

## TIGHT BOUNDS FOR ONLINE VECTOR SCHEDULING\*

SUNGJIN IM<sup>†</sup>, NATHANIEL KELL<sup>‡</sup>, JANARDHAN KULKARNI<sup>§</sup>, AND DEBMALYA PANIGRAHI<sup>‡</sup>

**Abstract.** Modern data centers face a key challenge of effectively serving user requests that arrive online. Such requests are inherently *multidimensional* and characterized by demand vectors over multiple resources such as processor cycles, storage space, and network bandwidth. Typically, different resources require different objectives to be optimized, and  $L_r$  norms of loads are among the most popular objectives considered. Furthermore, the server clusters are also often heterogeneous making the scheduling problem more challenging. To address these problems, we consider the *online vector scheduling* problem in this paper. Introduced by Chekuri and Khanna in 2006, vector scheduling is a generalization of classical load balancing, where every job has a vector load instead of a scalar load. The scalar problem, introduced by Graham in 1966, and its many variants (identical and unrelated machines, makespan and  $L_r$  norm optimization, offline and online jobs, etc.) have been extensively studied over the last 50 years. In this paper, we resolve the online complexity of the vector scheduling problem and its important generalizations—for all  $L_r$  norms and in both the identical and unrelated machines settings. For an instance with  $m$  machines and  $d$  dimensions, our main results are: For *identical machines*, we show that the optimal competitive ratio is  $\Theta(\log d / \log \log d)$  by giving an online lower bound and an algorithm with an asymptotically matching competitive ratio. The lower bound is technically challenging, and is obtained via an online lower bound for the minimum monochromatic clique problem using a novel online coloring game and randomized coding scheme. Our techniques also extend to asymptotically tight upper and lower bounds for general  $L_r$  norms. For *unrelated machines*, we show that the optimal competitive ratio is  $\Theta(\log m + \log d)$  by giving an online lower bound that matches a previously known upper bound. Unlike identical machines, however, extending these results, particularly the upper bound, to general  $L_r$  norms requires new ideas. In particular, we use a carefully constructed potential function that balances the individual  $L_r$  objectives with the overall (convexified) min-max objective to guide the online algorithm and track the changes in potential to bound the competitive ratio.

**Key words.** multidimensional, vector scheduling, makespan, identical machines, unrelated machines, various norms, competitive analysis

**AMS subject classification.** 68W27

**DOI.** 10.1137/17M111835X

**1. Introduction.** A key algorithmic challenge in modern data centers is the scheduling of online resource requests on the available hardware. Such requests are inherently *multidimensional* and simultaneously ask for multiple resources such as processor cycles, network bandwidth, and storage space [23, 27, 36] (see also multidimensional load balancing in virtualization [28, 33]). In addition to the multidimensionality of resource requests, another challenge is the *heterogeneity* of server clusters because

---

\*Received by the editors February 27, 2017; accepted for publication (in revised form) September 28, 2018; published electronically January 15, 2019. A preliminary version of this paper appeared in the Proceedings of the 56th IEEE Symposium on Foundations of Computer Science (FOCS) 2015 [32].

<http://www.siam.org/journals/sicomp/48-1/M111835.html>

**Funding:** The first author was supported in part by NSF Awards CCF-1409130 and CCF-1008065. The second and fourth authors were supported in part by NSF Awards CCF-1527084 and CCF-1535972, an NSF CAREER Award, a Google Faculty Research Award, and a Yahoo FREP Award. The third author was supported in part by NSF Awards CCF-0745761, CCF-1008065, and CCF-1348696.

<sup>†</sup>Electrical Engineering and Computer Science, University of California at Merced, Merced, CA 95343 (sim3@ucmerced.edu).

<sup>‡</sup>Department of Computer Science, Duke University, Durham, NC 27708 (Kell@cs.duke.edu, debmalya@cs.duke.edu).

<sup>§</sup>Microsoft Research, Redmond, WA 98052 (jakul@microsoft.com).

of incremental hardware deployment and the use of dedicated specialized hardware for particular tasks [1, 24, 46]. As a third source of nonuniformity, the objective of the load balancing exercise is often defined by the application at hand and the resource being allocated. In addition to the traditional goals of minimizing maximum ( $L_\infty$  norm) and total ( $L_1$  norm) machine loads, various intermediate  $L_r$  norms<sup>1</sup> are also important for specific applications. For example, the  $L_2$  norm of machine loads is suitable for disk storage [17, 20] while the  $L_r$  norm for  $r$  between 2 and 3 is used for modeling energy consumption [3, 40, 45].

In the algorithmic literature, the (single dimensional) load balancing problem, also called list scheduling, has a long history since the pioneering work of Graham in 1966 [26]. However, the multidimensional problem, introduced by Chekuri and Khanna [18] and called *vector scheduling* (vs), remains less understood. In the simplest version of this problem, each job has a vector load and the goal is to assign the jobs to machines so as to minimize the maximum machine load over all dimensions. As an example of our limited understanding of this problem, we note that the approximation complexity of this most basic version is not resolved yet—the current best approximation factor is  $O(\log d / \log \log d)$  (e.g., [30]), where  $d$  is the number of dimensions, while only an  $\omega(1)$  lower bound is known [18]. In this paper, we consider the online version of this problem, i.e., where the jobs appear in a sequence and have to be assigned irrevocably to a machine on arrival. Note that this is the most common scenario in the data center applications that we described earlier, and in other real world settings. In addition to the basic setting described above, we also consider more general scenarios to capture the practical challenges that we outlined. In particular, we consider this problem in both the identical and unrelated machines settings, the latter capturing the nonuniformity of servers. Furthermore, we also consider all  $L_r$  norm objectives of machine loads in addition to the makespan ( $L_\infty$ ) objective. In this paper, we completely resolve the online complexity of all these variants of the vector scheduling problem.

Formally, there are  $n$  jobs (denoted  $J$ ) that arrive online and must be immediately and irrevocably assigned on arrival to one among a fixed set of  $m$  machines (denoted  $M$ ). We denote the  $d$ -dimensional load vector of job  $j$  on machine  $i$  by  $p_{i,j} = \langle p_{i,j}(k) : k \in [d] \rangle$ , which is revealed on its online arrival. For identical machines, the load of job  $j$  in dimension  $k$  is identical for all machines  $i$ , and we denote it by  $p_j(k)$ . Let us denote the assignment function of jobs to machines by  $f : J \rightarrow M$ . An assignment  $f$  produces a *load* of  $\Lambda_i(k) = \sum_{j:f(j)=i} p_{i,j}(k)$  in dimension  $k$  of machine  $i$ ; we succinctly denote the machine loads in dimension  $k$  by an  $m$ -dimensional vector  $\Lambda(k)$ . (Note that for the scalar problem, there is only one such machine load vector.)

**The makespan norm.** The objective for the vs problem for the makespan norm (denoted VSMAX) is defined as follows.

DEFINITION 1.1. VSMAX: *The objective is the maximum load over all machines and all dimensions, i.e.,*

$$\max_k \|\Lambda(k)\|_\infty = \max_{k,i \in M} \Lambda_i(k).$$

We consider this problem in both the *identical machines* (denoted VSMAX-I) and the *unrelated machines* (denoted VSMAX-U) settings. First, we state our result for identical machines. Recall that we say an algorithm has competitive ratio  $\alpha$  if for

<sup>1</sup>Our  $L_r$  norms are typically referred to as  $p$ -norms or  $L_p$  norms. We use  $L_r$  norms to reserve the letter  $p$  for job processing times.

all instances, the objective of the algorithm is at most  $\alpha$  times the objective of the optimal solution.

**THEOREM 1.2.** *There is a lower bound of  $\Omega(\log d / \log \log d)$  on the competitive ratio of online algorithms for the VSMAX-I problem. Moreover, there is an online algorithm whose competitive ratio asymptotically matches this lower bound.*

The upper bound is a slight improvement over the previous best  $O(\log d)$  [8, 37], but the only lower bound known previously was NP-hardness of obtaining an  $O(1)$ -approximation for the offline problem [18]. We remark that while the offline approximability remains unresolved, the best offline algorithms currently known (see [8, 37], and this paper) are in fact online. Also, our lower bound is information-theoretic, i.e., relies on the online model instead of computational limitations.

For unrelated machines (VSMAX-U), an  $O(\log m + \log d)$ -competitive algorithm was given by Meyerson, Roytman, and Tagiku [37]. We show that this is the best possible.

**THEOREM 1.3.** *There is a lower bound of  $\Omega(\log m + \log d)$  on the competitive ratio of online algorithms for the VSMAX-U problem.*

**Extensions to other  $L_r$  norms.** As we briefly discussed earlier, there are many applications where an  $L_r$  norm (for some  $r \geq 1$ ) is more suitable than the makespan norm. First, we consider identical machines, where we aim to optimize *all norms* simultaneously (denoted VSALL-I). In particular, the objective for a particular  $L_r$  norm in dimension  $k$  is given as follows.

**DEFINITION 1.4.** VSALL-I: *For dimension  $k$  and norm  $L_r$ ,  $r \geq 1$ , the objective is*

$$\|\Lambda(k)\|_r = \left( \sum_{i \in M} \Lambda_i^r(k) \right)^{1/r}.$$

For VSALL-I, an algorithm is said to have competitive ratio  $\alpha_r$  under the  $L_r$  norm if  $\|\Lambda(k)\|_r \leq \alpha_r \|\Lambda_r^*(k)\|_r$  for every dimension  $k$  and for every possible load vector  $\Lambda_r^*$  resulting from a feasible assignment. The next theorem extends Theorem 1.2 to an all norms optimization. Namely, we obtain an algorithm that produces a single solution that is  $\alpha_r$ -competitive for all  $r \geq 1$  (i.e.,  $\alpha_r$  is a function of  $r$ ). We note that this result assumes that the algorithm is aware a priori of both the final volume of all jobs on each dimension and the largest load over all dimensions and jobs.

**THEOREM 1.5.** *There is an online algorithm for the VSALL-I problem that obtains a competitive ratio of  $O((\log d / \log \log d)^{1-1/r})$ , simultaneously for all  $L_r$  norms. Moreover, these competitive ratios are tight, i.e., there is a matching lower bound for every individual  $L_r$  norm.<sup>2</sup>*

For unrelated machines, there is a polynomial lower bound for simultaneously optimizing multiple  $L_r$  norms, even with scalar loads [35], which rules out an all norms approximation. Therefore, we focus on an *any norm* approximation, where the algorithm is given norms  $r_1, r_2, \dots, r_d$  (where  $1 \leq r_k \leq \log m$ ),<sup>3</sup> and the goal is to minimize the  $L_{r_k}$  norm for dimension  $k$ . The same lower bound also rules out

<sup>2</sup>We note that our upper bound holds even against a stronger benchmark, where  $\Lambda_r^*(k)$  is defined as the load vector in dimension  $k$  of the assignment that minimizes the  $L_r$  norm of dimension  $k$ .

<sup>3</sup>For any  $m$ -dimensional vector  $x$ ,  $\|x\|_\infty = \Theta(\|x\|_{\log m})$ . Therefore, for any  $r_k > \log m$ , an algorithm can instead use a  $L_{\log m}$  norm to approximate an  $L_{r_k}$  norm objective up to constant distortion. Thus, in both our upper and lower bound results we restrict  $1 \leq r_k \leq \log m$ .

the possibility of the algorithm being competitive against the optimal value of each individual norm in their respective dimensions. We use a standard trick in multi-objective optimization to circumvent this impossibility: we only require the algorithm to be competitive against any given feasible *target* vector  $T = \langle T_1, \dots, T_d \rangle$ .<sup>4</sup> Now, we are ready to define the VS problem with arbitrary  $L_r$  norms for unrelated machine—we call this problem VSANY-U.

DEFINITION 1.6. VSANY-U: *For dimension  $k$ , the objective is*

$$\|\Lambda(k)\|_{r_k} = \left( \sum_{i \in M} \Lambda_i^{r_k}(k) \right)^{1/r_k}.$$

An algorithm is said to be  $\alpha_{r_k}$ -competitive under the  $L_{r_k}$  norm if  $\|\Lambda(k)\|_{r_k} \leq \alpha_{r_k} T_k$  for every dimension  $k$ . Note the (necessary) difference between the definitions of VSALL-I and VSANY-U: in the former, the algorithm must be competitive in all norms in all dimensions simultaneously, whereas in VSANY-U, the algorithm only needs to be competitive against a single norm in each dimension that is specified in the problem input. We obtain the following result for the any norm problem.

