Sum Edge Coloring of Multigraphs via Configuration LP[*]

Magnús M. Halldórsson [†]   Guy Kortsarz [‡]   Maxim Sviridenko [§]

February 2, 2009

Abstract

We consider the scheduling of biprocessor jobs under sum objective (BPSMS). Given a collection of unit-length jobs where each job requires the use of two processors, find a schedule such that no two jobs involving the same processor run concurrently. The objective is to minimize the sum of the completion times of the jobs. Equivalently, we would like to find a sum edge coloring of a given multigraph, i.e. a partition of its edge set into matchings $M_1, \ldots, M_t$ minimizing $\sum_{i=1}^{t} i|M_i|$.

This problem is APX-hard, even in the case of bipartite graphs [M04]. This special case is closely related to the classic open shop scheduling problem. We give a 1.8298-approximation algorithm for BPSMS improving the previously best ratio known of 2 [BBH+98]. The algorithm combines a configuration LP with greedy methods, using non-standard randomized rounding on the LP fractions. We also give an efficient combinatorial 1.8886-approximation algorithm for the case of simple graphs.

Key-words: Edge Scheduling, Configuration LP, Approximation Algorithms

1   Introduction

We consider the problem of *biprocessor scheduling* unit jobs under sum of completion times measure (BPSMS). Given a collection of unit-length jobs where each job requires the use of two processors, find a schedule such that jobs using the same processor never run concurrently. The objective is to minimize the sum of the completion times of the jobs.

The problem can be formalized as an edge coloring problem on multigraphs, where the nodes correspond to processors and edges correspond to jobs. In each round, we schedule a matching

---

from the graph with the aim of scheduling the edges early on average. This is also known as *sum edge coloring*. Formally, the BPSMS problem is: given a multigraphs G, find an edge coloring, or a partition into a matchings $M_1, M_2, \ldots$, minimizing $\sum_i i \cdot |M_i|$.

Biprocessor scheduling problems have been studied extensively [AB$^+$00, KK85, GKMP02, CGJL85, Y05, GSS87, DK89, K96, FJP01]. Natural applications of biprocessor scheduling include, e.g., data migration problems, [GKMP02, CGJL85, Y05], running two (or more) programs on the same job in order to assure that the result is reliable [GSS87], mutual testing of processors in biprocessor diagnostic links [KK85], and batch manufacturing where jobs simultaneously require two resources for processing [DK89].

We survey some additional previous work on biprocessor scheduling. If jobs have varying length then biprocessor scheduling of trees is NP-hard both in preemptive [M05] and non-preemptive [K96] settings. Marx [M06] designed a polynomial-time approximation scheme (PTAS) for the case of preemptive biprocessor scheduling of trees. This PTAS was generalized in [M05'] to partial k-trees and planar graphs. Marx [M04] showed that BPSMS (unit jobs) is NP-hard even in the case of planar bipartite graphs of maximum degree 3 and APX-hard in general bipartite graphs. When the number of processors is constant, Afrati *et al.* [AB$^+$00] gave a polynomial time approximation scheme for minimizing the sum of completion times of biprocessor scheduling. Later, their results were generalized in [FJP01] to handle weighted completion times and release times.

Related problems: A problem related to ours is minsum edge coloring with total vertex completion time objective. This problem has a setting similar to ours but with different objective function: for each vertex compute the maximal color assigned to it and minimize the total sum of such colors over all vertices. This problem is NP-hard but admits an approximation ratio of $1 + \phi \approx 2.61$ [M08] on general graphs.

A problem more general than ours is the (vertex) *sum coloring* problem. Here, a feasible solution assigns a color (a positive integer) to each vertex of the input graph so that adjacent vertices get different colors. The goal is to minimize the sum of colors over all vertices. Thus, the BPSMS problem is the special case of sum coloring *line graphs* of general graphs.

We survey some of the work done on the sum coloring problem. This problem is NP-hard on interval graphs [M05], planar graphs [HK02], and is APX-hard on bipartite graphs [BK98]. It admits a 1.796-approximation on interval graphs [HKS01], a PTAS on planar graphs [HK02] and a 27/26-approximation on bipartite graphs [GJKM04].

The literature on scheduling with the total completion time objective is quite extensive. The approximation algorithms for such problems are usually based on completion time formulation [HSSW97] or time-indexed formulation [ScSk02]. We refer the reader to [CK04] for a thorough survey of this subject.

The only general approximation result for BPSMS is a 2-approximation algorithm in [BBH$^+$98]. Improved bounds of 1.796 [GHKS04] and $\sqrt{2}$ [GM06] have been given for the special case of bipartite graphs. The bipartite case lends additional importance to BPSMS as it captures a version of the classic open shop scheduling problem and can be denoted as $O|p_{ij} = 1|\sum C_{ij}$ using the standard scheduling notation [LLRS]. The minsum open shop scheduling with general processing times was studied in the series of papers [QS02, QS02', GHKS04].

Our main tool for the main result is the solution of a *configuration linear program (LP)*. A configuration LP is a linear program with a large (usually exponential) number of variables. Usually, it contains a variable corresponding to each "restricted" feasible solution, which in our case corresponds to a feasible assignment of edges to a particular color. See also configuration LP's for assignment problems [FGMS06, NBY06, BS06, COR01] and an interesting application of the technique giving a variable for any tree spanning a certain collection of terminals [CCG$^+$98]. Although we are not aware of the instances with the integrality gap close to our performance guarantee even for simpler linear program with assignment variables x_{et} for each edge e and color t it easy to see that the configuration LP is strictly stronger than the simple LP. For a multigraph consisting of three vertices and k parallel edges between each pair of vertices the simple LP has optimal value equal to $3k(k+1/2)$ (just spread each edge equally between $2k$ colors) while the value of an optimal integral solution and the optimal value of the configuration LP is equal to $3k(3k + 1)/2$.

1.1   Our results

Theorem 1.1 *The BPSMS problem admits an LP-based approximation algorithm with expected approximation ratio at most* 1.8298.

This results holds even if the graph has parallel edges.

This LP-based approach has high polynomial time complexity. Hence, it is of practical interest to study purely combinatorial algorithms.

Theorem 1.2 *The BPSMS problem on graphs with no parallel edges admits a combinatorial algorithm whose approximation ratio is at most* 1.8886.

This algorithm combines ideas from [HKS01] and [V64].

2   An Overview of the Main Algorithm

The main algorithm proceeds as follows. We formulate a configuration integer program that contains, for each matching M and time t, a variable x_{Mt} that represents whether M is to be scheduled at time t. We solve the linear programming relaxation by means of the ellipsoid algorithm. We then randomly round the variables in a way to be explained shortly. This results in a many-to-many

correspondence of matchings to time slots. This is resolved by first randomly reordering matchings assigned to the same time slot, and then scheduling the matching in the order of the rounding and the reordering. Finally, the edges left unscheduled after this phase are then scheduled by a simple greedy procedure that iteratively schedules the edges in the earliest possible time slot.

There are two twists to the randomized rounding. One is the introduction of a multiplier α, eventually chosen to be 0.9076. Namely, the probability that matching M is selected at time t is αx_{Mt}. This introduces a balance between the costs incurred by the edges of rounded matchings and edges scheduled in the greedy phase.

The other variation is that each x_{Mt} value is chopped up into many tiny fractions that are rounded independently. This ensures a much stronger concentration of the combined rounding, which is crucial in the greedy phase. The main reason for using this "chopping" operation is purely technical: we want the expected number of edges incident to a node v of degree d_v that are not covered during the first phase to be roughly $e^{-\alpha} d_v$ even for vertices of small degree. See Proposition 3.5. The minimum number of chopping required is to n^4 pieces (as the analysis indicates). Smaller chopping will not carry the inequalities through and larger chopping will only worsen the approximation ratio.

3   An Approximation Algorithm Using Configuration LP

Let n be the number of vertices in the graph and Δ be the maximum degree. Let $\mathcal{M}$ denote the collection of all (possibly non-maximal) matchings in the graph, including the empty matching. For an edge h, let $\mathcal{M}_h$ denote the set of matchings that contain h. We form an LP with a variable x_{Mt} for each matching M and each time slot t. Observe that the number of rounds in any reasonable schedule is at most $\Psi = 2\Delta - 1$, since each edge has at most $2\Delta - 2$ adjacent edges.

$$
\begin{aligned}
\text{Minimize} \qquad & opt^* = \sum_{t,M} t|M| x_{Mt} \\
\text{subject to:} \qquad & \sum_{M \in \mathcal{M}} x_{Mt} = 1 & \text{for each } t & \qquad (1) \\
& \sum_{M \in \mathcal{M}_h} \sum_t x_{Mt} = 1 & \text{for each edge } h & \qquad (2) \\
& x_{Mt} \geq 0 & \text{for each } t \text{ and } M &
\end{aligned}
$$

Equality (1) ensures that in an integral solution each round contains exactly one matching (if a round is not used by the optimum then it is assigned an empty matching). Equality (2) indicates that each edge belongs to exactly one scheduled matching. Since every valid solution must obey these constraints, the linear programming relaxation is indeed a relaxation of our problem.

