

Online Bipartite Matching with Decomposable Weights

Moses Charikar*

Monika Henzinger†

Huy L. Nguyễn‡

June 26, 2014

Abstract

We study a weighted online bipartite matching problem: $G(V_1, V_2, E)$ is a weighted bipartite graph where V_1 is known beforehand and the vertices of V_2 arrive online. The goal is to match vertices of V_2 as they arrive to vertices in V_1 , so as to maximize the sum of weights of edges in the matching. If assignments to V_1 cannot be changed, no bounded competitive ratio is achievable. We study the weighted online matching problem with *free disposal*, where vertices in V_1 can be assigned multiple times, but only get credit for the maximum weight edge assigned to them over the course of the algorithm. For this problem, the greedy algorithm is 0.5-competitive and determining whether a better competitive ratio is achievable is a well known open problem.

We identify an interesting special case where the edge weights are decomposable as the product of two factors, one corresponding to each end point of the edge. This is analogous to the well studied related machines model in the scheduling literature, although the objective functions are different. For this case of decomposable edge weights, we design a 0.5664 competitive randomized algorithm in complete bipartite graphs. We show that such instances with decomposable weights are non-trivial by establishing upper bounds of 0.618 for deterministic and 0.8 for randomized algorithms.

A tight competitive ratio of $1 - 1/e \approx 0.632$ was known previously for both the 0-1 case as well as the case where edge weights depend on the offline vertices only, but for these cases, reassignments cannot change the quality of the solution. Beating 0.5 for weighted matching where reassignments are necessary has been a significant challenge. We thus give the first online algorithm with competitive ratio strictly better than 0.5 for a non-trivial case of weighted matching with free disposal.

1 Introduction

In recent years, online bipartite matching problems have been intensely studied. Matching itself is a fundamental optimization problem with several applications, such as matching medical students to residency programs, matching men and women, matching packets to outgoing links in a router and so on. There is a rich body of work on matching problems, yet there are basic problems we don't understand and we study one such question in this work. The study of the online setting goes back to the seminal work of Karp, Vazirani and Vazirani [26] who gave an optimal $1 - 1/e$ competitive algorithm for the unweighted case. Here $G(V_1, V_2, E)$ is a bipartite graph where V_1

*Princeton University, USA, moses@cs.princeton.edu

†University of Vienna, Austria, monika.henzinger@univie.ac.at

‡Princeton University, USA, hlnghuyen@cs.princeton.edu

is known beforehand and the vertices of V_2 arrive online. The goal of the algorithm is to match vertices of V_2 as they arrive to vertices in V_1 , so as to maximize the size of the matching.

In the weighted case, edges have weights and the goal is to maximize the sum of weights of edges in the matching. In the application of assigning ad impressions to advertisers in display advertisement, the weights could represent the (expected) value of an ad impression to an advertiser and the objective function for the maximum matching problem encodes the goal of assigning ad impressions to advertisers to as to maximize total value. If assignments to V_1 cannot be changed and if edge weights depend on the *online* node to which they are adjacent, it is easy to see that no competitive ratio bounded away from 0 is achievable.

Feldman et al [18] introduced the *free disposal* setting for weighted matching, where vertices in V_1 can be assigned multiple times, but only get credit for the maximum weight edge assigned to them over the course of the algorithm. (On the other hand, a vertex in V_2 can only be assigned at the time that it arrives with no later reassignments permitted). [18] argues that this is a realistic model for assigning ad impressions to advertisers. The greedy algorithm is 0.5 competitive for the online weighted matching problem with free disposal. They study the weighted matching problem with capacities – here each vertex $v \in V_1$ is associated with a capacity $n(v)$ and gets credit for the largest $n(v)$ edge weights from vertices in V_2 assigned to v . They designed an algorithm with competitive ratio approaching $1 - 1/e$ as the capacities approach infinity. Specifically, if all capacities are at least k , their algorithm gets competitive ratio $1 - 1/e_k$ where $e_k = (1 + 1/k)^k$. If all capacities are 1, their algorithm is 1/2-competitive.

Aggarwal et al [1] considered the online weighted bipartite matching problem where edge weights are only dependent on the end point in V_1 , i.e. each vertex $v \in V_1$ has a weight $w(v)$ and the weight of all edges incident on v is $w(v)$. This is called the *vertex weighted setting*. They designed a $1 - 1/e$ competitive algorithm. Their algorithm can be viewed as a generalization of the Ranking algorithm of [26].

It is remarkable that some basic questions about a fundamental problem such as matching are still open in the online setting. Our work is motivated by the following tantalizing open problem; *Is it possible to achieve a competitive ratio better than 0.5 for weighted online matching?* Currently no upper bound better than $1 - 1/e$ is known for the setting of general weights – in fact this bound holds even for the setting of 0-1 weights. On the other hand, no algorithm with competitive ratio better than 0.5 (achieved by the greedy algorithm) is known for this problem. By the results of [18], the case where the capacities are all 1 seems to be the hardest case and this is what we focus on.

1.1 Our results

We identify an interesting special case of this problem where we have a complete graph between V_1 and V_2 and the edge weights are decomposable as the product of two factors, one corresponding to each end point of the edge. This is analogous to the well studied related machines model in the scheduling literature [5, 6, 9, 17] where the load of a job of size p on a machine of speed s is p/s although the objective functions are different. Scheduling problems typically involving minimizing the maximum machine load (makespan) or minimizing the ℓ_p norm of machine loads, where the load on a machine is the sum of loads of all jobs placed on the machine. By contrast, in the problem we study, the objective (phrased in machine scheduling terminology) is to maximize the sum of machine loads where the load of a machine is the load of the largest job placed on the machine. For this case of decomposable edge weights, we design a 0.5664 competitive algorithm (Section 3). For display

advertisement using a complete graph models the setting of a specific market segment (such as impressions for males between 20 and 30), where every advertiser is interested in every impression. The weight factor of the offline node u can model the value that a click has for advertiser u , the weight factor of the online node v can model the clickthrough probability of the user to which impression v is shown. Thus, the maximum weight matching in the setting we study corresponds to maximizing the sum of the expected values of all advertisers.

Our algorithm uses a now standard *randomized doubling technique* [8, 20, 12, 23]; however the analysis is novel and non-trivial. We perform a recursive analysis where each step proceeds as follows: We lower bound the profit that the algorithm derives from the fastest machine (i.e. the load of the largest job placed on it) relative to the difference between two optimum solutions - one corresponding to the original instance and the other corresponding to a modified instance obtained by removing this machine and all the jobs assigned to it. This is somewhat reminiscent of, but different from the local ratio technique used to design approximation algorithms. Finally, to exploit the randomness used by the algorithm we need to establish several structural properties of the worst case sequence of jobs – this is a departure from previous applications of this *randomized doubling technique*. While all previous online matching algorithms were analyzed using a *local*, step-by-step analysis, we use a *global* technique, i.e. we reason about the entire sequence of jobs at once. This might be useful for solving the case of online weighted matching for *general* weights. The algorithm and analysis is presented in Section 3 and an outline of the analysis is presented in Section 3.1.

A priori, it may seem that the setting of decomposable weights ought to be a much easier case of weighted online matching since it does not capture the well studied setting of 0-1 weights. We show that such instances with decomposable weights are non-trivial by establishing an upper bound of $(\sqrt{5} - 1)/2 \approx 0.618$ on the competitive ratios of deterministic algorithms (Section 4) and an upper bound of 0.8 on the competitive ratio of randomized algorithms (Section 5). The deterministic upper bound constructs a sequence of jobs that is the solution to a certain recurrence relation. Crucial to the success of this approach is a delicate choice of parameters to ensure that the solution of the recurrence is oscillatory (i.e. the roots are complex). In contrast to the setting with capacities, for which a deterministic algorithm with competitive ratio approaching $1 - 1/e \approx 0.632$ exists [18], our upper bound of $(\sqrt{5} - 1)/2 < 1 - 1/e$ for deterministic algorithms shows that no such competitive ratio can be achieved for the decomposable case with unit capacities. Note that the upper bound of $1 - 1/e$ for the unweighted case [26] is for randomized algorithms and does not apply to the setting of decomposable weights that we study here.

In contrast to the vertex weighted setting (and the special case of 0-1 weights) where reassignments to vertices in V_1 cannot improve the quality of the solution, any algorithm for the decomposable weight setting must necessarily exploit reassignments in order to achieve a competitive ratio bounded away from 0. For this class of instances, we give an upper bound approaching 0.5 for the competitive ratio of the greedy algorithm. This shows that for decomposable weights greedy’s performance cannot be better than for general weights, where it is 0.5-competitive (Section 2).

