WDM Split Cross-connects and 3-stage Clos Networks

Hung Q. Ngo [*]

July 9, 2002

Abstract

In this paper, we first show that a wavelength division multiplexed (WDM) split cross-connect is strictly, wide-sense, or rearrangeably non-blocking if and only if a corresponding 3-stage Clos network is strictly, wide-sense or rearrangeably non-blocking. This observation implies known results on strictly non-blocking WDM split cross-connects and gives rise to new results on wide-sense and rearrangeably non-blocking WDM split cross-connects. We next introduce a routing algorithm for the Clos networks which show that there are wide-sense non-blocking Clos networks which are not strictly non-blocking. Hence, the same holds for the WDM split cross-connects. We also consider another demand model for WDM split cross-connects and present associated results. Routing algorithms for WDM split cross-connects could be viewed as on-line and off-line algorithms to edge-color a class of bipartite multi-graphs.

1 Introduction

Wavelength division multiplexed (WDM) cross-connects are important in modern WDM optical networks. We can think of the cross-connects as black boxes with disjoint sets of input and output fibers, each of which is capable of carrying a set of signals on different wavelengths. We would like the cross-connect to be able to route any *demand* from an unused wavelength at an input fiber to an unused wavelength at an output fiber. Since the input and output wavelength of a demand might be different, *wavelength interchangers* (WIs) are needed to design the cross-connects. We assume the WIs to have the ability to convert any incoming wavelength to any outgoing wavelength. Since WIs are expensive components, an objective is to design cross-connects with as few WIs as possible. Obviously, the minimum number of WIs shall depend on which degree of non-blockingness (strictly, wide-sense, rearrangeably, to be defined) is desired.

The wavelength division multiplexed (WDM) split cross-connects (WSXC) were introduced by Rasala and Wilfong [13, 14] as an optimal architecture for WDM cross-connects in terms of number of WIs. An $n_1 \times n_2$ WSXC has n_1 input fibers and n_2 output fibers. The design is shown in Figure 1. A set of r_1 wavelengths $\{\lambda_1, \dots, \lambda_{r_1}\}$ are available at each input fiber, and a set of r_2 wavelengths $\{\gamma_1, \dots, \gamma_{r_2}\}$ are available at each output fiber. The input fibers are inputs to a switching fabric I. The fabric I has m outputs connected to m WIs. The output side consists of fibers from the WIs to the output fabric O of dimension $m \times r_2$. We use $S(n_1, r_1, m, n_2, r_2)$ to denote the WSXC. For convenience, we use $S(n, m, r)$ to denote $S(n, r, m, n, r)$. In [13] the authors show that $2k - 1$ WIs are necessary and sufficient for $S(n, m, r)$ to be strictly non-blocking (SNB), given that I and O are SNB. In [14], corresponding conditions were derived for the general case of $S(n_1, r_1, m, n_2, r_2)$. The authors called $S(n, m, r)$ the *homogeneous WSXC*,

[*]Computer Science and Engineering Department, 201 Bell Hall State University of New York at Buffalo, Amherst, NY 14260, USA. hungngo@cse.buffalo.edu

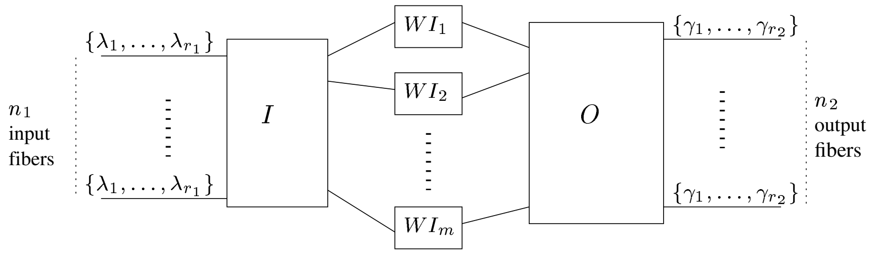

Figure 1: The WDM split cross-connect $S(n_1, r_1, m, n_2, r_2)$

and $S(n_1, r_1, m, n_2, r_2)$ the *heterogeneous WSXC*. The notations used in their paper were slightly different from ours. The new notations are used to be consistent with well-established notations for the 3-stage Clos networks.

