

An Approximate König's Theorem for Edge-Coloring Weighted Bipartite Graphs

[Extended Abstract]

José R. Correa
Operations Research Center
Massachusetts Institute of Technology
Cambridge, MA 02139
jcorrea@mit.edu

Michel X. Goemans
Department of Mathematics
Massachusetts Institute of Technology
Cambridge, MA 02139
goemans@math.mit.edu

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. 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. Our analysis is interesting in 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.

Categories and Subject Descriptors

F.2.m [Analysis of Algorithms and Problem Complexity]: Miscellaneous; G.2.2 [Discrete Mathematics]: Graph Theory—*Network Problems*

General Terms

Algorithms and Theory

Keywords

Bipartite Edge Coloring, Rearrangeability of 3-Stage Clos Networks, Bin Packing

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

STOC'04, June 13–15, 2004, Chicago, Illinois, USA.
Copyright 2004 ACM 1-58113-852-0/04/0006 ...\$5.00.

1. INTRODUCTION AND SUMMARY OF RESULTS

We consider a generalized bipartite edge coloring problem introduced by Du, Gao, Hwang and Kim [9]. In this natural extension of edge coloring, we are given an edge weighted bipartite (multi)graph $G = (V, E)$. We need to color the edges with as few colors as possible such that the total weight of edges of a color incident to a vertex is at most 1. If all weights are 1, this is the classical bipartite edge coloring problem, and, by König's theorem [8, 14], the number of colors needed is equal to the maximum degree of G . In case of general weights, an analogue to the maximum degree is the maximum total weight incident to any vertex, which we denote by n . The main striking question is the dependence of the number of colors needed on this parameter n in this generalized edge coloring problem. If all weights are $\frac{1}{2} + \epsilon$ for an arbitrarily small ϵ , the number of colors needed is clearly at least $2n - 1$. However, no instance requiring more than $2n - 1$ colors is currently known. One of the main results in this paper is to show that $2.557n + o(n)$ colors are sufficient and this improves upon a bound of $\frac{17n}{6} + O(1) \sim 2.8333n + O(1)$ due to Du et al. [9].

1.1 Rearrangeability of Clos Networks.

The motivation for studying this problem comes from the question of rearrangeability of multi-rate 3-stage Clos networks. Clos networks have been introduced in 1953 [5] and have been widely used for data communications and parallel computing systems (see e.g. [1, 11]). 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 the multi-rate environment, a connection request is a triple (i, j, w) where i is an inlet, j an outlet, and w the weight (which could for example represent the bandwidth consumed). 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 one. 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 have also 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 are violated. The Clos network $C(n, m, r)$ is said to be *multirate rearrangeable* (or just rearrangeable) if every request frame is routable. The main longstanding open question is to determine the minimum value m of middle switches such that $C(n, m, r)$ is multirate rearrangeable, and this minimum value is denoted by $m(n, r)$. Melen and Turner [18] initiated the research on multirate switching networks and proved that $m(n, r) \leq 2n - 1$ if all connection requests have weight no more than $1/2$. In 1991, Chung and Ross [4] conjectured that $m(n, r) \leq 2n - 1$ in the general setting; this conjecture remains open. The best bounds known so far on the function $m(n, r)$ were obtained by Ngo and Vu [19] (lower bound) and Du, Gao, Hwang, and Kim [9] (upper bound):

$$\frac{5n}{4} \leq m(n, r) \leq \frac{41n}{16} + O(1).$$

1.2 Problem Definition.

The question of rearrangeability 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 (with say $|A| = |B| = r$); in the sequel, all our graphs are multi-graphs. 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 $0 \leq w(e) \leq 1$. 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

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

Following the notation in [19], let \mathcal{B}_r^n be the collection of such edge-weighted bipartite multi-graphs. A Clos network $C(n, m, r)$ is then (multi-rate) 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 .

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

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

Here, \mathcal{D}_r^n denotes the natural counterpart of \mathcal{B}_r^n . The required number of colors in this case, denoted by $M(n, r)$, is clearly greater or equal to $m(n, r)$.

If all weights are forced to belong to a subset $I \subset [0, 1]$, $\mathcal{B}_r^n(I)$ (resp. $\mathcal{D}_r^n(I)$) denotes the natural extension of \mathcal{B}_r^n

(resp. \mathcal{D}_r^n). In this case $m_I(n, r)$ (resp. $M_I(n, r)$) is the smallest integer such that every graph in $\mathcal{B}_r^n(I)$ (resp. $\mathcal{D}_r^n(I)$) admits a proper coloring with $m_I(n, r)$ (resp. $M_I(n, r)$) colors.