THEOREM 1.7. *There is an online algorithm for the VSANY-U problem that simultaneously obtains a competitive ratio of  $O(r_k + \log d)$  for each dimension  $k$ , where the goal is to optimize the  $L_{r_k}$  norm in the  $k$ th dimension. Moreover, these competitive ratios are tight, i.e., there is a matching lower bound for any given set of  $L_{r_k}$  norms.*

**1.1. Our techniques.** First, we outline the main techniques used for the identical machines setting. At a high level, we leverage the connection with vertex coloring that was initially observed by Chekuri and Khanna [18]. However, the lower bound of Halldórsson and Szegedy [29] for online vertex coloring cannot be used to exploit this connection (we discuss why we need new approaches in detail in section 2.1.3), we therefore derive a lower bound of  $\Omega(\sqrt{t})$  for the online problem of minimizing the size of the largest monochromatic clique given a fixed set of colors, where  $t$  is the number of colors. To the best of our knowledge, this problem was not studied before and we believe this result should be of independent interest.<sup>5</sup> As is typical in establishing online lower bounds, the construction of the lower bound instance is viewed as a game between the online algorithm and the adversary. Our main goal is to force the online algorithm to grow cliques while guaranteeing that the optimal (offline) solution can color vertices in a way that limits clique sizes to a constant. The technical challenge is to show that the optimal solution does not form large cliques across the cliques that the algorithm has created. For this purpose, we develop a novel randomized code that dictates the choices of the optimal solution and restricts those of the online algorithm. Using the probabilistic method on this code, we are able to show the existence of codewords that always lead to a good optimal solution and an expensive algorithmic one. We also show that the same idea can be used to obtain a lower bound for any  $L_r$  norm.

<sup>4</sup>A target vector  $\langle T_1, \dots, T_d \rangle$  is feasible if there is an assignment such that for every dimension  $k$ , the value of the  $L_{r_k}$  norm in that dimension is at most  $T_k$ . Our results do not rely on the exact feasibility of the target vector; if there is a feasible solution that violates targets in all dimensions by at most a factor of  $\beta$ , then our results hold with an additional factor of  $\beta$  in the competitive ratio.

<sup>5</sup>In [39], the problem of coloring vertices without creating certain monochromatic subgraphs was studied, which is different from our goal of minimizing the largest monochromatic clique size. Furthermore, this previous work was only for random graphs and the focus was on whether the desirable coloring is achievable online depending on the parameters of the random graph.

We now turn our attention to our second main result which is in the unrelated machines setting: an upper bound for the VSANY-U problem. Our algorithm is greedy with respect to a potential function (as are algorithms for all special cases studied earlier [4, 6, 15, 37]), and the novelty lies in the choice of the potential function. For each individual dimension  $k$ , we use the  $L_{r_k}^{r_k}$  norm as the potential (following [4, 15]). The main challenge is to combine these individual potentials into a single potential. We use a weighted linear combination of the individual potentials for the different dimensions. This is somewhat counterintuitive since the combined potential can possibly allow a large potential in one dimension to be compensated by a small potential in a different one—indeed, a naïve combination only gives a competitive ratio of  $O(\max_k r_k + \log d)$  for all  $k$ . However, we observe that we are aiming for a competitive ratio of  $O(r_k + \log d)$  which allows some slack compared to scalar loads if  $r_k < \log d$ . Suppose  $q_k = r_k + \log d$ ; then we use weights of  $q_k^{-q_k}$  in the linear combination after changing the individual potentials to  $L_{r_k}^{q_k}$ . Note that as one would expect, the weights are larger for dimensions that allow a smaller slack. We show that this combined potential simultaneously leads to the asymptotically optimal competitive ratio on every individual dimension.

Finally, we briefly discuss our other results. Our slightly improved upper bound for the VS MAX-I problem follows from a simple random assignment and redistributing “overloaded” machines. We remark that derandomizing this strategy is relatively straightforward. Although this improvement is very marginal, we feel that this is somewhat interesting since our algorithm is simple and perhaps more intuitive yet gives the tight upper bound. For the VS ALL-I problem, we give a reduction to VS MAX-I by structuring the instance by “smoothing” large jobs and then arguing that for structured instances, a VS MAX-I algorithm is also optimal for other  $L_r$  norms.

**1.2. Related work.** Due to the large volume of related work, we will only sample some relevant results in online scheduling and refer the interested reader to more detailed surveys (e.g., [7, 41, 42]) and textbooks (e.g., [13]).

**Scalar loads.** Since the  $(2 - 1/m)$ -competitive algorithm by Graham [26] for online (scalar) load balancing on identical machines, a series of papers [2, 10, 34] have led to the current best ratio of 1.9201 [22]. On the negative side, this problem was shown to be NP-hard in the strong sense by Faigle, Kern, and Tuřan [21] and has since been shown to have a competitive ratio of at least 1.880 [2, 11, 25, 31]. For other norms, Avidor, Avar, and Sgall [5] obtained competitive ratios of  $\sqrt{4/3}$  and  $2 - O(\frac{\log r}{r})$  for the  $L_2$  and general  $L_r$  norms, respectively.

For unrelated machines, Aspnes et al., [4] obtained a competitive ratio of  $O(\log m)$  for makespan minimization, which is asymptotically tight [9]. Scheduling for the  $L_2$  norm was considered by [17, 20], and Awerbuch et al., [6] obtained a competitive ratio of  $1 + \sqrt{2}$ , which was shown to be tight [16]. For general  $L_r$  norms, Awerbuch et al., [6] (and Caragiannis [15]) obtained a competitive ratio of  $O(r)$  and showed that it is tight up to constants. Various intermediate settings such as related machines (machines have unequal but job-independent speeds) [4, 12] and restricted assignment (each job has a machine-independent load but can only be assigned to a subset of machines) [9, 16, 19, 43] have also been studied for the makespan and  $L_r$  norms.

**Vector loads.** The VS MAX-I problem was introduced by Chekuri and Khanna [18], who gave an offline approximation of  $O(\log^2 d)$  and observed that a random assignment has a competitive ratio of  $O(\frac{\log dm}{\log \log dm})$ . Azar et al., [8] and Meyerson, Roytman, and Tagiku [37] improved the competitive ratio to  $O(\log d)$  using deterministic online algorithms. An offline  $\omega(1)$  lower bound was also proved in [18], and it remains open

as to what the exact dependence of the approximation ratio on  $d$  should be. Our online lower bound asserts that a significantly sublogarithmic dependence would require a radically different approach from all the known algorithms for this problem.

For unrelated machines, Meyerson, Roytman, and Tagiku [37] noted that the natural extension of the algorithm of Aspnes et al., [4] to vector loads has a competitive ratio of  $O(\log dm)$  for makespan minimization; in fact, for identical machines, they used exactly the same algorithm but gave a tighter analysis. For the offline VSMAX-U problem, Harris and Srinivasan [30] recently showed that the dependence on  $m$  is not required by giving a randomized  $O(\log d / \log \log d)$  approximation algorithm.

**2. Identical machines.** First, we consider the online VS problem for identical machines. In this section, we obtain tight upper and lower bounds for this problem, both for the makespan norm (Theorem 1.2) and for arbitrary  $L_r$  norms (Theorem 1.5).

**2.1. Lower bounds for vsmax-i and vsall-i.** In this section, we will prove the lower bound in Theorem 1.2, i.e., show that any online algorithm for the VSMAX-I problem can be forced to construct a schedule such that there exists a dimension where one machine has load  $\Omega(\log d / \log \log d)$ , whereas the optimal schedule has  $O(1)$  load on all dimensions of all machines. This construction will also be extended to all  $L_p$  norms (VSALL-I) in order to establish the lower bound in Theorem 1.5.

We give our lower bound for VSMAX-I in two parts. First in section 2.1.1, we define a lower bound instance for an online graph coloring problem, which we call MONOCHROMATIC CLIQUE. Next, in section 2.1.2, we show how our lower bound instance for MONOCHROMATIC CLIQUE can be encoded as an instance for VSMAX-I in order to obtain the desired  $\Omega(\log d / \log \log d)$  bound. We then close this subsection by discussing in section 2.1.3 why the previous work is not sufficient towards establishing lower bounds.

**2.1.1. Lower bound for monochromatic clique.** The MONOCHROMATIC CLIQUE problem is defined as follows.

MONOCHROMATIC CLIQUE. We are given a *fixed* set of  $t$  colors. The input graph is revealed to an algorithm as an online sequence of  $n$  vertices  $v_1, \dots, v_n$  that arrive one at a time. When vertex  $v_j$  arrives, we are given all edges between vertices  $v_1, v_2, \dots, v_{j-1}$  and vertex  $v_j$ . The algorithm must then assign  $v_j$  to one of the  $t$  colors before it sees the next arrival. The objective is to minimize the size of the largest monochromatic clique in the final coloring.

The goal of the section will be to prove the following lemma, which we will use later in section 2.1.2 to establish our lower bound for VSMAX-I.

**THEOREM 2.1.** *The competitive ratio of any online algorithm for MONOCHROMATIC CLIQUE is  $\Omega(\sqrt{t})$ , where  $t$  is the number of available colors.*

More specifically, for any online algorithm  $A$ , there is an instance on which  $A$  produces a monochromatic clique of size  $\sqrt{t}$ , whereas the optimal solution can color the graph such that the size of the largest monochromatic clique is  $O(1)$ .

We will frame the lower bound as a game between an adversary and the online algorithm. At a high level, the instance is designed as follows. For each new arriving vertex  $v$  and color  $c$ , the adversary connects  $v$  to every vertex in some currently existing monochromatic clique of color  $c$ . Since we do this for every color, this ensures that regardless of the algorithm's coloring of  $v$ , some monochromatic clique grows by 1 in size (or the first vertex in a clique is introduced). Since this growth happens for every vertex, the adversary is able to quickly force the algorithm to create a

monochromatic clique of size  $\sqrt{t}$ .

The main challenge now is to ensure that the adversary can still obtain a good offline solution. Our choice for this solution will be naïve: the adversary will simply properly color the monochromatic cliques it attempted to grow in the algorithm's solution. Since the game stops once the algorithm has produced a monochromatic clique of size  $\sqrt{t}$ , and there are  $t$  colors, such a proper coloring of every clique is possible. The risk with this approach is that a large monochromatic clique may now form in the adversary's coloring from edges that cross these independently grown cliques (in other words, properly colored cliques in the algorithm's solution could now become monochromatic cliques for the adversary). This may seem hard to avoid since each vertex is connecting to some monochromatic clique for every color. However, in our analysis we show that if on each step the adversary selects which cliques to grow in a carefully defined random fashion, then with positive probability, all properly colored cliques in the algorithm's solution that hurt the adversary's naïve solution are of size  $O(1)$ .

**Instance construction.** We adopt the standard terminology used in online coloring problems (see, e.g., [29]). Namely, the algorithm will place each vertex in one of  $t$  bins to define its color assignments, whereas we will use *colors* to refer to the color assignment in the optimal solution (controlled by the adversary). For each vertex arrival, the game is defined by the following 3-step process:

1. The adversary issues a vertex  $v_j$  and defines  $v_j$ 's adjacencies with vertices  $v_1, \dots, v_{j-1}$ .
2. The online algorithm places  $v_j$  in one of the available  $t$  bins.
3. The adversary selects a color for the vertex.

We further divide each bin into  $\sqrt{t}$  slots  $1, 2, \dots, \sqrt{t}$ . These slots will only be used for the adversary's bookkeeping. Correspondingly, we partition the  $t$  colors into  $\sqrt{t}$  color sets  $C_1, \dots, C_{\sqrt{t}}$ , each of size  $\sqrt{t}$ . Each vertex will reside in a slot inside the bin chosen by the algorithm, and all vertices residing in slot  $i$  across all bins will be colored by the optimal solution using a color from  $C_i$ . The high-level goal of the construction will be to produce properly colored cliques inside each slot of every bin.

Consider the arrival of vertex  $v_j$ . Inductively assume the previous vertices  $v_1, \dots, v_{j-1}$  have been placed in the bins by the algorithm, and that every vertex within a bin lies in some slot. Further, assume that all of the vertices in any particular slot of a bin form a properly colored clique.

To specify the new adjacencies formed by vertex  $v_j$  for step 1, we will use a  $t$ -length  $\sqrt{t}$ -ary string  $s_j$ , where we connect  $v_j$  to every vertex in slot  $s_j[k]$  of bin  $k$ , for all  $k = 1, 2, \dots, t$ . Next, for step 2, the algorithm places  $v_j$  in some bin  $b_j$ . We say that  $v_j$  is then placed in slot  $q_j = s_j[b_j]$  in bin  $b_j$ . Finally for step 3, the adversary chooses an arbitrary color for  $v_j$  from the colors in  $C_{q_j}$  that have not yet been used for any vertex in slot  $q_j$  of bin  $b_j$ . The adversary will end the instance whenever there exists a slot in some bin that contains  $\sqrt{t}$  vertices. This ensures that as long as the game is running, there is always an unused color in every slot of every bin. Also observe that after this placement, the clique in slot  $q_j$  in bin  $b_j$  has grown in size by 1 but is still properly colored. So, this induction is well defined. This completes the description of the instance (barring our choice for each adjacency string  $s_j$ ). See Figure 2.1 for illustrations of the construction.