The Clos network has been widely used for data communications and parallel computing systems [1, 5, 6, 8]. The Clos network $C(n_1, r_1, m, n_2, r_2)$ is a 3-stage interconnection network, where the first stage consists of r_1 crossbars of size $n_1 \times m$, the last stage has r_2 crossbars of dimension $m \times n_2$, and the middle stage has m crossbars of dimension $r_1 \times r_2$ (see Figure 2). Each input crossbar I_i ($i = 1, \ldots, r_1$) is connected to each middle crossbar M_j ($j = 1, \ldots, m$).

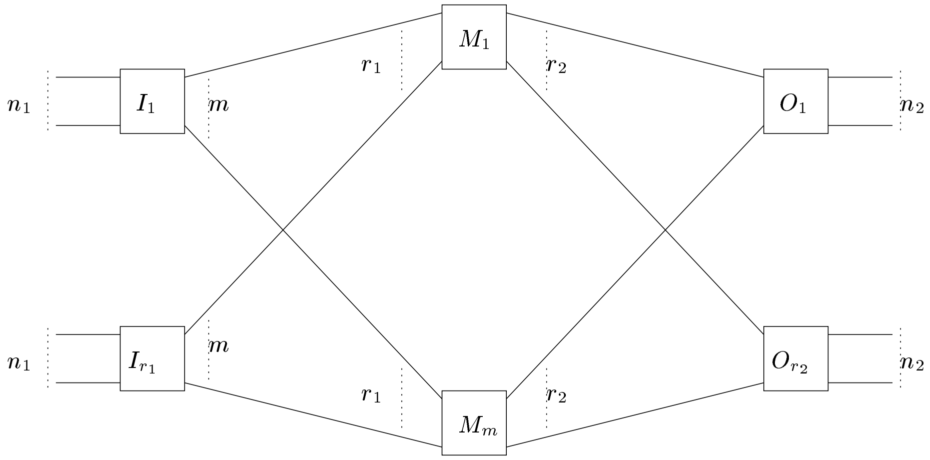

Figure 2: The 3-stage Clos network $C(n_1, r_1, m, n_2, r_2)$

Similarly, the middle stage and the last stage are fully connected. When $n_1 = n_2 = n$ and $r_1 = r_2 = r$, the network is called the *symmetric* 3*-stage Clos network*, denoted by $C(n, m, r)$.

In this paper, we show that $S(n_1, r_1, m, n_2, r_2)$ is "equivalent" to $C(n_1, r_1, m, n_2, r_2)$, and derive consequences of this fact. In particular, we get for free many results from [13, 14] on SNB WSXC. The observation also gives rise to a few results on wide-sense non-blocking (WSNB) and rearrangeably non-blocking (RNB) WSXC. Next, a routing algorithm for Clos networks and WSXCs called the *degree-counting* algorithm is presented in terms of an online algorithm for edge-coloring a class of bipartite multi-graphs. We also introduce a new demand model for the WSXC and present associated results under different blocking properties.

2 Preliminaries

2.1 The classical demand model

In this section, we present basic definitions of different degrees of non-blockingness for the WSXCs and Clos networks under the so-called *classical demand model*. In the WSXC case, this model assumes that each link can carry a number of channels each of which is on a different wavelength. As for Clos networks, this model assumes that each link can carry at most one channel. In the next section, another demand model is introduced, in which a wavelength in WSXC or a link in Clos networks can be further divided into sub-channels. Practically, the sub-channels could be thought of as representing time division multiplexed (TDM) channels within one wavelength.

