SIAM J. COMPUT.
Vol. 37, No. 3, pp. 870–894

IMPROVED BOUNDS ON NONBLOCKING 3-STAGE CLOS NETWORKS*

JOSÉ R. CORREA† AND MICHEL X. GOEMANS‡

Abstract. We consider a generalization of edge coloring bipartite graphs in which every edge has a weight in $[0, 1]$ and the coloring of the edges must satisfy that the sum of the weights of the edges incident to a vertex v of any color must be at most 1. For unit weights, König's theorem says that the number of colors needed is exactly the maximum degree. For this generalization, we show that $2.557n + o(n)$ colors are sufficient, where n is the maximum total weight adjacent to any vertex, improving the previously best bound of $2.833n + O(1)$ due to Du et al. Our analysis is interesting on its own and involves a novel decomposition result for bipartite graphs and the introduction of an associated continuous one-dimensional bin packing instance which we can prove allows perfect packing. This question is motivated by the question of the rearrangeability of 3-stage Clos networks. In that context, the corresponding parameter n of interest in the edge coloring problem is the maximum over all vertices of the number of unit-sized bins needed to pack the weights of the incident edges. In that setting, we are able to improve the bound to $2.5480n + o(n)$, also improving a bound of $2.5625n + O(1)$ of Du et al. We also consider the online version of this problem in which edges have to be colored as soon as they are revealed. In this context, we can show that $5n$ colors are enough. This contrasts with the best known lower bound of $3n - 2$ by Tsai, Wang, and Hwang but improves upon the previous best upper bound of $5.75n$ obtained by Gao and Hwang. Additionally, we show several improved bounds for more restricted versions of the problem. These online bounds are achieved by simple and easy-to-implement algorithms, inspired by the first fit heuristic for bin packing.

Key words. bipartite edge coloring, rearrangeability of 3-stage Clos networks, bin packing

AMS subject classifications. 68W40, 68R10, 90C27, 90C59

DOI. 10.1137/060656413

1. Introduction.

1.1. Clos networks.

Suppose we need to connect a set of inlets $i_1, \ldots, i_k$—which may represent telephone calls, parallel machines, or any kind of connection request—to a set of outlets $j_1, \ldots, j_k$, through an *interconnection network*, in such a way that any request permutation (i.e., a permutation of $\{1, \ldots, k\}$) can be routed simultaneously. More precisely, let us define a connection request as a pair, (i, j), where i is an inlet and j an outlet, and a *request frame* as any collection of requests such that every inlet and every outlet are associated with at most one request. The goal is to design a network that can route any request frame; such a network is called *nonblocking*.

Naturally, the simplest way to achieve this is to directly connect every inlet to every outlet by a different link. This solution, called *crossbar*, was already developed and implemented for telephone communications in the late 1930's by Western Electric (the Bell System). Despite the simplicity and nice properties of crossbar networks,

*Received by the editors April 5, 2006; accepted for publication (in revised form) February 16, 2007; published electronically August 15, 2007.
http://www.siam.org/journals/sicomp/37-3/65641.html

†School of Business, Universidad Adolfo Ibáñez, Avenida Presidente Errázuriz 3485, Santiago, Chile (correa@uai.cl). The research of this author was partially supported by grant FONDECYT 1060035.
‡Department of Mathematics, Massachusetts Institute of Technology, 77 Massachusetts Ave., Cambridge, MA 02139 (goemans@math.mit.edu). The research of this author was partially supported by NSF contract CCF-0515221 and ONR grant N00014-05-1-0148.

their main drawback is that they require too many links to achieve their goal: If we have k inlets and k outlets, they require k^2 links. In 1953, Clos [9] introduced a new type of interconnection network with the same property but that requires only $O(k^{3/2})$ links. These networks have been widely used for data communications and parallel computing systems (see, e.g., [3, 16]).

Formally, a 3-stage Clos network $C(n_1, r_1, m, n_2, r_2)$ is an 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 size $m \times n_2$, and the middle stage has m crossbars of size $r_1 \times r_2$. Moreover, each of the r_1 input switches is connected to each of the m middle switches. Similarly, the middle stage and the last stage are fully connected. We focus on the case in which $n_1 = n_2 = n$; i.e., the number of *inlets* or inputs of the input stage switches is equal to the number of *outlets* or outputs of the output stage switches. We also assume that $r_1 = r_2 = r$, even though all our results hold independently of what r_1 and r_2 are. The resulting Clos network is denoted by $C(n, m, r)$.

In a Clos network, a request frame is said to be *routable* if all requests can be routed through a middle switch so that no two requests share a link. The main question related to 3-stage Clos networks is to determine the number m of middle switches (crossbars) needed to route any request frame, i.e., for the network to be nonblocking. The answer, however, depends on the model we consider. Essentially there are three settings in which this question has been studied and used:

- *Rearrangeably nonblocking:* An interconnection network is *rearrangeably nonblocking* (or just rearrangeable) if every request frame is routable. This is the relevant question in an offline setting.
- *Strictly nonblocking:* An interconnection network is *strictly nonblocking* if any new connection request, compatible with a request frame, can be routed, independent of how the rest of the request frame is routed (i.e., independent of the state of the network). This is a relevant question in an online setting.
- *Wide-sense nonblocking:* If connection requests are revealed over time, an interconnection network is *wide-sense nonblocking* if any new connection request, compatible with a request frame, can be routed, provided that the rest of the request frame was routed according to a given routing algorithm. This question is the most important in practice, since it is motivated by the online environment, but it is less restrictive than the strictly nonblocking requirement. It is important to mention that several authors consider the more restrictive definition of wide-sense nonblocking in which the algorithm has to be able to route new connection requests even if previous connections terminate (see, e.g., Benes' original book [5]).

Clos himself noted that $C(n, 2n-1, r)$ is strictly nonblocking (which implies that it is wide-sense nonblocking as well), while, shortly after, Slepian [29] (see also [5]) proved that $C(n, n, r)$ is rearrangeable. Moreover, both results are best possible. It is then clear that if the total number of inlets is k and we choose $n = r = \sqrt{k}$, the number of links required in a 3-stage Clos network is $5 \times k^{3/2}$ or $8 \times k^{3/2}$, depending on whether we need it to be rearrangeably or strictly nonblocking.

Although our focus here will be the study of 3-stage Clos networks, let us briefly mention a few results for general interconnection networks. Shannon [27] showed that $\Omega(k \log k)$ links are needed for an interconnection network to even be rearrangeable. Surprisingly, this lower bound was matched by Benes [6] and Beizer [4], who designed rearrangeable networks of size $O(k \log k)$. Later, Bassalygo and Pinsker [2] constructively showed the existence of strictly nonblocking networks of size $O(k \log k)$.

The multirate environment. We just described the *classic* switching environment, in which connection requests fully use a link and have all the same bandwidth. However, in modern communications, different requests may have different bandwidths and may be combined in a given link if the "link capacity" is large enough to carry both requests. This setting is usually called the *multirate* environment. In such a setting, a connection request is a triple (i, j, w), where i is an inlet, j an outlet, and w the weight (thus, the classic environment corresponds to the special case in which all weights are 1). A *request frame* is a collection of requests such that the total weight of all requests in the frame involving a fixed inlet or outlet does not exceed 1. In a Clos network, all $r \times m$ links between the input switches and middle switches and all $m \times r$ links between the middle switches and the output switches also have capacity 1. A request frame is said to be *routable* if all requests can be routed through a middle switch so that none of the link capacities is violated. For a recent survey on multirate Clos networks, we refer the reader to the excellent survey by Turner and Melen [30], who also initiated the research on multirate switching networks [23]. As in the classic environment, the question is to determine the minimum value of m of middle switches such that any request frame can be routed; the network is then said to be multirate nonblocking. Again, the answer depends on whether the problem is considered online or offline. However, the questions are still wide open and need further investigations.

- *Rearrangeably nonblocking (offline):* An interconnection network is said to be *multirate rearrangeably nonblocking* (or just rearrangeable) if every request frame is routable. The question is thus to determine the minimum value of m of middle switches such that $C(n, m, r)$ is multirate rearrangeable, and this minimum value is denoted by $m(n, r)$. It is particularly interesting to obtain bounds that are independent of r.
- *Wide-sense nonblocking (online):* If connection requests are revealed over time (both the inlet-outlet pair and its weight), an interconnection network is said to be *wide-sense nonblocking* if any new connection request, compatible with a request frame, can be routed, provided that all the rest of the request frame was routed according to a given routing algorithm. Thus, the question is again to determine the minimum value m of middle switches such that $C(n, m, r)$ is wide-sense nonblocking, and this value is denoted by $m_W(n, r)$. Let us emphasize again that we assume that the requests never terminate, i.e., that we have no deletions during the execution; this is the same weaker setting as in [13], for example.
- *Strictly nonblocking:* In the multirate environment, we say that an interconnection network is *strictly nonblocking* if any new connection request, compatible with a request frame, can be routed, independent of how the rest of the request frame is routed (i.e., independent of the state of the network).

1.2. Problem definition. The question of rearrangeability and nonblocking properties of a 3-stage Clos network can be translated in graph-theoretic terms in the following way. We are given a bipartite (multi)graph $G = (V, E)$ with bipartition A, B (say with $|A| = |B| = r$); in what follows, all our graphs will be multigraphs. A and B represent the input and output switches, respectively. Edge $e = (i, j)$ represents a request between input switch i and output switch j and carries a weight $0 \le w(e) \le 1$. The assumption of the requests being a request frame can be translated into the assumption that the weights on the edges incident to $v \in V$ can be packed into n unit-sized bins. That is, for all $v \in V$, the set $\delta(v)$ of edges incident to v can be partitioned into n groups C_i^v, $i = 1, \dots, n$, satisfying

$$\sum_{e \in C_i^v} w(e) \leq 1 \quad \text{for all } i = 1 \dots, n. \tag{1}$$

Following the notation in [24], let $\mathcal{B}_r^n$ be the collection of such edge-weighted bipartite multigraphs.

A Clos network $C(n, m, r)$ is then (multirate) rearrangeable if, for every graph in $\mathcal{B}_r^n$, the edges can be colored with m colors so that the total weight of all edges of the same color incident to a vertex v is at most 1. The question is thus to determine the minimum number $m(n, r)$ of colors needed to properly color every weighted bipartite graph in $\mathcal{B}_r^n$. In the online setting, we know only a priori that the graph belongs to $\mathcal{B}_r^n$, but the edges and their weight are revealed over time. Similar to the rearrangeable case, a Clos network $C(n, m, r)$ is (multirate) wide-sense nonblocking if there exists an online algorithm A such that, for every graph in $\mathcal{B}_r^n$, the edges can be colored with m colors so that the total weight of all edges of the same color incident to a vertex v is at most 1. The question is thus to determine the minimum number $m_W(n, r)$ for which there is an online algorithm that properly colors every weighted bipartite graph in $\mathcal{B}_r^n$ using no more than $m_W(n, r)$ colors. In the same manner, $C(n, m, r)$ is (multirate) strictly nonblocking if, for any $G = (V, E) \in \mathcal{B}_r^n$ and any proper m-coloring of $(V, E \setminus \{e\})$, for any $e \in E$, edge e can be colored without changing the color of any already colored edge and using any extra color. Now $m_S(n, r)$ is the minimum number of colors such that $C(n, m_S(n, r), r)$ is strictly nonblocking.