**Instance analysis.** The following lemma follows directly from the construction.

**LEMMA 2.2.** *Any online algorithm produces a monochromatic clique of size  $\sqrt{t}$  by the end of the instance.*

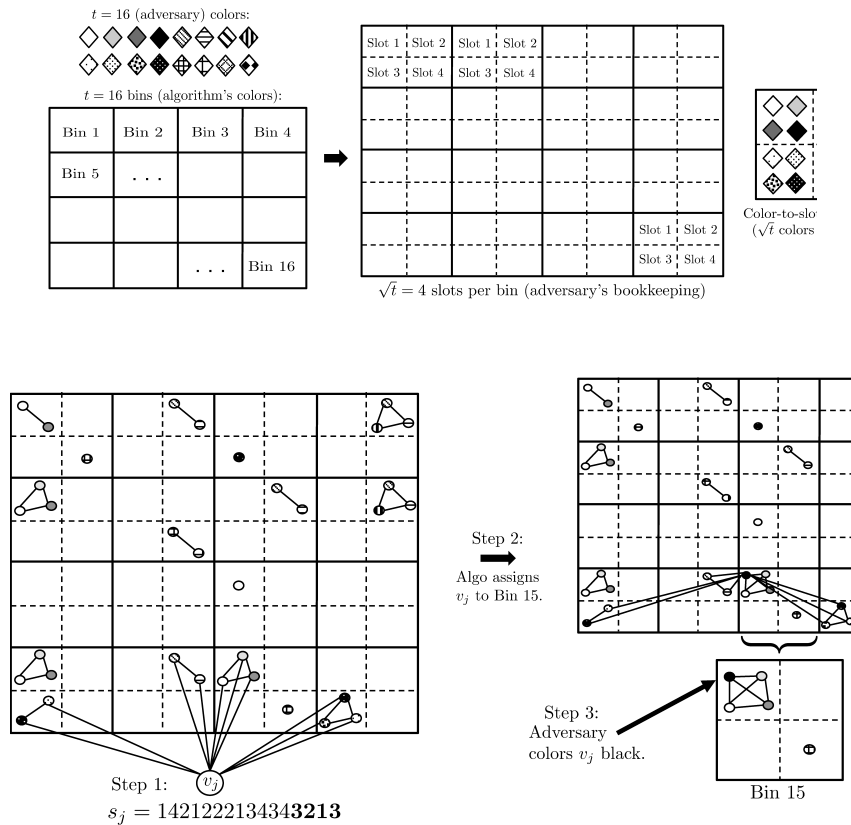


FIG. 2.1. Depiction of the three-step lower-bound game for Lemma 2.1. For simplicity, the only adjacencies shown for vertices issued before  $v_j$  are those between vertices in the same bin-slot pair (in reality, other adjacencies also exist). Also for simplicity, the only adjacencies shown for  $v_j$  are those it has with vertices in bins 13 through 16 (dictated by the bold substring “3213” in  $s_j$ ). Again note that in reality,  $v_j$  will also be adjacent to vertices in bins 1 through 12 due to the remaining prefix “142122213434.”

*Proof.* After  $t^2$  vertices are issued, there will exist some bin  $b$  containing at least  $t$  vertices, and therefore some slot in bin  $b$  containing at least  $\sqrt{t}$  vertices forming a clique of size  $\sqrt{t}$ . Since all of the vertices in the clique are in the same bin, there exists a monochromatic clique of size  $\sqrt{t}$  in the algorithm's solution.  $\square$

Thus, it remains to show that there exists a sequence of  $\sqrt{t}$ -ary strings of length  $t^2$  (recall that these strings define the adjacencies for each new vertex) such that the size of the largest monochromatic clique in the optimal coloring is  $O(1)$ . For brevity, we call such a sequence a *good sequence*.

First, observe that monochromatic edges (i.e., edges between vertices of the same color) cannot form between vertices in slots  $s$  and  $s' \neq s$  (in the same or in different bins) since the color sets used for the slots are disjoint. Moreover, monochromatic edges cannot form within the same slot in the same bin since these vertices always form a properly colored clique. Therefore, monochromatic edges can only form between two adjacent vertices  $v_j$  and  $v_{j'}$  such that  $q_j = q_{j'}$  and  $b_j \neq b_{j'}$ , i.e., vertices in the same slot but in different bins. Relating back to our earlier discussion, these are exactly the edges that are properly colored in the algorithm's solution that could potentially



form monochromatic cliques in the adversary's solution; we will refer to such edges as *bad edges*.

Thus, in order to define a good sequence of  $t^2$  strings, we need to ensure our adjacency strings do not induce large cliques of bad edges. To do this, we first need a handle on what structure must exist across the sequence in order for bad-edge cliques to form. This undesired structure is characterized by the following lemma.

LEMMA 2.3. Suppose  $K = \{u_{\phi(1)}, \dots, u_{\phi(w)}\}$  is a  $w$ -sized monochromatic clique of color  $c \in C_\ell$  that forms during the instance, where  $\phi : [w] \rightarrow [t^2]$  maps  $k \in [w]$  to the index of the  $k$ th vertex to join  $K$  (note, from the above discussion, that  $b_{\phi(j)}$  are different for all  $j \in [w]$ ). Then

$$s_{\phi(j)}[b_{\phi(i)}] = \ell \quad \forall j \in \{1, \dots, w\}, \forall i \in \{1, \dots, j-1\}.$$

*Proof.* Consider vertex  $u_{\phi(j)}$  (the  $j$ th vertex to join  $K$ ). Since  $K$  is a clique,  $u_{\phi(j)}$  must be adjacent to vertices  $u_{\phi(1)}, \dots, u_{\phi(j-1)}$ . Since all of these vertices are colored with  $c \in C_\ell$ , they must have been placed in slot  $\ell$  in their respective bins. Therefore, the positions in  $s_{\phi(j)}$  that correspond to these bins must also be  $\ell$ , i.e.,  $s_{\phi(j)}[b_{\phi(i)}] = \ell$  for all previous vertices  $u_{\phi(i)}$ .  $\square$

In the remainder of the proof, we show that the structure in Lemma 2.3 can be avoided with nonzero probability for constant sized cliques if we generate our strings uniformly at random, thus implying the existence of a good set of  $t^2$  strings.

Specifically, suppose the adversary picks each  $s_j$  uniformly at random, i.e., for each character in  $s_j$  we pick  $w \in [\sqrt{t}]$  with probability  $t^{-1/2}$ . We define the following notation:

- Let  $K_{20}$  be the event that the adversary creates a monochromatic clique of size 20 or greater.<sup>6</sup>
- Let  $K_{20}(S, c)$  be the event that a monochromatic clique  $K$  of color  $c$  and size 20 or greater forms such that the first ten vertices to join  $K$  are placed in the bins specified by the set of ten indices  $S$ .
- Let  $P_j(S, q)$  be a random variable that is 1 if  $s_j[i] = q \forall i \in S$  and 0 otherwise. Let  $P(S, q) = \sum_{j=1}^{t^2} P_j(S, q)$ .
- Let  $q(c) \in [\sqrt{t}]$  be the index of the color set to which color  $c$  belongs (i.e.,  $c \in C_{q(c)}$ ).
- Let  $[n]_k := \binom{[n]}{k}$  denote the set of all size- $k$  subsets of  $[n]$ .

The next lemma follows from standard Chernoff–Hoeffding bounds, which we state first for completeness.

THEOREM 2.4 (Chernoff–Hoeffding bounds (e.g., [38])). *Let  $X_1, X_2, \dots, X_n$  be independent binary random variables, and let  $a_1, a_2, \dots, a_n$  be coefficients in  $[0, 1]$ . Let  $X = \sum_i a_i X_i$ . Then,*

- For any  $\mu \geq \mathbb{E}[X]$  and any  $\delta > 0$ ,  $\Pr[X > (1 + \delta)\mu] \leq \left(\frac{e^\delta}{(1+\delta)^{(1+\delta)}}\right)^\mu$ .
- For any  $\mu \leq \mathbb{E}[X]$  and  $\delta > 0$ ,  $\Pr[X < (1 - \delta)\mu] \leq e^{-\mu\delta^2/2}$ .

We are now ready to state and prove the lemma.

LEMMA 2.5. *Let  $S \in [t]_{10}$ , and let  $r \in [\sqrt{t}]$ . If the adversary picks each  $s_j$  uniformly at random, then  $\Pr[P(S, r) \geq 10] < t^{-30}$ .*

<sup>6</sup>20 is an arbitrarily chosen large enough constant.

*Proof.* First, we observe that for any set  $S \in [t]_{10}$  and any  $r \in [\sqrt{t}]$ , we have  $\Pr[P_j(S, r) = 1] = (1/\sqrt{t})^{10} = t^{-5}$  for any string  $s_j$ . Therefore, by linearity of expectation, we have

$$(2.1) \quad \mathbb{E}[P(S, r)] = \mathbb{E} \left[ \sum_{j=1}^{t^2} P_j(S, r) \right] = t^2 \cdot t^{-5} = t^{-3}.$$

Applying Theorem 2.4 to  $P(S, r)$  with  $X_i = P_i(S, r)$ ,  $a_i = 1$ ,  $\delta = 10t^3 - 1$ , and  $\mu = t^{-3}$  from (2.1), we get

$$\Pr[P(S, r) \geq 10] \leq \left( \frac{e^{10t^3-1}}{(10t^3)^{10t^3}} \right)^{t^{-3}} \leq \left( \frac{e^{10}}{10^{10}} \right) \cdot \left( \frac{1}{t^{30}} \right) < t^{-30}.$$

This completes the proof.  $\square$

Using Lemmas 2.3 and 2.5, we argue that there exists an offline solution with no monochromatic clique of super constant size.

**LEMMA 2.6.** *There exists an offline solution where every monochromatic clique is of size  $O(1)$ .*

*Proof.* To show the existence of a good set of  $t^2$  strings, it is sufficient to show that  $\Pr[K_{20}] < 1$ . Using Lemma 2.5, we in fact show this event occurs with low probability. Observe that

$$(2.2) \quad \Pr[K_{20}] \leq \sum_{c \in [\sqrt{t}]} \sum_{S \in [t]_{10}} \Pr[K_{20}(S, c)] \leq \sum_{c \in [\sqrt{t}]} \sum_{S \in [t]_{10}} \Pr[P(S, q(c)) \geq 10].$$

The first inequality is a straightforward union bound. The second inequality follows Lemma 2.3. If the event  $K_{20}(S, c)$  occurs, then Lemma 2.3 implies  $s_j[b_i] = q(c)$  for  $j = 11, \dots, 20$ ,  $i \in S$ .

Since there are  $\sqrt{t}$  possible colors and  $|[t]_{10}| < t^{10}$ , applying both (2.2) and Lemma 2.5 we get

$$\begin{aligned} \Pr[K_{20}] &\leq \sum_{c \in [\sqrt{t}]} \sum_{S \in [t]_{10}} \Pr[P(S, q(c)) \geq 10] \\ &\leq \sum_{c \in [\sqrt{t}]} \sum_{S \in [t]_{10}} t^{-30} \leq t^{1/2} \cdot t^{10} \cdot t^{-30} = t^{-39/2} < 1 \end{aligned}$$

for all  $t > 1$ . Therefore, there is a coloring which has no monochromatic clique of size more than 20.  $\square$

Theorem 2.1 now follows directly from Lemmas 2.2 and 2.6.

### 2.1.2. vsmax-i and vsall-i lower bound using monochromatic clique.

We are now ready to use Theorem 2.1 to show an  $\Omega(\log d / \log \log d)$  lower bound for VSMAX-I. We will describe a lower bound instance for VSMAX-I whose structure is based on an instance of MONOCHROMATIC CLIQUE. This will allow us to use the lower bound instance from Theorem 2.1 as a black box to produce the desired lower bound for VSMAX-I.

We first set the problem definition of MONOCHROMATIC CLIQUE to be for  $m$  colors where  $m$  is also the number of machines used in the VSMAX-I instance. Let  $I_C$  be the

lower-bound instance for this problem given by Theorem 2.1. This produces a graph  $G$  of  $m^2$  vertices such that the algorithm forms a monochromatic clique of size  $\sqrt{m}$ , whereas the largest monochromatic clique in the optimal solution is of size  $O(1)$ . Let  $G_j = (V_j, E_j)$  be the graph in  $I_C$  after vertices  $v_1, \dots, v_j$  have been issued (and so  $G_n = G$ ). We define the corresponding lower bound instance for VSMAX-I as follows (see Figures 2.2 and 2.3 for an illustration):

- There are  $m^2$  jobs, which correspond to vertices  $v_1, \dots, v_{m^2}$  from  $I_C$ .
- Each job has  $d = \binom{m^2}{\sqrt{m}}$  dimensions, where each dimension corresponds to a specific  $\sqrt{m}$ -sized vertex subset of the  $m^2$  vertices. Let  $S_1, \dots, S_d$  be an arbitrary ordering of these subsets.
- Job vectors will be binary. Namely, the  $k$ th vector entry for job  $j$  is 1 if  $v_j \in S_k$  and the vertices in  $\{v_1, \dots, v_j\} \cap S_k$  form a clique in  $G_j$  (if  $\{v_1, \dots, v_j\} \cap S_k = \{v_j\}$ , then it is considered a 1-clique); otherwise, the  $k$ th entry is 0.
- Let  $c_1, \dots, c_m$  define an ordering on the available colors from  $I_C$ . We match each color from  $I_C$  to a machine in our scheduling instance. Therefore, when the VSMAX-I algorithm makes an assignment for a job, we translate this machine assignment as the corresponding color assignment in  $I_C$ . Formally, if job  $j$  is placed on machine  $i$  in the scheduling instance, then vertex  $v_j$  is assigned color  $c_i$  in  $I_C$ .