Proposition 3.1 *The LP can be solved exactly in polynomial time and we may assume the solution vector is a basic feasible solution.*

Proof: As the *LP* has an exponential number of variables, we consider the dual program. Equality (1) receives a variable y_t and Equality (2) receives a variable z_h. The dual LP is then:

$$\text{maximize} \qquad \sum_h z_h + \sum_t y_t \tag{3}$$

$$\text{subject to:} \qquad \sum_{h \in M} z_h + y_t \leq t \cdot |M| \qquad \text{for each round } t \text{ and matching } M \tag{4}$$

In spite of having an exponential number of inequalities, the dual LP can be approximated to an arbitrary precision by the ellipsoid algorithm when given a polynomial-time separation oracle. Such an oracle must answer, within arbitrary precision, whether a candidate fractional solution is $\{y_t\} \cup \{z_h\}$ is feasible, and in case it is not, output a violated inequality (see [Kh80]). A violated inequality implies that for some M and t:

$$\sum_{h \in M} z_h + y_t > t|M|.$$

To find a violated inequality we look for every t for the largest weighted matching with weights $z_h - t$. The value of the solution is $\sum_{h \in M} z_h - t|M|$ for some matching M. Since we are maximizing $\sum_{h \in M} z_h - t|M|$, if $\sum_{h \in M} z_h - t|M| \leq y_t$, this inequality is true to every M. Otherwise, we found a violated constraint. As long as we find separating hyperplanes we continue doing so until a feasible solution is found (as one clearly exists). We may restrict the variables in the primal program to those that correspond to violated dual constraints found (as these dual constraint suffice to define the resulting dual solution). Thus, the number of variables is polynomial and the primal program can be solved within arbitrary precision in polynomial time.

We may also assume that the solution is a basic feasible solution (see for example [J2001]).   ∎

3.1   The randomized rounding algorithm

For simplicity of the presentation the algorithm and analysis are for simple graphs. We point out later the changes needed (in only one proof) if the graph has parallel edges.   Let $\{x_{Mt}\}$ be the optimum fractional solution for the instance at hand. We assume that $\{x_{Mt}\}$ is a basic feasible solution. This means that the number of x_{Mt} variables with positive values is at most the number of variables in the dual *LP*; this is at most the number of rounds plus the number of edges, or at most $2n^2$.

Let $1/2 \leq \alpha \leq 1$ be a universal constant whose value will be optimized later. For technical reasons (to be clarified later) we need to do a *chopping operation* described as follows. Each x_{Mt} is replaced by a set of n^4 smaller fractions $y^1_{Mt}, \ldots, y^{n^4}_{Mt}$, each with a value of x_{Mt}/n^4. We think of y^j_{Mt} as associated with a matching $M_{j,t} = M$ and "responsible" for an x_{Mt}/n^4 fraction of x_{Mt}.

The algorithm has a randomized phase followed by a deterministic phase. The random phase first assigns to each time slot t a *bucket B_t* of matchings, all provisionally requesting to occupy that slot. Since bucket B_t may currently contain more than one matching, a procedure is needed to decide which unique matching to use in each round t. In this *spreading* procedure, the matchings are evenly distributed so that each *actual round t* receives at most one matching, thus forming a proper schedule. Because of the spreading, the actual unique matching assigned to round t may not even belong to B_t.

Formally, the algorithm is defined as follows:

(i) **The initial ordering step:**

   We choose an arbitrary *initial order* on all the chopped $M_{j,t}$ matchings. The order is deterministic and obeys time slots in that if $t < t'$ then any $M_{j,t}$ must come in the initial order *before* $M'_{j',t'}$ regardless of j', j, M', M.

(ii) **The randomized rounding step:** For each $M'_{j,t}$, $1 \leq t \leq \Psi$ and $1 \leq j \leq n^4$, place M' independently at random in bucket $B_{t'}$ with probability $\alpha \cdot y^j_{t',M'}$. Several matchings may be selected into the same bucket $B_{t'}$, indicating their provisional wish to be scheduled at time t'.

(iii) **The random reordering step:** Randomly reorder all matchings that reached bucket B_t.

   Remark: The random ordering step is completely independent of the randomized rounding step.

(iv) **The spreading step:** Spread the resulting matchings, in order of bucket number and the chosen order within each bucket. Thus, all matchings of bucket B_t will be scheduled before all matchings of bucket B_{t+1}, etc. For each i, let m_i denote the number of matchings in B_i. By increasing bucket numbers i, the matchings in B_i are assigned to rounds $\sum_{j=1}^{i-1} m_j + 1$ through $\sum_{j=1}^{i-1} m_j + m_i$, with the ordering of the m_i matchings that belong to B_i following the random order of the previous step.

(v) **The assignment step**: We assign each edge h to the first matching M that contains h (in the order of spreading). If M is the i-th matching in the order then the *finish time* of h is $f(h) = i$. We then remove copies of h from all other matchings.

Denote by $\mathcal{M}_1, \mathcal{M}_2, \ldots \mathcal{M}_\gamma$ the matchings chosen in this first phase. Let $E_c = \bigcup_{i=1}^{\gamma} \mathcal{M}_i$ be the set of edges covered in this first phase, and $E_u = E - E_c$ be the set of *uncovered* edges.

In the second phase, the edges of E_u are now scheduled in an arbitrary sequence by a straightforward greedy procedure. Namely, add edges one by one to the current schedule (collection of matchings) as early as possible, i.e. if an edge $(z, v) \in E_u$ is scheduled at time t then for every time step $t' < t$ either edge incident to node z or an edge incident to node v is scheduled at time t'. Let the final schedule be $\{\mathcal{M}_1, \mathcal{M}_2, \ldots, \mathcal{M}_\Phi\}$ where $\Phi \geq \gamma$.

3.2 Analysis

Recall that for each matching M and each bucket t, there are n^4 tries of adding M into t. Let A^j_{Mt} be the event the j-th try succeeds and Y^j_{Mt} the indicator random variable of A^j_{Mt}. Let $X_{Mt} = \sum_j Y^j_{Mt}$. Since $\Pr[A^j_{Mt}] = \alpha y^j_{Mt}$, by linearity of expectation, $\mathbf{E}[X_{Mt}] = \alpha \cdot x_{Mt}$.

Let $\mathcal{A}_h = \{A^j_{Mt} | j = 1, \dots, n^4, M \in \mathcal{M}_h, t = 1, \dots, \Psi\}$ be the set of events involving edge h. We order these events for h according to the initial order (see the first step of the algorithm), and denote $\mathcal{A}_h = \{A^h_1, A^h_2, \dots, A^h_{k_h}\}$ according to this order. The i-th event A^h_i in this order corresponds to some $A^j_{M',t'}$ that is the jth trial of some pair t', M'. Correspondingly, let y^h_i denote the chopped value $y^j_{M',t'}$, namely, $y^h_i = \Pr[A^j_{t'M'}]/\alpha$. Note that the order above depends for now only on the initial order and does not depend on the later random ordering in buckets. Let k_h be the total number of events related to h.

Let $deg(v)$ be the degree of vertex v, and $d_c(v)$ be the number of edges incident on v in E_c (as a random variable).

Let the *finish time* of an edge h, denoted $f(h)$, be the round i in which h was scheduled, i.e., if the unique matching containing h was $\mathcal{M}_i$. We split this quantity into two parts: its *covering contribution* $f_c(h)$ and its *uncovered contribution* $f_u(h)$. If $h \in E_c$, then $f_c(h) = f(h)$ and $f_u(h) = 0$, while if $h \in E_u$, then $f_c(h) = 0$ and $f_u(h) = f(h)$. Clearly, $f_c(h) + f_u(h) = f(h)$. Thus, $f_c(h)$ ($f_u(h)$) correspond to the amount contributed by h to the objective function from the first (second) phase, respectively.

Remark: Throughout, we assume, without explicit mention, that n is large enough (larger than any universal constant required in the analysis).

Outline of the proof We analyze the performance of the algorithm by bounding individually the expected covering and uncovered contribution of each edge. See Figure 1 for an overview of the structure of the proof.

The unconditional covering contribution is obtained in Lemma 3.7. For this, we use the initial ordering defined on the events involving the edge h and first evaluate the expected covering contribution conditioned on the edge being selected due to the first event and more generally i^{th} event in the ordering, respectively (Proposition 3.2). We use several propositions to manipulate these conditional contributions: 3.3, 3.8, 3.9. Additionally, it depends on Proposition 3.5 which bounds the probability that an edge remains uncovered after the randomized phase.