1.3 Our Approach and Results.

Our main contribution in this paper is to show 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 only improve 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 5.

For most of the paper, 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 to 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.

Before continuing to describe the approach and its analysis, it is first useful to understand the limitations of such an approach, independently of how the decomposition into matchings is performed. For this purpose, 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 $X = \{x_0, x_1, \dots, x_l\}$ with $x_0 > x_1 > \dots > x_l$, 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). If we end up coloring the big edges (i.e. those whose weight is greater or equal to α) with k colors then greedily coloring the edges of weight smaller than α (with any color that fits) does not create any new color provided that $k \geq \lceil \frac{2n}{1-\alpha} \rceil$ as was shown by Du et al. [9]. 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.

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 does not differ from the continuous bin packing instance with density $\frac{n}{x^2}$ for $x \in [\alpha, 1]$ by much more than the n additional items of size 1. More formally, we mean that the optimal number of bins of the original bin packing instance surpasses the number of bins of the continuous bin packing instance by only $n + o(n)$ (n accounts for the items of size 1). We should emphasize that our bin packing instance is independent of the given bipartite graph G ; it is only based 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 only deal with an arbitrarily fine discretization of it and consider discrete bin packing instances. Our decomposition of the graph into matchings is described in Section 2, while the construction of our bin packing instance is detailed in Section 3.

Once we have constructed this continuous bin packing instance with density $\frac{n}{x^2}$ for $x \in [\alpha, 1]$, we are then able 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 is described in Section 4 and relies on a result of Rhee and Talagrand [21]. The total number of bins used is thus

$$\left(1 + \frac{\alpha}{1-\alpha} + \ln \frac{1-\alpha}{\alpha}\right) n,$$

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 $\alpha = 0.217811\dots$, we obtain that the number of colors needed is less than $2.557n$.

Our decomposition of the bipartite graph into matchings is interesting in its own. One of the intermediate results we need in the proof concerns the following problem. Given a bipartite graph G and nonnegative numbers $\gamma_1, \dots, \gamma_l$ summing to 1, decompose the graph into F_1, \dots, 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.

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.

1.4 Discussion of Previous Work.

Let us review some existing results on this problem. The first important result we are aware of is due to Slepian

[23] (see also [2]). He used Hall's Theorem to prove that $m_{[1,1]}(n, r) = M_{[1,1]}(n, r) = n$. Melen and Turner [18] 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 $M_{[0,B]}(n, r) \leq \frac{n}{1-B}$. On the other hand it is easy to prove that $m_{[b,1]}(n, r) \leq n \lfloor \frac{1}{b} \rfloor$ and that $M_{[b,1]}(n, r) \leq \frac{n}{b}$.

As mentioned before, the best bounds known on $m(n, r)$ are $\frac{5n}{4} \leq m(n, r) \leq \frac{41n}{16} + O(1)$, and were obtained by Ngo and Vu [19] (lower bound) and Du, Gao, Hwang, and Kim [9] (upper bound). These last authors also obtained the previously best bounds for $M(n, r)$, namely $2n - 1 \leq M(n, r) \leq \frac{17n}{6} + O(1)$.

Another special case that has attracted attention is when all edge weights can only take k different values (known, in Clos networks terminology, as the bounded rate environment). For $k = 2$, one can easily verify the Chung-Ross conjecture, namely the $2n - 1$ bound holds in this case [4]. Moreover, Lin, Du, Wu and Yoo [16] proved that for $k = 3$, $\frac{9n}{4} + O(1)$ colors suffice. This bound is an improvement over the $\frac{7n}{3}$ bound obtained by Lin, Du, Hu, and Xue [15]. Unfortunately the proof of all bounds for the finite rate environment rely on rather tedious case analysis.

2. BALANCED DECOMPOSITIONS OF BIPARTITE GRAPHS

We start by showing that edges with small weights can be handled effectively in a greedy fashion.