Since assigning jobs to machines corresponds to colorings in  $I_C$ , it follows that the largest load in dimension  $k$  is the size of the largest monochromatic subclique in  $S_k$ .  $I_C$  is given by the construction in Theorem 2.1; therefore, at the end of the instance, there will exist a dimension  $k'$  such that the online algorithm colored every vertex in  $S_{k'}$  with some color  $c_i$ . Thus, machine  $i$  will have  $\sqrt{m}$  load in dimension  $k'$ . In contrast, Theorem 2.1 ensures that all of the monochromatic cliques in the optimal solution are of size  $O(1)$ , and therefore, the load on every machine in dimension  $k'$  is  $O(1)$ .

The relationship between  $m$  and  $d$  is given as follows.

**FACT 2.7.** *If  $d = \binom{m^2}{\sqrt{m}}$ , then  $\sqrt{m} = \Omega(\log d / \log \log d)$ .*

*Proof.* We will use the following well-known bounds on  $\binom{n}{k}$ : for integers  $0 \leq k \leq n$ ,  $\binom{n}{k}^k \leq \binom{n}{k} \leq \left(\frac{en}{k}\right)^k$ . First, we observe that

$$(2.3) \quad \log d = \log \binom{m^2}{\sqrt{m}} \leq \log \left(\frac{em^2}{\sqrt{m}}\right)^{\sqrt{m}} = \log(e^{\sqrt{m}} \cdot m^{(3/2)\sqrt{m}}) = \sqrt{m} \cdot (1 + (3/2) \log m).$$

We also have

$$(2.4) \quad \log \log d = \log \log \binom{m^2}{\sqrt{m}} \geq \log \log \left(\frac{m^2}{\sqrt{m}}\right)^{\sqrt{m}} \geq \log((3/2)\sqrt{m} \log m) \geq (1/2) \log m.$$

Hence, combining (2.3) and (2.4), we obtain

$$\sqrt{m} \geq \frac{\log d}{1 + (3/2) \log m} \geq \frac{\log d}{1 + 3 \log \log d},$$

which implies that  $\sqrt{m} = \Omega(\log d / \log \log d)$ , as desired.  $\square$

To end the section, we show that our lower bound for VSMAX-I extends to general  $L_r$  norms (Theorem 1.5). As before, our lower bound construction forces any algorithm to schedule jobs so that there exists a dimension  $k'$  where at least one machine

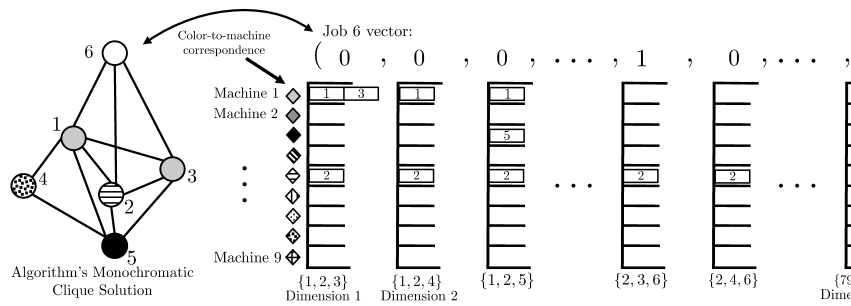


FIG. 2.2. Illustration of the lower bound construction for VSMAX-I using the MONOCHROMATIC CLIQUE lower bound (Theorem 2.1) for an instance where  $m = 9$  (and thus  $d = \binom{9^2}{\sqrt{9}} = \binom{81}{3}$  for the VSMAX-I instance and  $t = 9$  for the MONOCHROMATIC CLIQUE instance). Currently job 6 is being issued; its binary load vector, which is based on the current edge structure in the MONOCHROMATIC CLIQUE instance, is given above the machines/dimensions. Observe that job 6 has load 0 in the first three dimensions and the last dimension since 6 is not contained in any of the sets corresponding to these dimensions (indicated below). It does have load 1 in the dimension corresponding to set  $\{2, 3, 6\}$  since vertex 6 forms a clique with vertices 2 and 3 in the MONOCHROMATIC CLIQUE instance; however, it still has load 0 in the dimension corresponding to set  $\{2, 4, 6\}$  since vertex 6 does not form a clique with vertices 2 and 4.




FIG. 2.3. State of the construction after job 6 is assigned to machine 3. Since black is the color we associated with machine 3, this job assignment by the VSMAX-I algorithm is translated as coloring vertex 6 black in the MONOCHROMATIC CLIQUE instance.

has load at least  $\sqrt{m}$ , whereas the load on every dimension of every machine in the optimal solution is bounded by some constant  $C$ . Since any dimension has at most  $\sqrt{m}$  jobs with load 1, any assignment ensures that there are at most  $\sqrt{m}$  machines with nonzero load in a given dimension. Therefore, in the optimal solution, the  $L_r$  norm of the load vector for dimension  $k'$  is at most  $(C^r \cdot \sqrt{m})^{1/r} = C \cdot m^{1/(2r)}$ .

Thus, the ratio between the objective of the solution produced by the online algorithm and the optimal solution is at least  $m^{1/2}/(C \cdot m^{1/2r}) = (1/C) \cdot m^{(r-1)/(2r)}$ . Using Fact 2.7, we conclude the lower bound.

**2.1.3. Justification for our approach.** Finally, we justify our approach by discussing in detail why the previous work is not sufficient towards establishing a lower bound for VSMAX. As mentioned earlier, a natural starting point for lower bounds is the online *vertex coloring* (VC) lower bound of Halldórsson and Szegedy [29], for which connections to VSMAX-I [18] have previously been exploited. The basic idea is to encode a VC instance as a VSMAX-I instance where the number of dimensions  $d$  is (roughly)  $n^B$  and show that an approximation factor of (roughly)  $B$  for VSMAX-I implies an approximation factor of (roughly)  $n^{1-1/B}$  for VC. One may want to try to combine this reduction and the online lower bound of  $\Omega(n/\log^2 n)$  for VC [29] to get a better lower bound for VSMAX-I. However, the reduction crucially relies on the

fact that a graph with the largest clique size of at most  $k$  has a chromatic number of (roughly)  $O(n^{1-1/k})$ , and this does not imply that the graph can be colored online with a similar number of colors.

A second seemingly plausible approach is to explore the connection of VSMAX-I with online vector bin packing (VBP), where multidimensional items arriving online must be packed into a minimum number of identical multidimensional bins. Recently, Azar et al. [8] obtained strong lower bounds of  $\Omega(d^{1/B})$  where  $B \geq 1$  is the capacity of each bin in every dimension (the items have a maximum size of 1 on any dimension). It would be tempting to conjecture that the inability to obtain a constant approximation algorithm for the VBP problem unless  $B = \Omega(\log d)$  should yield a lower bound of  $\Omega(\log d)$  for the VSMAX-I problem. Unfortunately, this is false. The difference between the two problems is in the capacity of the bins/machines that the optimal solution is allowed to use: in VSMAX-I, this capacity is 1 whereas in VBP this capacity is  $B$ , and using bins with larger capacity can decrease the number of bins needed superlinearly in the increased capacity. Therefore, a lower bound for VBP does not imply any lower bound for VSMAX-I. On the other hand, an upper bound of  $O(d^{1/(B-1)} \log d)$  for the VBP problem is obtained in [8] via an  $O(\log d)$ -competitive algorithm for VSMAX-I. Improving this ratio considerably for VSMAX-I would have been a natural approach for closing the gap for VBP; unfortunately, our lower bound of  $\Omega(\log d / \log \log d)$  rules out this possibility.

**2.2. Upper bounds for vsmax-i and vsall-i.** In this section we prove the upper bounds in Theorem 1.2 (VSMAX-I) and Theorem 1.5 (VSALL-I). First, we give a randomized  $O(\log d / \log \log d)$ -competitive online algorithm for VSMAX-I (section 2.2.1) and then show how to derandomize it (section 2.2.2). Next, we give an  $O((\log d / \log \log d)^{\frac{r-1}{r}})$ -competitive algorithm for VSALL-I (section 2.2.3), i.e., for each dimension  $k$  and  $1 \leq r \leq \log m$ ,  $\|\Lambda(k)\|_r$  is competitive with the optimal schedule for dimension  $k$  under the  $L_r$  norm objective.

Throughout the section we assume that a priori the online algorithm is aware of both the final volume of all jobs on each dimension and the largest load over all dimensions and jobs. We note that the lower bounds claimed in Theorems 1.2 and 1.5 are robust against this assumption since the optimal makespan is always a constant and this knowledge does not help the online algorithm. Furthermore, these assumptions can be completely removed for our VSMAX-I algorithm by updating a threshold on the maximum job load on any dimension and the total volume of jobs that the algorithm has observed so far. However, in order to make our presentation more transparent and our notation simple, we present our results under these assumptions.

For each job  $j$  that arrives online, both our VSMAX-I and VSALL-I algorithms will perform the following transformation:

- **Transformation 1.** Let  $V = \langle V_1, \dots, V_d \rangle$  be the volume vector given to the algorithm a priori, where  $V_k$  denotes the total volume of all jobs for dimension  $k$ . For this transformation, we normalize  $p_j(k)$  by dividing it by  $V_k/m$  (for ease of notation, we will still refer to this normalized value as  $p_j(k)$ ).

Our VSMAX-I and VSALL-I algorithms will also perform subsequent transformations; however, these transformations will differ slightly for the two algorithms.

**2.2.1. Randomized algorithm for vsall-i.** We now present our randomized  $O(\log d / \log \log d)$ -competitive algorithm for VSMAX-I. Informally, our algorithm works as follows. For each job  $j$ , we first attempt to assign it to a machine  $i$  chosen uniformly at random; however, if the resulting assignment would result in a load larger than

$O(\log d / \log \log d)$  on machine  $i$ , then we dismiss the assignment and instead assign  $j$  greedily among other previously dismissed jobs. In general, a greedy assignment can be as bad as  $\Omega(d)$ -competitive; however, in our analysis we show that a job is dismissed by the random assignment with low probability. Therefore, in expectation, the total volume of these jobs is low enough to assign greedily and still remain competitive.

**Instance transformations.** Before formally defining our algorithm, we define additional online transformations and outline the properties that these transformations guarantee. Note that we perform these transformations for both the randomized algorithm presented in this section and the derandomized algorithm presented in section 2.2.2. These additional transformations are defined as follows (which are performed in sequence after Transformation 1):

- **Transformation 2.** Let  $T$  be the load of the largest job in the instance (given a priori). If for dimension  $k$  we have  $T \geq V_k/m$ , then for each job  $j$  we set  $p_j(k)$  to be  $(p_j(k) \cdot V_k)/(mT)$ . In other words, we normalize jobs in dimension  $k$  by  $T$  instead of  $V_k/m$ .
- **Transformation 3.** For each job  $j$  and dimension  $k$ , if

$$p_j(k) < (1/d) \max_{k'} p_j(k'),$$

then we increase  $p_j(k)$  to  $(1/d) \max_{k'} p_j(k')$ .

Observe that after we apply Transformations 1 and 2 to all jobs, we have  $\sum_j p_j(k) \leq m$  for all  $k \in [d]$  and  $0 \leq p_j(k) \leq 1$  for all jobs  $j$  and  $k \in [d]$ .

In Lemmas 2.8 and 2.9, we prove additional properties that Transformation 3 preserves. Since Transformations 1 and 2 are simple scaling procedures, an  $\alpha$ -competitive algorithm on the resulting scaled instance is also  $\alpha$ -competitive on the original instance, if we only apply the first two transformations. In Lemma 2.8, we prove that this property is still maintained after Transformation 3.

**LEMMA 2.8.** *After Transformations 1 and 2 have been applied, Transformation 3 increases the optimal makespan by a factor of at most 2.*

*Proof.* Fix a machine  $i$  and a dimension  $k$ . Let OPT denote the optimal assignment before Transformation 3 is applied. Let  $J^*(i)$  denote the jobs assigned to machine  $i$  in OPT, let  $\Lambda_i^*(k)$  be the load of OPT on machine  $i$  in dimension  $k$ , and let  $\Lambda^* = \max_{i,k} \Lambda_i^*(k)$  denote the makespan of OPT. We will show that Transformation 3 can increase the load on machine  $i$  in dimension  $k$  by at most  $\Lambda^*$ .