1.2 Related work

Goel and Mehta [21] and Birnbaum and Mathieu [10] simplified the analysis of the Ranking algorithm considerably. Devanur et al [15] recently gave an elegant randomized primal-dual interpretation of [26]; their framework also applies to the generalization to the vertex weighted setting by [1]. Haeupler et al [24] studied online weighted matching in the stochastic setting where vertices from V_2 are drawn from a known distribution. The stochastic setting had been previously studied in the

context of unweighted bipartite matching in a sequence of papers [19, 28]. Recent work has also studied the random arrival model (for unweighted matching) where the order of arrival of vertices in V_2 is assumed to be a random permutation: In this setting, Karande, et al [25] and Mahdian and Yan [27] showed that the Ranking algorithm of [26] achieves a competitive ratio better than $1 - 1/e$. A couple of recent papers analyze the performance of a randomized greedy algorithm and an analog of the Ranking algorithm for matching in general graphs [32, 22]. Another recent paper introduces a stochastic model for online matching where the goal is to maximize the number of successful assignments (where success is governed by a stochastic process) [30].

A related model allowing cancellation of previously accepted online nodes was studied in [13, 7, 4] and optimal deterministic and randomized algorithms were given. In their setting the weight of an edge depends *only* on the online node. Additionally in their model they decide in an online fashion only which online nodes to accept, *not* how to match these nodes to offline nodes. If a previously accepted node is later rejected, a non-negative cost is incurred. Since the actual matching is only determined after all online nodes have been seen, their model is very different from ours: Even if the cost of rejection of a previously accepted node is set to 0, the key difference is that they do not commit to a matching at every step and the intended matching can change dramatically from step to step. Thus, it does *not* solve the problem that we are studying.

A related problem that has been studied is online matching with preemption [29, 3, 16]. Here, the edges of a graph arrive online and the algorithm is required to maintain a subset of edges that form a matching. Previously selected edges can be rejected (preempted) in favor of newly arrived edges. This problem differs from the problem we study in two ways: (1) the graph is not necessarily bipartite, and (2) edges arrive one by one. In our (classic) case, vertices arrives online and all incident edges to a newly arrived vertex v are revealed when v arrives.

Another generalization of online bipartite matching is the Adwords problem [31, 14]. In addition, several online packing problems have been studied with applications to the Adwords and Display Advertisement problem [11, 21, 2].

1.3 Notation and preliminaries

We consider the following variant of the online bipartite matching problem. The input is a complete bipartite graph $G = (V_1 \cup V_2, V_1 \times V_2)$ along with two weight functions $s : V_1 \rightarrow \mathbb{R}_+$ and $w : V_2 \rightarrow \mathbb{R}_+$. The weight of each edge $e = (u, v)$ is the product $s(u) \cdot w(v)$. At the beginning, only s is given to the algorithm. Then, the vertices of V_2 arrive one by one. When a new vertex v arrives, $w(v)$ is revealed and the algorithm has to match it to a vertex in V_1 . At the end, the reward of each vertex $u \in V_1$ is the maximum weight assigned to u times $s(u)$. The goal of the algorithm is to maximize the sum of the rewards. To simplify the presentation we will call vertices of V_1 *machines* and vertices of V_2 *jobs*. The s -value of a machine u will be called the *speed* of the machines and the w -value of a job v is called the *size* of the job. Thus, the goal of the online algorithm is to assign jobs to machines. However, we are not studying the “classic” variant of the problem since we are using a different optimization function, motivated by display advertisements.

2 Upper bound for the greedy algorithm

We begin by addressing an obvious question, which is how well a greedy approach would solve our problem, and using the proof to provide some intuition for our algorithm in the next section. We

analyze here the following simple greedy algorithm: When a job v arrives, the algorithm computes for every machine u the difference between the weight of (u, v) and the weight (u, v') , where v' is the job currently assigned to u . If this difference is positive for at least one machine, the job is assigned to a machine with maximum difference.

Theorem 1. *The competitive ratio of the greedy algorithm is at most $\frac{1}{2-\epsilon}$ for any $\epsilon > 0$.*

Proof. Consider the following instance. V_1 consists of a vertex a with $s(a) = 1$ and $t = 1/\epsilon^2$ vertices b_1, \dots, b_t with $s(b_i) = \epsilon/2 \forall i$. V_2 consists of the following vertices arriving in the same order d_1, \dots, d_{1+t} where $w(d_i) = (1 - \epsilon/2)^{-i}$. We will prove by induction that all vertices d_i are assigned to a . When d_1 arrives, nothing is assigned so it is assigned to a . Assume that all the first t vertices are assigned to a when d_{t+1} arrives. The gain by assigning d_{t+1} to a is $(w(d_{t+1}) - w(d_t))s(a) = \epsilon(1 - \epsilon/2)^{-t-1}/2$. The gain by assigning d_{t+1} to some b_j is $w(d_{t+1})s(b_j) = \epsilon(1 - \epsilon/2)^{-t-1}/2$. Thus, the algorithm can assign d_{t+1} to a . The total reward of the algorithm is $(1 - \epsilon/2)^{-1-t}$. The optimal solution is to assign d_{1+t} to a and the rest to b_i 's, getting $(1 - \epsilon/2)^{-1-t} + (1 - \epsilon/2)^{-t} - 1 \geq (2 - \epsilon)(1 - \epsilon/2)^{-1-t}$. Thus, the competitive ratio is at most $\frac{1}{2-\epsilon}$. \square

The instance used in the proof above suggests some of the complications an algorithm has to deal with in the setting of decomposable weights: in order to have competitive ratio bounded away from 0.5, an online algorithm must necessarily place some jobs on the slow machines. In fact it is possible to design an algorithm with competitive ratio bounded away from 0.5 for the specific set of machines used in this proof (for any sequence of jobs). The idea is to ensure that a job is placed on the fast machine only if its size is larger than $(1 + \gamma)$ times the size of the largest job currently on the fast machine (for an appropriately chosen parameter γ). Such a strategy works for any set of machines consisting of one fast machine and several slow machines of the same speed. However, we do not know how to generalize this approach to an arbitrary set of machines. Still, this strategy (i.e. ensuring that jobs placed on a machine increase in size geometrically) was one of the motivations behind the design of the randomized online algorithm to be presented next.

3 Randomized algorithm

We now describe our randomized algorithm which uses a parameter c we will specify later: The algorithm picks values $x_i \in (0, 1]$ uniformly and at random, independently for each machine i . Each job of weight w considered by machine i is placed in the unique interval $w \in (c^{k+x_i}, c^{k+1+x_i}]$ where k ranges over all integers. When a new job w arrives, the algorithm checks the machines in the order of decreasing speed (with ties broken in an arbitrary but fixed way). For machine i it first determines the unique interval into which w falls, which depends on its choice of x_i . If the machine currently does not have a job in this or a bigger interval (with larger k), w is assigned to i and the algorithm stops, otherwise the algorithm checks the next machine.

The following function arises in our analysis:

Definition 2. Define $h(c) = 1 - \frac{1}{\beta} W\left(\frac{\beta e^\beta}{c}\right)$

where $\beta = \frac{c \ln(c)}{c-1} - 1$ and $W()$ is the Lambert W function (i.e. inverse of $f(x) = xe^x$).

We will prove the following theorem:

Theorem 3. For $c \geq e$, the randomized algorithm has competitive ratio $\min\left(\frac{c-1}{c \ln(c)}, h(c)\right)$. In particular, for $c = 3.55829$, the randomized algorithm has a competitive ratio 0.5664.

3.1 Analysis Outline

We briefly outline the analysis strategy before describing the details. An instance of the problem consists of a set of jobs and a set of machines. The (offline) optimal solution to an instance is obtained by ordering machines from fastest to slowest, ordering jobs from largest to smallest and assigning the i th largest job to the i th fastest machine. Say the machines are numbered $1, 2, \dots, n$, from fastest to slowest. Let OPT_i denote the value of the optimal solution for the instance seen by the machines from i onwards, i.e. the instance consisting of machines $i, i+1, \dots, n$, and the set of jobs passed by the $(i-1)$ st machine to the i th machine in the online algorithm. Then $OPT_1 = OPT$, the value of the optimal solution for the original instance. Even though we defined OPT_i to be the value of the optimal solution, we will sometimes use OPT_i to denote the optimal assignment, although the meaning will be clear from context. Define OPT_{n+1} to be 0. For $2 \leq i \leq n$, OPT_i is a random variable that depends on the random values $x_{i'}$ picked by the algorithm for $i' < i$. In the analysis, we will define random variables Δ_i such that $\Delta_i \geq OPT_i - OPT_{i+1}$ (see Lemma 4 later). Let A_i denote the profit of the online algorithm derived from machine i (i.e. the size of the largest job assigned to machine i times the speed of machine i). Let $A = \sum_{i=1}^n A_i$ be the value of the solution produced by the online algorithm. We will prove that for $1 \leq i \leq n$,

$$\mathbb{E}[A_i] \geq \alpha \mathbb{E}[\Delta_i] \geq \alpha(\mathbb{E}[OPT_i] - \mathbb{E}[OPT_{i+1}]) \quad (1)$$

for a suitable choice of $\alpha > 0.5$. The expectations in (1) are taken over the random choices of machine $1, \dots, i$. Note that $OPT_i - OPT_{i+1}$ is a random variable, but the sum of these quantities for $1 \leq i \leq n$ is $OPT_1 - OPT_{n+1} = OPT$, a deterministic quantity. Summing up (1) over $i = 1, \dots, n$, we get $\mathbb{E}[A] \geq \alpha \cdot OPT$, proving that the algorithm gives an α approximation.