To bound the uncovered contribution, we charge each edge h by some quantity $f'_u(h)$ (see Definition 3.1) so that $\sum_{h \in E_u} f'_u(h) = \sum_{h \in E_u} f_u(h)$. This is done because analyzing the expectation of $f'_u(h)$ is easier. We bound the expected number of covered edges incident on a vertex (Proposition 3.10) and then argue the uncovered contribution of each edge (Lemma 3.11) using $f'_u(h)$. Finally, summing over all the edges of the graph, we combine the two contributions to derive the
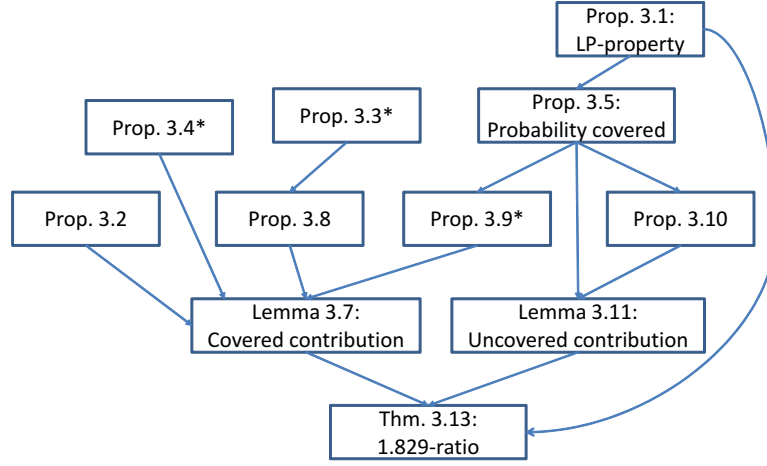
Figure 1: Structure of the proof of the main result. Asterisk denotes that proof is given in the appendix.

approximation in Theorem 3.13. We note that the part of the ratio corresponding to the covering contributions is in terms of the LP objective, while the uncovered contributions is in terms of a weaker previously studied lower bound on the optimal value (Definition 3.2).

Proof details We first bound $f_c(h)$, conditioned on $h \in E_c$. Recall that $\mathcal{A}_h = \{A_1^h, A_2^h, \ldots A_{k_h}^h\}$ is the initial ordering of matchings containing h. At least one of the events in this list has succeeded.

We compute the finish time of h under the assumption that events $A_1^h, \ldots, A_p^h$ failed and the first to succeed was A_{p+1}^h, for some $p \geq 0$. Let t_i^h be the slot number in the initial ordering of the event A_i^h.

Proposition 3.2 *The expected covering contribution of edge h, given that the $p+1$-th event in $\mathcal{A}_h$ was the first one to occur, is bounded by*

$$\mathbf{E}\left[f_c(h) \, \Big| \, \bigcap_{i \leq p} \overline{A_i^h} \cap A_{p+1}^h \right] \leq \alpha \cdot \left(t_{p+1}^h - \frac{1}{2} \sum_{i=1}^{p+1} y_i^h + \frac{1}{\alpha} - \frac{1}{2} \right) \, .$$

Proof: We bound the expected waiting time of an edge, or the expected number of matchings that precede its round. We first bound the number of those coming from earlier buckets, and then those coming from the same bucket as h.

The expected sum of fractions chosen per bucket is $\alpha \cdot \sum_{M \in \mathcal{M}} x_{Mt} = \alpha$. Thus the unconditional expected waiting time of h due to the $t_{p+1} - 1$ first rounds is $\alpha(t_{p+1}^h - 1)$. However, we are given that previous events involving h did not occur. Let $B = \{A_i^h, \ i \leq p \mid t_i^h < t_{p+1}^h\}$ be the set of

8

events concerning h that belong to earlier buckets. Then, h waits

$$\alpha \left(t_{p+1}^h - 1 - \sum_{A_i^h \in B} y_i^h \right) \tag{5}$$

in expectation for previous buckets.

We now consider the matchings belonging to the current bucket t_{p+1}^h. Let $W = \{A_i^h \mid i \leq p, t_i^h = t_{p+1}^h\}$ be the set of events involving h that precede it in the initial order but also concern the same bucket. The expected number of matchings in bucket t_{p+1}^h, conditioned on $\cap_{A \in W} \overline{A}$ (i.e. none of its preceding events involving h occurring), equals

$$\alpha \left(1 - \sum_{A_i^h \in W} y_i^h \right) \ .$$

This amount is independent of the random ordering step, but the matchings will be spread within the bucket randomly. Hence, taking the expectation also over the random orderings, the waiting time of h for this bucket is at most

$$\alpha \left(1/2 - \sum_{A_i^h \in W} y_i^h/2 - y_{p+1}^h/2 \right) \ . \tag{6}$$

We now add the waiting times of (5) and (6), observing that worst case occurs when $B = \emptyset$. In that case the expected waiting time is bounded by

$$\alpha \left(t_{p+1}^h - \frac{1}{2} - \sum_{i=1}^{p+1} y_h^i/2 \right) \ .$$

Adding the round of M_{p+1}^h itself yields the claim. ∎

Remark: The above number can indeed be strictly larger then the round of h. The round of h can be smaller if the following events happen. There is a matching containing h located *after* M_{p+1}^h in the initial ordering, this matching reaches slot t_{p+1}^h and is located before M_{p+1}^h by the random ordering. However, it seems hard to use this fact to improve the ratio.

The proof of the next claim appears in the appendix

Proposition 3.3 *Let* $y_1, y_2, \ldots, y_k$ *be non-negative numbers satisfying* $\sum_{i=1}^k y_i = 1$. *Then,*

$$\frac{y_1^2}{2} + y_2 \cdot \left(y_1 + \frac{y_2}{2} \right) + y_3 \left(y_1 + y_2 + \frac{y_3}{2} \right) + \ldots + y_k \cdot \sum_{i=1}^{k-1} y_i + \frac{y_k^2}{2} = \frac{1}{2} \ . \tag{7}$$

The next claim appeared many times in the approximation algorithms literature. Chudak and Shmoys [CS03] were the first ones to use it. A simpler proof from [FGMS06] is given in the appendix.

Proposition 3.4 Let $t_1, \ldots, t_k$ be positive numbers so that $t_1 \leq t_2 \leq \cdots \leq t_k$, and $x_1, \ldots, x_k$ be positive numbers such that $\sum_{i=1}^{k} x_i \leq 1$. Then:

$$t_1 \cdot x_1 + t_2 \cdot (1 - x_1) \cdot x_2 + \cdots + \prod_{i=1}^{k-1}(1 - x_i) \cdot x_k t_k \leq \frac{\left(1 - \prod_{i=1}^{k}(1 - x_i)\right) \cdot \sum_{i=1}^{k} t_i x_i}{\sum_{i=1}^{k} x_i}.$$

Proposition 3.5 The probability that an edge h is not covered is bounded by

$$\frac{1}{e^\alpha}\left(1 - \frac{3}{n^2}\right) \leq \Pr[h \in E_u] = \prod_{i=1}^{k_h}\left(1 - \alpha y_i^h\right) \leq \frac{1}{e^\alpha}.$$

Proof: Observe that

$$\sum_{i=1}^{k_h} y_i^h = \sum_t \sum_{M \in \mathcal{M}_h} \sum_{j=1}^{n^4} y_{Mt}^j = \sum_t \sum_{M \in \mathcal{M}_h} x_{Mt} = 1, \tag{8}$$

by constraint (2) of the LP. Therefore, the upper bound follows from the fact that $1 - x \leq e^{-x}$.

To get the lower bound we need to group the y_{Mt}^j variables according to their M, t values:

$$\Pr[h \in E_u] = \prod_{t, M \in \mathcal{M}_h} \prod_{i=1}^{n^4}\left(1 - \alpha \cdot y_{Mt}^i\right) = \prod_{t, M \in \mathcal{M}_h}\left(1 - \frac{\alpha \cdot x_{Mt}}{n^4}\right)^{n^4}.$$

The following inequality holds:

$$\left(1 - \frac{\alpha \cdot x_{Mt}}{n^4}\right)^{n^4} \geq \frac{1}{e^{\alpha x_{Mt}}} \cdot \left(1 - \frac{\alpha \cdot x_{Mt}}{n^4}\right)^{\alpha \cdot x_{Mt}}$$

because $(1 - 1/x)^{x-1} > 1/e$ holds for all $x > 1$. As $x_{Mt} \leq 1$ and $\alpha \leq 1$ we get:

$$\left(1 - \frac{\alpha \cdot x_{Mt}}{n^4}\right)^{n^4} \geq \frac{1}{e^{\alpha x_{Mt}}} \cdot \left(1 - \frac{1}{n^4}\right).$$

By assumption, $\{x_{Mt}\}$ is a basic feasible solution. Hence, the number of non-zero x_{Mt} is at most the number of edges plus the number of rounds, or at most $2n^2$. Thus

$$\prod_{i=1}^{k_h}(1 - y_i^h) \geq \left(1 - \frac{1}{n^4}\right)^{2n^2} \cdot \prod_{t, M \in \mathcal{M}_h} \frac{1}{e^{\alpha x_{Mt}}} \geq \frac{1}{e^\alpha}\left(1 - \frac{3}{n^2}\right).$$

We used Equality (2) in the LP and $(1 - 1/n^4)^{2n^2} \geq (1 - 3/n^2)$. ∎

Remark: More generally, if there are parallel edges, the fractions need to be chopped into at least $|E| \cdot n^2$ pieces.