2.1.1 WDM Split Cross-connects

Recall the general architecture of $S = S(n_1, r_1, m, n_2, r_2)$ as shown in Figure 1. A *demand* D is defined to be a 4-tuple (i, λ, o, γ), where $i \in \{1, \ldots, n_1\}$ is an input fiber to the fabric I, $o \in \{1, \ldots, n_2\}$ is an output fiber from the fabric O, $\lambda \in \{\lambda_1, \ldots, \lambda_{r_1}\}$, and $\gamma \in \{\gamma_1, \ldots, \gamma_{r_2}\}$. The fibers i and o are called *input fiber* and *output fiber* for demand D, respectively. Similarly, λ and γ are called the *input wavelength* and *output wavelength* of D. A route R in S is a path from an input of I to an output of I to the corresponding WI to an input of O and exits at an output of O. Along each fiber (or link), R is assigned a wavelength such that consecutive fibers in R are assigned the same wavelength unless the fibers are connected through a WI. Straightforwardly, a route R for a demand $D = (i, \lambda, o, \gamma)$ is a route from i to o such that the fibers i and o are assigned with wavelengths λ and γ, respectively. The WIs are able to simultaneously convert any incoming wavelength to any outgoing wavelength as long as there is no conflict at the input of the output.

A *valid demand set* is a set $\mathcal{D}$ of demands such that no two demands in $\mathcal{D}$ share both the input fiber and input wavelength, and no two demands share both the output fiber and output wavelength. A valid demand set $\mathcal{D}$ is *realizable* (or *satisfied*) by S if there is a set $\mathcal{R}$ of routes, one for each demand, such that any two routes which use the same fiber is assigned different wavelengths on that fiber. In this case, we say that $\mathcal{R}$ *realizes* $\mathcal{D}$. Given a valid demand set $\mathcal{D}$, a new demand $D \notin \mathcal{D}$ is said to be *compatible* with $\mathcal{D}$ iff $\mathcal{D} \cup \{D\}$ is valid.

A WSXC S is said to be *strictly non-blocking* if given a valid demand set $\mathcal{D}$ and a set $\mathcal{R}$ of routes realizing $\mathcal{D}$, for any new demand D compatible with $\mathcal{D}$ there exists a new route R so that $\mathcal{R} \cup \{R\}$ realizes $\mathcal{D} \cup \{D\}$.

As demands come and go, a strategy to pick a route for a new demand is called a *routing algorithm*. A WSXC is said to be *wide-sense non-blocking* (WSNB) with respect to a routing algorithm A if given a route set $\mathcal{R}$ realizing a valid demand set $\mathcal{D}$, for any new demand D compatible with $\mathcal{D}$ the algorithm A can choose a new route R so that $\mathcal{R} \cup \{R\}$ realizes $\mathcal{D} \cup \{D\}$. A WSXC is WSNB with respect to a class $\mathcal{A}$ of routing algorithms if it is WSNB with respect to every algorithm in $\mathcal{A}$. In general, a WSXC is WSNB if it is WSNB under some routing algorithm.

Lastly, a WSXC is *rearrangeably non-blocking* if for any valid demand set $\mathcal{D}$ there exists a set $\mathcal{R}$ of routes realizing $\mathcal{D}$.

Note that the switching fabrics I and O can also be strictly, wide-sense or rearrangeably non-blocking. Note also that SNB implies WSNB which in turn implies RNB. Unless mentioned otherwise, we will assume that I and O are SNB in order to simplify the discussion on minimizing the number of WIs. It is straightforward to extend the results in this paper to the cases when I and/or O are non-blocking in other senses.

2.1.2 Three-stage Clos Networks

The inlets (outlets) of the input (output) crossbars are the *inputs* (*outputs*) of the network. Inputs and outputs are referred to as *external links*, while links between crossbars are referred to as *internal links*. As in the case of WSXC, we shall assume that all crossbars are SNB to shorten our presentation.

A demand D for the Clos network $C = C(n_1, r_1, m, n_2, r_2)$ is a pair (i, o) where i is an inlet of an input crossbar and o is an outlet of an output crossbar. A set $\mathcal{D}$ of demands for C is *valid* if no two demands in $\mathcal{D}$ share the same inlet nor the same outlet. A *route* R consists of an inlet i, an input crossbar I one of whose inlets is i, a link from I to a middle crossbar M, M, a link from M to an output crossbar O, O, and an outlet o of O. A route R *realizes* a demand $D = (i, o)$ if R begins from i and ends at o. A set $\mathcal{R}$ of k routes *realizes* a valid demand set $\mathcal{D}$ of k demands if each route in $\mathcal{R}$ realizes a different demand in $\mathcal{D}$ so that no two routes in $\mathcal{R}$ share a link.