Inequality (1) applies to a recursive application of the algorithm to the subinstance consisting of machines i, \dots, n and the jobs passed from machine $i-1$ to machine i . The subinstance is a function of the random choices made by the first $i-1$ machines. We will prove that for any instance of the random choices made by the first $i-1$ machines,

$$\mathbb{E}[A_i] \geq \alpha \mathbb{E}[\Delta_i]. \quad (2)$$

Here, the expectation is taken over the random choice of machine i . (2) immediately implies (1) by taking expectation over the random choices made by the first $i-1$ machines.

We need to establish (2). In fact, it suffices to do this for $i = 1$ and the proof applies to all values of i since (2) is a statement about a recursive application of the algorithm. Wlog, we normalize so that the fastest machine has speed 1 and the largest job is c . Note that this is done by simply multiplying all machine speeds by a suitable factor and all job sizes by a suitable factor – both the LHS and the RHS of (2) are scaled by the same quantity.

In order to compare Δ_1 with the profit of the algorithm, we decompose the instance into a convex combination of simpler threshold instances in Lemma 4. Here, the speeds are either all the same or take only two different values, 0 and 1. It suffices to compare the profit of the algorithm to OPT on such threshold instances.

Intuitively, if there are so few fast machines that even a relatively large job (job of weight at least 1) got assigned to a slow machine in OPT, then the original instance is mostly comparable to

the threshold instance where only a few machines have speed 1 and the rest have speed 0. Even if the fastest machine gets jobs assigned to machines of speed 0 in OPT, this does not affect the profit of the algorithm relative to OPT because OPT does not profit from these jobs either. Thus we only care about jobs of weight at least 1. Because a single machine can get at most two jobs of value in the range $[1, c]$, handling this case only requires analyzing at most two jobs. The proof for this case is contained in Lemma 7.

On the other hand, if there are a lot of fast machines so that all large jobs are assigned to fast machines in OPT, then the original instance is comparable to the threshold instance where all machines have speed 1. In this case, the fastest machine can get assigned many jobs that all contribute to OPT. However, because all speeds are the same, we can deduce the worst possible sequence of jobs: after the first few jobs, all other jobs have weights forming a geometric sequence. The rest of the proof is to analyze the algorithm on this specific sequence. The detailed proof is contained in Lemma 9.

The proofs of both Lemmata 7 and 9 use the decomposable structure of the edge weights.

3.2 Analysis Details

Recall that OPT_1 is the value of the optimal solution for the instance, and OPT_2 is the value of the optimal solution for the subinstance seen by machine 2 onwards. Assume wlog that all job sizes are distinct (by perturbing job sizes infinitesimally). For $y \leq c$, let $j(y)$ be the size of the largest job $\leq y$ or 0, if no such job exists. Let $s(y)$ be the speed of the machine in the optimal solution that $j(y)$ is assigned to or 0 if $j(y) = 0$. If there is a job of size y then $s(y)$ is the speed of the machine in the optimal solution that this job is assigned to. Note that $s(y) \in [0, 1]$ is monotone increasing with $s(c) = 1$. We refer to the function s as the *speed profile*. Note that s is not a random variable. Let the *assignment sequence* $\mathbf{w} = (w, w_1, w_2, \dots)$ denote the set of jobs assigned to the fastest machine by the algorithm where $w > w_1 > w_2 > \dots$. Let $\max(\mathbf{w})$ denote the maximum element in the sequence \mathbf{w} , i.e. $\max(\mathbf{w}) = w$. In Lemma 3, we bound $OPT_1 - OPT_2$ by a function that depends *only* on \mathbf{w} , s , and c . Such a bound is possible because of the fact that any job can be assigned to any machine, i.e. the graph is a complete graph. The value we take for the aforementioned random variable Δ_1 turns out to be exactly this bound.

Lemma 4. $OPT_1 - OPT_2 \leq c - (c - w)s(w) + \sum_{k \geq 1} w_k \cdot s(w_k)$

Proof. Let I_1, I_2 be the instances corresponding to OPT_1 and OPT_2 . I_2 is obtained from I_1 by removing the fastest machine and the set of jobs that are assigned to the fastest machine by the algorithm. Let us consider changing I_1 to I_2 in two steps: (1) Remove the fastest machine and the largest job w assigned by the algorithm to the fastest machine. (2) Remove the jobs w_1, w_2, \dots . For each step, we will bound the change in the value of the optimal solution resulting in a feasible solution for I_2 and computing its value – this will be a lower bound for OPT_2 .

First we analyze Step 1: OPT_1 assigns the largest job c to the fastest machine, contributing c to its value. The algorithm assigns w to the fastest machine instead of c . In OPT_1 , w was assigned to a machine of speed $s(w)$. When we remove w and the fastest machine from I_1 , one possible assignment to the resulting instance is obtained by placing c on the machine of speed $s(w)$. The value of the resulting solution is lower by exactly $(c + w \cdot s(w)) - c \cdot s(w) = c - (c - w)s(w)$.

Next, we analyze Step 2: Jobs w_1, w_2, \dots were assigned to machines of speeds $s(w_1), s(w_2), \dots$ in OPT_1 . When we remove jobs w_1, w_2, \dots , one feasible assignment for the resulting instance is

simply not to assign any jobs to the machines $s(w_1), s(w_2), \dots$, and keep all other assignments unchanged. The value of the solution drops by exactly $\sum_{k \geq 1} w_k \cdot s(w_k)$.

Thus we exhibited a feasible solution to instance I_2 of value V where

$$OPT_1 - V = c - (c - w)s(w) + \sum_{k \geq 1} w_k \cdot s(w_k).$$

But $OPT_2 \geq V$. Hence, the lemma follows. \square

We define the random variable Δ_1 , a function of the assignment sequence \mathbf{w} and the speed profile s , to be

$$\Delta_1(\mathbf{w}, s) = c - (c - w)s(w) + \sum_{k \geq 1} w_k \cdot s(w_k).$$

As defined, $\Delta_1(\mathbf{w}, s) \geq OPT_1 - OPT_2$. We note that even though OPT_1 and OPT_2 are functions of all the jobs in the instance, Δ_1 only depends on the subset of jobs assigned to the fastest machine by the algorithm. Our goal is to show $\mathbb{E}[A_1] = E[\max(\mathbf{w})] \geq \alpha \mathbb{E}[\Delta_1]$.

First, we argue that it suffices to restrict our analysis to a simple set of step function speed profiles s_t , $0 \leq t \leq 1$: For $t \in (0, 1]$, $s_t(y) = 1$ for $y \in [c^t, c]$ and $s_t(y) = 0$ for $y < c^t$. For $t = 0$, $s_0(y) = 1$ for all $y \leq c$.

Lemma 5. *Suppose that for $t = 0$ and for all $t \in (0, 1]$ such that there exists a job of weight c^t , we have*

$$\mathbb{E}[\max(\mathbf{w})] \geq \alpha \mathbb{E}[\Delta_1(\mathbf{w}, s_t)] \tag{3}$$

Then, $\mathbb{E}[\max(\mathbf{w})] \geq \alpha(\mathbb{E}[OPT_1] - \mathbb{E}[OPT_2])$.

Proof. Consider function $s'(y)$ defined as follows: $s'(y) = s(y)$ for $y \in [1, c]$ and $s'(y) = s(1)$ for $y < 1$. Note that s' is not a random variable. We claim that $\Delta_1(\mathbf{w}, s') \geq \Delta_1(\mathbf{w}, s)$. Since the largest job assigned to the fastest machine is $w \in [1, c]$, the $s(w)$ term is unchanged in going from $\Delta_1(\mathbf{w}, s)$ to $\Delta_1(\mathbf{w}, s')$. Further, the $s(w_k)$ terms in (\mathbf{w}, s') are \geq the corresponding terms in $\Delta_1(\mathbf{w}, s)$.

It is easy to see that s' is a convex combination of the step functions s_t , $0 \leq t \leq 1$. More specifically, $s' = \sum_t p_t s_t$ for suitably chosen coefficients p_t such that (a) $\sum_t p_t = 1$ and (b) $p_t = 0$ if $t > 0$ and no job with weight c^t exists.

For a fixed assignment sequence \mathbf{w} , note that $\Delta_1(\mathbf{w}, s') = \Delta_1(\mathbf{w}, \sum_t p_t s_t) = \sum_t p_t \cdot \Delta_1(\mathbf{w}, s_t)$. Hence, for a distribution over assignment sequences \mathbf{w} ,

$$\mathbb{E}[\Delta_1(\mathbf{w}, s')] = \mathbb{E}[OPT_1(\mathbf{w}, \sum_t p_t s_t)] = \sum_t p_t \cdot \mathbb{E}[\Delta_1(\mathbf{w}, s_t)]$$

Now, suppose that for all $0 < t \leq 1$ such that there exists a job of weight c^t and for $t = 0$

$$\mathbb{E}[\max(\mathbf{w})] \geq \alpha \mathbb{E}[\Delta_1(\mathbf{w}, s_t)].$$

This implies that

$$\begin{aligned} \mathbb{E}[\max(\mathbf{w})] &\geq \alpha \sum_t p_t \cdot \mathbb{E}[\Delta_1(\mathbf{w}, s_t)] = \alpha \mathbb{E}[\Delta_1(\mathbf{w}, s')] \\ &\geq \alpha \mathbb{E}[\Delta_1(\mathbf{w}, s)] \geq \alpha(\mathbb{E}[OPT_1] - \mathbb{E}[OPT_2]) \end{aligned}$$

\square

Note that since we scaled job sizes, the thresholds (i.e interval boundaries) c^{k+x_1} should also be scaled by the same quantity (say γ). After scaling, let $x \in (0, 1]$ be such that c^x is the unique threshold from the set $\{\gamma c^{k+x_1}, k \text{ integer}\}$ in $(1, c]$. Since x_1 is uniformly distributed in $(0, 1]$, x is also uniformly distributed in $(0, 1]$. Having defined x thus, the interval boundaries picked by the algorithm for the fastest machine are c^{x+k} for integers k .