Notation 3.6 Let $b_i^h = \alpha \cdot \left(t_i^h - \left(\sum_{j=1}^{i} y_j^h\right)/2 + 1/\alpha - 1/2\right).$

10

Lemma 3.7 *The expected covering contribution of an edge h is bounded by*

$$\mathbf{E}\left[f_c(h)\right] \leq \left(1 + \frac{1}{n}\right) \cdot \left(1 - \frac{1}{e^\alpha}\right) \cdot \left(\left(1 - \frac{3\alpha}{4}\right) + \alpha \sum_{t, M \in \mathcal{M}_h} x_{Mt} \cdot t\right) .$$

Proof: By Proposition 3.2, $\mathbf{E}\left[f_c(h) \mid \bigcap_{j=1}^{i-1} \overline{A_j^h} \cap A_i^h\right] \leq b_i^h$. Thus,

$$\mathbf{E}\left[f_c(h)\right] = \mathbf{E}\left[f_c(h) \mid A_1^h\right] \cdot Pr[A_1^h] + \mathbf{E}\left[f_c(h) \mid \overline{A_1^h} \cap A_2^h\right] \cdot Pr[\overline{A_1^h} \cap A_2^h] + \dots$$

$$+ \mathbf{E}\left[f_c(h) \mid \bigcap_{i < k_h} \overline{A_i^h} \cap A_{k_h}^h\right] \cdot Pr[\cap_{i < k_h} \overline{A_i^h} \cap A_{k_h}^h] + 0 \cdot Pr[\bigcap_{i \leq k_h} \overline{A_i^h}]$$

$$\leq \alpha \cdot y_1^h \cdot b_1^h + \left(1 - \alpha \cdot y_1^h\right) \cdot \alpha \cdot y_2^h \cdot b_2^h + \left(1 - \alpha \cdot y_1^h\right) \cdot (1 - \alpha \cdot y_2^h) \cdot \alpha y_3^h \cdot b_3^h$$

$$+ \dots + \prod_{i=1}^{k_h - 1} \left(1 - \alpha \cdot y_i^h\right) \alpha \cdot y_{k_h}^h \cdot b_{k_h}^h$$

$$\leq \left(1 - \left(\prod_{i=1}^{k_h} \left(1 - \alpha \cdot y_i^h\right)\right)\right) \cdot \sum_{i=1}^{k_h} y_i^h b_i^h . \qquad \text{(By Proposition 3.4)}$$

The α term that multiplied every term before the last inequality was canceled because $\sum_{i=1}^{k_h} \alpha y_i^h = \alpha$.

The lemma now follows from the combination of the following two propositions. ∎

Proposition 3.8 $\displaystyle\sum_{i=1}^{k_h} y_i^h b_i^h \leq \alpha \sum_{t, M \in \mathcal{M}} x_{Mt} \cdot t + (1 - 3\alpha/4).$

Proof: Recall that

$$y_i^h b_i^h = \alpha y_i^h t_i^h + \alpha \cdot \left(y_i^h \cdot \left(-\frac{1}{2} \sum_{j=1}^{i} y_j^h + \frac{1}{\alpha} - \frac{1}{2}\right)\right) .$$

We break the sum into three parts, and analyze first the total contribution of the two last terms $\left(-\alpha \cdot y_i^h \left(\sum_{j=1}^{i} y_j^h\right)\right)/2$ and $\alpha \cdot y_i^h(1/\alpha - 1/2))$ when summing over all i. By Proposition 3.3,

$$-\alpha \sum_{i=1}^{k_h} y_i^h \cdot \left(\frac{1}{2} \sum_{j=1}^{i} y_j^h\right) \leq -\frac{\alpha}{2} \sum_{i=1}^{k_h} y_i^h \cdot \left(\sum_{j=1}^{i-1} y_j^h + \frac{y_i^h}{2}\right) = -\alpha/4 . \tag{9}$$

Since $\sum_{i=1}^{k_h} y_i^h = 1$ we have that

$$\alpha \left(\frac{1}{\alpha} - \frac{1}{2}\right) \cdot \sum_{i=1}^{k_h} y_i^h = \left(1 - \frac{\alpha}{2}\right) . \tag{10}$$

11

Finally, consider the sum of the terms $\alpha y_i^h t_i^h$. By re-indexing, we have that

$$\sum_{i=1}^{k_h} \alpha y_i^h \cdot t_i^h = \alpha \sum_{t,M \in \mathcal{M}_h} \sum_{i=1}^{n^4} y_{Mt}^i \cdot t = \alpha \sum_{t,M \in \mathcal{M}_h} x_{Mt} \cdot t . \tag{11}$$

Adding (9), (10) and (11) now yields the claim. ∎

The proof of the following lemma is given in the appendix.

Proposition 3.9 $\left(1 - \left(\prod_{i=1}^{k_h} \left(1 - \alpha \cdot y_i^h \right) \right) \right) \leq \left(1 + \frac{1}{n} \right) \cdot \left(1 - \frac{1}{e^\alpha} \right) .$

We now turn to bounding the expected contribution of uncovered edges.

Proposition 3.10 *The expected number of covered edges incident on a vertex v is bounded by*

$$\mathbf{E}\left[d_c(v)\right] \leq \left(1 + \frac{1}{n} \right) \left(1 - \frac{1}{e^\alpha} \right) deg(v) .$$

Proof: By Proposition 3.5, we have for each $h \in E(v)$ that

$$\Pr[h \in E_c] = 1 - \Pr[h \in E_u] \leq 1 - \frac{1}{e^\alpha} \left(1 - \frac{3}{n^2} \right) = \left(1 - \frac{1}{e^\alpha} \right) + \frac{3}{n^2 \cdot e^\alpha}.$$

Using $\alpha \geq 1/2$, we have that, for large enough n,

$$\Pr[h \in E_c] \leq \left(1 + \frac{1}{n} \right) \left(1 - \frac{1}{e^\alpha} \right) .$$

By linearity of expectation,

$$\mathbf{E}\left[d_c(v)\right] \leq \left(1 + \frac{1}{n} \right) \left(1 - \frac{1}{e^\alpha} \right) deg(v) .$$

∎

Remark: It appears difficult to bound the expectation of $d_c(v)$ as above without the chopping operation, especially when degrees are small.

We are ready to bound from above the expectation of the sum of finish times in the greedy phase. Consider an uncovered edge $(u, v) \in E_u$. This edge would first have to wait a total of at most $d_c(v) + d_c(u)$ rounds until all covered edges incident on v and u are scheduled. After that, each time the edge (u, v) is not selected, at least one of u or v is matched. Thus, each time the edge (u, v) waits can be charged to an edge in E_u that is incident on either u or v. Thus, the contribution of the greedy step to the sum of finish times is at most

$$\sum_{v \in V} \sum_{i=d_c(v)+1}^{deg(v)} i = \sum_{v \in V} \left(\sum_{i=1}^{deg(v)} i - \sum_{i=1}^{d_c(v)} i \right) = \sum_{w \in V} \left(\binom{deg(w) + 1}{2} - \binom{d_c(w) + 1}{2} \right) . \tag{12}$$

Now, divide the term $\binom{deg(v)+1}{2} - \binom{d_c(v)+1}{2}$, equally among the edges in $E_u(v)$. Then, the edge $h = (z, v)$ is charged

$$\frac{\binom{deg(z)+1}{2} - \binom{d_c(z)+1}{2}}{deg(z) - d_c(z)} + \frac{\binom{deg(v)+1}{2} - \binom{d_c(v)+1}{2}}{deg(v) - d_c(v)} = \frac{deg(z) + d_c(z)}{2} + \frac{deg(v) + d_c(v)}{2} + 1 \ .$$

Definition 3.1 *Define*

$$f_u'(h) = \begin{cases} \frac{deg(z)+d_c(z)}{2} + \frac{deg(v)+d_c(v)}{2} + 1, & \text{if } h \in E_u \\ 0, & \text{otherwise.} \end{cases}$$

Observe that $\sum_{h \in E} f_u'(h) \geq \sum_{h \in E} f_u(h)$. Hence, for the purpose of evaluating the sum of the expected values, it is enough to bound the expectation of $f_u'(h)$.