Let  $V_i^* = \sum_{j \in J^*(i)} \sum_{k' \in [d]} p_j(k')$  denote the total volume of jobs that OPT assigns to machine  $i$ . Observe that by a simple averaging argument, we have  $V_i^*/d \leq \max_{k' \in [d]} \Lambda_i^*(k')$ . Since Transformation 3 can increase the load of a job  $j$  in a fixed dimension by at most  $(1/d) \max_{k'} p_j(k')$ , we can upper bound the total increase in load on machine  $i$  in dimension  $k$  as follows:

$$(2.5) \quad \sum_{j \in J^*(i)} (1/d) \max_{k'} p_j(k') \leq V_i^*/d \leq \max_{k' \in [d]} \Lambda_i^*(k') \leq \Lambda^*,$$

as desired. Note that the first inequality follows from the fact that the sum of maximum loads on a machine is at most the total volume of its jobs.  $\square$

Recall that after Transformations 1 and 2,  $\sum_j p_j(k) \leq m$  for all  $k \in [d]$ . In Lemma 2.9, we show that this property is preserved within a constant factor after Transformation 3.

LEMMA 2.9. *After performing Transformation 3,  $\sum_j p_j(k) \leq 2m$  for all  $k \in [d]$ .*

*Proof.* Consider any fixed dimension  $k \in [d]$ . After Transformation 3, each job  $j$ 's load on dimension  $k$  increases by at most  $(1/d) \max_{k'} p_j(k')$ . Hence the total increase in load from jobs in dimension  $k$  is at most

$$\sum_j (1/d) \max_{k'} p_j(k') \leq (1/d) \sum_j \sum_{k' \in [d]} p_j(k') \leq (1/d)md \leq m,$$

where the second inequality and the lemma follow from the fact that  $\sum_j p_j(k) \leq m$  before Transformation 3.  $\square$

In summary, the properties that we collectively obtain from these transformations are as follows:

- **Property 1.** For all  $k \in [d]$ ,  $\sum_j p_j(k) \leq 2m$ .
- **Property 2.** For all  $j$  and  $k \in [d]$ ,  $0 \leq p_j(k) \leq 1$ .
- **Property 3.** For all  $j$  and  $k \in [d]$ ,  $(1/d) \max_{k'} p_j(k') \leq p_j(k) \leq \max_{k'} p_j(k')$ .
- **Property 4.** The optimal makespan is at least 1.

Property 1 is a restatement of Lemma 2.9. Property 2 was true after the first two transformations, and Transformation 3 has no effect on this property. Property 3 is a direct consequence of Transformation 3.

To see why Property 4 is true, let  $j$  be the job with the largest load  $T$  in the instance, and let  $k = \arg \max_{k'} p_j(k')$  (i.e.,  $\max_{k'} p_j(k') = T$ ). If Transformation 2 is applied to dimension  $k$ , then  $p_j(k) = 1$  afterwards, which immediately implies Property 4. Otherwise, only Transformations 1 and 3 are applied to dimension  $k$  and we have  $\sum_{j'} p_{j'}(k) \geq m$ , which again leads to Property 4 by a simple volume argument. Thus, by Property 4 and Lemma 2.8, it is sufficient to now show that the makespan of the algorithm's schedule is  $O(\log d / \log \log d)$ .

**Algorithm definition.** As discussed earlier, our algorithm consists of two procedures: a random assignment and greedy packing. It will be convenient to assume that the algorithm has two disjoint sets  $M_1, M_2$  of  $m$  identical machines that will be used independently by the two procedures, respectively. Each machine in  $M_1$  is paired with an arbitrary distinct machine in  $M_2$ , and the actual load on a machine will be evaluated as the sum of the loads on the corresponding pair of machines. In other words, to show competitiveness it is sufficient to prove that all machines in both  $M_1$  and  $M_2$  have load  $O(\log d / \log \log d)$ .

Define the parameter  $\alpha := \frac{10 \log d}{\log \log d}$ . Our two procedures are formally defined as follows.

- **First procedure (random assignment).** Assign each job to one of the machines in  $M_1$  uniformly at random. Let  $J_j^1(i)$  denote the subset of the first  $j$  jobs  $\{1, 2, \dots, j\}$  that are assigned to machine  $i$  in this procedure, and let  $\Lambda_{i,j}^1(k)$  denote the resulting load on machine  $i$  on dimension  $k$  due to jobs in  $J_j^1(i)$ . If  $\Lambda_{i,j}^1(k) \geq 2\alpha + 1$  for some  $k \in [d]$ , then we pass job  $j$  to the second procedure. (However, note that *all* jobs are still scheduled by the first procedure; so even if a job  $j$  is passed to the second procedure after being assigned to machine  $i$  in the first procedure,  $j$  still contributes to the final load of  $\Lambda_{i,n}^1(k)$  for all  $k$ ).
- **Second procedure (greedy packing).** This procedure is only concerned with the jobs  $J^2$  that are passed from the first procedure. It allocates each job in  $J^2$  (in the order that the jobs arrive in) to one of the machines in  $M_2$  such that the resulting makespan,  $\max_{i \in M_2, k \in [d]} \Lambda_{i,j}^2(k)$  is minimized;  $\Lambda_{i,j}^2(k)$

is analogously defined for this second procedure as above.

This completes the description of the algorithm. We will let  $J^1(i) := J_n^1(i)$  and  $\Lambda_i^1(k) := \Lambda_{i,n}^1(k)$ , and define  $J^2(i)$  and  $\Lambda_i^2(k)$  similarly. We emphasize again that jobs in  $J^2$  are scheduled only on machines  $M_2$ ; all other jobs are scheduled on  $M_1$  machines.

**Algorithm analysis.** Before beginning our analysis, recall that by Property 4, the cost of the optimal solution must be at least 1; thus, in order to establish the upper bound in Theorem 1.2, it suffices to show that the algorithm's objective is  $O(\log d / \log \log d)$ .

First, observe that it follows directly from the definition of the algorithm that the loads on machines in  $M_1$  are at most  $2\alpha + 1 = O(\log d / \log \log d)$ . Therefore, we are only left with bounding the loads on machines in  $M_2$ . The following lemma shows that the second procedure receives only a small fraction of the total volume, which then allows us to argue that the greedy assignment in the second procedure is  $\alpha$ -competitive.

LEMMA 2.10. *The probability that a job  $j$  is passed to the second procedure is at most  $1/d^3$ , i.e.,  $\Pr[j \in J^2] \leq 1/d^3$ .*

*Proof.* Fix a machine  $i$ , job  $j$ , and dimension  $k$ . Suppose job  $j$  was assigned to machine  $i$  by the first procedure and is passed to the second procedure because we would have had  $\Lambda_{i,j}^1(k) \geq 2\alpha + 1$ . Since  $p_j(k) \leq 1$  due to Property 2, it follows that  $\Lambda_{i,j-1}^1(k) \geq 2\alpha$ . Therefore, we will show

$$(2.6) \quad \Pr[\Lambda_{i,j-1}^1(k) \geq 2\alpha] \leq 1/d^4,$$

where the probability space is over the random choices of jobs  $1, 2, \dots, j-1$ . Once inequality (2.6) is established, the lemma follows from a simple union bound over all dimensions.

To show (2.6), we use standard Chernoff–Hoeffding bounds (stated in Theorem 2.4 earlier). Note that  $\mathbb{E}[\Lambda_{i,j-1}^1(k)] \leq 2$  due to Property 1 and the fact that jobs are assigned to machines uniformly at random. To apply the inequality, we define random variables  $X_1, X_2, \dots, X_{j-1}$ , where  $X_{j'} = 1$  if job  $j'$  is assigned to machine  $i$ ; otherwise,  $X_{j'} = 0$ . Set the parameters of Theorem 2.4 as follows:  $a_{j'} = p_{j'}(k)$ ,  $\mu = 2$ , and  $\delta = \alpha - 1$ . Thus we have

$$\begin{aligned} \Pr[\Lambda_{i,j-1}^1(k) \geq 2\alpha] &= \Pr \left[ \sum_{j' \in [j-1]} a_{j'} X_{j'} \geq \alpha\mu \right] = \Pr \left[ \sum_{j' \in [j-1]} a_{j'} X_{j'} \geq (1 + \delta)\mu \right] \\ &\leq \left( \frac{e^\delta}{(1 + \delta)(1 + \delta)} \right)^\mu \leq \frac{e^\delta}{(1 + \delta)^{(1 + \delta)}} \\ &\leq 1 / (5 \log d / \log \log d)^{(5 \log d / \log \log d)} \\ &\leq 1/d^4 \quad (\text{for sufficiently large } d), \end{aligned}$$

as desired.  $\square$

Next, we upper bound the makespan of the second procedure in terms of its total volume of jobs  $V^2$ , i.e.,  $V^2 = \sum_{j \in J^2, k \in [d]} p_j(k)$ .

LEMMA 2.11.  $\max_{i \in M_2, k \in [d]} \Lambda_i^2(k) \leq V^2/m + 1$ .

*Proof.* For the sake of contradiction, suppose that at the end of the instance there exists a dimension  $k$  and machine  $i$  such that  $\Lambda_i^2(k) > V^2/m + 1$ . Let  $j$  be the job



that made machine  $i$  first cross this  $V^2/m + 1$  threshold in dimension  $k$ . For each machine  $i'$ , let  $k_{i'} = \arg \max_{k'} \Lambda_{i',j-1}^2(k')$  denote the dimension with maximum load on machine  $i'$  before  $j$  was assigned.

By Property 2 and the greediness of the algorithm, we have that  $\Lambda_{i',j-1}^2(k_{i'}) > V^2/m$  for all  $i'$ . Otherwise,  $j$  would have been assigned to a machine other than  $i$  resulting in a makespan less than  $V^2/m + 1$  (since  $\max_{k,j} p_j(k) \leq 1$ ). However, this implies that every machine in  $M_2$  has a dimension with more than  $V^2/m$  load. Clearly, this contradicts the definition of  $V^2$ .  $\square$

We are now ready to complete the analysis. From Lemma 2.10 and linearity of expectation, we know that

$$(2.7) \quad \mathbb{E}[V^2] \leq \frac{1}{d^3} \sum_{j,k \in [d]} p_j(k) \leq \frac{1}{d^3} \cdot 2dm = \frac{m}{d^2},$$

where the second inequality follows from Property 1. Hence, inequality (2.7) along with Lemma 2.11 imply that the second procedure yields an expected makespan of  $O(1)$ , which completes our analysis.

**2.2.2. Derandomized algorithm for vsmax-i.** Our derandomization borrows the technique developed in [14]. To derandomize the algorithm, we replace the first procedure—a uniformly random assignment—with a deterministic assignment guided by the following potential  $\Phi$ . Let  $f(x) := \alpha^x$  for notational simplicity. Recall that  $\alpha := 10 \log d / \log \log d$ :

$$\Phi_{i,k}(j) := f \left( \Lambda_{i,j}^1(k) - \frac{\alpha}{m} \sum_{j' \in [j]} p_{j'}(k) \right) \quad \forall i \in M_1, j \in [n], k \in [d],$$

$$\Phi(j) := \sum_{i \in M} \sum_{k=1}^d \Phi_{i,k}(j).$$

- **(New deterministic) first procedure.** Each job  $j$  is assigned to a machine  $i$  such that  $\Phi(j)$  is minimized. If  $\Lambda_{i,j}^1(k) \geq 3\alpha + 1$ , then  $j$  is added to queue  $J^2$  so that it can be scheduled by the greedy second procedure.

Again, as mentioned at the beginning of section 2.2.1, our derandomized algorithm also performs Transformations 2 and 3, and therefore, Properties 1 through 4 still hold; therefore, it again suffices to show that the algorithm's objective is  $O(\log d / \log \log d)$ . We begin by showing the following lemma.

LEMMA 2.12.  $\Phi(j)$  is nonincreasing in  $j$ .