We prove (3) for $\alpha = \min\left(\frac{c-1}{c \ln(c)}, h(c)\right)$ in two separate lemmata, one for the case $t > 0$ (Lemma 7) and the other for the case $t = 0$ (Lemma 9). Recall that the expression for Δ_1 only depends on the subset of jobs assigned to the fastest machine. We call a job a *local maximum* if it is larger than all jobs preceding it. Since the algorithm assigns a new job to the fastest machine if and only if it falls in a larger interval than the current largest job, it follows that any job assigned to the fastest machine must be a local maximum.

Define $m_S(y)$ to be the minimum job in the sequence of all local maxima in the range $(y, cy]$, i.e., the first job larger than y and at most cy , if such a job exists and 0 otherwise. We use $m_S(y)$ in two ways. (1) We define $u_0 = m_S(1)$. Note that u_0 is not a random variable. We use u_0 in Lemma 7 to prove the desired statement for $t > 0$. Specifically, we use u_0 to compute (i) a lower bound for $\mathbb{E}[w]$ as a function of u_0 (and not of any other jobs) and (ii) an upper bound for $\mathbb{E}[\Delta_1]$ as a function of u_0 . Combining (i) and (ii) we prove that the desired inequality holds for all u_0 . (2) In Lemma 9 we bound $\mathbb{E}[\sum_{k \geq 1} w_k]$ by a sum of $m_S(y)$ over suitable values of y . This simplifies the analysis since the elements in the subsequence of *all* local maxima are *not* random variables, while the values in \mathbf{w} are random variables.

We first prove some simple properties of u_0 that we will use:

Claim 6. (1) $u_0 \leq w$ and (2) $u_0 \geq w_1$.

Proof. $u_0 \leq w$ as u_0 is the minimum element in the sequence of all local maxima in $(1, c]$ and w is the element from the interval $(1, c]$ picked by the algorithm.

w_1 is the minimum element in the sequence of local maxima in the range $(c^{x-k-1}, c^{x-k}]$ for $x \in (0, 1]$ and k a non-negative integer. Either $u_0 \geq c^{x-k} \geq w_1$, or u_0 also falls into $(c^{x-k-1}, c^{x-k}]$ and $u_0 \geq w_1$ follows from the fact that w_1 is the smallest local maximum in this range, while u_0 is an arbitrary local maximum in this range. \square

The next lemmata conclude our algorithm analysis.

Lemma 7. For $c \geq e$, $t \in (0, 1]$ such that there exists a job of weight c^t , and $\alpha = \min\left(\frac{c-1}{c \ln(c)}, h(c)\right)$, we have

$$\alpha \mathbb{E}[\Delta_1(\mathbf{w}, s_t)] \leq \mathbb{E}[\max(\mathbf{w})]$$

Proof. Because there is a job with weight c^t , it must be the case that $u_0 \leq c^t$. As w is the job placed by the algorithm on the fastest machine, w is in the same interval as c for any choice of the random value $x \in (0, 1]$. Thus, $w \geq c^x$. As $c \geq w > cw_2$ it follows that $w_k < 1 \leq c^t$ for all $k > 1$ and, thus, $s_t(w_k) = 0$ for all $k > 1$. Hence $\sum_{k > 1} w_k \cdot s_t(w_k) = 0$. To analyze $\mathbb{E}[\Delta_1]$ we have to consider two cases, depending on whether $u_0 = c^t$ (and hence w_1 might contribute to $\mathbb{E}[\Delta_1]$) or whether $u_0 < c^t$ (and, thus, $s_t(w_1) = 0$ and w_1 does not contribute to $\mathbb{E}[\Delta_1]$).

Case 1: $u_0 = c^t$. Since $w \geq u_0$ it holds that $s_t(w) = 1$ for all choices of x . Thus we have

$$\mathbb{E}[c - (c - w)s_t(w)] = \mathbb{E}[w]$$

As discussed above, $\sum_{k>1} w_k \cdot s_t(w_k) = 0$ and, thus, the only contribution to $\mathbb{E}[\sum_{k \geq 1} w_k \cdot s_t(w_k)]$ is from w_1 . Additionally $s_t(w_1) = 1$ only if $w_1 = u_0 = c^t$, and this only happens when x is chosen such that $x \geq t$. Thus,

$$\mathbb{E}\left[\sum_{k \geq 1} w_k \cdot s_t(w_k)\right] \leq (1-t)c^t$$

Note that $w \geq \max(c^x, c^t)$. Thus we have

$$\mathbb{E}[w] \geq \int_t^1 c^x dx + \int_0^t c^t dx \geq \frac{c-c^t}{\ln c} + tc^t$$

In this case, we want to show

$$\alpha \mathbb{E}[\Delta_1] \leq \mathbb{E}[w].$$

This holds if

$$\alpha \left(c^t + \frac{c-c^t}{\ln c} \right) \leq \frac{c-c^t}{\ln c} + tc^t$$

Since $u_0 = c^t$, this inequality follows for all $\alpha \leq h(c)$ from Inequality 5 below.

Case 2: $u_0 < c^t$. As $w_k \leq u_0$ for $k \geq 1$, in this case, for all choices of x , all speeds $s_t(w_1) = s_t(w_2) = \dots = 0$ so $\mathbb{E}[\sum_{k \geq 1} w_k \cdot s_t(w_k)] = 0$. Thus, it suffices to show that $\mathbb{E}[\alpha(c - (c-w)s_t(w))] \leq \mathbb{E}[w]$, or equivalently that $\alpha \mathbb{E}[c(1 - s_t(w))] \leq \mathbb{E}[w - \alpha w s_t(w)]$.

Let c^z be the greatest local maximum that is smaller than c^t . If $x > z$, then $w \geq c^t$ and, thus, $s_t(w) = 1$. If $x \leq z$, then w is the first local maximum greater than c^x , while c^z is a local maximum greater than c^x . Thus, it is either equal to w or a later local maximum, which by the definition of local maximum implies that it is larger than w . Hence, $w \leq c^z < c^t$ and thus, $s_t(w) = 0$. Therefore,

$$\mathbb{E}[c(1 - s_t(w))] \leq c \int_0^z 1 dx = cz$$

We also have

$$\begin{aligned} \mathbb{E}[(1 - \alpha s_t(w))w] &\geq (1-\alpha) \int_t^1 c^x dx + (1-\alpha) \int_z^t c^t dx + \int_{\log_c u_0}^z c^x dx + \int_0^{\log_c u_0} u_0 dx \\ &= (1-\alpha) \frac{c-c^t}{\ln c} + (1-\alpha) \frac{c^t(t-z)}{\ln c} + \frac{c^z - u_0}{\ln c} + u_0 \log_c u_0 \\ &=: V(t, u_0) \end{aligned}$$

Thus it suffices to show that $\alpha \mathbb{E}[c(1 - s_t(w))] \leq V(t, u_0) \forall z \in [0, 1], t \in [z, 1], u_0 \in [1, c^z]$. For any fixed t and z , the value of u_0 minimizing $V(t, u_0)$ is $u_0 = 1$. After fixing $u_0 = 1$, we have

$$V(t, 1) = (1-\alpha) \frac{c-c^t}{\ln c} + (1-\alpha) \frac{c^t(t-z)}{\ln c} + \frac{c^z - 1}{\ln c}$$

Therefore,

$$\frac{\partial V(t, 1)}{\partial t} = (1-\alpha) \frac{c^t + tc^t/\ln c - c^t/\ln c - zc^t/\ln c}{\ln c}$$

Notice that $\partial V(t, 1)/\partial t$ is non-negative for all $t \in [z, 1]$ if $c \geq e$. Therefore, for any $c \geq e$, it suffices to consider only $u_0 = 1, t = z$ and prove that

$$\alpha \left(ct + \frac{c-c^t}{\ln c} \right) \leq \frac{c-1}{\ln c}$$

for α as large as possible. The following claim shows that this inequality holds for $\alpha \leq \frac{c-1}{c \ln c}$.