LEMMA 2 (DU ET AL. [9]). *Consider $G = (V, E) \in \mathcal{D}_r^n$ with bipartition $V = A \cup B$ and assume that we have colored all edges $e \in E$ with $w(e) \geq \alpha$ using at least $\frac{2n}{1-\alpha}$ colors. Then, we can greedily color the remaining edges without using any additional color.*

In other words, Lemma 2 says that if $M_{(\alpha,1)}(n, r) \leq \lceil 2n/(1-\alpha) \rceil$ then $M(n, r) \leq \lceil 2n/(1-\alpha) \rceil$.

From now on we fix a parameter $0 < \alpha < 1$ and work with graphs in $\mathcal{D}_r^n(\alpha, 1)$. 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 [10]). *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. \square

The next theorem is an extension of Hoffman's result stated in Lemma 3.

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

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

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

PROOF. Let $L = \{1, \dots, l\}$. 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.

To 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 $\deg_x(v)$) and an index $i \in L$. Let $\{i\} = A_0 \subset A_1 \subset \dots \subset A_k = L$ 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

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

where $e_i(v) = 1 + \frac{\gamma(A_0)}{\gamma(A_1)} + \frac{\gamma(A_0)}{\gamma(A_2)} + \dots + \frac{\gamma(A_0)}{\gamma(A_{k-1})}$. 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

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

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

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

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 [7] 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 $i = 1, \dots, 13$, 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 only used in the proof 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 et al. [3] 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 [22] 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 [13]. 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.

The decomposition we need to construct our bin packing instance is given below.

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

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

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

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

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

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

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

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

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

$$\begin{aligned} \left(\frac{1}{\alpha_2} - \frac{1}{\alpha_1}\right) \alpha_2 \deg_{D_2}(v) + \left(\frac{1}{\alpha_2} - \frac{1}{\alpha_1}\right) \alpha_3 \deg_{D_3}(v) + \\ \dots + \left(\frac{1}{\alpha_2} - \frac{1}{\alpha_1}\right) \alpha_p \deg_{D_p}(v) \leq \left(\frac{1}{\alpha_2} - \frac{1}{\alpha_1}\right) n \end{aligned}$$

⋮

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

Note that, for all $i = 1, \dots, p$, 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

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

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

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

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 $F_k = D_k^k \cup D_{k+1}^k \cup \dots \cup D_p^k$ for all $k = 1, \dots, p$. Thus letting $a_k(v) = \sum_{i=k}^p e_k^i(v) \leq 3(p-k+1)$, we have the following

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

□

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, \dots, 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 $k = 2, \dots, p$ 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 $\left(\frac{1}{\alpha_k} - \frac{1}{\alpha_{k-1}}\right) n$ matchings (or, more formally the floor of this quantity) in the decomposition of F_k , for $k = 2, \dots, p$ 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_1 . Let \mathcal{M} be those matchings selected in F_1, \dots, F_p and, by construction, we have an item for each element of \mathcal{M} . Our bin packing instance is thus:

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

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 only depends on n and the values of α_i selected.

Given any solution to this bin packing instance with say 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

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

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 on 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 to for fractionally opened bins). The argument above regarding discarded items show 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. [9], 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 than $1 - \alpha$ 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 $\alpha_1, \dots, \alpha_p$).

3.1 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 $g : [\alpha, 1 - \alpha] \rightarrow \mathbb{R}$ be defined by $g(x) = 1/x^2$. As $n \int_{\alpha_k}^{\alpha_{k-1}} g(x) dx = \left(\frac{1}{\alpha_k} - \frac{1}{\alpha_{k-1}} \right) n$ 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

$$\begin{aligned} \sum_{k=2}^p \alpha_{k-1} \left(\frac{1}{\alpha_k} - \frac{1}{\alpha_{k-1}} \right) n &\geq n \int_{\alpha}^{1-\alpha} xg(x) dx \\ &= \int_{\alpha}^{1-\alpha} \left(\frac{n}{x} \right) dx = n \ln \frac{1-\alpha}{\alpha}. \end{aligned}$$

Therefore, from Lemma 2, we derive that our analysis can not 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 $\alpha_p = \alpha$.

3.2 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 $a = x_1 < \dots < x_q = b$ 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 $\lim_{q \rightarrow \infty} \lim_{n \rightarrow \infty} \frac{Q_n^q}{n}$. 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:

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

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 [21] 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 $\left(\frac{1}{\alpha_{k-1}} - \frac{1}{\alpha_{k-2}} \right) n$ 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 be packed each 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.