If all weights are forced to belong to a subset $I \subset [0, 1]$, let $\mathcal{B}_r^n(I)$ denote the natural extension of $\mathcal{B}_r^n$. In this case, $m_I(n, r)$ is the smallest integer such that every graph in $\mathcal{B}_r^n(I)$ admits a proper coloring with $m_I(n, r)$ colors. The quantities $m_{W\,I}(n, r)$ and $m_{S\,I}(n, r)$ are the natural counterparts of $m_I(n, r)$ in the wide-sense and strictly nonblocking setting.

Another special case that has attracted attention is when all edge weights can take only k different values (known, in Clos network terminology, as the bounded rate environment, or k-rate environment). We denote by $m^k(n, r)$ the minimum number of middle switches so that $C(m^k(n, r), n, r)$ is multirate rearrangeable when all request frames have weights with only k different values. Similarly, $m_W^k(n, r)$ is the corresponding counterpart of $m^k(n, r)$.

In section 3, we focus on a generalized bipartite edge-coloring problem, very similar to the one just described, except that we require only the weights incident to any vertex to add up to at most n. That is, condition (1) is replaced by the following weaker condition:

$$\sum_{e \in \delta(v)} w(e) \leq n \quad \text{for all } v \in V. \tag{2}$$

Here $\mathcal{D}_r^n$ denotes the natural counterpart of $\mathcal{B}_r^n$. As $\mathcal{B}_r^n \subseteq \mathcal{D}_r^n$, the required number of colors in this case, denoted by $M(n, r)$, is clearly greater than or equal to $m(n, r)$. If all weights are forced to belong to a subset $I \subset [0, 1]$, $\mathcal{D}_r^n(I)$ and $M_I(n, r)$ denote the natural counterparts of $\mathcal{B}_r^n$(I) and $m_I(n, r)$.

1.3. Discussion of previous work. Let us review some existing results on this problem. We start by giving the most relevant results on rearrangeability, and later we focus on wide-sense and strictly nonblocking properties.

Rearrangeability. The first important result was proved shortly after the introduction of 3-stage Clos networks and is due to Slepian [29] (see also [5]). He used König's edge-coloring theorem [19] (see also [11]) to prove that $m_{[1,1]}(n,r) = n$. Melen and Turner [23] initiated the research on multirate switching networks and proved that $m_{[0,1/2]}(n,r) \leq M_{[0,1/2]}(n,r) \leq 2n-1$. More generally, they proved that

$$m_{[0,B]}(n,r) \leq M_{[0,B]}(n,r) \leq \frac{n}{1-B}.$$

On the other hand, it is easy to prove that $m_{[b,1]}(n,r) \leq n\lfloor \frac{1}{b} \rfloor$ and that $M_{[b,1]}(n,r) \leq \frac{n}{b}$.

Previous to this work, the best bounds known on $m(n,r)$ in the general setting are $\frac{5n}{4} \leq m(n,r) \leq \frac{41n}{16} + O(1)$ and were obtained by Ngo and Vu [24] (lower bound) and Du et al. [12] (upper bound). The latter authors also obtained the previously best bounds for $M(n,r)$, namely $2n-1 \leq M(n,r) \leq \frac{17n}{6} + O(1)$.

In the k-rate environment, better bounds have been proved. For $k=2$, one can actually verify the Chung–Ross conjecture, namely, that the $2n-1$ bound holds in this case [8]. Moreover, Lin et al. [21] proved that

$$m^3(n,r) \leq \frac{9n}{4} + O(1) \quad \text{ and } \quad m^3_{(\frac{1}{5},1]}(n,r) \leq 2n.$$

The first bound is an improvement over the $\frac{7n}{3}$ bound obtained by Lin et al. [20]. Unfortunately, the proofs of all bounds for the finite rate environment rely on rather tedious case analysis.

Wide-sense and strictly nonblocking. Let us now survey some of the most relevant results concerning nonblocking properties of 3-stage Clos networks. In the classical environment, Clos [9] proved that $C(n, 2n-1, r)$ is strictly nonblocking. Unfortunately, as first noted in [23], in the multirate environment, $C(n,m,r)$ cannot be strictly nonblocking unless m is infinity. Indeed, consider the network $C(n,m,1)$ and assume that there is a connection request of weight 1 and $(n-1)/\varepsilon$ connection requests of weight $\varepsilon = (n-1)/m$ between the only input and output switch pair in the network. A possible current state for the network is that each small connection request is routed along a different middle switch, and thus the large request cannot be routed, implying that the network is in a blocking state. However, if connection requests are restricted to have weights within some interval, finite bounds can be obtained. Indeed, Melen and Turner [23] proved that $m_{S[b,1]}(n,r) \leq 2\lfloor (n-1)/b \rfloor + 3$, which was further improved by Chung and Ross [8] to $m_{S[b,1]}(n,r) \leq 2\lfloor 1/b \rfloor (n-1) + 1$. The latter authors also proved that $m_{S(0,B]} \leq 2\lceil \frac{n-B}{1-B} \rceil + 1$.

The bad example above motivated the algorithmic concept of wide-sense nonblocking. Indeed, already in [23] it was noted that $8n$ middle switches are enough to ensure the wide-sense nonblocking condition, i.e., $m_W(n,r) \leq 8n$. Later, Chung and Ross [8] used their bounds on $m_{S[b,1]}(n,r)$ and $m_{S(0,B]}(n,r)$ to improve the bound. Indeed, their algorithm would split connection requests according to their weight: the smaller than or equal to 1/2 and those strictly larger than 1/2. The bound is therefore

$$m_W(n,r) \leq m_{S(1/2,1]}(n,r) + m_{S(0,1/2]}(n,r) \leq 2n - 2 + 1 + 4n + 1 = 6n.$$

The best known bound previous to our result was obtained by Gao and Hwang [13]. They used a *quota scheme*, which consists of reserving some middle switches for large connections while letting the rest carry any connection request. This approach led them to the bound $m_{W[0,1/2]} \leq 3.75n$, implying, in the same manner as above, that

$$m_W(n,r) \leq 5.75n.$$

The study of lower bounds for wide-sense nonblocking properties has been much more recent. Bar-Noy, Motwani, and Naor [1] were the first to prove that in the classical setting $m_{W[1,1]}(n, r) \geq 2n-1$ for exponentially large r. This surprising result essentially says that in the classical single-rate environment the strictly nonblocking and wide-sense nonblocking conditions are the same. Moreover, recent work by Haxell et al. [14] shows that this lower bound holds even for $r = \Omega(n^2)$. In the multirate environment, there is only a recent improvement on the previous bound. Tsai, Wang, and Hwang [31] proved that $m_W(n, r) \geq 3n - 2$, and their proof also works in the more restricted 2-rate environment.

1.4. Overview and main results in the paper. The main goal of this paper is to present bounds on $M(n, r)$, $m(n, r)$, and $m_W(n, r)$. Indeed, we will show that $2.557n$, $2.548n$, and $5n$ are, respectively, upper bounds on these numbers.

We start in section 2 by showing a result on balanced decomposition of bipartite graphs into matchings. In the context of Clos networks, this result becomes useful only in section 3; however, we believe it is interesting on its own, and so we have decided to present it in a separate section. The question that is addressed is as follows: Given a bipartite graph G and nonnegative numbers $\gamma_1, \ldots, \gamma_l$ summing to 1, decompose the graph into $F_1, \ldots, F_l$ such that the degree of any vertex v in F_i is approximately γ_i times the degree of v in G. We show that the decomposition can be done such that for all i and all vertices, the degree of v in F_i differs from its required value by an additive constant less than 3. The question whether this constant can be decreased to 1 is to the best of our knowledge open.

Our main contribution in this paper, proved in section 3, is the following result.

THEOREM 1. *The number of colors required to properly color every weighted bipartite graph in $\mathcal{D}_r^n$ is at most $2.557n + o(n)$. In other words,*

$$M(n, r) \leq 2.557n + o(n).$$

Observe that this does not improve only upon Du et al.'s bound of $\frac{17}{6}n + O(1)$ on $M(n, r)$ but even slightly upon their bound of $\frac{41}{16}n + O(1) = 2.5625n + O(1)$ on $m(n, r)$. In fact, our approach can also be applied to bounding $m(n, r)$ directly, and this gives us a slightly improved bound of $m(n, r) \leq 2.5480n + o(n)$. The latter improvement is sketched in section 3.7.

For most of section 3, we consider the generalized bipartite edge coloring problem in which the weights on edges incident to any vertex sum to at most n, i.e., graphs in $\mathcal{D}_r^n$. The approach we consider to attack this problem associates a bin packing instance with every such generalized edge coloring instance. For this purpose, we first decompose the edge weighted bipartite graph $G = (V, E)$ into a union of matchings. We then create a bin packing instance in which all bins have size 1. We create an item of our bin packing instance for each matching in our decomposition, and we set its size to be the maximum weight of any edge in the matching. A packing with k bins immediately leads to a valid k-coloring by simply coloring the edges of all matchings (items) placed in the same bin with the same unique color. As we shall see in section 3.1, this approach needs that we first discard all edges whose weight is less than some parameter α (to be determined). This can be done using the following result implicit in Du et al. [12].

LEMMA 2. *Consider $G = (V, E) \in \mathcal{D}_r^n$ with bipartition $V = A \cup B$ and assume that we have used at least $\frac{2n}{1-\alpha}$ colors to color all edges except some edges e with $w(e) \leq \alpha$. Then we can greedily color these remaining edges without using any additional color. In particular, if $M_{(\alpha,1]}(n, r) \leq \lceil 2n/(1-\alpha) \rceil$, then $M(n, r) \leq \lceil 2n/(1-\alpha) \rceil$.*

Proof. If $e = (u, v) \in E$ with $w(e) \le \alpha$ cannot be colored then the total weight of edges of a given color i incident to either u or v is greater than $1 - \alpha$. Summing over all $\frac{2n}{1-\alpha}$ colors, we get a contradiction with sum of conditions (2) for u and v. ∎

Therefore, we can focus on instances in which all weights are in $[\alpha, 1]$, provided that we are willing to use $\lceil \frac{2n}{1-\alpha} \rceil$ colors. As the proof of this last result used only any greedy algorithm, the result also holds in the wide-sense nonblocking setting.

Our main contribution is to show that, for *any* generalized edge coloring problem with weights in $[\alpha, 1]$, we can decompose the bipartite graph into matchings in such a way that the corresponding bin packing instance can be packed into at most $n + o(n)$ bins plus the number of bins required to pack a continuous bin packing instance with density $\frac{n}{x^2}$ for $x \in [\alpha, 1]$ (i.e., the number of items with size in the interval $(x, x + dx)$ is $\frac{n}{x^2} dx$). We should emphasize that our bin packing instance is independent of the given bipartite graph G; it is based only on the fact that $G \in \mathcal{D}_r^n$. Although it is easier to refer in the statements here to the continuous bin packing instance, we actually deal only with an arbitrarily fine discretization of it and consider discrete bin packing instances. Our decomposition of the graph into matchings relies on the result of section 2 and is described in section 3.2, while the construction of our bin packing instance is detailed in section 3.3.