Lemma 3.11 *The contribution of an uncovered edge $h = (z, v)$ is bounded by*

$$\mathbf{E}\left[f_u'(h) \right] \leq \left(1 + \frac{1}{n} \right) \cdot \frac{1}{e^\alpha} \left[\left(2 - \frac{1}{e^\alpha} \right) \cdot \left(\frac{deg(z) + deg(v)}{2} + 1 \right) - \left(1 - \frac{1}{e^\alpha} \right) \right]$$

Proof: The probability that h is uncovered is at most $1/e^\alpha$ (see Proposition 3.5). Thus,

$$\mathbf{E}\left[f_u'(h) \right] \leq \frac{1}{e^\alpha} \cdot \mathbf{E}\left[\frac{deg(z) + d_c(z)}{2} + \frac{deg(v) + d_c(v)}{2} + 1 \right]$$
$$\leq \frac{1}{e^\alpha} \left(1 + \frac{1}{n} \right) \left(2 - \frac{1}{e^\alpha} \right) \left(\frac{deg(z) + deg(v)}{2} + 1 \right)$$
$$- \frac{1}{e^\alpha} \left(1 + \frac{1}{n} \right) \left(1 - \frac{1}{e^\alpha} \right) \ . \qquad \text{(By Proposition 3.10)}$$

∎

The following measure has been useful for lower bounding the cost of the optimal solution.

Definition 3.2 *Let* $q(G) = \sum_{v \in V} \binom{deg(v) + 1}{2} \ .$

Let $opt(G)$ denote the cost of the optimal sum edge coloring of G, and recall that opt^* denotes the value of the linear program. It was shown in [BBH$^+$98] that $opt(G) \geq q(G)/2$.

Observation 3.12 $opt(G) \geq q(G)/2 = \sum_{(z',v') \in E} \left(\frac{deg(z') + deg(v')}{4} + \frac{1}{2} \right)$

Proof:

$$q(G)/2 = \sum_{v \in V} \frac{deg(v)(deg(v) + 1)}{4} = \sum_{v \in V} \frac{deg^2(v)}{4} + \frac{|E|}{2} = \sum_{(z',v') \in E} \left(\frac{deg(z') + deg(v')}{4} + \frac{1}{2} \right) \ .$$

∎

Theorem 3.13 *The expected approximation ratio of the algorithm is at most 1.8298.*

13

Proof: Each edge h' contributes two terms that depend *only on n and α*. There is the positive contribution (see Lemma 3.7) of

$$\left(1 + \frac{1}{n}\right) \cdot \left(1 - \frac{1}{e^\alpha}\right) \cdot \left(1 - \frac{3\alpha}{4}\right) \tag{13}$$

and the negative term from Lemma 3.11 of

$$-\frac{1}{e^\alpha} \cdot \left(1 + \frac{1}{n}\right) \cdot \left(1 - \frac{1}{e^\alpha}\right) \ . \tag{14}$$

This amounts to

$$\left(1 + \frac{1}{n}\right) \cdot \left(1 - \frac{1}{e^\alpha}\right) \cdot \left(1 - \frac{3\alpha}{4} - \frac{1}{e^\alpha}\right) \ . \tag{15}$$

For $\alpha \geq 0.606$, the above term is negative and only makes the upper bound smaller.

Ignoring the terms (13) and (14) we get:

$$E\left[\sum_h f(h)\right] = \sum_h \left(\mathbf{E}\left[f_c(h)\right] + \mathbf{E}\left[f_u(h)\right]\right) = \sum_h \left(\mathbf{E}\left[f_c(h)\right] + \mathbf{E}\left[f'_u(h)\right]\right)$$

$$\leq \sum_h \left(\left(1 + \frac{1}{n}\right)\alpha \cdot \left(1 - \frac{1}{e^\alpha}\right) \sum_{t, M \in \mathcal{M}_h} x_{Mt} \cdot t\right) \quad \text{(By Lemma 3.7)}$$

$$+ \sum_{h=(z,v)} \left(\left(1 + \frac{1}{n}\right)\frac{2}{e^\alpha}\left(2 - \frac{1}{e^\alpha}\right) \cdot \left(\frac{deg(z) + deg(v)}{4} + \frac{1}{2}\right)\right) \quad \text{(By Lemma 3.11)}$$

$$= \left(1 + \frac{1}{n}\right)\alpha\left(1 - \frac{1}{e^\alpha}\right)opt^* + \left(1 + \frac{1}{n}\right)\frac{2}{e^\alpha}\left(2 - \frac{1}{e^\alpha}\right) \cdot \frac{q(G)}{2}$$

$$\leq \left(1 + \frac{1}{n}\right) \cdot \left(\alpha \cdot \left(1 - \frac{1}{e^\alpha}\right) + \frac{2}{e^\alpha}\left(2 - \frac{1}{e^\alpha}\right)\right)opt(G) \ . \quad \text{(By Observation 3.12)}$$

For $\alpha = 0.9076$ and large enough n, the right hand side evaluates to $1.8298 opt(G)$.   ∎

4   Combinatorial Approach

We give a combinatorial algorithm that applies to simple graphs. We use the algorithm ACS of [HKS01]. This algorithm is designed to operate in rounds, finding in each round a sub-solution of maximum throughput.

A *b-matching* is a subgraph where each node has at most b incident edges. A maximum weight b-matching can be found in polynomial time by a reduction to matching (cf. [CCPS98]). Denote this algorithm by b-Mat.

In each round i, the algorithm ACS finds a maximum k_i-matching, where $k_1, k_2, \ldots$ forms a geometric sequence. The base q of the sequence is a parameter to the algorithm, while the offset

α is selected uniformly at random from the range $[0,1)$. We show later how to derandomize this strategy. The b-matching is then turned into a collection of matchings, using Vizing's algorithm, and those are then scheduled in the natural non-decreasing order of size. See Fig. 2.

ACS(G**,**q**)**
   $\alpha = \mathbf{U}[0,1);\ i \leftarrow 0$
   while $(G \neq \emptyset)$ do
      $k_i = \lfloor q^{i+\alpha} \rfloor$
      $G_i \leftarrow b\text{-Matching}(G, k_i).$
      Color the edges of G_i using Vizing's algorithm
      Schedule the colors in non-increasing order of cardinality
      $G \leftarrow G - G_i$
      $i \leftarrow i + 1;$
   end

Figure 2: Combinatorial algorithm ACS for simple graphs

This algorithm attains a performance ratio of 1.796 when an algorithm is available for obtaining an maximum k-(edge)-colorable subgraph [HKS01]. This leads to a 1.796-approximation of BPSMS in bipartite graphs [GHKS04]. The main issue for analysis is to assess the extra cost per round due to the application of Vizing's algorithm. We use ideas from [EHLS08] for simplifying parts of the analysis.

Analysis   The analysis proceeds as follows. After defining a simple lower bound on the cost of edges in an optimal solution, we introduce a series of notation, followed by an expression that we show forms an upper bound on the amortized cost of each edge in the algorithm's solution. We simplify it and break it into parts that can be viewed as the asymptotic cost on one hand, and the nearly-additive incremental cost due to the application of Vizing's algorithm. The main effort lies in evaluating the latter precisely. In particular, we utilize that when the first block involves a 1-matching, then we already have a single matching and avoid paying the extra cost incurred by Vizing's algorithm. The performance ratio is then derived from these upper and lower bounds.

For any natural number r, denote by d_r the minimum number of colors needed to color at least r vertices from V. Observe the following easy bound.

Observation 4.1 $\sum_{r=1}^{n} d_r \leq opt(G)$.

Let the edges be indexed $r = 1, 2, \ldots, n$ in agreement with the schedule formed by ACS. Let the *block* containing edge r denote the iteration i of ACS (starting at 0) in which the edge is scheduled. Let δ be characteristic function of the event that the first block is of size 1, i.e. $\delta = 1$ if $\alpha < \log_q 2$ and 0 otherwise. Let b_r be the minimum value i such that $d_r \leq k_i$. Let β_r be the characteristic

function of the event that b_r is positive, i.e. $\beta_r = 1$ if $\alpha < \min(\log_q d_r, 1)$ and 0 otherwise. We analyze the following amortized upper bound on the algorithm's cost of coloring edge r:

$$\pi_r = \sum_{i=0}^{b_r-1}(k_i + 1) + \frac{k_{b_r} + 2}{2} - \frac{\beta_r + 1}{2}\delta \ .$$

Lemma 4.2 $ACS(G) \leq \sum_{r=1}^{n} \pi_r$.

Proof: Define b'_r as the block number of edge r. Recall that a maximum b-matching contains at least as many edges as can be colored with b colors. Thus, $b'_r \leq b_r$, for all r. We first argue that $ACS(G) \leq \sum_{r=1}^{n} \pi'_r$, where

$$\pi'_r = \begin{cases} \sum_{i=0}^{b'_r-1}(k_i + 1) + \frac{k_{b'_r} + 2}{2} - \delta & \text{if } b_{r'} > 0 \\ \frac{k_0 + 2 - \delta}{2} & \text{if } b_{r'} = 0. \end{cases} \tag{16}$$

First, let us consider nodes r that fall into block later than the first one, i.e. $b'_r > 0$. The first two terms of the first case of (16) bound the amortized cost of coloring the edge r, under the assumption that all blocks use an additional color due to the application of Vizing's theorem. The first term represents the number of colors used on blocks preceding that of r's. Each block is formed by a k_i-matching and is colored into $k_i + 1$ matchings. The second term represents an upper bound on the contribution of the last block. This bound is the mean over two orderings of the block: some arbitrary ordering and its reverse. The cost of the order obtained by the algorithm is at most the cost of the smaller of these two.