Claim 8. For $c \geq e$,

$$\frac{\frac{c-1}{\ln c}}{ct + \frac{c-c^t}{\ln c}} \geq \frac{c-1}{c \ln c} \quad \forall t \in [0, 1]$$

Proof. Consider $f(t) = ct + \frac{c-c^t}{\ln c}$. We have $f'(t) = c - \frac{c^t}{\ln^2 c} \geq 0 \quad \forall t \in [0, 1]$. Thus, the maximum $f(t)$ is achieved when $t = 1$ and $f(1) = c$. Therefore,

$$\frac{\frac{c-1}{\ln c}}{ct + \frac{c-c^t}{\ln c}} \geq \frac{c-1}{c \ln c} \quad \forall t \in [0, 1]$$

□

Thus, altogether the lemma holds for $\alpha = \min\left(\frac{c-1}{c \ln(c)}, h(c)\right)$. □

Lemma 9. For $s_0(x) \equiv 1$ and $\alpha = h(c)$, we have $\mathbb{E}[\max(\mathbf{w})] \geq \alpha \mathbb{E}[\Delta_1(\mathbf{w}, s_0)]$.

Proof. Since $s_0(w) = 1$ for all choices of x , it holds that $\mathbb{E}[c - (c-w)s_t(w)] = \mathbb{E}[w]$. Thus we need to show $\alpha \mathbb{E}[\sum_{k \geq 1} w_k] \leq (1-\alpha) \mathbb{E}[w]$. As $w \geq \max(u_0, c^x)$ we have the following lower bound for $\mathbb{E}[w]$:

$$\mathbb{E}[w] \geq \int_{\log_c u_0}^1 c^x dx + \int_0^{\log_c u_0} u_0 dx = \frac{c - u_0}{\ln c} + u_0 \log_c u_0$$

Now, to prove the inequality, we only need to bound from above $\mathbb{E}[\sum_{k \geq 1} w_k]$ for a fixed $u_0 \in [1, c]$. We can write $\mathbb{E}[\sum_{k \geq 1} w_k]$ in terms of $m_S(x)$ as follows.

$$\mathbb{E}[\sum_{k \geq 1} w_k] \leq \sum_{i=1}^{\infty} \int_0^1 m_S(c^{-i+x}) dx = \int_{-\infty}^0 m_S(c^x) dx =: B_S$$

The following claims analyze the structure of the jobs smaller than u_0 in the worst case, i.e., if S maximizes B_S .

Claim 10. For any sequence S of all local maxima where there are 2 consecutive local maxima $u_0 \geq w'_u \geq w'_{u+1}$ with $w'_u > cw'_{u+1}$ there is a sequence S' with $B_{S'}$ at least as large and no such pair of consecutive local maxima.

Proof. Add a new local maximum of weight w'_u/c to S to form S' . Notice that $m_S(x) \leq m_{S'}(x) \quad \forall x$. This argument can be repeated until there is no pair of consecutive local maxima with ratio greater than c . □

Claim 11. Consider a sequence of all local maxima S with 3 consecutive local maxima $u_0 \geq w'_u \geq w'_{u+1} \geq w'_{u+2}$ where $w'_u \leq cw'_{u+2}$. After removing w'_{u+1} , the resulting sequence S' has $B_{S'} \geq B_S$.

Proof. For all $y \notin [w'_{u+2}, w'_u]$, we have $m_{S'}(y) = m_S(y)$. For all $y \in [w'_{u+2}, w'_u]$, we have $m_{S'}(y) \geq m_S(y)$. Thus, $B_{S'} \geq B_S$. □

Claim 12. Consider a sequence of all local maxima S containing $w'_u > w'_{u+1} > \dots > w'_v$ satisfying

- (1) $u_0 \geq w'_u$
- (2) $w'_{u+i} = c^{1-i}w'_{u+1} \quad \forall 1 \leq i < v - u$
- (3) $w'_u/w'_{u+1} \leq c$, and
- (4) $w'_{v-1}/w'_v \leq c$

Then either one of the following conditions applies

- (1) $w'_u/w'_{u+1} = c$, or
- (2) $w'_{v-1}/w'_v = c$, or

there is a sequence S' with at most the same number of local maxima and $B_{S'} > B_S$.

Proof. Assume that none of the conditions applies. We will show it is possible to move the jobs to form a sequence S' with $B_{S'} \geq B_S$.

We consider the effect of moving $z = w'_{u+1}$ while maintaining the relation $w'_{u+j} = c^{1-j}w'_{u+1} \quad \forall 1 \leq j < v - u$. We have

$$\begin{aligned}
B_S &= \int_{-\infty}^{\log_c w'_v} m_S(x) dx + \int_{\log_c w'_u}^1 m_S(x) dx \\
&\quad + \int_{\log_c w'_v}^{\log_c w'_u} m_S(x) dx \\
&= T + \int_{\log_c z}^{\log_c w'_u} w'_u dx + \sum_{j=0}^{v-u-3} \int_{\log_c c^{-j-1}z}^{\log_c c^{-j}z} c^{-j} z dx \\
&\quad + \int_{\log_c w'_v}^{\log_c c^{u+2-v}z} c^{u+2-v} z dx
\end{aligned}$$

where T is a function that does not depend on z . Furthermore, we have

$$\begin{aligned}
\frac{\partial B_S}{\partial z} &= -w'_u \frac{1}{z \ln c} + \frac{(1 - c^{u+2-v})}{1 - c^{-1}} \\
&\quad + c^{u+2-v} \left(u + 2 - v + \frac{1 + \ln z}{\ln c} - \frac{\ln w'_v}{\ln c} \right)
\end{aligned}$$

Notice that $\partial B_S / \partial z$ is monotonically increasing so the maximum of B_S is achieved at an extreme point, which is either w'_u , $c^{v-u-2}w'_v$, w'_u/c , $c^{v-u-1}w'_v$. When $z = w'_u$ or $z = c^{v-u-2}w'_v$, there are 2 jobs of the same weight and we can remove one without changing B_S . If $z = w'_u/c$, the first condition in the lemma holds. If $z = c^{v-u-1}w'_v$, then the second condition ($w'_{v-1}/w'_v = c$) holds as $z = c^{v-u}w'_{v-1}$ by the assumptions of the lemma. Thus, the conclusion follows from an inductive argument on the number of jobs. \square

By the above claims, the only sequences we need to consider to prove Lemma 9 are of the form

$$u_0, u_0/c, \dots, u_0/c^m, v/c^{m+1}, v/c^{m+2}, \dots$$

where $u_0 \leq v$, i.e., all pairs of consecutive jobs have ratio exactly c except for possibly one pair.

Thus it holds that

$$\begin{aligned}
B_S &= \int_{\log_c u_0}^1 u_0 dx + \sum_{i=1}^m \int_0^1 u_0 c^{-i} dx \\
&\quad + \int_{\log_c v}^{\log_c u_0} u_0 c^{-m} dx + \sum_{i=m+1}^{\infty} \int_0^1 v c^{-i} dx \\
&= (1 - \log_c u_0)u_0 + \frac{(1 - c^{-m})u_0}{c - 1} \\
&\quad + u_0 c^{-m} \log_c(u_0/v) + \frac{v c^{-m}}{c - 1}
\end{aligned}$$

Notice that $\frac{\partial B_S}{\partial v} = -\frac{u_0 c^{-m}}{v \ln c} + \frac{c^{-m}}{c-1}$ is monotonically increasing so the choice of v maximizing B_S is either $v = u_0$ or $v = c$. The following lemma proves that the value of B_S when $v = u_0$ is larger than the value of B_S when $v = c$. Thus B_S is maximized when $v = u_0$, i.e., *all* pairs of consecutive jobs less than u_0 have ratio exactly c .

Claim 13.

$$\frac{u_0}{c-1} \geq u_0 \log_c(u_0/c) + \frac{c}{c-1} \quad \forall u_0 \in [1, c]$$

Proof. Let $f(x) = \frac{x}{c-1} - x \log_c(x/c) + \frac{c}{c-1}$. We have $f(1) \geq 0$ and $f(c) \geq 0$. Also notice that

$$f'(x) = \frac{1}{c-1} - \frac{\ln x + 1}{\ln c}$$

is monotonically decreasing in x so the minimum of $f(x)$ is achieved at the extreme points. In other words, $f(x) \geq 0 \quad \forall x \in [1, c]$. \square

It follows that $B_S = (1 - \log_c u_0)u_0 + \frac{u_0}{c-1}$. Thus, we need to show for $\alpha = h(c)$

$$\alpha \left((1 - \log_c u_0)u_0 + \frac{u_0}{c-1} \right) \leq (1 - \alpha) \left(\frac{c - u_0}{\ln c} + u_0 \log_c u_0 \right) \quad (4)$$

or equivalently

$$\alpha \left(\frac{u_0 c}{c-1} + \frac{c - u_0}{\ln c} \right) \leq \frac{c - u_0 + u_0 \ln u_0}{\ln c} \quad (5)$$

Let $\beta = \frac{c \ln c}{c-1} - 1$. We can rewrite the above inequality as

$$\alpha(\beta u_0 + c) \leq c - u_0 + u_0 \ln u_0$$

Claim 14. For $\alpha = h(c)$,

$$\alpha(\beta u_0 + c) \leq c - u_0 + u_0 \ln u_0$$

Proof. Let $f(u_0) = c - u_0 + u_0 \ln u_0$ and $g(u_0) = \beta u_0 + c$. We need to show that $f(u_0)/g(u_0) \geq \alpha$. Note that $(f(u_0)/g(u_0))' = 0$ if and only if

$$\begin{aligned}
0 &= f'(u_0)g(u_0) - g'(u_0)f(u_0) \\
&= \ln u_0(\beta u_0 + c) - \beta(c - u_0 + u_0 \ln u_0) \\
&= c \ln u_0 - \beta c + \beta u_0
\end{aligned}$$

