n_1 = 1$	$n_1 = 2$	$n_1 \geq 3$
$n_2 = 1$	trivial	linear	linear
$n_2 = 2$	linear	linear	NP-hard
$n_2 \geq 3$	NP-hard	NP-hard	NP-hard

Let us comment briefly on the case of $n_1 = n_2 = 2$. In this case the line graph of such a bipartite hypergraph is of degree at most 3. If such a line graph is K_4-free then, in virtue of Brooks' theorem [1], it has a vertex-coloring with 3 colors which can be found in linear time. Since such a coloring is equivalent to 3-coloring the edges of H, it follows that there exists an efficient algorithm for rearranging 2-cast Clos networks $C(2_1,r_1,3,2,r_2)$, unless $L(H)$ contains K_4. If $n_2 = 1$ then the line graph $L(H)$ is a collection of at most r_1 complete graphs each of size at most n_1. If, however, $n_1 = 1$ and $n_2 = 2$ then $L(H)$ is a collection of paths and cycles. Of course, Clos networks $C(n_1,r_1,m,n_2,r_2)$ with $n_1 = 1$ or $n_2 = 1$ have no practical significance.

Below we give a similar complexity table for $f = 3$. Notice the difference in the middle of the table (i.e. for $n_1 = n_2 = 2$). This is so because $L(H)$ is no longer of degree 3 and, consequently, the previous linear coloring algorithm does not work. Also, note that if $n_1 = 1$, $n_2 = 2$, the corresponding line graphs are of degree 3, which implies linear solvability in this case.

Table 2. The complexity of edge coloring depending on the in- and out-degrees for $f = 3$.

	$n_1 = 1$	$n_1 = 2$	$n_1 \geq 3$
$n_2 = 1$	trivial	linear	linear
$n_2 = 2$	linear	NP-hard	NP-hard
$n_2 \geq 3$	NP-hard	NP-hard	NP-hard

The last Table 3 gives the complexity status of control algorithms for f-cast Clos networks, where $f \geq 4$. Notice the difference for $n_1 = 1$, $n_2 = 2$. The NP-hardness of this case follows from the fact that $L(H)$ is no longer of degree 3.

Table 3. The complexity of edge coloring depending on the in- and out-degrees for $f \geq 4$.

	$n_1 = 1$	$n_1 = 2$	$n_1 \geq 3$
$n_2 = 1$	trivial	linear	linear
$n_2 = 2$	NP-hard	NP-hard	NP-hard
$n_2 \geq 3$	NP-hard	NP-hard	NP-hard

6 Conclusions

We have introduced a bound for the number of middle stage crossbars and showed that the bound is tight. However, we believe that some further steps could be done towards enhancing it. In the last section we have discussed the complexity status for Clos networks control. This is just a preliminary sketch and we are going to expand this topic in further publications.

References:

[1] R. L. Brooks, On colouring the nodes of a network, *Proc. Cambridge Philosophical Society, Math. Phys. Sci.*, Vol.37, 1941, pp. 194-197.

[2] H-B Chen, F. K. Hwang, On multicast rearrangeable 3-stage Clos networks without first-stage fan-out, *SIAM J. Disc. Math.*, Vol.20, 2006, pp. 287-290.

[3] F. K. Hwang, *The Mathematical Theory of Nonblocking Switching Networks*, 2nd ed., World Scientific, Hackensack, NJ, 2004.

[4] A. Jajszczyk, Rearrangeability of multicast Clos networks with middle-stage fan-out, *Communications Letters, IEEE,* Vol.7, 2004, pp. 599 – 600.

[5] M. G. Kirkpatrick, M. Klawe, and N. Pippenger, Some graph-coloring theorems with applications to genera-lized connection networks, *SIAM J. Algebraic Discrete Methods*, Vol.6, 1985, pp. 576-582.

[6] P. Obszarski, A. Jastrzębski, and M. Kubale, Towards the boundary between easy and hard control problems in multicast Clos networks, *Bulletin of the Polish Academy of Sciences, Technical Sciences,* Vol.63, 2015, pp. 739-744.