*Proof.* Consider the arrival of job  $j$ . To structure our argument, we assume the algorithm still assigns  $j$  to a machine in  $M_1$  uniformly at random. Our goal now is to show that  $\mathbb{E}[\Phi(j)] \leq \Phi(j-1)$ , which implies the existence of a machine  $i$  such that assigning job  $j$  to the machine  $i$  leads to  $\Phi(j) \leq \Phi(j-1)$  (and such an assignment is actually found by the algorithm since its assignment maximizes the decrease in

potential). We bound  $\mathbb{E}[\Phi_{i,k}(j)]$  as follows.

$$\begin{aligned}
 E[\Phi_{i,k}(j)] &= \frac{1}{m} f \left( \Lambda_{i,j-1}^1 + p_j(k) - \frac{\alpha}{m} p_j(k) - \frac{\alpha}{m} \sum_{j' \in [j-1]} p_{j'}(k) \right) \\
 &\quad + \left( 1 - \frac{1}{m} \right) f \left( \Lambda_{i,j-1}^1 - \frac{\alpha}{m} p_j(k) - \frac{\alpha}{m} \sum_{j' \in [j-1]} p_{j'}(k) \right) \\
 &= \Phi_{i,k}(j-1) \cdot \alpha^{-\frac{\alpha}{m} p_j(k)} \cdot \left( \frac{1}{m} (\alpha^{p_j(k)} - 1) + 1 \right) \\
 (2.8) \quad &\leq \Phi_{i,k}(j-1) \cdot \alpha^{-\frac{\alpha}{m} p_j(k)} \left( \frac{p_j(k)}{m} (\alpha - 1) + 1 \right) \\
 (2.9) \quad &\leq \Phi_{i,k}(j-1) \cdot \exp \left( -(\alpha \log \alpha) \cdot \frac{p_j(k)}{m} \right) \exp \left( \frac{p_j(k)}{m} \cdot (\alpha - 1) \right) \\
 &\leq \Phi_{i,k}(j-1).
 \end{aligned}$$

Inequality (2.8) follows since  $\alpha^x - 1 \leq (\alpha - 1)x$  for  $x \in [0, 1]$ , and  $p_j(k) \leq 1$  due to Property 2. Inequality (2.9) follows from the fact that  $x + 1 \leq e^x$ . Therefore, by linearity of expectation, we have  $\mathbb{E}[\Phi(j)] \leq \Phi(j-1)$ , thereby proving the lemma.  $\square$

The next corollary follows from Lemma 2.12 and the simple observation that  $\Phi(0) = md$ .

COROLLARY 2.13.  $\Phi(n) \leq md$ .

As in section 2.2.1, it is straightforward to see that the algorithm forces machines in  $M_1$  to have makespan  $O(\alpha)$ , so we again focus on the second procedure of the algorithm. Here, we need a deterministic bound on the total volume  $V^2 = \sum_{j \in J^2} \sum_{k \in [d]} p_j(k)$  that can be scheduled on machines in  $M_2$ . Lemma 2.14 provides us with such a bound.

LEMMA 2.14.  $V^2 \leq m/d$ .

*Proof.* Consider a job  $j \in J_2$  that was assigned to machine  $i$  in the first procedure. Let  $k(j)$  be an arbitrary dimension  $k$  with  $\Lambda_{i,j}^1 \geq 3\alpha + 1$  (such a dimension exists since  $j \in J^2$ ). Let  $J_i^2(k) = \{j : j \in J^1(i) \cap J^2 \text{ and } k(j) = k\}$  denote the set of jobs  $j \in J^2$  that were assigned to machine  $i$  by the first procedure and are associated with dimension  $k$ . We upper bound  $V^2$  as follows:

$$\begin{aligned}
 V^2 &= \sum_{j \in J^2} \sum_{k' \in [d]} p_j(k') \\
 &= \sum_{i \in M_1, k \in [d]} \sum_{j \in J_i^2(k)} \sum_{k' \in [d]} p_j(k') \\
 &\quad \text{(since we associate job } j \in J^2 \text{ with a unique dimension } k(j)) \\
 &\leq \sum_{i \in M_1, k \in [d]} \sum_{j \in J_i^2(k)} d^2 p_j(k) \quad \text{(by Property 3)} \\
 (2.10) \quad &\leq d^2 \sum_{i \in M_1} \sum_{k \in [d]} (\Lambda_i^1(k) - 3\alpha)_+.
 \end{aligned}$$

To see why the last inequality holds, recall that  $\Lambda_{i,j}^1(k) \geq 3\alpha + 1$  when  $j \in J^1(i)$  and  $k = k(j)$ . This can happen only when  $\Lambda_{i,j-1}^1(k) \geq 3\alpha$  since  $p_j(k) \leq 1$  due to Property 2. Since  $\Lambda_{i,j'}^1(k)$  is nondecreasing in  $j'$ , the sum of  $p_j(k)$  over all such jobs  $j$  is at most  $(\Lambda_i^1(k) - 3\alpha)_+$ ; here  $(x)_+ := \max\{0, x\}$ .

We claim that for all  $i \in M_1, k \in [d]$ ,

$$(2.11) \quad \Phi_{i,k}(n) \geq \alpha^\alpha (\Lambda_{i,j}^1(k) - 3\alpha)_+. \quad \square$$

If  $\Lambda_{i,j}^1(k) - 3\alpha \leq 0$ , then the claim is obviously true since  $\Phi_{i,k}(n)$  is always nonnegative. Otherwise, we have

$$\Phi_{i,k}(n) \geq \alpha^{\Lambda_{i,n}^1(k) - 2\alpha} \geq \alpha^\alpha (\Lambda_{i,j}^1(k) - 3\alpha),$$

where the first inequality follows from Property 1. So in either case, (2.11) holds.

By combining (2.10), (2.11), Corollary 2.13, and recalling  $\alpha = \frac{10 \log d}{\log \log d}$ , we have

$$V^2 \leq d^2 \sum_i \sum_k (\Lambda_i^1(k) - 3\alpha)_+ \leq \frac{d^2}{\alpha^\alpha} \sum_i \sum_k \Phi_{i,k}(n) \leq \frac{d^3}{\alpha^\alpha} m \leq \frac{m}{d}.$$

By Lemma 2.11, we have  $\max_{i \in M_2, k \in [d]} \Lambda_i^2(k) \leq \frac{1}{m} V^2 + 1 = O(1)$ . Thus, we have shown that each of the two deterministic procedures yields a makespan of  $O(\alpha) = O(\log d / \log \log d)$ , thereby proving the upper bound.

**2.2.3. Algorithm for vsall-i.** We now present our  $O((\log d / \log \log d)^{\frac{r-1}{r}})$ -competitive algorithm for VSALL-I. Throughout this section, let  $A$  denote the  $O(\log d / \log \log d)$ -competitive algorithm for VSMAX-I defined in section 2.2.2. Our VSALL-I algorithm essentially works by using  $A$  as a black box; however, we will perform a smoothing transformation on large loads before scheduling jobs with  $A$ .

**Algorithm definition.** We will apply the following transformation to all jobs  $j$  that arrive online after Transformation 1 has been performed (note that this is in replacement of Transformations 2 and 3 defined in section 2.2.1).

- **Transformation 2.** If  $p_j(k) > 1$ , we reduce  $p_j(k)$  to be 1. If this load reduction is applied in dimension  $k$  for job  $j$ , we say  $j$  is *large* in  $k$ ; otherwise,  $j$  is *small* in dimension  $k$ . On this transformed instance, our algorithm simply schedules jobs using our VSMAX-I algorithm  $A$ .

**Algorithm analysis.** Let  $\alpha = O(\log d / \log \log d)$  be the competitive ratio of algorithm  $A$ . For our analysis in this section, let  $p_j(k)$  denote the load of job  $j$  in dimension  $k$  after only Transformation 1 is applied (in other words, Transformation 2 is only used by the algorithm). Call this instance the *scaled instance*. Observe that because of Transformation 1, the scaled instance has the following property:

- **Property 1.**  $\sum_{j \in J} p_j(k) = m$  for all  $k \in [d]$ .

Let  $\text{OPT}'(k, r)$  be the cost of the optimal solution of the scaled instance in dimension  $k$ . In Lemma 2.15, we establish two lower bounds on  $\text{OPT}'(k, r)^r$ .

LEMMA 2.15. *The following inequality holds:*

$$\text{OPT}'(k, r)^r \geq \max \left( \sum_{j \in J} p_j(k)^r, m \cdot \left( \sum_{j \in J} p_j(k) / m \right)^r \right) = \max \left( \sum_{j \in J} (p_j(k)^r), m \right).$$

*Proof.* Consider any fixed assignment of jobs, and let  $J'(i) \subseteq J$  be the set of jobs assigned to machine  $i$ . Consider any fixed  $k$ . The first lower bound (within the max in the statement of the lemma) follows since

$$\sum_{i \in M} \left( \sum_{j \in J'(i)} p_j(k) \right)^r \geq \sum_{i \in M} \sum_{j \in J'(i)} p_j(k)^r = \sum_{j \in J} p_j(k)^r.$$

The second lower bound is due to the convexity of  $x^r$  when  $r \geq 1$ , and the last equality follows from Property 1.  $\square$

Let  $J(i) \subseteq J$  be the set of jobs assigned to machine  $i$  by the online algorithm. Let  $\ell(i, k)$  and  $s(i, k)$  be the set of jobs assigned to machine  $i$  that are large and small in dimension  $k$ , respectively. For brevity, let  $\sigma_\ell(i, k) = \sum_{j \in \ell(i, k)} p_j(k)$  and  $\sigma_s(i, k) = \sum_{j \in s(i, k)} p_j(k)$ . Observe that since algorithm  $A$  is  $\alpha$ -competitive on an instance of VSMAX-1 (which ensures no machine has more than  $\alpha$  load), we obtain the following additional two properties for the algorithm's schedule:

- **Property 2.**  $|\ell(i, k)| \leq \alpha$  for all  $i \in M, k \in [d]$ .
- **Property 3.**  $\sigma_s(i, k) \leq \alpha$  for all  $i \in M, k \in [d]$ .

Using these additional properties, the next two lemmas will bound the contribution of both large and small loads to the objective; namely, we need to bound both  $\sigma_\ell(i, k)^r$  and  $\sum_i \sigma_s(i, k)^r$  in terms of  $\alpha$ . Lemma 2.16 provides this bound for large loads, while Lemma 2.17 will be used to bound small loads.

LEMMA 2.16. *The following inequality holds:*

$$\sigma_\ell(i, k)^r = \left( \sum_{j \in \ell(i, k)} p_j(k) \right)^r \leq \alpha^{r-1} \sum_{j \in \ell(i, k)} p_j(k)^r.$$

*Proof.* Let  $h = |\ell(i, k)|$ . Then, it follows that

$$\begin{aligned} \left( \sum_{j \in \ell(i, k)} p_j(k) \right)^r &= \left( \frac{1}{h} \sum_{j \in \ell(i, k)} (p_j(k) \cdot h) \right)^r \leq \frac{1}{h} \sum_{j \in \ell(i, k)} (p_j(k) \cdot h)^r \quad (\text{convexity of } x^r) \\ &= h^{r-1} \sum_{j \in \ell(i, k)} p_j(k)^r \leq \alpha^{r-1} \sum_{j \in \ell(i, k)} p_j(k)^r \quad (\text{Property 2}). \quad \square \end{aligned}$$

Recall that by Property 1, we have that  $\sigma_s(i, k) \leq m$ . Using this fact and along with Property 3, the general statement shown in Lemma 2.17 will immediately provide us with the desired bound on  $\sum_i \sigma_s(i, k)^r$  (stated formally in Corollary 2.18).

LEMMA 2.17. *Let  $f(x) = x^r$  for some  $r \geq 1$  whose domain is defined over a set of variables  $x_1, \dots, x_n \in [0, \alpha]$  where  $\alpha \geq 1$ . If  $\sum_{i=1}^m x_i \leq m$ , then*

$$\sum_{i=1}^m f(x_i) \leq 2m \alpha^{r-1}.$$

*Proof.* Let  $\tilde{f} = \sum_{i=1}^m f(x_i)$ . We claim that  $\tilde{f}$  is maximized when  $0 < x_i < \alpha$  for at most one  $i \in [m]$ . If there are two such variables  $x_i$  and  $x_j$  with  $0 < x_i \leq x_j < \alpha$ , it is easy to see that we can further increase  $\tilde{f}$  by decreasing  $x_i$  and increasing  $x_j$  by an infinitesimal equal amount (i.e.,  $x_i \leftarrow x_i - \epsilon$  and  $x_j \leftarrow x_j + \epsilon$ ) due to convexity of  $f$ .

Hence,  $\tilde{f}$  is maximized when the multiset  $\{x_i : i \in [m]\}$  has  $\lfloor m/\alpha \rfloor$  copies of  $\alpha$ , and one copy of  $m - \alpha \lfloor m/\alpha \rfloor$  (which is at most  $\alpha$ ), which gives,

$$(2.12) \quad \tilde{f} \leq \lfloor m/\alpha \rfloor f(\alpha) + f(m - \alpha \lfloor m/\alpha \rfloor).$$

If  $\lfloor m/\alpha \rfloor \geq 1$ , then it follows that

$$\begin{aligned} \sum_{i=1}^m f(x_i) = \tilde{f} &\leq (\lfloor m/\alpha \rfloor + 1)f(\alpha) && \text{(by (2.12) and since } m - \alpha \lfloor m/\alpha \rfloor \leq \alpha) \\ &\leq 2(m/\alpha)\alpha^r = 2m \alpha^{r-1} && \text{(since } \lfloor m/\alpha \rfloor \geq 1). \end{aligned}$$

In the case where  $m < \alpha$ ,  $\tilde{f}$  is maximized by setting  $x_i = m$  for a single  $i$  and  $x_{i'} = 0$  for all other  $i' \neq i$ . Therefore,  $\tilde{f} \leq f(m) = m^r \leq m\alpha^{r-1}$ .  $\square$

**COROLLARY 2.18.** *For all dimensions  $k$ ,  $\sum_{i \in M} \sigma_s(i, k)^r \leq 2m \alpha^{r-1}$ .*

We are now ready to bound  $\|\Lambda(k)\|_r$  against  $\text{OPT}'(k, r)$ .