Once the continuous bin packing instance with density $\frac{n}{x^2}$ for $x \in [\alpha, 1]$ is constructed, in sections 3.4, 3.5, and 3.6, we turn to compute the number of bins it requires. First, we observe that all items of size greater than $1 - \alpha$ need to be placed alone in bins; they therefore require $\int_{1-\alpha}^{1} \frac{n}{x^2} dx = \frac{\alpha}{1-\alpha} n$ bins. For the remaining items with density $\frac{n}{x^2}$ for $x \in [\alpha, 1 - \alpha]$, we prove that they can be *perfectly* packed. This means that the number of bins they require is simply their total size, up to lower-order terms (accounting for the discretization). This means that they require $\int_{\alpha}^{1-\alpha} x \frac{n}{x^2} dx = n \ln \frac{1-\alpha}{\alpha}$ additional bins. This relies on a result of Rhee and Talagrand [26]. The total number of bins used is thus $(1 + \frac{\alpha}{1-\alpha} + \ln \frac{1-\alpha}{\alpha})n$, and we choose α so that this equals $\frac{2}{1-\alpha} n$ in order to be able to greedily color the edges with weight lower than α. For $\alpha = 0.217811\ldots$, we obtain that the number of colors needed is less than $2.557n$.

It is worth mentioning that our main result can be done algorithmically. Indeed, the continuous bin packing instance is independent of the input; therefore, a discretization of it can be solved optimally a priori by exhaustive search (or by using any good algorithm for bin packing). The matching decomposition, for edges with weight in $[\alpha, 1]$, can be efficiently done using network flows techniques (see Lemma 5). Finally, the edges with weight in $(0, \alpha)$ can be greedily colored as in Lemma 2.

In section 4, we will use simple adaptations of the first fit (FF) heuristic for the classical bin packing problem to obtain improved bounds on the wide-sense nonblocking properties of 3-stage Clos networks. In the bin packing setting, FF places a new item in the first bin that has space available for it; in our online setting, it will simply color an edge with the smallest possible color (under some arbitrary order on the colors) as long as it does not violate condition (1). Our main result, proved in section 4.1, is to show that $C(n, 5n, r)$ is wide-sense nonblocking, i.e., $m_W(n, r) \le 5n$. Later, in section 4.2, we show that $m_{W(0,1/2]} \le 3.601n + 3$. Both bounds improve upon the bounds obtained by Gao and Hwang [13] of $5.75n$ and $3.75n$, respectively.

Additionally, in section 4.3 we are able to show that in the 2-rate environment, there is an online algorithm that uses no more than $3n$ middle switches to schedule any request frame. This not only improves the previous best known bound of $4n$ [13] but also almost matches the lower bound on $m_W(n, r)$ of $3n - 2$ obtained by

Tsai, Wang, and Hwang [31] which is also valid in the 2-rate case. We can therefore conclude that

$$3n - 2 \leq m^2_W(n, r) \leq 3n.$$

Finally, in section 5 we prove that using an analogue of the FF decreasing heuristic for bin packing, no more than $\frac{8n}{3}$ middle switches are needed to route any request frame. As sorting is needed, this bound holds only in the offline setting, and it does not improve upon the bound of $m(n, r) \leq 2.548n + o(n)$ given in section 3. However, it has the following advantages: (i) it is a nonasymptotic result; (ii) it is a very simple to implement algorithm; and (iii) it can be implemented to run in time $O(n \log n)$.

2. Balanced decompositions of bipartite graphs. Given a subset of edges F of a graph G and a vertex v, we let $\deg_F(v)$ denote the degree of vertex v in F, that is, $|\delta(v) \cap F|$, where $\delta(v)$ is the set of edges incident to v in the graph. The following result follows easily from network flow theory.

LEMMA 3 (Hoffman [15]). *Consider a bipartite graph $G = (V, E)$ and let $0 \leq \mu_1, \mu_2$ with $\mu_1 + \mu_2 = 1$. Then there exists a partition of E into E_1 and E_2 such that*

$$\lfloor \mu_i \deg_E(v) \rfloor \leq \deg_{E_i}(v) \leq \lceil \mu_i \deg_E(v) \rceil$$

for $i = 1, 2$ and all $v \in V$.

Proof. Let A, B be the bipartition of the bipartite graph G. Orient all edges from A to B. Add a source with arcs to all vertices in A and a sink with arcs from all vertices in B. Set the capacity of all the arcs in E to be 1, and set upper and lower capacities on the arcs adjacent to the source and sink to be $\lceil \mu_1 \deg_E(v) \rceil$ and $\lfloor \mu_1 \deg_E(v) \rfloor$, where v is the corresponding adjacent vertex. As a feasible flow can be obtained by setting the flow on every arc in E to be μ_1, there exists an integer feasible flow, and this flow corresponds to the edge set E_1. The remaining edges E_2 also satisfy the required property. ◻

The next theorem is an extension of Hoffman's result.

THEOREM 4. *Consider a bipartite graph $G = (V, E)$ and let $\gamma_1, \ldots, \gamma_l \in (0, 1)$ such that $\sum_{i=1}^{l} \gamma_i = 1$. Then there exists a partition $E_1, \ldots, E_l$ of E such that for all $v \in V$ and all $i = 1, \ldots, l$,*

$$\gamma_i \deg_E(v) - e_i(v) < \deg_{E_i}(v) < \gamma_i \deg_E(v) + e_i(v).$$

Here $e_i(v) < 3$, and $\sum_{i=1}^{l} e_i(v) \leq 2(l - 1)$.

Proof. Let $L = \{1, \ldots, l\}$. We construct a binary tree T with $l-1$ internal nodes and l leaves, each node being labelled by a subset of L. The root is labelled with L, and the l leaves are labelled by a distinct singleton subset of L. If an internal node is labelled with N, then its two children are labelled with I and $N \setminus I$, where $I, N \setminus I$ is the most balanced number partition of N; i.e., I is such that $\max\{\gamma(I), \gamma(N \setminus I)\}$ is minimized (for a set S, $\gamma(S)$ denotes $\sum_{i \in S} \gamma_i$).

With every node with label I, we also associate an edge set $E(I)$. We first set $E(L) = E$. Given $E(N)$ for an internal node N, we obtain $E(I)$ and $E(N \setminus I)$ for its children by applying Lemma 3 to the graph with edge set $E(N)$ and with $\mu_1 = \gamma(I)/\gamma(N)$ and $\mu_2 = 1 - \mu_1$. The leaves are thus associated with subgraphs $E(\{i\})$ which make a partition of E. We claim that $E(\{i\})$ satisfies the required properties for E_i.

Fix a vertex $v \in V$ (for simplicity, we just drop v when writing $\deg_*(v)$) and an index $i \in L$. Let $\{i\} = A_0 \subset A_1 \subset \cdots \subset A_k = L$ be the labels on the path from the leaf $\{i\}$ to the root. We now derive an upper bound on $\deg_{E_i}(v)$ (and we could proceed similarly for the lower bound). From Lemma 3, we have that

$$\begin{aligned}
\deg_{E_i}(v) = \deg_{E(A_0)} &< \frac{\gamma(A_0)}{\gamma(A_1)} \deg_{E(A_1)} + 1 \\
&< \frac{\gamma(A_0)}{\gamma(A_1)} \left(\frac{\gamma(A_1)}{\gamma(A_2)} \deg_{E(A_2)} + 1 \right) + 1 \\
&< \frac{\gamma(A_0)}{\gamma(A_1)} \left(\frac{\gamma(A_1)}{\gamma(A_2)} \left(\cdots \left(\frac{\gamma(A_{k-1})}{\gamma(A_k)} \deg_E + 1 \right) \cdots \right) + 1 \right) + 1 \\
&= \frac{\gamma(A_0)}{\gamma(A_k)} \deg_E + 1 + \frac{\gamma(A_0)}{\gamma(A_1)} + \frac{\gamma(A_0)}{\gamma(A_2)} + \cdots + \frac{\gamma(A_0)}{\gamma(A_{k-1})} \\
&= \gamma_i \deg_E + e_i(v),
\end{aligned}$$

where $e_i(v) = 1 + \frac{\gamma(A_0)}{\gamma(A_1)} + \frac{\gamma(A_0)}{\gamma(A_2)} + \cdots + \frac{\gamma(A_0)}{\gamma(A_{k-1})}$. Let $\eta = \min_{i \in A_1} \gamma_i$ and let j be the arg min. Let $a = \gamma(A_0) \geq \eta$. Thus we have $\gamma(A_1) \geq a + \eta$. In general, when considering A_k, we split it into A_{k-1} and $A_k \setminus A_{k-1}$, while we could have split it into $A_{k-1} \setminus \{j\}$ and the rest. This implies that $\gamma(A_{k-1}) - \eta \leq \gamma(A_k) - \gamma(A_{k-1})$, i.e., $\gamma(A_k) \geq 2\gamma(A_{k-1}) - \eta$. Using this repeatedly, we get $\gamma(A_2) \geq 2a + \eta$, $\gamma(A_3) \geq 4a + \eta$, and, generally, $\gamma(A_{k-1}) \geq 2^k a + \eta$. Thus the bound becomes

$$\begin{aligned}
e_i(v) &\leq 1 + \frac{a}{a+\eta} + \frac{a}{2a+\eta} + \frac{a}{4a+\eta} + \frac{a}{8a+\eta} + \cdots \\
&\leq 1 + \frac{a}{a} + \frac{a}{2a} + \frac{a}{4a} + \frac{a}{8a} + \cdots < 3.
\end{aligned}$$

Finally, in order to get a bound on $\sum_i e_i(v)$, observe that

$$\sum_{i=1}^{l} e_i(v) = \sum_{(\text{all labels } N \text{ except the root})} \sum_{i \in N} \frac{\gamma_i}{\gamma(N)} = 2(l-1),$$

since there are $2l-1$ nodes in the binary tree. A proof of the lower bound on $\deg_{E_i}(v)$ is identical. $\square$

We suspect that the bound can be further improved. If $\gamma_i = 1/l$ for every i, de Werra [10] has shown that we can impose $\lfloor \gamma_i \deg_E(v) \rfloor \leq \deg_{E_i}(v) \leq \lceil \gamma_i \deg_E(v) \rceil$ for every i, while Theorem 4 implies $\lfloor \gamma_i \deg_E(v) \rfloor - 2 \leq \deg_{E_i}(v) \leq \lceil \gamma_i \deg_E(v) \rceil + 2$ for every i (without making assumptions on the γ_i's). We do not know whether the tighter condition (without the $+2$) can be imposed in the general case. The proof technique used here, however, cannot even improve the $+2$ term into a $+1$ term. Indeed, for $\gamma_i = 1/13$ for $i = 1, \ldots, 13$, one can see that no partitioning scheme would give a bound on $e_i(v)$ (using the analysis in the proof of Theorem 4) better than $1 + \frac{1}{2} + \frac{1}{4} + \frac{1}{5} + \frac{1}{13} = 2 + \frac{7}{260}$ (and this can be shown to be the worst when all γ_i's are equal).