The definitions of *compatibility*, *routing algorithms*, *SNB*, *WSNB*, and *RNB* are similar to those in the previous section.

2.2 The multi-rate demand model

The multi-rate demand models were initially introduced in [10]. Other works on these models include those in [3, 4, 9, 11, 12]. In the multi-rate demand model, each demand D for a WSXC is a 5 tuple $(i, \lambda, o, \gamma, w)$, where $w \in (0, 1]$ is the requested *rate* (or *weight*) of the demand. Other parameters remain the same as before. For simplicity, we shall assume that all wavelengths have unit capacity. A rate w can be thought of as the fraction of bandwidth on a particular wavelength. A set $\mathcal{D}$ of demands is *valid* if for each fiber i and wavelength λ, the total weight of demands having i as the input fiber and λ as the input wavelength does not exceed unity. The same must also hold for every pair of output fiber and output wavelength.

All the definitions are similar to the classical demand model except that now many routes can share the same wavelength on the same fiber as long as the total weight of the corresponding demands does not exceed unity. The concepts of SNB, WSNB, and RNB are now called *multi-rate SNB*, *multi-rate WSNB*, and *multi-rate RNB*.

The corresponding definitions for the Clos networks is similar, and thus omitted here.

The multi-rate model for the Clos networks given here is often referred to as the *continuous* multi-rate model [8]. It is more general than the *discrete* multi-rate model where the weights are required to be chosen from a given set W of k weights where all weights are integer multiples of the smallest weight. The smallest weight is meant to model the smallest granularity of a TDM channel.

3 The equivalence of WSXCs and Clos networks

3.1 The classical demand model

Lemma 3.1. *Assume that switching fabrics I and O of $S = S(n_1, r_1, m, n_2, r_2)$ and the crossbars of $C = C(n_1, r_1, m, n_2, r_2)$ are are strictly non-blocking, Then, S is strictly non-blocking, wide-sense non-blocking (under any algorithm), and rearrangeably non-blocking if and only if C is strictly non-blocking, wide-sense non-blocking, and rearrangeably non-blocking, respectively.*

Proof. We first establish a one-to-one correspondence between the following two sets:

$$\begin{aligned} A &= \{(D, R) \mid R \text{ realizes } D \text{ in } S\} \\ B &= \{(D', R') \mid R' \text{ realizes } D' \text{ in } C\}. \end{aligned}$$

Consider $(D, R) \in A$ where $D = (i, \lambda_x, o, \gamma_y)$. Note that $1 \le x \le r_1$, $1 \le y \le r_2$, $1 \le i \le n_1$, and $1 \le o \le n_2$. Further assume that the route R routes D through the wavelength interchanger WI_z. The corresponding pair $(D', R') \in B$ is defined as follows. The demand D' is from i' to o', where i' is the ith input to input crossbar I_x, and o' is the oth output from output crossbar O_y. The route R' is the sequence $i' \to I_x \to M_z \to O_y \to o'$. Thus, the middle-stage crossbars of C can be thought of as representing the WIs of S. It is easy to see that the construction is invertible, and hence the mapping is bijective.

Next, observe that two routes R_1 and R_2 for S share a fiber and the same wavelength on that fiber if and only if the corresponding two routes for C share the same link. In essence, any routing of S induces a routing of C and vice versa. The equivalence of S and C now readily follows. □

We are now ready to derive some results on WSXCs for free. All the corollaries follow from the corresponding results about 3-stage Clos networks. There are numerous results on 3-stage Clos networks (see [8] and [1]). We will quote here a few more important corollaries.

The following corollary includes all the results about SNB WSXCs in [13, 14].