The last term, δ, takes into account the fact that the first block, when of unit size, has no additive cost. Thus, we subtract the additive 1 for the first block when the first block is of unit size, i.e. when $\alpha < \log_q 2$.

Consider now nodes r that fall into the first block, i.e. $b'_r = 0$. The size of that block is $k_0 = \lfloor q^\alpha \rfloor$. If $\alpha < \log_q 2$, the block is of unit size and the node will be colored $\pi'_r = \frac{1+2-1}{2} = 1$. If $\alpha \geq \log_q 2$, the block is of non-unit size, and $k_0 + 1$ colors will be used. The average cost of nodes in the block is again at most the average of two opposite orderings, or $(k_0 + 2)/2 = \pi'_r$.

Finally, we claim that $\sum_{r=1}^{n} \pi'_r \leq \sum_{r=1}^{n} \pi_r$. Recall that $b'_r \leq b_r$. Observe that when $b'_r = b_r$, for all r, then $\pi'_r = \pi_r$. We argue that π'_r is monotone with the values of $\{b'_r\}_r$. Namely, it is easy to verify that if one b'_r value increases, the cost of some π'_r values will strictly increase and none will decrease. Hence, the claim. ∎

We split our treatment into two terms that intuitively represent the asymptotic contribution of the r-th edge and the nearly-additive incidental cost. The following characterization is obtained by rewriting k_i as $(\lfloor q^\alpha \rfloor - q^\alpha) + q^\alpha$.

Lemma 4.3 *For each edge r,*

$$\pi_r \leq \varphi_r \left[\frac{1}{q-1} + \frac{1}{2} \right] + b_r + 1 - \frac{q^\alpha}{q-1} + \gamma_r \ .$$

16

where $\varphi_r = q^{b_r + \alpha}$ and

$$\gamma_r = \frac{\beta_r + 1}{2} \left(\lfloor q^\alpha \rfloor - q^\alpha - \delta \right) \ .$$

The following bound on the asymptotic costs of π_r was given in [HKS01] (see also [EHLS08]).

Lemma 4.4 $\mathbf{E}[\varphi_r] = \frac{q-1}{\ln q} d_r$.

Proof: Let s be a real number such that $d_r = q^{s+\alpha}$. Recall that b_r is the smallest integer i such that $\lfloor q^{s+\alpha} \rfloor = d_r \le \lfloor q^{i+\alpha} \rfloor$. Thus, $b_r = \lceil s \rceil$. Hence, $\varphi_r = q^{b_r - s} d_r = q^{\lceil s \rceil - s} d_r$. Since $\alpha \sim \mathbf{U}[0,1)$, so is $\lceil s \rceil - s$. Thus,

$$\frac{\mathbf{E}[\varphi_r]}{d_r} = \mathbf{E}\left[\frac{\varphi_r}{d_r} \right] = \mathbf{E}\left[q^{\lceil s \rceil - s} \right] = \mathbf{E}[q^\alpha] = \int_0^1 q^x dx = \frac{q-1}{\ln q} \ .$$

∎

We evaluate the incidental costs in the following lemma, whose somewhat tedious proof is deferred to the appendix.

Lemma 4.5 *For* $q \in [3, 4)$,

$$\mathbf{E}[\gamma_r] = \begin{cases} \frac{1}{2}\left(3 - \log_q 12 - \frac{q-1}{\ln q} \right) & \text{if } d_r = 1 \\ \frac{1}{2}\left(3 - \log_q 12 - \frac{q}{\ln q} \right) & \text{if } d_r = 2 \\ \frac{1}{2}\left(3 - \log_q 16/3 - \frac{q+1}{\ln q} \right) & \text{if } d_r = 3, \ and \\ 3 - \log_q 12 - \frac{q-1}{\ln q} & \text{if } d_r \ge 4. \end{cases}$$

Theorem 4.6 *The performance ratio of ACS is at most* 1.8886.

Proof: By Lemma 4.2 and Observation 4.1 it suffices to show that $\mathbf{E}[\sum_{r=1}^n \pi_r] \le 1.8886 \sum_{r=1}^n d_r$. We use $q = 3.666$, which is slightly larger than the value used in [HKS01] when there was no additive overhead. By Lemmas 4.3, 4.4 and 4.5, it then holds for any r that

$$\mathbf{E}[\pi_r] \le \Pi_{d_r} \doteq \begin{cases} 1.54361 & \text{if } d_r = 1 \\ 1.74407 & \text{if } d_r = 2 \\ 1.84111 & \text{if } d_r = 3 \\ \log_q d_r - .73474 & \text{if } d_r \ge 4. \end{cases}$$

The function Π_d/d, plotted in Figure 3, attains a maximum of 1.90488 at $d = 7$, establishing that $\mathbf{E}[\pi_r] \le 1.90488 d_r$, for any r.

A slightly better performance ratio can be argued as follows. We can view the d_r-values (that intuitively represent the costs in the optimal solution) as a histogram; see Figure 4. Formally,

17

Figure 3: Plots of performance ratios. The upper line is the per-edge bound of Π_d/d, while the lower line is an amortized bound.

(a) An example graph, with a numbering on the edges

(b) Representation of the d_r-values as a histogram

Figure 4: A graph and its d_r-values

define $f_x = 1 + |\{r : d_r = x\}|$ as the number of items mapped by d to value x, or equivalently, the height of column x. Similarly, $g_y = \max\{x : f_x \geq y\}$ is the length of row y. Observe that the lower bound on the optimal solution can be expressed as $\sum_{r=1}^{n} d_r = \sum_{x=1} x \cdot f_x = \sum_{y=1} \binom{g_y+1}{2}$.

Let $\psi_g = \sum_{d=1}^{g} \Pi_d$, for $g = g_y$. We bound the algorithm's cost by

$$\sum_r \Pi_{d_r} = \sum_{y=1} \psi_g.$$

It suffices therefore to bound the ratio of ψ_g to $\binom{g}{2}$, maximized over g. This ratio can easily be seen to be a concave function of g with a single maxima. We show this function in Figure 3 along with the bound Π_r/d_r, observing that it has a maximum of 1.8886 at $g = 14$. ∎

Derandomization and weighted graphs The algorithm can be derandomized as described in [HKS01] and improved in [EHLS08]. We describe here the latter approach.

By Theorem 4.6, there exists a value for α that results in the proven expected performance ratio of ACS. Once the sequence $\lfloor q^{i+\alpha} \rfloor_{i=0}$ is fixed, the operation of the algorithm is deterministic. Each

value d_r can belong to at most two ranges of the form $[\lfloor q^{i+\alpha} \rfloor, \ldots, \lfloor q^{i+1+\alpha} \rfloor]$ as α varies. Thus, there are at most n threshold values for α when a value moves from one range to another. Hence, it suffices to test the at most n solutions to find the best value for α.

The analysis above is given for unweighted graphs. It is most easily extended to weighted graphs by viewing each weighted edge as being represented by multiple units all of which are scheduled in the same color. The analysis holds when we view each value d_r (or π_r) as corresponding to one of these small units. By making these units arbitrarily small, we obtain a matching performance ratio.

Finally, we remark that the algorithm obtains better approximation on non-sparse graphs (albeit, again it is restricted to graph with no parallel edges). Namely, when the average degree is high, the cost-per-edge in the optimal solution is linear in the average degree, while the additional charges will be only logarithmic. Hence, in the case of non-sparse simple graphs, we obtain an improvement over the LP-based approach. This approach, however, fails to work in the case of multigraphs.

Theorem 4.7 *For a simple graph, ACS obtains a solution of cost* $1.796 opt(G) + O(\log opt(G))$. *Hence, it achieves a 1.796-factor approximation for simple graphs of average degree* $\omega(1)$.

References

[AB$^+$00] F. Afrati, E. Bampis, A. Fishkin, K. Jansen and C. Kenyon. Scheduling to minimize the average completion time of dedicated tasks. In Proc. *20th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*, LNCS Volume 1974, pages 454–464, Springer, 2000.

[BBH$^+$98] A. Bar-Noy, M. Bellare, M. M. Halldórsson, H. Shachnai and T. Tamir. On chromatic sums and distributed resource allocation. *Inf. Comp.* 140:183–202, 1998.

[BK98] A. Bar-Noy and G. Kortsarz. The minimum color-sum of bipartite graphs. *J. of Algorithms* 28(2):339–365, 1998.

[BS06] N. Bansal and M. Sviridenko. The Santa Claus problem. In *Proc. 38th Ann. ACM Symp. Theory of Comput. (STOC)*, pages 31–40, 2006.