The above expression is monotonically increasing in u_0 . We show next that it has a root. Thus, $f(u_0)/g(u_0)$ is minimized at this root. Let $y = u_0\beta/c$. The above equation is equivalent to

$$\beta + \ln(\beta/c) = \ln y + y$$

The solution of this equation is $y = W\left(\frac{\beta e^\beta}{c}\right)$, where W is the Lambert W function. Thus, $u_0 = \frac{c}{\beta}W\left(\frac{\beta e^\beta}{c}\right)$. Substituting the identity $\ln u_0 = \beta - \beta u_0/c$ into $f(u_0)/g(u_0)$, we get

$$\begin{aligned} \frac{f(u_0)}{g(u_0)} &= \frac{c - u_0 + u_0(\beta - \beta u_0/c)}{\beta u_0 + c} = 1 - \frac{u_0}{c} \\ &= 1 - \frac{1}{\beta}W\left(\frac{\beta e^\beta}{c}\right) = h(c) \end{aligned}$$

□

Thus, inequality (5) holds for $\alpha = h(c)$. This completes the proof of Lemma 9. □

4 Upper bound for deterministic algorithms

To prove an upper bound of a , we construct an instance such that any deterministic algorithm has competitive ratio at most a for some prefix of the request sequence. The instance has one *fast* machine of speed $r > 1$ and n *slow* machines of speed 1. The request sequence has non-decreasing job sizes satisfying a certain oscillatory recurrence relation.

Theorem 15. *The competitive ratio of any deterministic algorithm is at most $(\sqrt{5} - 1)/2 + \epsilon \approx 0.618034 + \epsilon$ for any $\epsilon > 0$.*

Proof. Our construction to establish this bound uses parameters $r > 1$ and $a < 1$ that we will fix later. We construct an instance such that any deterministic algorithm has competitive ratio at most a for some prefix of the request sequence. The instance has one *fast* machine of speed $r > 1$ and n *slow* machines of speed 1. The request sequence has non-decreasing job sizes: $1 = w_0 \leq w_1 \leq \dots \leq w_{n-1} \leq w_n$.

The instance will satisfy the following properties:

$$a \cdot r \geq 1 \tag{6}$$

$$\forall k = 1 \dots n, \quad a \cdot \left(r \cdot w_k + \sum_{j=0}^{i-1} w_j \right) \geq r \cdot w_{k-1} + w_k \tag{7}$$

$$a \cdot \left(r \cdot w_n + \sum_{j=0}^{n-1} w_j \right) \geq r \cdot w_n \tag{8}$$

Lemma 16. *Properties (6) - (8) imply that any deterministic algorithm has competitive ratio at most a .*

Proof. Consider the action of a deterministic algorithm A on the instance we construct. If w_0 is placed on a slow machine, (6) implies that the competitive ratio is at most a . For $k = 1, \dots, n$, note

that the optimum value for the prefix of the request sequence ending at w_k is $r \cdot w_k + \sum_{j=0}^{i-1} w_j$. We consider two cases:

Case 1: Suppose A places some job on a slow machine. Let k be the first job placed on a slow machine. The the value of A for the prefix ending at w_i is $r \cdot w_{k-1} + w_k$. Now (7) implies that the competitive ratio is at most a .

Case 2: All jobs are placed on the fast machine. Then, once the entire sequence is processed upto w_n , (8) implies that the competitive ratio is at most a . \square

In order to produce the instance, we focus on satisfying the following properties instead, which imply the previous properties (6) - (8).

$$a \cdot r \geq 1 \tag{9}$$

$$a \cdot (r \cdot w_1 + w_0) = r \cdot w_0 + w_1 \tag{10}$$

$$\begin{aligned} \forall i = 2 \dots n, \quad & a \cdot (r \cdot w_k - r \cdot w_{k-1} + w_{k-1}) \\ & = r \cdot w_{k-1} + w_k - (r \cdot w_{k-2} + w_{k-1}) \end{aligned} \tag{11}$$

$$r \cdot w_{n-1} + w_n \geq r \cdot w_n \tag{12}$$

Note that (10) and (11) together imply (7) with equality. Further, (12) and (11) for $k = n$ implies (8).

Next, we rewrite (11) as a recurrence relation for the sequence $\{w_k\}$.

$$(a \cdot r - 1)w_k - (a + 1)(r - 1)w_{k-1} + r \cdot w_{k-2} = 0 \tag{13}$$

Note that the initial conditions are $w_0 = 1$ and from (10), $w_1 = (r - a)/(a \cdot r - 1)$.

Let $\delta > 0$ be a sufficiently small constant.

Lemma 17.

$$\text{For } a = \frac{1 + \sqrt{5 + 12\delta + 4\delta^2}}{3 + \sqrt{5} + 2\delta} \text{ and } r = \frac{1 + \sqrt{5 + 12\delta + 4\delta^2}}{3 - \sqrt{5} + 2\delta},$$

the roots of the characteristic equation of recurrence relation (13) are $\frac{1 + \sqrt{5}}{2} \pm i\sqrt{\delta}$

Proof. We will choose parameters a, r such that the roots of the characteristic equation of the recurrence relation (13) are of the form $z \pm i\sqrt{\delta}$ for some small δ . Note that i here is the complex square root of -1 . The reason for this choice of roots will become clear later.

First, we derive relationships between a, r and z, δ . Using the standard formula for the roots of a quadratic equation, we get

$$\frac{(a + 1)(r - 1)}{2(r \cdot a - 1)} = z \tag{14}$$

$$\left(\frac{(a + 1)(r - 1)}{2(r \cdot a - 1)} \right)^2 - \frac{r}{r \cdot a - 1} = -\delta \tag{15}$$

$$(15) \Rightarrow \frac{r}{r \cdot a - 1} = z^2 + \delta \tag{16}$$

$$(14) \Rightarrow \frac{r - a}{r \cdot a - 1} = 2z - 1 \tag{17}$$

$$(16) - (17) \Rightarrow \frac{a}{r \cdot a - 1} = (z - 1)^2 + \delta \tag{18}$$

An easy calculation shows that the system of equations

$$\frac{r}{r \cdot a - 1} = x, \quad \frac{a}{r \cdot a - 1} = y$$

has the solution $a = \frac{1 + \sqrt{1 + 4xy}}{2x}$, $r = \frac{1 + \sqrt{1 + 4xy}}{2y}$. Substituting $x = z^2 + \delta$ and $y = (z - 1)^2 + \delta$, we get

$$a = \frac{1 + \sqrt{1 + 4(z^2 + \delta)((z - 1)^2 + \delta)}}{2(z^2 + \delta)}$$

$$r = \frac{1 + \sqrt{1 + 4(z^2 + \delta)((z - 1)^2 + \delta)}}{2((z - 1)^2 + \delta)}$$

Recall that a is the upper bound on the competitive ratio that we establish and we would like to minimize a to get the best bound possible. In fact, a is minimized for $z = (1 + \sqrt{5})/2$ and $\delta = 0$ for which we get $a = (\sqrt{5} - 1)/2$. In fact, our construction will need $\delta > 0$, but we fix $z = (1 + \sqrt{5})/2$ to minimize the upper bound a we obtain from this argument. For this value of z we get

$$a = \frac{1 + \sqrt{5 + 12\delta + 4\delta^2}}{3 + \sqrt{5} + 2\delta} \quad r = \frac{1 + \sqrt{5 + 12\delta + 4\delta^2}}{3 - \sqrt{5} + 2\delta}. \quad \square$$

Note that the upper bound a on the competitive ratio is of the form $(\sqrt{5} - 1)/2 + \epsilon$ with ϵ a suitable function of δ (and $\epsilon \rightarrow 0$ as $\delta \rightarrow 0$).