As stated, the proof of Theorem 4 is not algorithmic, since we need to solve number partition as a subroutine. However, we used in the proof only the fact that the partitioning of N used is locally optimum in the sense that no item can be moved to the other side of the partition while making it more balanced. A locally optimum number partition can be obtained in polynomial time in several ways. Brucker, Hurink, and Werne [7] show (in the context of scheduling parallel machines) that iteratively improving the partition until a local optimum is reached takes $O(|N|^2)$ iterations. Schuurman and Vredeveld [28] noted that iteratively finding the best local improvement requires $O(|N|)$ iterations, which implies an overall running time of $O(|N| \log |N|)$. One can also use the differencing method of Karmarkar and Karp [18]. This differencing method, which also runs in $O(|N| \log |N|)$ time, consists of repeatedly replacing the largest two items by one new item whose size (i.e., γ value) equals the difference in sizes of these largest two items until only one item of size say Δ remains. By inverting the process, one can easily obtain a partition $(I, N \setminus I)$ with $\gamma(I) = \gamma(N \setminus I) + \Delta$. A simple inductive argument shows that all items in I have $\gamma_i \geq \Delta$, and therefore the partition obtained is locally optimum. Using any of these algorithms to find a local optimum, a partition of the edge set satisfying the conditions of Theorem 4 can be obtained in polynomial time.

3. Rearrangeably nonblocking Clos networks. In this section, we consider the generalized bipartite edge coloring problem in which the weights on edges incident to any vertex sum to at most n, i.e., graphs in $\mathcal{D}_r^n$. As described earlier, we associate a bin packing instance with every such generalized edge coloring instance by decomposing the edge weighted bipartite graph $G = (V, E)$ into a union of matchings. We then create a bin packing instance in which all bins have size 1. We create an item of our bin packing instance for each matching in our decomposition, and we set its size to be the maximum weight of any edge in the matching. A packing with k bins immediately leads to a valid k-coloring by simply coloring the edges of all matchings (items) placed in the same bin with the same unique color.

3.1. Limitations. Consider the following trivial instance of our generalized edge coloring problem. Let X be a finite subset of $(0, 1]$ and create a vertex in A and in B for each element $x \in X$ and $\lfloor \frac{n}{x} \rfloor$ edges between them. In this case, $2n - 1$ colors are sufficient (and needed if $\frac{1}{2} + \epsilon \in X$ for some small ϵ). No matter what decomposition into matchings we consider, our bin packing instance has at least $\lfloor \frac{n}{x} \rfloor$ items (matchings) of size at least x for every $x \in X$. If $X = \{x_0, x_1, \ldots, x_l\}$ with $x_0 > x_1 > \cdots > x_l$, this bin packing instance requires no fewer bins than another bin packing instance with $\lfloor \frac{n}{x_i} \rfloor - \lfloor \frac{n}{x_{i-1}} \rfloor$ items of size x_i for every $i \geq 1$ and $\lfloor \frac{n}{x_i} \rfloor$ items of size x_0. As X gets denser in $(0, 1]$, this bin packing instance tends to a continuous bin packing instance with density $\frac{n}{x^2}$ (i.e., the number of items of size in $(x, x + dx)$ is $\frac{n}{x^2} dx$) after having removed the n items of size 1. Now the number of bins required is at least the total size of all items $n + \int_0^1 x \frac{n}{x^2} dx$, which is unbounded!

To overcome this problem, we first discard all edges whose weight is less than some parameter α (to be determined) by using Lemma 2. This said, we can turn to proving the graph partitioning result in which our work is based.

3.2. Partitioning the graph. From now on we fix a parameter $0 < \alpha < 1$ and work with graphs in $\mathcal{D}_r^n(\alpha, 1)$. The decomposition we need to construct our bin packing instance is given below.

LEMMA 5. *Consider the sequence $\alpha_0 = 1 > \alpha_1 > \alpha_2 > \cdots > \alpha_p = \alpha \geq 0$. Let $G = (V, E) \in \mathcal{D}_r^n(\alpha, 1)$. Then there exist sets $F_1, \ldots, F_p$ partitioning E such that the following hold:*

(i) $\max_{e \in F_k} w(e) \leq \alpha_{k-1}$.

(ii) *For all vertices* $v \in V$,

$$\deg_{F_k}(v) \leq \left(\frac{1}{\alpha_k} - \frac{1}{\alpha_{k-1}}\right) n + a_k(v) \quad \text{for all } 2 \leq k \leq p,$$

$$\deg_{F_1}(v) \leq \frac{n}{\alpha_1} + a_1(v),$$

where $a_k(v) \leq 3(p-k+1)$.

Proof. Consider an instance $G = (L, R, E)$ with weight function w and let

$$D_i = \{e \in E \ : \ w(e) \in (\alpha_i, \alpha_{i-1}]\}$$

for $i = 1, \ldots, p$. From inequality (2) we can easily deduce that for all $v \in L \cup R$,

$$\sum_{i=k}^{p} \alpha_i \deg_{D_i}(v) \leq n \quad \text{for all } k = 1 \ldots, p. \tag{3}$$

If we divide the inequality (3) corresponding to $k = 1$ by α_1 and multiply the kth inequality (3) by $(\frac{1}{\alpha_k} - \frac{1}{\alpha_{k-1}})$, we obtain the following set of inequalities:

$$\left(\frac{1}{\alpha_1}\right) \alpha_1 \deg_{D_1}(v) + \left(\frac{1}{\alpha_1}\right) \alpha_2 \deg_{D_2}(v) + \cdots + \left(\frac{1}{\alpha_1}\right) \alpha_p \deg_{D_p}(v) \leq \left(\frac{1}{\alpha_1}\right) n$$

$$\left(\frac{1}{\alpha_2} - \frac{1}{\alpha_1}\right) \alpha_2 \deg_{D_2}(v) + \left(\frac{1}{\alpha_2} - \frac{1}{\alpha_1}\right) \alpha_3 \deg_{D_3}(v) + \cdots + \left(\frac{1}{\alpha_2} - \frac{1}{\alpha_1}\right) \alpha_p \deg_{D_p}(v)$$

$$\leq \left(\frac{1}{\alpha_2} - \frac{1}{\alpha_1}\right) n$$

$$\vdots$$

$$\left(\frac{1}{\alpha_p} - \frac{1}{\alpha_{p-1}}\right) \alpha_p \deg_{D_p}(v) \leq \left(\frac{1}{\alpha_p} - \frac{1}{\alpha_{p-1}}\right) n.$$

Note that, for all $i = 1, \ldots, p$, the coefficients in front of $\deg_{D_i}(v)$ over the above inequalities sum to 1. Therefore, for each D_i we can apply Theorem 4 with

$$\gamma_1^i = \frac{1}{\alpha_1}\alpha_i, \gamma_2^i = \left(\frac{1}{\alpha_2} - \frac{1}{\alpha_1}\right) \alpha_i, \ldots, \gamma_i^i = \left(\frac{1}{\alpha_i} - \frac{1}{\alpha_{i-1}}\right) \alpha_i$$

to partition D_i into sets $D_i^1, \ldots, D_i^i$ such that for all $k = 1, \ldots, i$ and all $v \in V$,

$$\gamma_k^i \deg_{D_i}(v) - e_k^i(v) < \deg_{D_i^k}(v) < \gamma_k^i \deg_{D_i}(v) + e_k^i(v),$$

where $e_k^i(v) \leq 3$ and $\sum_{k=1}^{i} e_k^i(v) \leq 2(i-1)$.

We are now ready to finish the proof. Define $F_k = D_k^k \cup D_{k+1}^k \cup \cdots \cup D_p^k$ for all

$k = 1, \ldots, p$. Thus letting $a_k(v) = \sum_{i=k}^{p} e_k^i(v) \le 3(p-k+1)$, we have the following:

$$\begin{aligned} \deg_{F_1}(v) &\le \sum_{i=1}^{p} \left(\gamma_1^i \deg_{D_i}(v) + e_1^i(v)\right) \\ &\le \left(\frac{1}{\alpha_1}\right) n + a_1(v); \\ \deg_{F_k}(v) &\le \sum_{i=k}^{p} \left(\gamma_k^i \deg_{D_i}(v) + e_k^i(v)\right) \\ &\le \left(\frac{1}{\alpha_k} - \frac{1}{\alpha_{k-1}}\right) n + a_k(v), \qquad 2 \le k \le p. \qquad \square \end{aligned}$$

3.3. The associated bin packing problem. Let us now consider a bipartite graph $G = (V, E) \in \mathcal{D}_r^n(\alpha, 1)$ and $F_1, \ldots, F_p$ as in Lemma 5. By König's theorem, F_k can be decomposed into no more than

$$\left(\frac{1}{\alpha_k} - \frac{1}{\alpha_{k-1}}\right) n + \max_{v \in V} a_k(v)$$

matchings, for all $k = 2, \ldots, p$ and F_1 can be decomposed into $\frac{n}{\alpha_1} + \max_{v \in V} a_1(v)$ matchings. We now construct an instance of the one-dimensional bin packing problem with unit-sized bins. Arbitrarily select $(\frac{1}{\alpha_k} - \frac{1}{\alpha_{k-1}})n$ matchings (or, more formally, the floor of this quantity) in the decomposition of F_k for $k = 2, \ldots, p$ and assign each of them an item of size α_{k-1}. Similarly, arbitrarily select $\frac{n}{\alpha_1}$ matchings in the decomposition of F_1 and assign each of them an item of size α_0. Let $\mathcal{M}$ be those matchings selected in $F_1, \ldots, F_p$, and, by construction, we have an item for each element of $\mathcal{M}$. Our bin packing instance is thus the following:

Input: $\frac{n}{\alpha_1}$ items of size 1 and $(\frac{1}{\alpha_k} - \frac{1}{\alpha_{k-1}})n$ items of size α_{k-1} for $k = 2, \ldots, p$.

Output: A packing of the items into the minimum number of bins.

Observe that this bin packing instance is *independent* of $G = (V, E) \in \mathcal{D}_r^n(\alpha, 1)$ and depends only on n and the values of α_i selected.

Given any solution to this bin packing instance, say with k opened bins, we can easily obtain a coloring of all the edges in the union of the matchings in $\mathcal{M}$ using just k colors. Indeed, we can simply color an edge belonging to a matching by a color representing the bin in which the corresponding item is packed. In constructing the bin packing instance, we have discarded at most

$$\sum_{k=1}^{p} \max_{v \in V} a_k(v) \le 3 \sum_{k=1}^{p} (p-k+1) = \frac{3}{2} p(p+1)$$

matchings, and they can be colored with a new color for each of them. In summary, the number of colors we need is at most the optimal number of bins of our bin packing instance plus $\frac{3}{2}p(p+1)$. An interesting feature of the results on the previous section is that they do not assume any conditions on p. We will see later that the optimal value for p is $\Theta(n^{1/3})$, which implies that the number of additional colors we need to accommodate the matchings not in $\mathcal{M}$ is $\frac{3}{2}p(p+1) = O(n^{2/3}) = o(n)$ and hence negligible.

As an example of the associated bin packing instance, consider the case with $p = 3$ and $\alpha_1 = \frac{1}{2}$, $\alpha_2 = \frac{1}{3}$, and $\alpha_3 = \alpha = \frac{1}{4}$. The bin packing instance then consists of $2n$ items of size 1, n items of size $\frac{1}{2}$, and n items of size $\frac{1}{3}$, and these items can be packed into $2n + \frac{n}{2} + \frac{n}{3} = \frac{17}{6}n$ bins (plus $O(1)$ bins for fractionally opened bins). The argument above regarding discarded items shows that we need $O(p^2) = O(1)$ additional bins. Using Lemma 2, we then obtain that $M(n, r) \leq \frac{17}{6}n + O(1)$. This derivation is essentially identical to the result of Du et al. [12], and the approach taken here can be viewed as an extension of it.