Corollary 3.2 (Rasala-Wilfong [13, 14]). *The WDM split cross-connect $S(n_1, r_1, m, n_2, r_2)$ is strictly non-blocking if and only if the number of wavelength interchangers m is at least m_0 where*

$$m_0 = \begin{cases} n_1 r_1 & \text{if } n_1 r_1 \le n_2 \\ n_2 r_2 & \text{if } n_2 r_2 \le n_1 \\ n_1 + n_2 - 1 & \text{otherwise.} \end{cases}$$

Corollary 3.3. *The WDM split cross-connect $S(n_1, r_1, m, n_2, r_2)$ is rearrangeably non-blocking if and only if the number of wavelength interchangers m is at least* $\max\{n_1, n_2\}$.

Unlike the cases of SNB and RNB, very little is known on WSNB properties of the Clos networks. Particularly, the difficulty comes from the fact that WSNB is algorithm dependent, and that there are very few known routing algorithms for the network except some common-sense strategies. We first give a negative result, which basically says that when r_1 and r_2 are large, $S(n_1, r_1, m, n_2, r_2)$ can not be WSNB without being SNB. The argument given is from an argument by Moore credited by Kurshan and Beneš.

Corollary 3.4. *If one of the following two conditions holds:*

$$r_1 \ge (n_2 - 1)\binom{n_1 + n_2 - 2}{n_1 - 1} + 1, \text{ and } r_2 \ge \frac{n_1 - 1}{n_2} r_1 + 1 \tag{1}$$

$$r_2 \ge (n_1 - 1)\binom{n_1 + n_2 - 2}{n_2 - 1} + 1, \text{ and } r_1 \ge \frac{n_2 - 1}{n_1} r_2 + 1 \tag{2}$$

Then, the WDM split cross-connect $S(n_1, r_1, m, n_2, r_2)$. is wide-sense non-blocking if and only if it is strictly non-blocking.

Proof. Due to symmetry, we only need to show the corresponding version on the Clos network $C = C(n_1, r_1, m, n_2, r_2)$ assuming condition (1).

When (1) holds, $n_1 r_1 > n_2$ and $n_2 r_2 > n_1$. Thus, C is SNB iff $m \geq n_1 + n_2 - 1$. To complete the proof, we show that $m \geq n_1 + n_2 - 1$ is also necessary for C to be WSNB. The second part of (1) says $r_1(n_1 - 1) \leq (r_2 - 1)n_2$. Consequently, there exists a state of the network where each input crossbars has $n_1 - 1$ demands to a subset of the first $r_2 - 1$ output crossbars. Intentionally, we want a state where the last output crossbar O_{r_2} to be free of demands. For each input crossbar I_i, identify a set $M(i)$ of precisely $(n_1 - 1)$ middle-stage crossbars which are carrying the $(n_1 - 1)$ requests from I_i. Suppose $m < n_1 + n_2 - 1$. There can be at most $\binom{n_1+n_2-2}{n_1-1}$ subsets of middle-stage crossbars of size $(n_1 - 1)$. Hence, when r_1 satisfies (1) there are n_2 identical $M(i)$. Call that identical $M(i)$ set X. Now, add n_2 demands from the n_2 input crossbars corresponding to the n_2 identical $M(i)$ to the last output crossbar. These new demands require n_2 crossbars other than those in X, making $m \geq |X| + n_2 = n_1 + n_2 - 1$. Contradiction! □

Thus, when r_1 and r_2, the number of input or output fibers are large, a WSNB WSXC can do no better than a SNB WSXC in terms of the number of WIs. The only other known results on general WSNB for Clos networks yield

Corollary 3.5. *The WDM split cross-connect $S(n, m, 2)$ is wide-sense non-blocking if and only if the number m of wavelength interchangers is at least $\lfloor 3n/2 \rfloor$. On the other hand, $S(n, m, 3)$ is wide-sense non-blocking only if m is at least $\lfloor 7n/4 \rfloor$.*

Note that the $r = 3$ case has only a necessary condition as opposed to the necessary and sufficiency in the $r = 2$ case. The previous corollary shows that there is a class of WSXCs which are WSNB but not SNB. A few common-sense routing strategies are as follows.