Claim 18. For a and r chosen as in Lemma 17, the term w_k of the solution to the recurrence relation is given by

$$w_k = \left(\frac{1}{2} - \frac{\sqrt{5} - 1}{4\sqrt{\delta}} \cdot i \right) \left(\frac{1 + \sqrt{5}}{2} + i\sqrt{\delta} \right)^k$$

$$+ \left(\frac{1}{2} + \frac{\sqrt{5} - 1}{4\sqrt{\delta}} \cdot i \right) \left(\frac{1 + \sqrt{5}}{2} - i\sqrt{\delta} \right)^k \quad (19)$$

Proof. A general term w_k of the sequence is given by the following expression:

$$(b + c \cdot i) \left(\frac{1 + \sqrt{5}}{2} + i\sqrt{\delta} \right)^k + (d + e \cdot i) \left(\frac{1 + \sqrt{5}}{2} - i\sqrt{\delta} \right)^k$$

where b, c, d, e are real numbers to be determined. Recall the initial conditions $w_0 = 1$ and $w_1 = (r - a)/(r \cdot a - 1) = 2z - 1 = \sqrt{5}$. This gives the following equations:

$$(b + c \cdot i) \cdot 1 + (d + e \cdot i) \cdot 1 = 1$$

$$(b + c \cdot i) \left(\frac{1 + \sqrt{5}}{2} + i\sqrt{\delta} \right)$$

$$+ (d + e \cdot i) \left(\frac{1 + \sqrt{5}}{2} - i\sqrt{\delta} \right) = \sqrt{5}$$

Solving these, we get $b = d = 1/2$, $c = -(\sqrt{5} - 1)/4\sqrt{\delta}$ and $e = (\sqrt{5} - 1)/4\sqrt{\delta}$. Hence the term w_k of the sequence is given by

$$\begin{aligned} & \left(\frac{1}{2} - \frac{\sqrt{5}-1}{4\sqrt{\delta}} \cdot i \right) \left(\frac{1+\sqrt{5}}{2} + i\sqrt{\delta} \right)^k \\ & + \left(\frac{1}{2} + \frac{\sqrt{5}-1}{4\sqrt{\delta}} \cdot i \right) \left(\frac{1+\sqrt{5}}{2} - i\sqrt{\delta} \right)^k \end{aligned} \quad \square$$

Lemma 19. For a and r chosen as in Lemma 17 and $\delta > 0$ sufficiently small, the solution w_k of the recurrence relation satisfies the conditions (9)-(12).

Proof. The choice of recurrence relation (13) ensures that (11) is satisfied, and the initial condition for w_1 ensures that (10) is satisfied. For $\delta = 0$, $a \cdot r = (6 + 2\sqrt{5})/4$. Hence for δ sufficiently small, $a \cdot r \geq 1$ and hence (9) is satisfied.

Rewriting condition (12), we need to show that $w_n/w_{n-1} \leq r/(r-1)$. Note that $r = \frac{1+\sqrt{5}}{3-\sqrt{5}} + O(\delta)$ and $r/(r-1) \geq (3 + \sqrt{5})/4 - O(\delta) \geq 1.309$. Examining the solution (19) of the recurrence, we see that w_k is twice the real part of the first term in (19). In other words, w_k is of the form

$$w_k = 2\Re \left((r_1 e^{i\Phi_1})(r_2 e^{i\Phi_2})^k \right) \quad (20)$$

where $\Phi_1 = \tan^{-1} \left(-\frac{\sqrt{5}-1}{2\sqrt{\delta}} \right)$, $\Phi_2 = \tan^{-1} \left(\frac{2\sqrt{\delta}}{1+\sqrt{5}} \right)$, and $r_2 = \frac{1+\sqrt{5}}{2} + O(\delta)$. As $\delta \rightarrow 0$,

$$\Phi_1 = -\frac{\pi}{2} + \frac{2\sqrt{\delta}}{\sqrt{5}-1} - O(\delta^{3/2})$$

and $\Phi_2 = \frac{2\sqrt{\delta}}{1+\sqrt{5}} - E$ for $E = O(\delta^{3/2})$. Note that

$$w_k = 2r_1 (r_2)^k \cos(\Phi_1 + k\Phi_2) \quad (21)$$

$$\begin{aligned} & = 2r_1 \left(\frac{1+\sqrt{5}}{2} + O(\delta) \right)^k \\ & \cos \left(-\frac{\pi}{2} + \frac{2\sqrt{\delta}}{\sqrt{5}-1} + k \frac{2\sqrt{\delta}}{1+\sqrt{5}} - O(\delta^{3/2}) - kE \right) \end{aligned} \quad (22)$$

$$\begin{aligned} \frac{w_n}{w_{n-1}} & = \left(\frac{1+\sqrt{5}}{2} + O(\delta) \right) \\ & \frac{\cos \left(-\frac{\pi}{2} + \frac{2\sqrt{\delta}}{\sqrt{5}-1} + n \frac{2\sqrt{\delta}}{1+\sqrt{5}} - O(\delta^{3/2}) - nE \right)}{\cos \left(-\frac{\pi}{2} + \frac{2\sqrt{\delta}}{\sqrt{5}-1} + (n-1) \frac{2\sqrt{\delta}}{1+\sqrt{5}} - O(\delta^{3/2}) - (n-1)E \right)} \end{aligned} \quad (23)$$

Note that we hope to show $w_n/w_{n-1} \leq 1.309 \leq (3 + \sqrt{5})/4 - O(\delta)$. The term $(1 + \sqrt{5})/2 + O(\delta) < 1.619$, so it is critical that the ratio of cosines be small enough to give us the condition we want (but not too small to ensure that $w_n/w_{n-1} \geq 1$).

First we show that the smallest ratio of consecutive terms is w_n/w_{n-1} . Consider $f(x) = \frac{\cos(x+\Phi_2)}{\cos(x)} = \cos(\Phi_2) - \tan(x) \sin(\Phi_2)$.

We have $f'(x) = -\frac{\sin(\Phi_2)}{\cos^2(x)} < 0$ so the ratio gets smaller as n increases, as long as $\Phi_1 + (n-1)\Phi_2 < \pi/2$.

Next we show there exists n such that $13/20 \leq \frac{\cos(\Phi_1+n\Phi_2)}{\cos(\Phi_1+(n-1)\Phi_2)} \leq 3/4$. For $n = 1$, the ratio is greater than 1 for sufficiently small δ . For the smallest n such that $\Phi_1 + n\Phi_2 \geq \pi/2$, the ratio is smaller than 0. Furthermore, if $\frac{\cos(\Phi_1+n\Phi_2)}{\cos(\Phi_1+(n-1)\Phi_2)} > 3/4$ then

$$\begin{aligned} & \frac{\cos(\Phi_1 + (n + 1)\Phi_2)}{\cos(\Phi_1 + n\Phi_2)} \\ &= \frac{2 \cos(\Phi_1 + n\Phi_2) \cos(\Phi_2) - \cos(\Phi_1 + (n - 1)\Phi_2)}{\cos(\Phi_1 + n\Phi_2)} \\ &\geq 2 \cos(\Phi_2) - \frac{4}{3} \geq 13/20 \quad \text{if } \cos(\Phi_2) \geq \frac{119}{120} \end{aligned}$$

This is true for sufficiently small δ . □

□

5 Upper bound for randomized algorithms

To establish the bound for randomized algorithms, we use Yao's principle and show an upper bound on the expected competitive ratio of any deterministic algorithm on a distribution of instances. The construction uses one fast machine of speed 1 and n slow machines of speed $1/4$. The request sequence has non-decreasing sizes 2^i . The prefix of this sequence ending with size 2^i is presented to the algorithm with probability $c/2^i$, where c is a normalizing constant. We show that the best algorithm for this sequence achieves at most $cn + 1$ while the optimal algorithm achieves roughly $5nc/4$.

Theorem 20. *The competitive ratio of any randomized algorithm against an oblivious adversary is at most $0.8 + \epsilon$ for any $\epsilon > 0$.*

Proof. In order to establish the bound for randomized algorithms, we use Yao's principle and show an upper bound on the expected competitive ratio of any deterministic algorithm on a distribution of instances. The construction uses a set of machines with one *fast* machine of speed 1 and n *slow* machines of speed $1/4$. The request sequence has non-decreasing sizes $w_i = 2^i$ for $i = 1, \dots, n$. Our construction uses a probability distribution over prefixes of this sequence: the prefix ending at w_i is presented to the algorithm with probability $p_i = c/2^i$, where the normalizing constant $c = 1/(1 - 1/2^n)$.

Let OPT_i denote the optimal solution for the length i prefix of the input. It is easy to see that OPT_i places w_i on the fastest machine and jobs w_1, \dots, w_{i-1} on the slow machines. Hence $OPT_i = 2^i + (2^i - 1)/4 = (5/4)2^i - 1/4$. We will compute the expected value of the optimal solution for the distribution on inputs specified above.

$$\begin{aligned} \sum_{i=1}^n \frac{c}{2^i} OPT_i &= c \sum_{i=1}^n \frac{(5/4)2^i - 1/4}{2^i} = c \sum_{i=1}^n \frac{5}{4} - \frac{1}{4 \cdot 2^i} \\ &= \left(1 - \frac{1}{2^n}\right)^{-1} \left(\frac{5n}{4} - \frac{1}{4}(1 - 1/2^n)\right) \end{aligned}$$