[CCPS98] W.J. Cook, W.H. Cunningham, W.R. Pulleyblank and A. Schrijver. Combinatorial Optimization. *John Wiley and Sons*, 1998.

[CCG$^+$98] M. Charikar, C. Chekuri, A. Goel, S. Guha and S. Plotkin. Approximating a Finite Metric by a Small Number of Tree Metrics. In Proc. *39th Ann. IEEE Symp. Found. of Comp. Sci. (FOCS)*, pages 379–388, 1998.

[CK04] C. Chekuri and S. Khanna. Approximation Algorithms for Minimizing Average Weighted Completion Time. Ch. 11 in *Handbook of Scheduling: Algorithms, Models, and Perf. Anal.*. J. Y.-T. Leung ed. Chapman & Hall/CRC, 2004.

[COR01] J. Chuzhoy, R. Ostrovsky and Y. Rabani. Approximation Algorithms for the Job Interval Selection Problem and Related Scheduling Problems. In Proc. *42nd Ann. IEEE Symp. Found. of Comp. Sci. (FOCS)*, pages 348–356, 2001.

[CS03] F. Chudak and D. Shmoys. Improved approximation algorithms for the uncapacitated facility location problem. *SIAM J. Comput.* 33(1):1–25, 2003.

[CGJL85] E. G. Coffman, Jr., M. R. Garey, D. S. Johnson and A. S. LaPaugh. Scheduling file transfers. *SIAM J. Comput.* 14:744–780, 1985.

[DK89] G. Dobson and U. Karmarkar. Simultaneous resource scheduling to minimize weighted flow times. *Oper. Res.* 37:592–600, 1989.

[EHLS08] L. Epstein, M. M. Halldórsson, A. Levin, and H. Shachnai. Weighted sum coloring in batch scheduling of conflicting jobs. *Algorithmica*, to appear, 2008.

[FJP01] A. V. Fishkin, K. Jansen and L. Porkolab. On minimizing average weighted completion time of multiprocessor tasks with release dates. In *Proc. 28th International Colloquium on Automata, Languages and Programming (ICALP)*. LNCS 2076, pages 875–886, 2001.

[FGMS06] L. Fleischer, M. Goemans, V. Mirrokni and M. Sviridenko. Tight approximation algorithms for maximum general assignment problems. In *Proc. 17th Ann. ACM-SIAM Symp. Disc. Algor. (SODA)*, 611–620, 2006.

[GJKM04] K. Giaro, R. Janczewski, M. Kubale, M. Małafiejski. Sum coloring of bipartite graphs with bounded degree. *Algorithmica* 40, 235-244, 2004.

[GHKS04] R. Gandhi, M. M. Halldórsson, G. Kortsarz and H. Shachnai. Approximating non-preemptive open-shop scheduling and related problems. In *Proc. 31st International Colloquium on Automata, Languages and Programming (ICALP)*. LNCS 3142, pp. 658–669, 2004.

[GM06] R. Gandhi and J. Mestre. Combinatorial algorithms for data migration. In *Proc. 9th Intl. Workshop on Approx. Algor. for Combin. Optimiz. Problems (APPROX 2006)* LNCS # 4110, Springer, 2006.

[GKMP02] K. Giaro, M. Kubale, M. Małafiejski and K. Piwakowski. Dedicated scheduling of biprocessor tasks to minimize mean flow time. In Proc. *Fourth Intl. Conf. Parallel Proc. and Appl. Math. (PPAM)* LNCS 2328, pages 87–96, 2001.

[GSS87] E. F. Gehringer, D. P. Siewiorek and Z. Segall. Parallel Processing: The Cm* Experience. *Digital Press*, Bedford, 1987.

[HK02] M. Halldórsson and G. Kortsarz. Tools for Multicoloring with applications to Planar Graphs and Partial k-trees. *J. Algorithms* 42:334-366, 2002.

[HKS01] M. M. Halldórsson, G. Kortsarz and H. Shachnai. Sum coloring interval graphs and k-claw free graphs with applications for scheduling dependent jobs. *Algorithmica* 37:187–209, 2001.

[HKS08] M. M. Halldórsson, G. Kortsarz, and M. Sviridenko. Min sum edge coloring in multigraphs via configuration LP. In *Proc. 13th Conf. Integer Prog. Combin. Optimiz. (IPCO)*, 2008.

[HSSW97] L. A. Hall, A. Schulz, D. B. Shmoys and J. Wein. Scheduling to minimize average completion time: Off-line and on-line approximation algorithms. *Math. Oper. Res.* 22:513–544, 1997.

[J2001] K. Jain. A factor 2 approximation algorithm for the generalized Steiner network problem. *Combinatorica* 21:39–60, 2001.

[K96] M. Kubale. Preemptive versus non-preemptive scheduling of biprocessor tasks on dedicated processors. *European J. Operational Research* 94:242–251, 1996.

[Kh80] L. Khachiyan. Polynomial-time algorithm for linear programming. *USSR Comput. Math. and Math. Physics* 20(1):53–72, 1980

[KK85] H. Krawczyk and M. Kubale. An approximation algorithm for diagnostic test scheduling in multicomputer systems. *IEEE Trans. Comput.* 34:869–872, 1985

[LLRS] E.L. Lawler, J.K. Lenstra, A.H.G. Rinnooy Kan and D.B. Shmoys. Sequencing and scheduling: Algorithms and complexity. In: *Handbook in Operations Research and Management Science*, Vol. 4, North-Holland, 1993, 445-522.

[M04] D. Marx. Complexity results for minimum sum edge coloring. *Manuscript*, `http://www.cs.bme.hu/~dmarx/publications.php`, 2004.

[M05] D. Marx. A short proof of the NP-completeness of minimum sum interval coloring. *Oper. Res. Lett.* 33(4):382–384, 2005.

[M06] D. Marx. Minimum sum multicoloring on the edges of trees. *Theoret. Comput. Sci.* 361(2–3):133–149, 2006.

[M05'] D. Marx. Minimum sum multicoloring on the edges of planar graphs and partial k-trees. In *Proc. 2nd Int'l Workshop on Approx. and Online Algorithms (WAOA 2004)* pages 9–22, Lecture Notes in Comput. Sci., 3351, Springer, Berlin, 2005.

[M08] J. Mestre. Adaptive Local Ratio. In *Proc. 19th Ann. ACM-SIAM Symp. Disc. Algor. (SODA)*, ACM Press, pages 152–161, 2008.

[NBY06] Z. Nutov, I. Beniaminy and R. Yuster. A $(1 - 1/e)$-approximation algorithm for the generalized assignment problem. *Oper. Res. Lett.* 34(3):283–288, 2006.

[QS02] M. Queyranne and M. Sviridenko. A (2+epsilon)-approximation algorithm for generalized preemptive open shop problem with minsum objective. *J. Algorithms* 45:202–212, 2002

[QS02'] M. Queyranne and M. Sviridenko. Approximation Algorithms for Shop Scheduling Problems with Minsum Objective, *J. Scheduling* 5:287–305, 2002.

[ScSk02] A. Schulz and M. Skutella. Scheduling Unrelated Machines by Randomized Rounding, *SIAM J. Disc. Math.* 15(4):450-469, 2002.

[V64] V. G. Vizing. On an estimate of the chromatic class of a p-graph. *Diskrete Analiz* 3:23–30, 1964.

[Y05] Y.-A. Kim. Data migration to minimize the total completion time. *J. Algorithms* 55:42–57, 2005.

A Proofs of Some Lemmas

Proposition 3.3 *Let $y_1, y_2, \ldots, y_k$ be non-negative numbers satisfying $\sum_{i=1}^{k} y_i = 1$. Then,*

$$\frac{y_1^2}{2} + y_2 \cdot \left(y_1 + \frac{y_2}{2} \right) + y_3 \left(y_1 + y_2 + \frac{y_3}{2} \right) + \ldots + y_k \cdot \sum_{i=1}^{k-1} y_i + \frac{y_k^2}{2} = \frac{1}{2} \ . \tag{17}$$

Proof: Denote the left hand side by S and rearrange its terms to obtain

$$S = y_1 \left(\frac{y_1}{2} + y_2 + \cdots + y_k \right) + y_2 \left(\frac{y_2}{2} + y_3 + \cdots + y_k \right) + \cdots + y_k \left(\frac{y_k}{2} \right) .$$

Applying the equality $\sum_{i=1}^{k} y_i = 1$ to each of the parenthesized sums, we have that

$$S = y_1 \left(1 - \frac{y_1}{2} \right) + y_2 \left(1 - y_1 - \frac{y_2}{2} \right) + \cdots + y_k \left(1 - y_1 - y_2 - \cdots - y_{k-1} - \frac{y_k}{2} \right)$$

$$= \sum_{i=1}^{k} y_i - S = 1 - S \ .$$

Hence, S is $1/2$, as claimed. ∎

Proposition 3.4 *Let $t_1, \ldots, t_k$ be positive numbers so that $t_1 \leq t_2 \leq \cdots \leq t_k$, and $x_1, \ldots, x_k$ be positive numbers such that $\sum_{i=1}^{k} x_i \leq 1$. Then:*

$$t_1 \cdot x_1 + t_2 \cdot (1 - x_1) \cdot x_2 + \cdots + \prod_{i=1}^{k-1} (1 - x_i) \cdot x_k t_k \leq \frac{\left(1 - \prod_{i=1}^{k} (1 - x_i) \right) \cdot \sum_{i=1}^{k} t_i x_i}{\sum_{i=1}^{k} x_i} \ .$$

Proof: ∎

Proposition 3.9 $\left(1 - \left(\prod_{i=1}^{k_h} \left(1 - \alpha \cdot y_i^h\right)\right)\right) \leq \left(1 + \frac{1}{n}\right) \cdot \left(1 - \frac{1}{e^\alpha}\right)$.