**LEMMA 2.19.** *For all dimensions  $k$ ,  $\|\Lambda(k)\|_r = O(\alpha^{(r-1)/r})\text{OPT}'(k, r)$ , i.e., the  $L_r$  norm of the vector load is at most  $O(\alpha^{(r-1)/r})$  times the  $L_r$  norm of the vector load of the optimal solution.*

*Proof.* Using Lemmas 2.15, 2.16, and Corollary 2.18, we have the following bound for  $\|\Lambda(k)\|_r^r = \sum_{i \in M} \left( \sum_{j \in J(i)} p_j(k) \right)^r$ :

$$\begin{aligned} &\sum_{i \in M} \left( \sum_{j \in J(i)} p_j(k) \right)^r \\ &= \sum_{i \in M} \left( \sum_{j \in \ell(i, k)} p_j(k) + \sum_{j \in s(i, k)} p_j(k) \right)^r \\ &\leq \sum_{i \in M} \left( 2 \max \left( \sum_{j \in \ell(i, k)} p_j(k), \sum_{j \in s(i, k)} p_j(k) \right) \right)^r \\ &\leq 2^r \sum_{i \in M} \left( \left( \sum_{j \in \ell(i, k)} p_j(k) \right)^r + \left( \sum_{j \in s(i, k)} p_j(k) \right)^r \right) \\ &\leq 2^r \left( \alpha^{r-1} \sum_{j \in \ell(i, k)} p_j(k)^r + 2m \cdot \alpha^{r-1} \right) && \text{(by Lemma 2.16 and Corollary 2.18)} \\ &\leq (2^r \cdot 3\alpha^{r-1})\text{OPT}'(k, r)^r && \text{(by Lemma 2.15),} \end{aligned}$$

which, raising both the left-hand side (LHS) and right-hand side (RHS) to  $1/r$ , gives us

$$\|\Lambda(k)\|_r \leq \left( 2 \cdot 3^{1/r} \alpha^{(r-1)/r} \right) \text{OPT}'(k, r) = O(\alpha^{(r-1)/r})\text{OPT}'(k, r).$$

This completes the proof.  $\square$

The upper bound in Theorem 1.5 now follows immediately from Lemma 2.19.

**3. Unrelated machines.** Now, we consider the online vector scheduling problem for unrelated machines. In this section, we obtain tight upper and lower bounds for this problem, both for the makespan norm (Theorem 1.3) and for arbitrary  $L_r$  norms (Theorem 1.7).

**3.1. Lower bound for vsany-u.** In this section we prove the lower bound in Theorem 1.7, i.e., we show that we can force any algorithm to make an assignment where there exists a dimension  $k$  that has cost at least  $\Omega(\log d + r_k)$  where  $1 \leq r_k \leq \log m$ .

Our construction is an adaptation of the lower bounds in [15] and [6] but for a multidimensional setting. Informally, the instance is defined as follows. We set  $m = d$  and then associate the  $i$ th machine with the  $i$ th dimension, i.e., machine  $i$  only receives load in the  $i$ th dimension. We then issue jobs in a series of  $\log d + 1$  phases. In a given phase, there will be a current set of *active machines*, which are the only machines that can be loaded in the current phase and for the rest of the instance (so once a machine is inactivated it stays inactive). More specifically, in a given phase we arbitrarily pair off the active machines and then issue one job for each pair, where each job has unit load but is defined such that it must be assigned to a unique machine in its pair. When a phase completes, we inactivate all of the machines that did not receive load (so we cut the number of active machines in half). This process eventually produces a load of  $\log d + 1$  on some machine, whereas reversing the decisions of the algorithm gives an optimal schedule with an  $L_{r_k}$  norm of 1 for all  $k \in [d]$ .

More formally let  $d = 2^h$ . The adversary sets the instance target parameters to be  $T_k = 1$  for all  $k \in [d]$  (it will be clear from our construction that these targets are feasible). For each job  $j$ , let  $m_1(j), m_2(j) \in [m]$  denote the machine pair the adversary associates with job  $j$ . We define  $j$  to have unit load on machines  $m_1(j), m_2(j)$  in their respective dimensions and arbitrarily large load on all other machines. Formally,  $p_{i,j}(k)$  is defined to be

$$p_{i,j}(k) = \begin{cases} 0 & \text{if } i \neq k \text{ and } i \in \{m_1(j), m_2(j)\}, \\ 1 & \text{if } i = k \text{ and } i \in \{m_1(j), m_2(j)\}, \\ \infty & \text{otherwise.} \end{cases}$$

As discussed above, the adversary issues jobs in  $h + 1$  phases. Phases 1 through  $h$  will work as previously specified (we describe how the final  $(h + 1)$ th phase works shortly). Let  $S_\ell$  denote the active machines in phase  $\ell$ . In the  $\ell$ th phase, we issue a set of jobs  $J_\ell$  where  $|J_\ell| = 2^{h-\ell}$ . We then pair off the machines in  $S_\ell$  and use each machine pair as  $m_1(j)$  and  $m_2(j)$  for a unique job  $j \in J_\ell$ . Clearly the algorithm must schedule  $j$  on  $m_1(j)$  or  $m_2(j)$ , and thus  $2^{h-\ell}$  machines accumulate an additional load of 1 in phase  $\ell$ . Machines that receive jobs in phase  $\ell$  remain active in phase  $\ell + 1$ ; all other machines are set to be inactive. In the final phase  $h + 1$ , there will be a single remaining active machine  $i'$ ; thus, we issue a single job  $j'$  with unit load that must be scheduled on  $i'$  (note that this final phase is added to the instance only to make our target vector feasible).

Based on this construction, there will exist a dimension  $k'$  at the end of the instance that has load  $h + 1$  on machine  $k'$  and 0 on all other machines. Observe that the optimal schedule is obtained by reversing the decisions of the algorithm, which places a unit load on one machine in each dimension. Namely, if  $j$  was assigned to  $m_1(j)$ , then the optimal schedule assigns  $j$  to  $m_2(j)$  (and vice versa), with the exception that  $j'$  is assigned to its only feasible machine.

In the case that  $\log d \geq r_{k'}$ , the adversary stops. Since  $T_{k'} = 1$  and the  $L_{r_{k'}}$  norm of the algorithm in dimension  $k'$  is  $h + 1 = \log d + 1$ , we have that the competitive ratio of the algorithm is  $\Omega(\log d + r_k)$ . If  $\log d < r_{k'}$ , then the adversary stops the current instance and begins a new instance. In the new instance, we simply simulate the lower bound from [6] in dimension  $k'$  (i.e., the only dimension that receives load is dimension  $k'$ ; the adversary also resets the target vectors accordingly). Here, the adversary forces the algorithm to be  $\Omega(r_{k'})$ -competitive, which, since  $\log d < r_{k'}$ , gives us the desired bound of  $\Omega(\log d + r_{k'})$ .

**3.2. Upper bound.** Our goal is to prove the upper bound in Theorem 1.7. Recall that we are given targets  $T_k$ , and we have to show that  $\|\Lambda(k)\|_{r_k} = O(\log d + r_k) \cdot T_k$  for all  $k \in [d]$ . ( $\Lambda(k)$  is the load vector in dimension  $k$  and  $r_k$  is the norm that we are optimizing.) First, we normalize  $p_{i,j}(k)$  to  $p_{i,j}(k)/T_k$  for all dimensions  $k$ ; to keep the notation simple, we will also denote this normalized load  $p_{i,j}(k)$ . This ensures that the target objective is 1 in every dimension. (We assume without loss of generality that  $T_k > 0$ . If  $T_k = 0$ , the algorithm discards all assignments that put nonzero load on dimension  $k$ ).

**3.2.1. Description of the algorithm.** As described in the introduction, our algorithm is greedy with respect to a potential function defined on modified  $L_{r_k}$  norms. Let  $L_k = \|\Lambda(k)\|_{r_k}$  denote the  $L_{r_k}$  norm of the machine loads in the  $k$ th dimension, and  $q_k = r_k + \log d$  denote the desired competitive ratio; all logs are base 2. We define the potential for dimension  $k$  as  $\Phi_k = L_k^{q_k}$ . The potentials for the  $d$  different dimensions are combined using a weighted linear combination, where the weight of dimension  $k$  is  $\alpha_k = (3q_k)^{-q_k}$ . Note that dimensions that allow a smaller slack in the competitive ratio are given a larger weight in the potential. We denote the combined potential by  $\Phi = \sum_{k=1}^d \alpha_k \cdot \Phi_k$ . The algorithm assigns job  $j$  to the machine that minimizes the increase in potential  $\Phi$ .

**3.2.2. Competitive analysis.** Let us fix a solution satisfying the target objectives, and call it the optimal solution. Let  $\Lambda_i(k)$  and  $\Lambda_i^*(k)$  be the load on the  $i$ th machine in the  $k$ th dimension for the algorithmic solution and the optimal solution, respectively. We also use  $L_k^*$  to denote the  $L_{r_k}$  norm in the  $k$ th dimension for the optimal solution; we have already asserted that by scaling,  $L_k^* = 1$ .

Similar to [4, 15], we compare the actual assignment made by the algorithm (starting with zero load on every machine in every dimension) to a hypothetical assignment made by the optimal solution starting with the final algorithmic load on every machine (i.e., load of  $\Lambda_i(k)$  on machine  $i$  in dimension  $k$ ).

We will need the following fact for our analysis, which follows by observing that all parameters are positive, the function is continuous in the domain, and its derivative is nonnegative.

**FACT 3.1.** *The function  $f(x_1, x_2, \dots, x_m) = (\sum_i (x_i + a_i)^w)^z - (\sum_i x_i^w)^z$  is non-decreasing if for all  $i \in [m]$  we restrict the domain of  $x_i$  to be  $[0, \infty)$ ,  $w \geq 1$ ,  $z \geq 1$ , and  $a_i \geq 0$ .*

Using greediness of the algorithm and convexity of the potential function, we argue in Lemma 3.2 that the change in potential in the former process is upper bounded by that in the latter process.

**LEMMA 3.2.** *The total change in potential in the online algorithm satisfies*

$$\sum_{k=1}^d \alpha_k L_k^{q_k} = \Phi(n) - \Phi(0) \leq \sum_{k=1}^d \alpha_k \left( \sum_{i=1}^m \left( \Lambda_i(k) + \Lambda_i^*(k) \right)^{r_k} \right)^{q_k/r_k} - \sum_{k=1}^d \alpha_k L_k^{q_k}.$$

*Proof.* Let  $y_{i,j} = 1$  if the algorithm assigns job  $j$  to machine  $i$ ; otherwise,  $y_{i,j} = 0$ . Define  $y_{i,j}^*$  similarly but for the optimal solution's assignments. We can express the resulting change in potential from scheduling job  $j$  as follows.

$$\begin{aligned} & \Phi(j) - \Phi(j-1) \\ &= \sum_{k=1}^d \alpha_k (L_k^{q_k}(j) - L_k^{q_k}(j-1)) = \sum_{k=1}^d \alpha_k \left( \left( \sum_{i=1}^m \Lambda_{i,j}^{r_k}(k) \right)^{q_k/r_k} - \left( \sum_{i=1}^m \Lambda_{i,j-1}^{r_k}(k) \right)^{q_k/r_k} \right) \\ (3.1) \quad &= \sum_{k=1}^d \alpha_k \left( \left( \sum_{i=1}^m \left( \Lambda_{i,j-1}(k) + p_{i,j}(k) \cdot y_{i,j} \right)^{r_k} \right)^{q_k/r_k} - \left( \sum_{i=1}^m \Lambda_{i,j-1}^{r_k}(k) \right)^{q_k/r_k} \right). \end{aligned}$$

Since the online algorithm schedules greedily based on  $\Phi(j)$ , using the optimal schedule's assignment for job  $j$  must result in a potential increase that is at least as large. Therefore, by (3.1) we have

$$(3.2) \quad \Phi(j) - \Phi(j-1) \leq \sum_{k=1}^d \alpha_k \left( \left( \sum_{i=1}^m \left( \Lambda_{i,j-1}(k) + p_{i,j}(k) \cdot y_{i,j}^* \right)^{r_k} \right)^{q_k/r_k} - \left( \sum_{i=1}^m \Lambda_{i,j-1}^{r_k}(k) \right)^{q_k/r_k} \right).$$