Our goal now is to focus on our general bin packing instance and analyze the number of bins it requires. Since all items in the bin packing instance have size at least $\alpha = \alpha_p$, it is clear that items whose size is more that $1-\alpha$ are forced to use a full bin in any feasible packing. Hence, without loss of generality, we can let $\alpha_1 = 1 - \alpha$. With this, an optimal packing always needs $n/(1 - \alpha)$ bins to pack items of size 1 plus a certain number of bins to pack the remaining items (of size $\alpha_1, \ldots, \alpha_p$).

3.4. A lower bound. A trivial lower bound on the number of unit bins required to pack our discrete instance is $n/(1-\alpha)$ bins (for the items of size greater than $1-\alpha$) plus the total size of the remaining items:

$$\frac{n}{1-\alpha} + \sum_{k=2}^{p} \alpha_{k-1} \left(\frac{1}{\alpha_k} - \frac{1}{\alpha_{k-1}} \right) n.$$

This can be lower bounded in the following way. Let $g : [\alpha, 1 - \alpha] \to \mathbb{R}$ be defined by $g(x) = 1/x^2$. As $n \int_{\alpha_k}^{\alpha_{k-1}} g(x)dx = (\frac{1}{\alpha_k} - \frac{1}{\alpha_{k-1}})n$ is the number of items of size α_{k-1} and $\alpha_{k-1} \geq x$ for any $x \in [\alpha_k, \alpha_{k-1}]$, we have that

$$\sum_{k=2}^{p} \alpha_{k-1} \left(\frac{1}{\alpha_k} - \frac{1}{\alpha_{k-1}} \right) n \geq n \int_{\alpha}^{1-\alpha} x g(x) dx$$

$$= \int_{\alpha}^{1-\alpha} \left(\frac{n}{x} \right) dx = n \ln \frac{1-\alpha}{\alpha}.$$

Therefore, from Lemma 2, we derive that our analysis cannot give an upper bound on $M_{[\alpha,1]}(n, r)$ better than

$$\min_{\alpha \in (0,1]} \max \left\{ \frac{2n}{1-\alpha}, \frac{n}{1-\alpha} + n \ln \frac{1-\alpha}{\alpha} \right\} = M \cdot n,$$

with $2.5569 \leq M \leq 2.5570$. The term $\frac{2n}{1-\alpha}$ comes from Lemma 2, while the other term is the bound just obtained. The value of α for which the minimum is attained is $\alpha \approx 0.2178117$. From now on, we fix α to be the argmin of the above expression. In what follows, we show that this lower bound is actually achievable by relating the number of bins required by our bin packing instance to a continuous bin packing instance and analyzing it. For this purpose, we assume that the α_i's in the definition of our bin packing instance are equally spaced in $[\alpha, 1-\alpha]$, i.e., $\alpha_{k-1} - \alpha_k = \Delta = \frac{1-2\alpha}{p-1}$ with $\alpha_1 = 1 - \alpha$ and $\alpha_p = \alpha$.

3.5. The continuous packing problem. We round our bin packing instance to a continuous bin packing problem for which packing strategies with sublinear waste exist. We first define what we mean by a continuous bin packing instance. Consider a finite positive measure μ with density g defined over $[a, b]$ (with $0 \leq a \leq b \leq 1$) and, for any integer q, consider a uniform discretization $a = x_1 < \cdots < x_q = b$ of the interval $[a, b]$. Let Q_n^q be the optimal number of bins needed to pack the bin packing instance in which, for all $1 \leq i < q$, there are $\lceil n\mu([x_i, x_{i+1}))\rceil$ items of size x_{i+1}. The *value* of our bin packing instance is then defined as $\lim_{q\to\infty} \lim_{n\to\infty} \frac{Q_n^q}{n}$. By simply considering the total size of the items, we see that the value of a continuous instance is never smaller than

$$\int_a^b x d\mu(x) = \int_a^b xg(x)dx.$$

We say that μ admits a perfect packing if we have equality

$$\lim_{q\to\infty} \lim_{n\to\infty} \frac{Q_n^q}{n} = \int_a^b x d\mu(x) = \int_a^b xg(x)dx.$$

The lower bound in the previous section suggests that we consider the continuous bin packing instance with the continuous density $g(x) = \frac{1}{x^2}$ over $x \in [\alpha, 1-\alpha]$. In the next section, we show that a result of Rhee and Talagrand [26] can be applied to prove that g actually admits a perfect packing. What we show now is that the difference between the number of bins we need in our discrete instance and the value of this continuous instance times n is $O(\frac{n}{p})$ and hence sublinear whenever p grows with n. For this purpose, we show that we can discard $O(n/p)$ items in our discrete instance and obtain an instance which is dominated by discrete realizations of our continuous instance. Indeed, as $\int_{\alpha_{k-1}}^{\alpha_{k-2}} g(x)dx = \frac{1}{\alpha_{k-1}} - \frac{1}{\alpha_{k-2}}$, the continuous instance would dominate the discrete instance if we had only $(\frac{1}{\alpha_{k-1}} - \frac{1}{\alpha_{k-2}})n$ items of size α_{k-1}. We therefore need to discard a number of items of size α_{k-1} equal to

$$\begin{aligned}
&\left(\frac{1}{\alpha_k} - \frac{1}{\alpha_{k-1}}\right)n - \left(\frac{1}{\alpha_{k-1}} - \frac{1}{\alpha_{k-2}}\right)n \\
&= \frac{2\Delta^2}{\alpha_k\alpha_{k-1}\alpha_{k-2}} n \leq \frac{2\Delta^2}{\alpha^3} n.
\end{aligned}$$

Over all values of k, this amounts to discarding $p\frac{2\Delta^2}{\alpha^3}n = \Theta(\frac{n}{p})$ items, and they can each be packed in a separate bin.

As announced, we show in the next section that g admits a perfect packing. This implies that the total number of colors needed to color any graph $G \in \mathcal{D}_r^n(\alpha, 1)$ is at most $M \cdot n + O(p^2) + O(n/p)$, which is optimized choosing $p = \Theta(n^{1/3})$. For the optimal choice of α, which is approximately 0.2178117, the previous quantity becomes

$$M \cdot n + O(n^{\frac{2}{3}}) < 2.557 \cdot n + O(n^{\frac{2}{3}}),$$

concluding the proof of Theorem 1.

3.6. Perfect packing. Consider the positive measure μ defined over the interval $[\alpha, 1-\alpha]$ with density $g(x) = 1/x^2$ for the optimal parameter α just obtained. To show that a perfect packing exists, we decompose g as the sum of three other positive functions, f_1, f_2 and f_3, all of which allow perfect packing. Furthermore, all bins used for the items corresponding to f_i will contain exactly $i+1$ items. With this, μ is a mixture of the corresponding measures μ_1, μ_2 and μ_3. The decomposition is depicted in Figure 1.

Consider the following functions:

1. $f_1(x) = \begin{cases} g(1-x) & \text{if } x \in [\alpha, 1/2), \\ g(x) & \text{if } x \in [1/2, 1-\alpha], \\ 0 & \text{otherwise,} \end{cases}$
2. $f_2(x) = \begin{cases} g(x) - f_1(x) - c & \text{if } x \in [1/4, \beta), \\ d & \text{if } x \in [\beta, \delta), \\ g(x) - f_1(x) & \text{if } x \in [\delta, 1/2), \\ 0 & \text{otherwise,} \end{cases}$
3. $f_3(x) = \begin{cases} g(x) - f_1(x) & \text{if } x \in [\alpha, 1/4), \\ c & \text{if } x \in [1/4, \beta), \\ g(x) - f_1(x) - d & \text{if } x \in [\beta, \delta), \\ 0 & \text{otherwise.} \end{cases}$

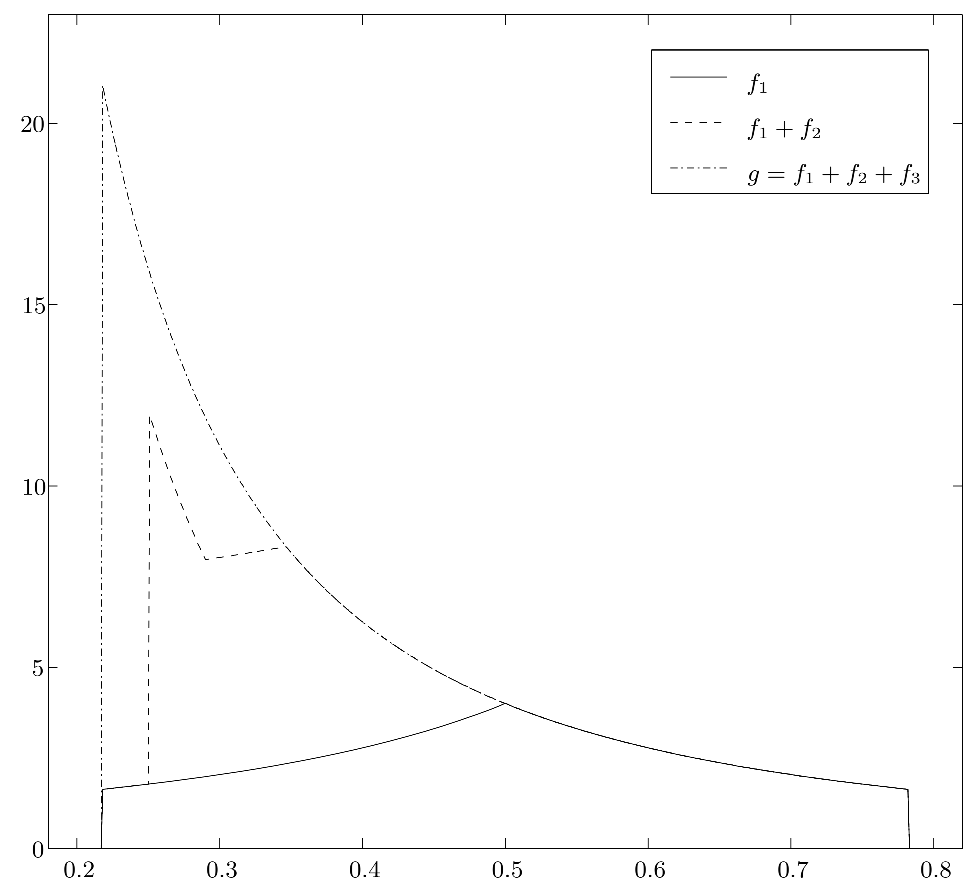

FIG. 1. *Decomposition of g into f_1, f_2, and f_3.*

Here $c = g(\beta) - f_1(\beta) - d$ and $d = g(\delta) - f_1(\delta)$ (so that f_2 is continuous). Clearly, for all $x \in [\alpha, 1-\alpha]$, $g(x) = f_1(x) + f_2(x) + f_3(x)$. The values of β and δ are uniquely determined by imposing that the average value of f_2 is $1/3$ and that of f_3 is $1/4$.