- Save the unused ($\mathcal{STU}$): do not route a new demand through an unused middle stage switch unless there is no other choice.
- Minimum index ($\mathcal{MI}$): route the new demand through the middle switch M_j whose index j is as small as possible.
- Packing ($\mathcal{P}$): route the new demand through a busiest middle stage switch. It is a folklore that packing is a very good routing algorithm.

(Note that these routing algorithms also apply to the WDM split cross-connects. We simply think of the WIs as the middle stage switches.)

In [8], it was shown that when $r \geq 3$, packing also requires $2n - 1$ middle-stage crossbars, implying the following:

Corollary 3.6. *Under the packing rule, the WDM split cross-connect $S(n, m, r)$ is wide-sense non-blocking if and only if it is strictly non-blocking.*

This negative result was obtained by constructing a sequence of demands establishments and releases which do not violate the packing rule but require $2n - 1$ middle-stage crossbars. What happens is that the packing rule alone is too general, and "bad" sequences of connection establishments and releases cause the rule to behave badly. Lastly, we cite another negative result:

Corollary 3.7. *Under $\mathcal{STU}, \mathcal{MI}, \mathcal{P}$, the WDM split cross-connect $S(n, m, r)$ is wide-sense non-blocking only if $m \geq 2n - \lceil \frac{n}{2^{r-1}} \rceil$.*

We now give a class of routing algorithms, a rarity in the literature, which seems to work better than the ones listed above. At least, in the case when $r = 3$ we know $S(n, m, 3)$ is WSNB but not SNB.

3.2 A routing algorithm

Routing algorithms for the Clos networks and WSXCs can be viewed as online edge-coloring algorithms of a class of bipartite multi-graphs. For simplicity, let us consider the symmetric Clos network $C(n, m, r)$. The network initially is at the empty state, namely there is no demand yet. Let $G = (\mathcal{I}, \mathcal{O}, E)$ be an $r \times r$ bipartite graphs, where $\mathcal{I} = \{I_1, \dots, I_r\}$ and $\mathcal{O} = \{O_1, \dots, O_r\}$. Each time there is a demand $D = (i, o)$ with i an inlet of I_x and o an outlet of O_y, we add an edge (I_x, O_y) into G. If D is routed through middle-stage crossbar M_z, then color this new edge z. Hence, the set of colors is $\{1, \dots, m\}$. As edges come and go, we are to assign a color to a new-coming edge using one of the m colors such that two edges incident to the same vertex are colored differently. Since there can only be at most n demands out of any I_x or to any O_y, the multi-graph G always has maximum degree at most n.

Our routing algorithm shall be casted in terms of this online coloring version. The gist of any coloring algorithm is to decide which color to pick amongst available colors. Without loss of generality, suppose the new "arriving" edge for G is (I_1, O_1). Let $G' = G - \{I_1, O_1\}$. Let C be the set of colors incident to either I_1 or O_1, and C' be the set of colors for edges of G'. The colors which are reusable are in $C' - C$. In the case $C' - C = \emptyset$, we have no choice but to use a new color, which could be the least-indexed M_z not in C.

When there are some colors in $C' - C$ to choose from, we employ the following rule:

(a) Order the colors in $C' - C$ in decreasing order of the number of occurrences (multiplicity) of the color in G'.

(b) We shall pick a color with highest multiplicity.

(c) In case of ties, suppose $Z = \{z_1, \dots, z_k\}$ is the set of colors with the same highest multiplicity. Let $E_1, \dots, E_k$ be the sets of edges whose colors are $z_1, \dots, z_k$, respectively. Define the function $c : Z \to \mathbb{N}$ as follows.

$$c(z_i) = \sum_{(I_x, O_y) \in E_i} \left(deg_{G'}(I_x) + deg_{G'}(O_y) \right).$$

Now pick color z_i whose "degree-counting" function $c(z_i)$ is highest.

(d) There might be further ties, in which case we pick the color randomly.