4. 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 functions:

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

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

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

To prove that all f_1 , f_2 and f_3 allow perfect packing we use a perfect packing result proved by Karmarkar [12] and Loulou [17] and a powerful theorem by Rhee and Talagrand [21]. 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 [21]). *Consider a decreasing measure μ defined over $[a, b]$ (with $0 \leq a \leq b \leq 1$) and an integer $p \geq 3$ such that $1/p \in [a, b]$. Then, μ allows perfect packing if the following are satisfied:*

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

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

Let us mention however, that although the previous result was proved in a probabilistic setting (namely under the definition: μ allows perfect packing if and only if the expected number of bins needed to pack n i.i.d. 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 $g : [\alpha, 1-\alpha] \rightarrow \mathbb{R}$ with $g(x) = 1/x^2$ allows perfect packing.*

PROOF. As $g = f_1 + f_2 + f_3$, we only need to 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 $(4-1)\alpha + \delta \approx 0.9999607$).

□

For the sake of completeness, in a forthcoming paper [6], we give a more constructive proof of Theorem 6, and the analysis is carried out using the definition of perfect packing needed in the setting of the present paper. The proof, however, still goes through under the definition in [21].

In what follows, we briefly outline this result. Let $0 \leq a \leq b \leq c \leq 1$ be such that $(p-1)a + c \leq 1$ 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.,

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

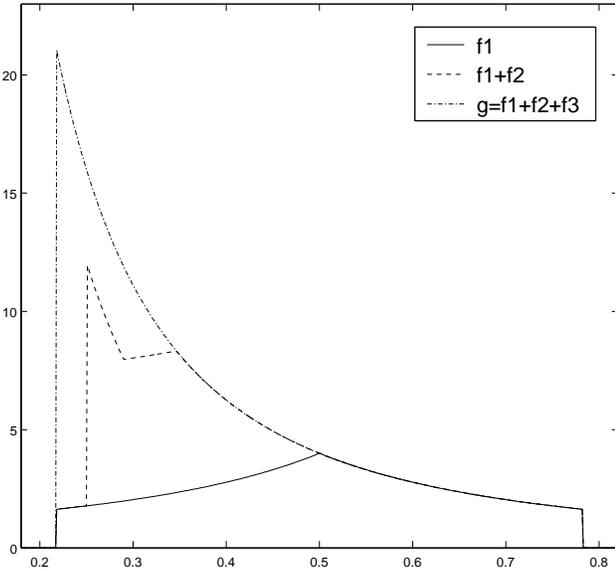
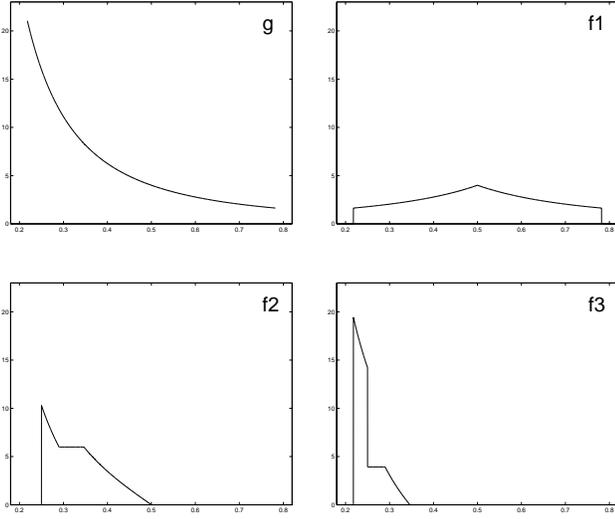


Figure 1: Decomposition of g into f_1 , f_2 and f_3

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

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

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

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, so they overcame the problem using a perfect packing characterization by Rhee [20] together with a complicated (and implicit) “exhaustion method”, that decomposes an L-shaped function into possibly uncountably many perfectly packable functions. In our approach [6], we exhibit a somewhat simple (albeit tedious) packing strategy that perfectly packs L-shaped functions. This is done by an exhaustion procedure that decomposes L-shaped functions into not too many L-shaped functions, all of which allow simple packing strategies.