Next we compute the expected value of the best deterministic algorithm on this distribution. Notice that for the setting we specified, a deterministic algorithm is completely specified by n choices of whether to put w_i on the fast machine or to put it on an unoccupied slow machine. Let c_i be the indicator variable of whether the algorithm puts w_i on the fast machine. Let $a_1 \leq \dots \leq a_k$ be the indices of the jobs the algorithm puts on the fast machine. Let m_i be the maximum a_l such that $a_l \leq i$. In other words, m_i is the index of the largest job on the fastest machine if the sequence of jobs ends at the i th job. For notational convenience, assume that we already put a job of size $w_0 = 0$ on the fast machine and $a_0 = 0$. The expected value of the algorithm is

$$\begin{aligned}
& \sum_{i=1}^n p_i \left(w_{m_i} + \sum_{j: j \notin \{a_1, \dots, a_k\} \wedge j \leq i} \frac{w_j}{4} \right) \\
&= \sum_{i=1}^n p_i \left(\sum_{j: a_j \leq i} (w_{a_j} - w_{a_{j-1}}) + \sum_{j: \bar{c}_j \wedge j \leq i} \frac{w_j}{4} \right) \\
&= \sum_{j=1}^n \left(c_j (w_j - w_{m_{j-1}}) + (1 - c_j) \frac{w_j}{4} \right) \sum_{i \geq j} p_i \\
&= \sum_{j=1}^n x_j
\end{aligned}$$

where $x_j = \left(c_j (w_j - w_{m_{j-1}}) + (1 - c_j) \frac{w_j}{4} \right) \sum_{i \geq j} p_i$. Let $f(t, j)$ be the maximum value of $\sum_{l=1}^t x_l$ over all choices of c_1, \dots, c_t with the restriction that $m_t = j$. First, if $j < t$ then $c_{j+1} = \dots = c_t = 0$ and we have

$$f(t, j) = f(t-1, j) + \frac{w_t}{4} \sum_{i \geq t} p_i \leq f(t-1, j) + \frac{c}{2} \quad (24)$$

Next, if $j = t$ then

$$\begin{aligned}
f(t, t) &= \max_{m_{t-1}} \left(f(t-1, m_{t-1}) + (w_t - w_{m_{t-1}}) \sum_{i \geq t} p_i \right) \\
&\leq \max_{m_{t-1}} \left(f(m_{t-1}, m_{t-1}) + (w_t - w_{m_{t-1}}) \frac{2c}{2^t} \right. \\
&\quad \left. + \frac{(t-1 - m_{t-1})c}{2} \right)
\end{aligned} \quad (25)$$

The above inequality follows from $\sum_{i \geq j} p_i \leq \frac{2c}{2^j}$ and Equation 24.

We now prove by induction that $f(i, i) \leq ci + 1$. The base cases $f(0, 0) = 0$, $f(1, 1) = 2$ and $f(2, 2) = 4 - c$ are obvious. Assume that the claim holds up to $i = t-1$ and we want to prove it for $i = t$. Consider 3 cases for m_{t-1} . First, if $m_{t-1} = t-1$ then $f(t, t) = f(t-1, t-1) + (2^t - 2^{t-1}) \frac{2c}{2^t} \leq c(t-1) + 1 + c = ct + 1$. Next, if $m_{t-1} = t-2$ then $f(t, t) = f(t-2, t-2) + (2^t - 2^{t-2}) \frac{2c}{2^t} + c/2 \leq ct + 1$. Lastly, if $m_{t-1} \leq t-3$ then

$$\begin{aligned}
f(t, t) &\leq f(m_{t-1}, m_{t-1}) + w_t \frac{2c}{2^t} + (t-1 - m_{t-1})c/2 \\
&\leq m_{t-1}c + 1 + 2c + (t-1 - m_{t-1})c/2 \leq ct + 1
\end{aligned}$$

Thus, we have proved the inductive case. By Equation 24 and the fact that $f(i, i) \leq ci + 1$, we have $f(n, j) \leq cn + 1 \forall j$. However, $\max_j f(n, j)$ is exactly the expected value of the best algorithm so the competitive ratio of any randomized algorithm on the specified instance is at most $\frac{cn+1}{5nc/4-1/4} \rightarrow 4/5$ as n goes to infinity. □

5.0.1 Acknowledgments.

MC was supported by NSF awards CCF 0832797, AF 0916218 and a Google research award. MH's support: The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement no. 340506 and the Austrian Science Fund (FWF) grant P23499-N23. HN was supported by NSF awards CCF 0832797, AF 0916218, a Google research award, and a Gordon Wu fellowship.

References

- [1] G. Aggarwal, G. Goel, C. Karande, and A. Mehta. Online vertex-weighted bipartite matching and single-bid budgeted allocations. In *SODA*, pages 1253–1264, 2011.
- [2] S. Agrawal, Z. Wang, and Y. Ye. A dynamic near-optimal algorithm for online linear programming. *CoRR*, abs/0911.2974, 2009.
- [3] B. V. Ashwinkumar. Buyback problem-approximate matroid intersection with cancellation costs. *Automata, Languages and Programming*, pages 379–390, 2011.
- [4] B. V. Ashwinkumar and R. Kleinberg. Randomized online algorithms for the buyback problem. In *WINE*, pages 529–536, 2009.
- [5] J. Aspnes, Y. Azar, A. Fiat, S. Plotkin, and O. Waarts. On-line routing of virtual circuits with applications to load balancing and machine scheduling. *Journal of the ACM (JACM)*, 44(3):486–504, 1997.
- [6] Y. Azar. On-line load balancing. *Online Algorithms*, pages 178–195, 1998.
- [7] M. Babaioff, J. D. Hartline, and R. D. Kleinberg. Selling ad campaigns: online algorithms with cancellations. In *EC*, pages 61–70, 2009.
- [8] A. Beck and D. Newman. Yet more on the linear search problem. *Israel journal of mathematics*, 8(4):419–429, 1970.
- [9] P. Berman, M. Charikar, and M. Karpinski. On-line load balancing for related machines. *Journal of Algorithms*, 35(1):108–121, 2000.
- [10] B. E. Birnbaum and C. Mathieu. On-line bipartite matching made simple. *SIGACT News*, 39(1):80–87, 2008.
- [11] N. Buchbinder, K. Jain, and J. Naor. Online primal-dual algorithms for maximizing ad-auctions revenue. In *ESA*, pages 253–264, 2007.

- [12] S. Chakrabarti, C. Phillips, A. Schulz, D. Shmoys, C. Stein, and J. Wein. Improved scheduling algorithms for minsum criteria. *Automata, Languages and Programming*, pages 646–657, 1996.
- [13] F. Constantin, J. Feldman, S. Muthukrishnan, and M. Pál. An online mechanism for ad slot reservations with cancellations. In *SODA*, pages 1265–1274, 2009.
- [14] N. R. Devanur and T. P. Hayes. The adwords problem: Online keyword matching with budgeted bidders under random permutations. In *EC*, pages 71–78, 2009.
- [15] N. R. Devanur, K. Jain, and R. Kleinberg. Randomized primal-dual analysis of ranking for online bipartite matching. In *SODA*, 2013.
- [16] L. Epstein, A. Levin, D. Segev, and O. Weimann. Improved bounds for online preemptive matching. *CoRR*, abs/1207.1788, 2012.
- [17] L. Epstein and J. Sgall. A lower bound for on-line scheduling on uniformly related machines. *Operations Research Letters*, 26(1):17–22, 2000.
- [18] J. Feldman, N. Korula, V. S. Mirrokni, S. Muthukrishnan, and M. Pál. Online ad assignment with free disposal. In *WINE*, pages 374–385, 2009.
- [19] J. Feldman, A. Mehta, V. Mirrokni, and S. Muthukrishnan. Online stochastic matching: Beating $1-1/e$. In *Foundations of Computer Science, 2009. FOCS'09. 50th Annual IEEE Symposium on*, pages 117–126. IEEE, 2009.
- [20] S. Gal. Search games, volume 149 of mathematics in science and engineering, 1980.
- [21] G. Goel and A. Mehta. Online budgeted matching in random input models with applications to adwords. In *SODA*, pages 982–991, 2008.
- [22] G. Goel and P. Tripathi. Matching with our eyes closed. In *FOCS*, pages 718–727, 2012.
- [23] M. Goemans and J. Kleinberg. An improved approximation ratio for the minimum latency problem. *Mathematical Programming*, 82(1):111–124, 1998.
- [24] B. Haeupler, V. S. Mirrokni, and M. Zadimoghaddam. Online stochastic weighted matching: Improved approximation algorithms. In *WINE*, pages 170–181, 2011.
- [25] C. Karande, A. Mehta, and P. Tripathi. Online bipartite matching with unknown distributions. In *STOC*, pages 587–596, 2011.
- [26] R. M. Karp, U. V. Vazirani, and V. V. Vazirani. An optimal algorithm for on-line bipartite matching. In *STOC*, pages 352–358, 1990.
- [27] M. Mahdian and Q. Yan. Online bipartite matching with random arrivals: an approach based on strongly factor-revealing LPs. In *STOC*, pages 597–606, 2011.
- [28] V. Manshadi, S. Gharan, and A. Saberi. Online stochastic matching: Online actions based on offline statistics. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1285–1294. SIAM, 2011.

- [29] A. McGregor. Finding graph matchings in data streams. *Approximation, Randomization and Combinatorial Optimization. Algorithms and Techniques*, pages 611–612, 2005.
- [30] A. Mehta and D. Panigrahi. Online matching with stochastic rewards. In *FOCS*, pages 728–737, 2012.
- [31] A. Mehta, A. Saberi, U. V. Vazirani, and V. V. Vazirani. Adwords and generalized on-line matching. In *FOCS*, pages 264–273, 2005.
- [32] M. Poloczek and M. Szegedy. Randomized greedy algorithms for the maximum matching problem with new analysis. In *FOCS*, pages 708–717, 2012.