We shall refer to this algorithm the *degree counting* algorithm (or *DC algorithm*) for convenience. The rationale behind line (b) is the same as that of packing, namely we pick a busiest color, with the hope of saving lightly loaded colors for later use. The intuition behind line (c) is as follows. If the total degree of vertices of G' which are incident to the chosen color z_i is large, then there might be less a chance that some of them are part of new coming edges, since the maximum degree of our graph is n. Hence, the color z_i might even be further re-used if a new coming edge does not involve these vertices.

Note that since this is an online algorithm, a straightforward data structure can be maintained so that no sorting needs to be done at each step. We simply re-adjust the data structure at a connection establishment or release.

The analysis of the DC algorithm is much more difficult, just as it is the case with packing or any other routing algorithm. One reason is that the number of possible states of the network is exponentially large. Another reason, which is the main reason, is that there is no good method for analyzing this kind of algorithm, except the one given in [2] for the $r = 2$ case. We shall employ that technique here to give a sketch of the analysis of our DC algorithm when $r = 3$, due to space limit and to the fact that the actual tedious details are not that entertaining.

Theorem 3.8. *Under the degree-counting algorithm, the Clos network $C(n, m, 3)$ is wide-sense non-blocking if $m \geq \frac{11n}{6} + 3$.*

Sketch of proof. Let $\mathcal{M}$ be the set of all non-empty matchings (not necessarily perfect) of $K_{r,r}$. Consider the graph G at some instance. The set of all edges of G with the same color z form a matching $\mu(z) \in \mathcal{M}$. At any moment, let $y(G, \mu)$ be the number of colors z whose induced matching is μ, namely $y(G, \mu) = |\{z \mid \mu(z) = \mu\}|$. As new edges come and go, the variables $y(G, \mu)$ change their values. The total number of colors used by an online coloring algorithm is

$$t = \max_G \sum_{\mu \in \mathcal{M}} y(G, \mu),$$

where the maximum goes over all possible edge-colored graphs G created by a sequence of edges coming and leaving, and the way the algorithm colors them.

When $r = 3$, the proof shows inductively that $t \leq 11n/6 + 3$ by considering any sequence of additions and deletions of edges, starting from the empty graph. The fact that $t \leq 11n/6 + 3$ is obtained by also inducting simultaneously that a set of 9 other inequalities are always satisfied. This proof is in much the same way as the Beneš proof that $C(n, 3n/2, 2)$ is WSNB, plus several small twists to take into account the DC algorithm. The details are long and not very interesting. I have not been able to generalize the proof to larger values of r. □

Corollary 3.9. *Under the degree-counting algorithm, the homogeneous WDM split cross-connect $S(n, m, r)$ is wide-sense non-blocking if $m \geq \frac{11n}{6} + 3$.*

3.3 The multi-rate demand model

The proof of the following lemma is similar to that of Lemma 3.1, and hence omitted.

Lemma 3.10. *Assume that switching fabrics I and O of $S = S(n_1, r_1, m, n_2, r_2)$ and the crossbars of $C = C(n_1, r_1, m, n_2, r_2)$ are are multi-rate strictly non-blocking, Then, S is multi-rate strictly non-blocking, wide-sense non-blocking (under any algorithm), and rearrangeably non-blocking if and only if C is multi-rate strictly non-blocking, wide-sense non-blocking, and rearrangeably non-blocking, respectively.*

A result in [10] now yields

Corollary 3.11. *Suppose the rates are all in the interval $[b, B]$ ($0 < b \leq B \leq 1$), then $S(n, m, r)$ is multi-rate strictly non-blocking if*

$$m \leq 2 \max_{b \leq w \leq B} \left\lfloor \frac{n - w}{s(w)} \right\rfloor,$$

where $s(w) = \max\{1 - w, b\}$.