Namely, if $\beta \approx 0.2900708$ and $\delta \approx 0.3465256$, then

$$\frac{\int_\alpha^{1-\alpha} x f_1(x)dx}{\int_\alpha^{1-\alpha} f_1(x)dx} = \frac{1}{2},$$

$$\frac{\int_{1/4}^{1/2} x f_2(x)dx}{\int_{1/4}^{1/2} f_2(x)dx} = \frac{1}{3},$$

$$\frac{\int_\alpha^{\delta} x f_3(x)dx}{\int_\alpha^{\delta} f_3(x)dx} = \frac{1}{4}.$$

To prove that all f_1, f_2, and f_3 allow perfect packing, we use a perfect packing result proved by Karmarkar [17] and by Loulou [22] and a powerful theorem by Rhee and Talagrand [26]. The former result says that measures that are symmetric around $1/2^k$ for some integer k allow perfect packing. The latter can be stated as follows.

THEOREM 6 (Rhee and Talagrand [26]). *Consider a decreasing measure μ defined over $[a, b]$ (with $0 \le a \le b \le 1$) and an integer $p \ge 3$ such that $1/p \in [a, b]$. Then μ allows perfect packing if the following are satisfied:*

(i) $(p-1)a + b \le 1$.

(ii) $\int_a^b x d\mu(x) = \frac{1}{p}\int_a^b d\mu(x)$.

In what follows, we briefly outline this result. Let $0 \le a \le b \le c \le 1$ be such that $(p-1)a + c \le 1$ and $a + b < 2/p < a + c$. The *L-shaped function*, denoted by $L(a, b, c)$, is the unique (up to a multiplicative constant) nondecreasing real function defined over $[a, c]$, which is constant on $[a, b]$ and constant on $(b, c]$, and whose average value is $1/p$, i.e.,

$$\frac{\int_a^c xL(a,b,c)(x)dx}{\int_a^c L(a,b,c)(x)dx} = \frac{1}{p}.$$

In order to prove Theorem 6, Rhee and Talagrand first showed how to decompose a density satisfying the assumptions of the theorem as the limit of sum of L-shaped functions with the above properties. Then the central part of their work was to show that all such L-shaped functions do allow perfect packing. Unfortunately, they did not find a simple perfect packing strategy, and so they overcame the problem using a perfect packing characterization by Rhee [25], together with a complicated (and implicit) "exhaustion method," that decomposes an L-shaped function into possibly uncountably many perfectly packable functions.

Let us mention, however, that although the previous result was proved in a probabilistic setting (namely, under the following definition: μ allows perfect packing if and only if the expected number of bins needed to pack n independent and identically distributed random variables drawn according to μ divided by n approaches the expected size of an item), the proof also applies to our setting here.

LEMMA 7. *The measure μ with density function $g : [\alpha, 1-\alpha] \to \mathbb{R}$ with $g(x) = 1/x^2$ allows perfect packing.*

Proof. As $g = f_1 + f_2 + f_3$, we need only show that each f_i, $i = 1, 2, 3$, allows perfect packing. The result follows immediately for f_1. Indeed, f_1 is symmetric around $1/2$. It remains to prove that both f_2 and f_3 satisfy the conditions of the previous theorem.

(1) The density f_2 is clearly decreasing in $[1/4, 1/2]$. Moreover,

$$\int_{1/4}^{1/2} x f_2(x) dx = \frac{1}{3} \int_{1/4}^{1/2} f_2(x) dx.$$

Finally, $(3-1)\frac{1}{4} + \frac{1}{2} = 1$. Thus all conditions are satisfied.

(2) Again, the density f_3 is decreasing in $[\alpha, \delta]$. In this case,

$$\int_{\alpha}^{\delta} x f_3(x) dx = \frac{1}{4} \int_{\alpha}^{\delta} f_3(x) dx,$$

and $(4-1)\alpha + \delta < 1$ (indeed, $(4-1)\alpha + \delta \approx 0.9999607$). ∎

3.7. Improved analysis for the rearrangeability of 3-stage Clos networks. In this section, we briefly discuss how a slight improvement of the $2.557n$ bound can be achieved when considering graphs belonging to $\mathcal{B}_r^n$. Specifically, we establish that $m(n,r) \leq 2.5480n + o(n)$. The analysis is essentially the same as the one for the bound on $M(n,r)$; therefore, we give only the main differences.

Let $G = (V,E) \in \mathcal{B}_r^n$. Since the weights satisfy condition (1), we can strengthen the main inequality used in Lemma 5 to be $\sum_{i=k}^{p} \deg_{D_i}(v) \leq 4n$ whenever $\alpha_p > 1/5$ (this is a strengthening only for $\alpha_k \leq 1/4$). This inequality, combined with the ideas in Lemma 5, can be used to prove the following result.

LEMMA 8. *Let $G = (V,E) \in \mathcal{B}_r^n(\alpha, 1)$ and consider a sequence $\alpha_0 = 1 > \alpha_1 > \cdots > \alpha_l = 1/4 > \cdots > \alpha_p = \alpha > 1/5$. Then there exist sets $F_1, \ldots, F_p$ partitioning E such that the following hold:*

(i) $\max_{e \in F_k} w(e) \leq \alpha_{k-1}$.

(ii) *For all vertices $v \in V$,*

$$\deg_{F_1}(v) \leq \frac{n}{\alpha_1} + a_1(v),$$

$$\deg_{F_k}(v) \leq \left(\frac{1}{\alpha_k} - \frac{1}{\alpha_{k-1}}\right) n + a_k(v), \qquad 2 \leq k \leq l,$$

$$\deg_{F_k}(v) \leq 16\,(\alpha_{k-1} - \alpha_k)\, n + a_k(v), \qquad l+1 \leq k \leq p,$$

where $a_k(v) \leq 3(p-k+1)$.

By mimicking the analysis in section 3.3, the problem now translates into packing the function $g : [\alpha, 1-\alpha] \to \mathbb{R}$ such that $g(x) = 16$ if $x \in [\alpha, 1/4]$, and $g(x) = 1/x^2$ otherwise. The value of α now has to be taken a bit smaller than it used to be: $\alpha \approx 0.2151$ is the optimal choice. For that value of α, a decomposition of g very similar to that in section 3.6 can be found. Applying again the result in [26], such decomposition amounts to concluding that g allows perfect packing. The total number of colors needed is therefore

$$n \int_{\alpha}^{1-\alpha} x g(x) dx + \frac{n}{1-\alpha} + o(n) = \frac{2n}{1-\alpha} + o(n)$$

$$< 2.5480 \cdot n + o(n),$$

where the inequality comes from the choice of α.

4. Wide-sense nonblocking Clos networks. In what follows, we consider the online coloring formulation of the problem, assuming that edge weights satisfy condition (1). We start by describing two variants of the FF heuristic in the context of wide-sense nonblocking 3-stage Clos networks. Let $G = (V, E)$ be the bipartite graph with bipartition $V = A \cup B$ such that edge weights satisfy (1). Let $\{1, \ldots, M\}$ be the colors with which we attempt to find a valid coloring of G. Assume all edges in $\tilde{E} \subset E$ have been revealed and colored so far, and a new edge $e = (u, v) \notin \tilde{E}$ is revealed. The *first-fit-min (FF-Min)* heuristic assigns $c(e) = j$ (i.e., colors e with color j), where j is the smallest color for which adding e does not violate the valid coloring condition. In other words,

$$c(e) = j = \min \Bigg\{ 1 \leq i \leq M : w(e) + \sum_{f:\, f \in \delta(v) \cap \tilde{E},\, c(f) = i} w(f) \leq 1, \quad w(e) + \sum_{f:\, f \in \delta(u) \cap \tilde{E},\, c(f) = i} w(f) \leq 1 \Bigg\}.$$

In the first-fit-max (FF-Max) heuristic, the above minimization is replaced by a maximization.

The main issue now is to determine the smallest M such that FF-Min can always assign a color to a given edge. In order to establish our main result, we need a preliminary definition regarding the *blocking number* of $\mathcal{B}_r^n([0,1])$ under algorithm FF-Min. For an interval $I \subset [0, 1]$, we define the blocking number of $\mathcal{B}_r^n(I)$ under algorithm A (or simply the blocking number of algorithm A) as the maximum over all vertices $v \in V$ and over any graph $G = (V, E) \in \mathcal{B}_r^n(I)$ of the number of colors whose total weight adjacent to v is more than $1/2$. We denote it by $B_I(A)$:

$$B_I(A) = \max_{G=(V,E) \in \mathcal{B}_r^n(I)} \max_{v \in V} \left\{ \text{number of colors } i : \sum_{f:\, f \in \delta(v)\; c(f) = i} w(f) > \frac{1}{2} \text{ under A} \right\}.$$

By definition of $\mathcal{B}_r^n(I)$, we have that, for any algorithm A and for any $I \subseteq [0, 1]$, $B_I(A) \leq 2n - 1$. This bound, for interval $[0, 1]$, is exactly what we need to establish our main result.

We remark that our results do not hold for the more restrictive definition of wide-sense nonblocking in which the algorithm has to be able to route new connection requests *even if* previous connections terminate. In terms of the graph coloring problem, the more restrictive condition allows not only additions but also deletions of edges over time.

4.1. Wide-sense nonblocking for general connection requests. We now give the main result of this section, namely the bound on $m_W(n, r)$ in the general case.

THEOREM 9. *The number of colors needed to color any graph in $\mathcal{B}_r^n$ using algorithm FF-Min is at most $5n$, i.e.,*

$$m_W(n, r) \leq 5n.$$

Proof. Consider algorithm FF-Min, with $M = 5n$, applied to $G \in \mathcal{B}_r^n$. Let A, B be the bipartition of V, i.e., $V = A \cup B$. Let us say that an edge e is *large* if $w(e) > 1/2$ and *small* if $w(e) \leq 1/2$.

Consider iteration k of the algorithm and assume $e_k = e = (u, v)$ for some $u \in A$, $v \in B$. To see that the algorithm indeed works, we prove that edge e can be colored with some of $5n$ available colors. For this we consider two cases:

- Edge $e = e_k$ is small ($w(e) \leq 1/2$). Since FF-Min is a greedy-type heuristic, by Lemma 2, $\frac{2n}{1-1/2} = 4n$ colors are enough; and thus a color smaller than or equal to $4n$ is assigned to e.
- Edge $e = e_k$ is large ($w(e) > 1/2$). In this case, assume e cannot be colored. Let S^k_{uv} be the set of colors $1 \leq i \leq 5n$ such that there exists an edge e_t with $t < k$ satisfying that
  - e_t is colored with i,
  - $w(e_t) > 1/2$, and
  - e_t is adjacent to either u or v.