5. IMPROVED ANALYSIS FOR THE RE-ARRANGEABILITY 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 only give 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: $\sum_{i=k}^p \deg_{D_i}(v) \leq 4n$, 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 $\alpha_0 = 1 > \alpha_1 > \dots > \alpha_l = 1/4 > \dots > \alpha_p = \alpha > 1/5$. Then, there exist sets F_1, \dots, F_p partitioning E , such that:*

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

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

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

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

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

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

By mimicking the analysis in Section 3, the problem now translates into packing the function $g: [\alpha, 1-\alpha] \rightarrow \mathbb{R}$ such that $g(x) = 16$ if $x \in [\alpha, 1/4]$, and $g(x) = 1/x^2$ otherwise. The value of α has now 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 on Section 4 can be found. Applying again the result in [21], such decomposition amounts to conclude that g allows perfect packing. The total number of colors needed is therefore

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

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

where the inequality comes from the choice of α .

Acknowledgments

The authors would like to thank Andreas Schulz for many useful discussions. The research was supported in part by NSF contracts CCR-0098018 and ITR-0121495.

6. REFERENCES

- [1] J. Beetem, M. Denneau & D. Weingarten, “The GF11 supercomputer,” In *Proceedings of the 12th Annual International Symposium on Computer Architecture*, 108–115, 1985.
- [2] V.E. Benes, “Mathematical Theory of Connecting Networks and Telephone Traffic,” Academic Press, New York, 1965.
- [3] P. Brucker, J. Hurink and F. Werner, “Improving local search heuristics for some scheduling problems — I,” *Discrete Applied Mathematics*, 65, 97–122, 1996.
- [4] S-P. Chung & K.W. Ross, “On nonblocking multirate interconnection networks,” *SIAM Journal on Computing* 20, 726–736, 1991.
- [5] C. Clos, “A study of nonblocking switching networks,” *Bell Systems Tech. Journal* 32, 406–424, 1953.
- [6] J.R. Correa & M.X. Goemans, in preparation.
- [7] D. de Werra, “On some combinatorial problems arising in scheduling,” *Operations Research Society Journal* 8, 165–175, 1970.
- [8] R. Diestel, “Graph Theory,” *Springer Verlag Graduate Texts in Mathematics* 173, New York, 1997.
- [9] D. Du, B. Gao, F. Hwang & J. Kim, “On multirate rearrangeable Clos networks,” *SIAM Journal on Computing* 28, 463–470, 1998.
- [10] A. J. Hoffman, “Generalization of a theorem of König,” *Journal of the Washington Academy of Science* 46, 211–212, 1956.
- [11] A. Itoh, W. Takahashi, H. Nagano, M. Kurisaka & S. Iwasaki, “Practical implementation and packaging technologies for a large-scale ATM switching system,” *Journal of Selected Areas in Communications* 9, 1280–1288, 1991.
- [12] N. Karmarkar, “Probabilistic analysis of some bin-packing algorithms,” In *Proceedings of the 23rd Symposium on Foundations of Computer Science*, 107–111, 1982.
- [13] N. Karmarkar and R.M. Karp, “The differencing method of set partitioning,” Technical Report UCB/CSD 82/113, University of California, Berkeley, 1982.
- [14] D. König, “Graphok és alkalmazásuk a determinánsok és a halmazok elméletére,” *Mathematikai és Természettudományi Értesítő* 34, 104–119, 1916.
- [15] G. Lin, D. Du, X Hu & G. Xue, “On rearrangeability of multirate Clos networks,” *SIAM Journal on Computing* 28, 1225–1231, 1999.
- [16] G. Lin, D. Du, W. Wu & K. Yoo, “On 3-rate rearrangeability of Clos networks,” *DIMACS Series in Discrete Mathematics and Theoretical Computer Science* Vol 42, 315–333, 1998.
- [17] R. Loulou, “Probabilistic behavior of optimal bin-packing solutions,” *Operations Research Letters* 3, 129–135, 1984.
- [18] R. Melen & S. Turner, “Nonblocking multirate networks”, *SIAM Journal on Computing* 18, 301–313,

1989.

- [19] H.Q. Ngo & V.H. Vu, “Multirate rearrangeable Clos networks and a generalized edge coloring problem on bipartite graphs,” In *Proceedings of the 14th ACM-SIAM Symposium on Discrete Algorithms (SODA 2003)*, 834–840, 2003.
- [20] W. Rhee, “Optimal bin packing with items of random sizes,” *Mathematics of Operations Research* 13, 140–151, 1988.
- [21] W. Rhee & M. Talagrand, “Some distributions that allow perfect packing,” *Journal of the ACM* 35, 564–578, 1988.
- [22] 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)*, 370–382, 2001.
- [23] D. Slepian, “Two theorems on a particular crossbar switching,” unpublished manuscript, 1958.