Note that when $b \to 0$ and $B \to 1$, the value of m tends to infinity. Thus, strictly non-blocking in the multi-rate environment is too restrictive. Fortunately, there are algorithms to make the network wide-sense non-blocking with relatively few WIs. See [7] for the Clos network versions of the following result

Corollary 3.12. $S(n, \lceil 5.75n \rceil, r)$ *is multi-rate wide-sense non-blocking.*

In terms of rearrangeability, there are a few interesting results [4, 9, 11, 12]. The problem of rearrangeablity is equivalent to a generalized bipartite multi-graph edge-coloring problem which somehow escaped the literature. Chung and Ross [3] conjectured that $C(n, 2n - 1, r)$ is multi-rate rearrangeably non-blocking, which is equivalent to saying a class of edge-weighted bipartite graphs is $2n - 1$ edge-colorable (see [12]). The following summarizes some of the known results:

Corollary 3.13. *(i)* $S(n,m,r)$ *is multi-rate rearrangeable when* $m \geq 41n/16 + O(1)$ *[4].*

(ii) $S(n,m,r)$ *is multi-rate rearrangeable when* $m \geq 2n - 1 + (r-1)/2$ *[12].*

(iii) $S(2,m,r)$ *is multi-rate rearrangeable if and only if* $m = 3$ *[12]. (This is a rarity since it is both necessary and sufficient.)*

(iv) $S(n,m,2)$ *is multi-rate rearrangeable when* $m \geq 3n/2$*. [12].*

In fact, Ngo and Vu [12] conjectured that $S(n,m,r)$ is multi-rate rearrangeable when $m \geq 2n(1 - 1/2^r)$. This is certainly a very interesting open problem.

4 Closing Remarks

Beside the conjecture mentioned at the end of the last section, an obvious open problem would be to have a good analysis of the DC algorithm. I do not expect it to work very well when r is relatively large (neither does any other known algorithm). However, analysis of the algorithm in terms of blocking probability would be of great practical interest. Another direction for further investigation might be a randomized-kind of routing algorithm, which may make the blocking probability analysis easier.

References

[1] V. E. Beneš, *Mathematical theory of connecting networks and telephone traffic*, Academic Press, New York, 1965. Mathematics in Science and Engineering, Vol. 17.

[2] ——, *Applications of group theory to connecting networks*, Bell System Tech. J., 54 (1975), pp. 407–420.

[3] S.-P. Chung and K. W. Ross, *On nonblocking multirate interconnection networks*, SIAM J. Comput., 20 (1991), pp. 726–736.

[4] D. Z. Du, B. Gao, F. K. Hwang, and J. H. Kim, *On multirate rearrangeable Clos networks*, SIAM J. Comput., 28 (1999), pp. 464–471 (electronic).

[5] H. Hinton, *An Introduction to Photonic Switching Fabrics*, Plenum Press, New York, NY, 1993.

[6] J. H. Hui, *Switching and traffic theory for integrated broadband networks*, Kluwer Academic Publishers, Boston/Dordrecht/London, 1990.

[7] F. K. Hwang, *Rearrangeability of multiconnection three-stage Clos networks*, Networks, 2 (1972), pp. 301–306.

[8] F. K. Hwang, *The mathematical theory of nonblocking switching networks*, World Scientific Publishing Co. Inc., River Edge, NJ, 1998.

[9] G.-H. Lin, D.-Z. Du, X.-D. Hu, and G. Xue, *On rearrangeability of multirate Clos networks*, SIAM J. Comput., 28 (1999), pp. 1225–1231 (electronic).

[10] R. Melen and J. S. Turner, *Nonblocking multirate networks*, SIAM J. Comput., 18 (1989), pp. 301–313.

[11] H. Q. Ngo, *A new routing algorithm for multirate rearrangeable Clos networks*, Theoretical Computer Science, to appear, (2002).

[12] H. Q. Ngo and V. H. Vu, *Multirate rearrangeable Clos networks and a generalized bipartite graph edge coloring problem*, (2002). Preprint.

[13] A. RASALA AND G. WILFONG, *Strictly non-blocking WDM cross-connects*, in Proceedings of the Eleventh Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'2000, San Francisco, CA), New York, 2000, ACM, pp. 606–615.

[14] ——, *Strictly non-blocking WDM cross-connects for heterogeneous networks*, in Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing (STOC'2000, Portland, OR), New York, 2000, ACM, pp. 513–524.