Proof: The term is maximized when $\prod_{i=1}^{k_h} \left(1 - \alpha \cdot y_i^h\right)$ is minimized. Recall the first inequality of Proposition 3.5:

$$\frac{1}{e^\alpha}\left(1 - \frac{3}{n^2}\right) \leq \prod_{i=1}^{k_h}\left(1 - \alpha y_i^h\right) \ .$$

We then get

$$\left(1 - \left(\prod_{i=1}^{k_h}\left(1 - \alpha \cdot y_i^h\right)\right)\right) \leq \left(1 - \frac{1}{e^\alpha}\cdot\left(1 - \frac{3}{n^2}\right)\right) \leq \left(1 + \frac{1}{n}\right)\cdot\left(1 - \frac{1}{e^\alpha}\right) \ .$$

The above is easily reduced to $3/(e^\alpha - 1) \leq n$ which holds for large enough n as $\alpha \geq 1/2$.   ∎

Lemma 4.3 *For each edge r,*

$$\pi_r \leq \varphi_r\left[\frac{1}{q-1} + \frac{1}{2}\right] + b_r + 1 - \frac{q^\alpha}{q-1} + \gamma_r \ .$$

where $\varphi_r = q^{b_r + \alpha}$ and

$$\gamma_r = \frac{\beta_r + 1}{2}\left(\lfloor q^\alpha\rfloor - q^\alpha - \delta\right) \ .$$

Proof: For r for which $b_r > 0$, and thus $\beta_r = 1$, we have

$$\pi_r = \sum_{i=0}^{b_r - 1}\lfloor q^{i+\alpha} + 1\rfloor - \delta + \frac{\lfloor q^{b_r+\alpha}\rfloor + 2}{2}$$

$$\leq \left(\lfloor q^\alpha\rfloor - q^\alpha\right) + \sum_{i=0}^{b_r - 1}\left(q^{i+\alpha} + 1\right) - \delta + \frac{q^{b_r + \alpha} + 2}{2}$$

$$= q^{b_r + \alpha}\left[\frac{1}{q-1} + \frac{1}{2}\right] - \frac{q^\alpha}{q-1} + (b_r + 1) + \left(\lfloor q^\alpha\rfloor - q^\alpha\right) - \delta$$

$$= \varphi_r\left[\frac{1}{q-1} + \frac{1}{2}\right] + b_r + 1 - \frac{q^\alpha}{q-1} + \gamma_r \ .$$

For r with $b_r = 0$, and thus $\beta_r = 0$, we have

$$\pi_r = \frac{\lfloor q^\alpha\rfloor + 2 - \delta}{2} = q^\alpha\left[\frac{1}{q-1} + \frac{1}{2}\right] - \frac{q^\alpha}{q-1} + 1 + \frac{\lfloor q^\alpha\rfloor - q^\alpha - \delta}{2} = \varphi_r\left[\frac{1}{q-1} + \frac{1}{2}\right] + b_r + 1 - \frac{q^\alpha}{q-1} + \gamma_r \ .$$

∎

23

Lemma 4.5 *For $q \in [3, 4)$,*

$$\mathbf{E}[\gamma_r] = \begin{cases} \frac{1}{2}\left(3 - \log_q 12 - \frac{q-1}{\ln q}\right) & \text{if } d_r = 1 \\ \frac{1}{2}\left(3 - \log_q 12 - \frac{q}{\ln q}\right) & \text{if } d_r = 2 \\ \frac{1}{2}\left(3 - \log_q 16/3 - \frac{q+1}{\ln q}\right) & \text{if } d_r = 3, \text{ and} \\ 3 - \log_q 12 - \frac{q-1}{\ln q} & \text{if } d_r \geq 4. \end{cases}$$

Proof: Let s be a real number such that $d_r = q^{s+\alpha}$. Recall from the proof of Lemma 4.4 that $b_r = \lceil s \rceil = \lceil \log_q d_r - \alpha \rceil$. This can be rewritten as $b_r = \lfloor \log_q d_r \rfloor + t_r$, where t_r is a Bernoulli variable with probability $\log_q d_r - \lfloor \log_q d_r \rfloor$. Thus,

$$\mathbf{E}[b_r] = \log_q d_r . \tag{18}$$

Note that

$$\mathbf{E}\left[\lfloor q^\alpha \rfloor\right] = \sum_{i=1}^{\lfloor q \rfloor} i \cdot \int_{\alpha = \log_q i}^{\min(1, \log_q i + 1)} d\alpha = \lfloor q \rfloor - \sum_{i=2}^{\lfloor q \rfloor} \log_q i = 3 - \log_q 6, \tag{19}$$

whenever $3 \leq q < 4$.

For r with $d_r \geq 4$, we have that $b_r > 0$, and thus $\beta = 1$. Hence, using (19), we have that

$$\mathbf{E}[\gamma_r] = \mathbf{E}\left[\lfloor q^\alpha \rfloor\right] - \mathbf{E}\left[q^\alpha\right] - \mathbf{E}[\delta] = (3 - \log_q 6) - \frac{q-1}{\ln q} - \log_q 2 = 3 - \log_q 12 - \frac{q-1}{\ln q} .$$

For r with $d_r = 1$ it holds that $b_r = \beta_r = 0$. Hence,

$$\mathbf{E}[\gamma_r] = \frac{1}{2}\left(\mathbf{E}\left[\lfloor q^\alpha \rfloor\right] - \mathbf{E}\left[q^\alpha\right] - \mathbf{E}[\delta]\right) = \frac{1}{2}\left(3 - \log_q 12 - \frac{q-1}{\ln q}\right) .$$

For the remaining two cases, we need to consider three regions for the value of α, with breakpoints at $\log_q 2$ and $\log_q 3$. The variables $\lfloor q^\alpha \rfloor$, δ and β_r are constant in each of these regions. Namely,

$$\lfloor q^\alpha \rfloor = \begin{cases} 1 & \alpha \in [0, \log_q 2) \\ 2 & \alpha \in [\log_q 2, \log_q 3) \\ 3 & \alpha \in [\log_q 3, 1), \end{cases} \qquad \delta = \begin{cases} 1 & \alpha \in [0, \log_q 2) \\ 0 & \alpha \in [\log_q 2, 1), \end{cases} \qquad \beta_r = \begin{cases} 1 & \alpha \in [0, \min(\log_q d_r, 1)) \\ 0 & \text{otherwise.}) \end{cases}$$

When $d_r = 2$,

$$\mathbf{E}[\gamma_r] = \int_{x=0}^{\log_q 2} (1 - q^x - 1)dx + \int_{x=\log_q 2}^{\log_q 3} \frac{1}{2}(2 - q^x)dx + \int_{x=\log_q 3}^{1} \frac{1}{2}(3 - q^x)dx$$

$$= -\left[\frac{q^x}{\ln q}\right]_{x=0}^{\log_q 2} + [x]_{x=\log_q 2}^{\log_q 3} + \left[\frac{3x}{2}\right]_{x=\log_q 3}^{1} - \left[\frac{q^x}{2\ln q}\right]_{x=\log_q 2}^{1}$$

$$= \frac{1}{2}\left(3 - \log_q 12 - \frac{q}{\ln q}\right) .$$

24

while when $d_r = 3$,

$$\mathbf{E}[\gamma_r] = \int_{x=0}^{\log_q 2} (1 - q^x - 1)dx + \int_{x=\log_q 2}^{\log_q 3} (2 - q^x)dx + \int_{x=\log_q 3}^{1} \frac{1}{2}(3 - q^x)dx$$

$$= -\left[\frac{q^x}{\ln q}\right]_{x=0}^{\log_q 3} + [2x]_{\log_q 2}^{\log_q 3} + \frac{1}{2}\left[3x - \frac{q^x}{\ln q}\right]_{\log_q 3}^{1}$$

$$= \frac{1}{2}\left(3 - \log_q 16/3 - \frac{q+1}{\ln q}\right) \ .$$

∎