  The bound $B_I(A) \leq 2n-1$ for any algorithm implies that $|S^k_{uv}| \leq 2(2n-1) = 4n-2$. Now consider s, the smallest color in S^k_{uv}, such that $i > s$ implies that $i \in S^k_{uv}$. Since $|S^k_{uv}| \leq 4n-2$, we have that $s \geq n+3$. By definition, there is a small edge, say $f = (u, t)$, colored with $s-1$. The fact that f is small amounts to concluding that for all $i < s-1$, either

$$\sum_{e:\, e\in\delta(u), c(e)=i} w(e) > \frac{1}{2} \quad \text{or} \quad \sum_{e:\, e\in\delta(t), c(e)=i} w(e) > \frac{1}{2}.$$

  The latter can happen for at most $2n-1$ colors (see the bound on the blocking number above), and therefore the former holds for at least $s-2-(2n-1) = s-2n-1$ colors. (As the former can happen only for at most $2n-1$ colors, this actually also implies $s \leq 4n$.) On the other hand, the number of large edges adjacent to u or v which are colored with $j \geq s$ is at least $5n-s+1$. Since, by condition (1), at most n of these can be adjacent to v, at least $5n-s-n+1 = 4n-s+1$ are adjacent to u.
  Overall we have that

$$\sum_{e:\, e\in\delta(u)} w(e) > \frac{s-2n-1}{2} + \frac{4n-s+1}{2} = n,$$

  which contradicts (1). $\square$

4.2. Improved bounds for the case of small connection requests. We now turn to the case in which all connection requests have weights in $[0, 1/2]$. In terms of our graph coloring problem, this means considering graphs in $\mathcal{B}^n_r([0,1/2])$. Gao and Hwang [13] have proved that $m_{W[0,1/2]}(n, r) \leq 3.75n$. Let us now see how an improvement of this result can be obtained.

LEMMA 10. *The number of colors needed to color any graph in $\mathcal{B}^n_r([0,1/2])$ using algorithm FF-Min is at most $3.601n+3$, i.e.,*

$$m_{W[0,1/2]}(n, r) \leq 3.601n + 3.$$

Proof. Observe first that from Lemma 2, if e is an edge with weight $w(e)$, FF-Min actually assigns to it a color $c(e)$ satisfying $c(e) \leq \frac{2n}{1-w(e)} + 1$, or

$$w(e) \geq 1 - \frac{2n}{c(e)-1}. \tag{4}$$

This immediately implies that edges with weight below $\frac{1}{4}$ are assigned to the first $\lceil \frac{8n}{3} \rceil$ colors.

Let M be the number of colors needed by FF-Min and let $e = (u, v)$ be an edge that could not be assigned to any of the first $M-1$ colors. Consider a color $j \le M-1$, since e was not assigned to j (and $w(e) \le 1/2$); then $\sum_{f\colon f\in\delta(u),c(f)=j} w(f) > 1/2$ or $\sum_{f\colon f\in\delta(v),c(f)=j} w(f) > 1/2$. Assume, without loss of generality, that the latter holds. Since all weights are at most $1/2$, at least two edges in $\delta(v)$ are colored j, and thus

$$\sum_{f\colon f\in\delta(v),c(f)=j} w(f) \ge 2\left(1 - \frac{2n}{j-1}\right).$$

We can now compute $\sum_{g\in\delta(v)} w(g) + \sum_{g\in\delta(u)} w(g)$ using the previous equation, the fact that for a color $j \le \lceil \frac{8n}{3} \rceil$ either $\sum_{f\colon f\in\delta(u),c(f)=j} w(f) > 1/2$ or $\sum_{f\colon f\in\delta(v),c(f)=j} w(f) > 1/2$, and (4):

$$\begin{aligned}
\sum_{g\in\delta(v)} w(g) + \sum_{g\in\delta(u)} w(g) &> 2w(e) + \left\lceil \frac{8n}{3} \right\rceil \frac{1}{2} + \sum_{j=\lceil \frac{8n}{3} \rceil+1}^{M-1} 2\left(1 - \frac{2n}{j-1}\right) \\
&\ge \left\lceil \frac{8n}{3} \right\rceil \frac{1}{2} + \sum_{j=\lceil \frac{8n}{3} \rceil+1}^{M} 2\left(1 - \frac{2n}{j-1}\right) \\
&\ge 2M - \frac{3}{2}\left\lceil \frac{8n}{3} \right\rceil - 4n \int_{\lceil 8n/3 \rceil -1}^{M-1} \frac{1}{x} dx \\
&\ge 2M - 4n - \frac{3}{2} - 4n \int_{8n/3-1}^{M-1} \frac{1}{x} dx \\
&= 2M - 4n - \frac{3}{2} + 4n \ln\left(\frac{8n-3}{3M-3}\right) \\
&= 2M - 4n - \frac{3}{2} + 4n \ln\left(\frac{8n}{3M-9}\right) \\
&\quad + 4n \ln\left(\frac{(3M-9)(8n-3)}{(3M-3)8n}\right).
\end{aligned}$$

However, for $M \ge 3.601n + 3$, the above quantity surpasses $2n$, leading to a contradiction. Indeed, for this choice of M, the last term is greater than -4, and so the previous quantity is greater than $2 \cdot 3.601n - 4n + 4n \ln(8/10.803) > 2n$. ∎

4.3. The 2-rate environment. We now prove the bound on $m_W(n, r)$ when connection requests can take only two values, which are known beforehand. As mentioned before, this result almost closes the gap with the best known lower bound in this environment. Indeed, the result in this section, together with results in [13, 31], implies that

$$3n - 2 \le m_W^2(n, r) \le 3n.$$

In what follows, we denote by b and B the two rates (or edge weights) and assume that $0 < b < B \le 1$. Gao and Hwang [13] already proved the bound in the case $B \le 1/2$.

Lemma 11 (Gao and Hwang [13]). *If $0 < b < B \leq 1/2$ are the two rates, then $m^2_W(n, r) \leq 3n$.*

We complete Gao and Hwang's result by proving a slightly better bound when $B > 1/2$. Of course, we may assume $b \leq 1/2$, for otherwise condition (1) allows us to reason that every vertex has degree at most n, and thus an even stronger bound of $2n$ holds for any online algorithm. Let k be the largest integer such that $B + kb \leq 1$ and ℓ be the largest integer such that $\ell b \leq 1$. Let us associate a *height* of 1 with every edge of weight b and a height of $(\ell - k)$ to every edge of weight B; and denote by $h(e)$ the height of an edge. As at most one item of size B can fit into a bin, (1) implies that the height of edges in any bin is at most l, and thus $\sum_{e \in \delta(v)} h(e) \leq n\ell$.

The algorithm we need to consider is the following.

Algorithm FF-Min-Max
(1) Assume the edges are revealed in the order $\{e_1, \dots, e_m\}$. (2) For $p = 1$ to m do: (a) If $w(e_p) = B$, assign a color $1 \leq i \leq 3n-1$ to e_p using FF-Max. (b) If $w(e_p) = b$: ∗ Assign any color $1 \leq i \leq 3n - 1$ to e_p such that at most $k - 1$ small edges adjacent to u have been colored i, and at most $k - 1$ small edges adjacent to v have been colored i, if one such color exists. ∗ Otherwise, assign a color $1 \leq i \leq 3n - 1$ to e_p using FF-Min.

Lemma 12. *The number of colors needed to color any graph in $\mathcal{B}^n_r(\{b, B\})$, with $0 < b \leq 1/2 < B \leq 1$, using Algorithm FF-Min-Max is at most $3n - 1$.*

Proof. Consider the graph $G = (V, E) \in \mathcal{B}^n_r(\{b, B\})$ and assume the set of colors is $\{1, \dots, 3n - 1\}$. For the purpose of this proof, let us say that an edge e is *large* if $w(e) = B$ and *small* if $w(e) = b$.

Consider step (1) of the algorithm and let $e_i = (u, v) \in E$ be the edge currently considered.

Let us first see that in step (2)(b) of the algorithm, FF-Min does not attempt to color e_i using a color larger than $2n - 1$ (assuming thus that no color could be found such that either u or v has at most $k - 1$ edges adjacent to it of that color). Indeed, it is enough to observe that FF-Min is a greedy-type algorithm. With this in mind, assume that FF-Min could not color e_i (which is a small edge) with a color $j \leq 2n-1$. Then, for any color $j \leq 2n - 1$, one of the following is satisfied:

- There are ℓ small edges $f_1 \dots, f_\ell$ such that $c(f_r) = j$ for all $r = 1, \dots, \ell$, and either $\{f_1, \dots, f_\ell\} \subseteq \delta(u)$ or $\{f_1, \dots, f_\ell\} \subseteq \delta(u)$.
- There are k small edges $f_1, \dots, f_k$ and one large edge f such that $c(f_r) = c(f) = j$ for all $r = 1 \dots, k$, and either $\{f_1, \dots, f_k, f\} \subseteq \delta(u)$ or $\{f_1, \dots, f_k, f\} \subseteq \delta(u)$.

We can therefore assume that there is a set of n colors $c_1, \dots, c_n$ such that for each c_r, ℓ small edges or k small and one large edge are colored with c_r and are adjacent to u (otherwise, the property is true with v). Since also e_i is adjacent to u, we obtain that the total height of edges adjacent to u is at least $n\ell + 1 > n\ell$, a contradiction with condition (1).

We conclude that the algorithm will always find a feasible color for a small edge.

Suppose now that the algorithm cannot assign any color to $e_i = (u, v)$ because e_i is a large edge ($w(e_i) = B$). Then Algorithm FF-Min-Max attempted to color e_i in step (2)(a) using FF-Max, but no color was found. We see in what follows that this is impossible.

Let j be the largest color having at least $k + 1$ small edges adjacent to either u or v. From the analysis above, no color i with $i \geq 2n$ has more than k small edges adjacent to it, and thus $j < 2n$. In addition, for any color i with $i > j$, there must already be a large edge colored i adjacent to either u or v, and as this can happen to at most $2n - 2$ edges, we have $j > n$. Thus, $n < j < 2n$. Without loss of generality, assume that at least $k + 1$ small edges of color j are adjacent to u and let $f = (u, t)$ be one of those edges that was colored using FF-min (there is at least one such f, since there are at least $k+1$ edges in total). Since f was colored by FF-min, for every color $1 \leq r \leq 3n - 1$, at least k small edges of color r are adjacent to u or t. Also, by definition of FF-min, for every color $1 \leq r < j$, there is a set of edges colored with r, of total height ℓ, adjacent to u or t (such a set consists of either ℓ small edges or k small and one large edge). Overall we have that

$$\sum_{e \in \delta(u) \cup \delta(t) : w(e) = b} h(e) + \sum_{e \in \delta(u) \cup \delta(t) : w(e) = B \text{ and } c(e) < j} h(e) \geq (3n-1)k + (j-1)(\ell - k) + 1,$$

where the final $+1$ comes from the fact that $k+1$ small edges adjacent to u are colored with j. Thus, since $n\ell$ is the maximum total height of edges adjacent to t,

$$\sum_{e \in \delta(u) : w(e) = b} h(e) + \sum_{e \in \delta(u) : w(e) = B \text{ and } c(e) < j} h(e) \geq (3n-1)k + (j-1)(\ell - k) + 1 - n\ell. \tag{5}$$

On the other hand, for every color r with $j < r \leq 3n - 1$, there is a large edge of color r adjacent to u or v. Since at most $n - 1$ (already colored) large edges can be adjacent to v, at least $3n - 1 - j - (n - 1) = 2n - j$ large edges of color larger than j are adjacent to u. Thus,

$$\sum_{e \in \delta(u) :\, w(e) = B \text{ and } c(e) > j} h(e) \geq (2n - j)(\ell - k).$$

Combining this with (5) and the fact that e_i is large, we conclude that the total height of edges adjacent to u is at least

$$[\ell - k] + [(3n - 1)k + (j - 1)(\ell - k) + 1 - nl] + [(\ell - k)(2n - j)],$$

where the first term corresponds to the height of e_i, the second to inequality (5), and the third to large edges adjacent to u of color $r > j$. The above quantity equals $(3n - 1)k - n\ell + 2n(\ell - k) + 1 = (n - 1)k + n\ell + 1 > n\ell$, which is impossible. ∎

5. A simple algorithm for multirate rearrangeability. In this section, we reconsider the offline setting and present a simple algorithm that is multirate rearrangeably nonblocking and that uses no more than $8n/3$ colors. In comparison, the algorithm of section 3 uses $2.548n + o(n)$ colors but is more complex. The algorithm we consider is the following.

Algorithm FF-Min-Decreasing

(1) Sort the edges according to their weight such that $w(e_1) \geq w(e_2) \geq \cdots \geq w(e_m)$. Let $k = 1$.
(2) While $k \leq m$ do:
   (a) Assign a color $1 \leq i \leq \lceil 8n/3 \rceil$ to e_k using FF-Min.
   (b) $k = k + 1$.

As is the case for FF decreasing for bin packing, Algorithm FF-Min-Decreasing can be implemented in time $O(n \log n)$. As it involves sorting, it is applicable only to the offline setting. We have the following result.

THEOREM 13. *The number of colors needed to color any graph in $\mathcal{B}_r^n$ using Algorithm FF-Min-Decreasing is at most $\lceil 8n/3 \rceil$.*

Proof. Consider the graph $G = (V, E) \in \mathcal{B}_r^n$ and let k be the smallest index such that $w(e_k) \leq \frac{1}{4}$. As FF-Min-Decreasing is a greedy algorithm, we know from Lemma 2 that our algorithm will always be able to color $\{e_k, \ldots, e_m\}$. Thus we can assume that $w(e_i) > \frac{1}{4}$ for all $i = 1 \ldots, m$; i.e., we can assume $G = (V, E) \in \mathcal{B}_r^n(]1/4, 1])$.

Let $e = e_\ell = (u, v)$ be the first edge that could not be colored by FF-Min-Decreasing. We distinguish two cases:

(i) $w(e) = \alpha > 1/3$. In this case, we will prove that FF-Min-Decreasing colors e with a color no larger than $2n$. We define the function $g_\alpha : [0, 1] \to [0, 1]$ as

$$g_\alpha(x) = \begin{cases} 1 & \text{if } 1 - \alpha < x, \\ 1/2 & \text{if } \alpha \leq x \leq 1 - \alpha, \\ 0 & \text{if } x < \alpha \end{cases}$$

and consider the modified edge weights $w'(e_i) = g_\alpha(w(e_i))$. We know that $G_\ell = (V, \{e_1, \ldots, e_\ell\})$ together with w satisfies condition (1), and from the sorting step $w(e_i) \geq \alpha$ for all $i = 1, \ldots, \ell$. Thus, G, together with w', also satisfies condition (1). Now, as e could not be colored using the first $2n$ colors, for all $1 \leq i \leq 2n$ either

$$\sum_{e: e \in \delta(u), c(e) = i} w(e) > 1 - \alpha \quad \text{or} \quad \sum_{e: e \in \delta(v), c(e) = i} w(e) > 1 - \alpha.$$

We can then assume that for a set $B \subset \{1, \ldots, 2n\}$ with $|B| \geq n$ the first inequality holds. For $i \in B$, the previous condition implies that the edges in $\delta(u)$ colored with i are either one edge f with $w(f) > 1 - \alpha$ or two edges f, g with $w(f) \geq \alpha$ and $w(g) \geq \alpha$. In both cases,

$$\sum_{e: e \in \delta(u), c(e) = i} w'(e) = 1.$$

It follows that

$$\sum_{e: e \in \delta(u)} w'(e) > \sum_{e: e \in \delta(u), c(e) \in B} w'(e) \geq n,$$

a contradiction with condition (1).

(ii) $1/3 \geq w(e) = \alpha > \frac{1}{4}$. We define the function $f_\alpha : [0,1] \to [0,1]$ as

$$f_\alpha(x) = \begin{cases} 1 & \text{if } 1-\alpha < x, \\ 1/2 & \text{if } \frac{1-\alpha}{2} < x \leq 1-\alpha, \\ 1/4 & \text{if } \alpha \leq x \leq \frac{1-\alpha}{2}, \\ 0 & \text{if } x < \alpha \end{cases}$$

and consider the modified edge weights $w'(e_i) = f_\alpha(w(e_i))$. As in the previous case, $G_\ell = (V, \{e_1, \ldots, e_\ell\})$, together with w, satisfies condition (1). Thus, from the sorting step $w(e_i) \geq \alpha$ for all $i = 1, \ldots, \ell$, and it is easy to check that G_ℓ, together with w', also satisfies condition (1).
Additionally, as e could not be colored using the $\lceil \frac{8n}{3} \rceil$ colors, for all $1 \leq i \leq \lceil \frac{8n}{3} \rceil$ either

$$\sum_{e: e \in \delta(u), c(e)=i} w(e) > 1 - \alpha \quad \text{or} \quad \sum_{e: e \in \delta(v), c(e)=i} w(e) > 1 - \alpha.$$

We can then assume that for a set $B \subset \{1, \ldots, \lceil 8n/3 \rceil\}$ with $|B| \geq 4n/3$ the first inequality holds. For $i \in B$, the previous condition implies that the edges in $\delta(u)$ colored with i are either one edge f with $w(f) > 1 - \alpha$; two edges f, g with at least one of them (the largest), say f, satisfying $w(f) > \frac{1-\alpha}{2}$; or three edges f, g, h (all with weights greater than α). In any case,

$$\sum_{e: e \in \delta(u), c(e)=i} w'(e) \geq \frac{3}{4}.$$

It follows that

$$\sum_{e: e \in \delta(u)} w'(e) > \sum_{e: e \in \delta(u), c(e) \in B} w'(e) \geq |B| \cdot \frac{3}{4} \geq n. \qquad \square$$

REFERENCES

[1] A. Bar-Noy, R. Motwani, and J. Naor, *The greedy algorithm is optimal for online edge coloring*, Inform. Process. Lett., 44 (1992), pp. 251–253.

[2] L. A. Bassalygo and M. S. Pinsker, *Complexity of optimum nonblocking switching networks without reconnections*, Problems Inform. Transmission, 9 (1974), pp. 64-66.

[3] J. Beetem, M. Denneau, and D. Weingarten, *The GF11 supercomputer*, in Proceedings of the 12th ACM Annual International Symposium on Computer Architecture, 1985, pp. 108–115.

[4] B. Beizer, *The analysis and synthesis of signal switching networks*, in Proceedings of the Symposium on Mathematical Theory of Automata, 1962, pp. 563–576.

[5] V. E. Benes, *Mathematical Theory of Connecting Networks and Telephone Traffic*, Academic Press, New York, 1965.

[6] V. E. Benes, *Optimal rearrangeable multistage connecting networks*, Bell System Tech. J., 43 (1964), pp. 1641–1656.

[7] P. Brucker, J. Hurink, and F. Werner, *Improving local search heuristics for some scheduling problems*—I, Discrete Appl. Math., 65 (1996), pp. 97–122.

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

[9] C. Clos, *A study of nonblocking switching networks*, Bell Systems Tech. J., 32 (1953), pp. 406–424.

[10] D. de Werra, *On some combinatorial problems arising in scheduling*, CORS. J., 8 (1970), pp. 165–175.

[11] R. Diestel, *Graph Theory*, Grad. Texts in Math. 173, Springer-Verlag, New York, 1997.
[12] D. Z. Du, B. Gao, F. K. Hwang, and J. H. Kim, *On multirate rearrangeable Clos networks*, SIAM J. Comput., 28 (1998), pp. 463–470.
[13] B. Gao and F. Hwang, *Wide-sense nonblocking for multirate* 3*-stage Clos networks*, Theoret. Comput. Sci., 182 (1997), pp. 171–182.
[14] P. Haxell, A. Rasala, G. Wilfong, and P. Winkler, *Wide-sense nonblocking WDM cross-connects*, in Proceedings of the 10th European Symposium on Algorithms (ESA 2002), Lecture Notes in Comput. Sci. 2461, Springer-Verlag, Berlin, 2002, pp. 538–550.
[15] A. J. Hoffman, *Generalization of a theorem of König*, J. Washington Acad. Sci., 46 (1956), pp. 211–212.
[16] A. Itoh, W. Takahashi, H. Nagano, M. Kurisaka, and S. Iwasaki, *Practical implementation and packaging technologies for a large-scale ATM switching system*, IEEE J. Selected Areas in Comm., 9 (1991), pp. 1280–1288.
[17] N. Karmarkar, *Probabilistic analysis of some bin-packing algorithms*, in Proceedings of the 23rd Annual IEEE Symposium on Foundations of Computer Science, 1982, pp. 107–111.
[18] N. Karmarkar and R. M. Karp, *The Differencing Method of Set Partitioning*, Technical report UCB/CSD 82/113, University of California, Berkeley, CA, 1982.
[19] D. König, *Graphok és alkalmazásuk a determinánsok és a halmazok elméletére*, Mathematikai és Természettudományi Értesitö, 34 (1916), pp. 104–119.
[20] 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.
[21] G. Lin, D. Du, W. Wu, and K. Yoo, *On* 3*-rate rearrangeability of Clos networks*, in Advances in Switching Networks, DIMACS Ser. Discrete Math. Theoret. Comput. Sci. 42. AMS, Providence, RI, 1998, pp. 315–333.
[22] R. Loulou, *Probabilistic behavior of optimal bin-packing solutions*, Oper. Res. Lett., 3 (1984), pp. 129–135.
[23] R. Melen and J. S. Turner, *Nonblocking multirate networks*, SIAM J. Comput., 18 (1989), pp. 301–313.
[24] H. Q. Ngo and V. H. Vu, *Multirate rearrangeable Clos networks and a generalized edge coloring problem on bipartite graphs*, in Proceedings of the Fourteenth ACM-SIAM Symposium on Discrete Algorithms (SODA 2003), 2003, pp. 834–840.
[25] W. Rhee, *Optimal bin packing with items of random sizes*, Math. Oper. Res., 13 (1988), pp. 140–151.
[26] W. Rhee and M. Talagrand, *Some distributions that allow perfect packing*, J. Assoc. Comput. Math., 35 (1988), pp. 564–578.
[27] C. E. Shannon, *Memory requirements in a telephone exchange*, Bell Systems Tech. J., 29 (1950), pp. 343–349.
[28] P. Schuurman and T. Vredeveld, *Performance guarantees of local search for multiprocessor scheduling*, in Proceedings of the 8th Conference on Integer Programming and Combinatorial Optimization (IPCO 2001), Springer-Verlag, Berlin, 2001, pp. 370–382.
[29] D. Slepian, *Two Theorems on a Particular Crossbar Switching*, manuscript, 1958.
[30] J. S. Turner and R. Melen, *Multirate Clos networks*, IEEE Comm. Mag., 41 (2003), pp. 38–44.
[31] K.-H. Tsai, D.-W. Wang, and F. Hwang, *Lower bound for wide-sense nonblocking Clos network*, Theoret. Comput. Sci., 261 (2001), pp. 323–328.