As loads are nondecreasing,  $\Lambda_i(k) \geq \Lambda_{i,j-1}(k)$ . Also note that  $r_k \geq 1$  and

$$q_k/r_k = (r_k + \log d)/r_k > 1.$$

Thus, we can apply Fact 3.1 to (3.2) (setting  $w = r_k$ ,  $z = q_k/r_k$ , and  $a_i = p_{i,j}(k) \cdot y_{i,j}^*$ ) to obtain

$$(3.3) \quad \Phi(j) - \Phi(j-1) \leq \sum_{k=1}^d \alpha_k \left( \left( \sum_{i=1}^m \left( \Lambda_i(k) + p_{i,j}(k) \cdot y_{i,j}^* \right)^{r_k} \right)^{q_k/r_k} - \left( \sum_{i=1}^m \Lambda_i^{r_k}(k) \right)^{q_k/r_k} \right).$$

We can again use Fact 3.1 to further bound the potential increase (using the same



values of  $a_i$ ,  $w$ , and  $z$ , but now  $\Delta x_i = \Lambda_{i,j-1}^*(k)$ :

(3.4)

$$\begin{aligned} & \Phi(j) - \Phi(j-1) \\ & \leq \sum_{k=1}^d \alpha_k \left( \left( \sum_{i=1}^m \left( \Lambda_i(k) + \Lambda_{i,j-1}^*(k) + p_{i,j}(k) \cdot y_{ij}^* \right)^{r_k} \right)^{q_k/r_k} \right. \\ & \quad \left. - \left( \sum_{i=1}^m \left( \Lambda_i(k) + \Lambda_{i,j-1}^*(k) \right)^{r_k} \right)^{q_k/r_k} \right) \end{aligned}$$

(3.5)

$$= \sum_{k=1}^d \alpha_k \left( \left( \sum_{i=1}^m \left( \Lambda_i(k) + \Lambda_{i,j}^*(k) \right)^{r_k} \right)^{q_k/r_k} - \left( \sum_{i=1}^m \left( \Lambda_i(k) + \Lambda_{i,j-1}^*(k) \right)^{r_k} \right)^{q_k/r_k} \right).$$

Observe that for a fixed  $k$ , the RHS of (3.5) is a telescoping series if we sum over all jobs  $j$ :

$$\begin{aligned} (3.6) \quad & \sum_{j=1}^n \alpha_k \left( \left( \sum_{i=1}^m \left( \Lambda_i(k) + \Lambda_{i,j}^*(k) \right)^{r_k} \right)^{q_k/r_k} - \left( \sum_{i=1}^m \left( \Lambda_i(k) + \Lambda_{i,j-1}^*(k) \right)^{r_k} \right)^{q_k/r_k} \right) \\ & = \alpha_k \left( \left( \sum_{i=1}^m \left( \Lambda_i(k) + \Lambda_i^*(k) \right)^{r_k} \right)^{q_k/r_k} - \left( \sum_{i=1}^m \left( \Lambda_i(k) \right)^{r_k} \right)^{q_k/r_k} \right). \end{aligned}$$

We have

$$\sum_{j=1}^n (\Phi(j) - \Phi(j-1)) = \Phi(n) - \Phi(0),$$

since this is also a telescoping series. By definition,  $\Phi(0) = 0$  and  $\Phi(n) = \sum_{k=1}^d \alpha_k L_k^{q_k}$ . Using these facts along with (3.5) and (3.6), we establish the lemma:

$$\begin{aligned} \sum_{k=1}^d \alpha_k L_k^{q_k} &= \sum_{j=1}^n (\Phi(j) - \Phi(j-1)) \\ & \left( \text{since } \Phi \text{ telescopes, } \Phi(0) = 0, \text{ and } \Phi(n) = \sum_{k=1}^d \alpha_k L_k^{q_k} \right) \\ & \leq \sum_{k=1}^d \alpha_k \left( \sum_{i=1}^m \left( \Lambda_i(k) + \Lambda_i^*(k) \right)^{r_k} \right)^{q_k/r_k} - \sum_{k=1}^d \alpha_k L_k^{q_k} \quad (\text{by (3.5) and (3.6)}). \end{aligned}$$

This completes the proof.  $\square$

We proceed by applying the Minkowski inequality (e.g., [44]), which states that for any two vectors  $v_1$  and  $v_2$ , we have  $\|v_1 + v_2\|_r \leq \|v_1\|_r + \|v_2\|_r$ . Applying this

inequality to the RHS in Lemma 3.2, we obtain

$$\begin{aligned}
 \sum_{k=1}^d \alpha_k L_k^{q_k} &\leq \sum_{k=1}^d \alpha_k \left( \left( \sum_{i=1}^m \Lambda_i^{r_k}(k) \right)^{1/r_k} + \left( \sum_{i=1}^m (\Lambda_i^*(k))^{r_k} \right)^{1/r_k} \right)^{q_k} - \sum_{k=1}^d \alpha_k L_k^{q_k} \\
 (3.7) \quad &= \sum_{k=1}^d \alpha_k (L_k + L_k^*)^{q_k} - \sum_{k=1}^d \alpha_k L_k^{q_k}.
 \end{aligned}$$

Next, we prove a simple lemma that we will apply to inequality (3.7).

LEMMA 3.3.  $(L_k + L_k^*)^{q_k} \leq e^{1/2} L_k^{q_k} + (3q_k \cdot L_k^*)^{q_k}$  for all  $k \in [d]$ .

*Proof.* First, consider the case  $L_k \geq 2q_k \cdot L_k^*$ . Then it follows,

$$\begin{aligned}
 (L_k + L_k^*)^{q_k} &\leq (1 + 1/(2q_k))^{q_k} \cdot L_k^{q_k} \\
 (3.8) \quad &\leq (e^{1/(2q_k)})^{q_k} \cdot L_k^{q_k} = e^{1/2} L_k^{q_k}.
 \end{aligned}$$

Otherwise,  $L_k < 2q_k \cdot L_k^*$ , and then we have

$$(L_k + L_k^*)^{q_k} \leq (3q_k \cdot L_k^*)^{q_k}.$$

Combining these two upper bounds completes the proof. □

Thus, we can rearrange (3.7) and bound  $2 \sum_{k=1}^d \alpha_k L_k^{q_k}$  as follows:

$$\begin{aligned}
 2 \sum_{k=1}^d \alpha_k L_k^{q_k} &\leq \sum_{k=1}^d \alpha_k (L_k + L_k^*)^{q_k} \\
 &\leq e^{1/2} \sum_{k=1}^d \alpha_k L_k^{q_k} + \sum_{k=1}^d \alpha_k (3q_k \cdot L_k^*)^{q_k} \quad (\text{by Lemma 3.3}) \\
 (3.9) \quad &= e^{1/2} \sum_{k=1}^d \alpha_k L_k^{q_k} + \sum_{k=1}^d (L_k^*)^{q_k}.
 \end{aligned}$$

Note that the last equality is due to the fact that  $\alpha_k^{-1} = (3q_k)^{q_k}$ . By our initial scaling,  $L_k^* = 1$  for all  $k$ . Therefore, after rearranging (3.9), we obtain

$$(2 - e^{1/2}) \sum_{k=1}^d \alpha_k L_k^{q_k} \leq \sum_{k=1}^d (L_k^*)^{q_k} \leq d,$$

which for any fixed  $k$  implies

$$\begin{aligned}
 L_k &\leq \frac{1}{(2 - e^{1/2})^{1/q_k}} \cdot \left( \frac{d}{\alpha_k} \right)^{1/q_k} \leq \frac{1}{2 - e^{1/2}} \cdot \left( \frac{d}{\alpha_k} \right)^{1/q_k} \\
 &= \frac{3}{2 - e^{1/2}} \cdot \left( d^{\frac{1}{r_k + \log d}} \right) q_k < 10 \cdot d^{\frac{1}{\log d}} \cdot q_k = 20q_k = O(r_k + \log d),
 \end{aligned}$$

where the second inequality uses  $q_k \geq 1$  and  $2 - e^{1/2} < 1$ . This completes the proof of the upper bound claimed in Theorem 1.7.

**4. Conclusions.** In this paper, we presented algorithms with tight competitive ratios for the online vector scheduling problem in both the identical and unrelated machines settings. To conclude the paper, we briefly mention some natural directions for future work.

In both the unrelated and identical machines settings, our bounds are tight in an asymptotic sense, i.e., our results apply to problem instances with a large number of dimensions. However, in many real world settings with multidimensional jobs, the number of dimensions is likely to be a small constant. Therefore, an interesting direction for future work would be to obtain improved competitive ratios when the number of dimensions is small, e.g.,  $d = 2, 3, 4$ , etc. In fact, for problem instances with a small number of dimensions, our algorithms have worse competitive ratios than even naïve greedy algorithms. More specifically, a simple greedy strategy yields a  $(d + 1)$ -competitive algorithm for identical machines with  $d$  dimensions, and thus improving on this bound is a very interesting question, even for the case when  $d = 2$ .

For another natural direction, recall that we made a technical assumption in our VSALL-I algorithm; in particular, we assumed that in each dimension, the cumulative volume across all jobs, along with the load value of the job with maximum load, are given to the algorithm a priori at the start of the instance. Such assumptions are easy to remove in single dimensional problems by using the so-called “guess-and-double” strategy; however, it is unclear how to remove these assumptions when scheduling vector jobs. Therefore, an interesting technical question is determine whether we can design an online algorithm for VSALL-I that does not need to know the cumulative volume of all jobs, nor the value of the largest load, in each dimension.

## REFERENCES

- [1] F. AHMAD, S. T. CHAKRADHAR, A. RAGHUNATHAN, AND T. N. VIJAYKUMAR, *Tarazu: Optimizing mapreduce on heterogeneous clusters*, in ACM SIGARCH Computer Architecture News, Vol. 40, ACM, New York, 2012, pp. 61–74.
- [2] S. ALBERS, *Better bounds for online scheduling*, SIAM J. Comput., 29 (1999), pp. 459–473, <https://doi.org/10.1137/S0097539797324874>.
- [3] S. ALBERS, *Energy-efficient algorithms*, Comm. ACM, 53 (2010), pp. 86–96.
- [4] J. ASPNES, Y. AZAR, A. FIAT, S. A. PLOTKIN, AND O. WAARTS, *On-line routing of virtual circuits with applications to load balancing and machine scheduling*, J. ACM, 44 (1997), pp. 486–504.
- [5] A. AVIDOR, Y. AZAR, AND J. SGALL, *Ancient and new algorithms for load balancing in the  $l_p$  norm*, Algorithmica, 29 (2001), pp. 422–441.
- [6] B. AWERBUCH, Y. AZAR, E. F. GROVE, M.-Y. KAO, P. KRISHNAN, AND J. S. VITTER, *Load balancing in the  $l_p$  norm*, in IEEE Symposium on Foundations of Computer Science, 1995, pp. 383–391.
- [7] Y. AZAR, *On-line load balancing*, in Online Algorithms, Lecture Notes in Comput. Sci. 1442, Springer, Berlin, 1998, pp. 178–195.
- [8] Y. AZAR, I. R. COHEN, S. KAMARA, AND F. B. SHEPHERD, *Tight bounds for online vector bin packing*, in ACM Symposium on Theory of Computing, ACM, New York, 2013, pp. 961–970.
- [9] Y. AZAR, J. NAOR, AND R. ROM, *The competitiveness of on-line assignments*, J. Algorithms, 18 (1995), pp. 221–237.
- [10] Y. BARTAL, A. FIAT, H. J. KARLOFF, AND R. VOHRA, *New algorithms for an ancient scheduling problem*, J. Comput. System Sci., 51 (1995), pp. 359–366.
- [11] Y. BARTAL, H. J. KARLOFF, AND Y. RABANI, *A better lower bound for on-line scheduling*, Inform. Process. Lett., 50 (1994), pp. 113–116.
- [12] P. BERMAN, M. CHARIKAR, AND M. KARPINSKI, *On-line load balancing for related machines*, J. Algorithms, 35 (2000), pp. 108–121.
- [13] A. BORODIN AND R. EL-YANIV, *Online Computation and Competitive Analysis*, Cambridge University Press, New York, 1998.

- [14] N. BUCHBINDER AND J. NAOR, *Online primal-dual algorithms for covering and packing problems*, in European Symposium on Algorithms, Lecture Notes in Comput. Sci. 3669, Springer, Berlin, 2005, pp. 689–701.
- [15] I. CARAGIANNIS, *Better bounds for online load balancing on unrelated machines*, in Proceedings of the ACM-SIAM Symposium on Discrete Algorithms, ACM, New York, 2008, pp. 972–981.
- [16] I. CARAGIANNIS, M. FLAMMINI, C. KAKLAMANIS, P. KANELLOPOULOS, AND L. MOSCARDELLI, *Tight bounds for selfish and greedy load balancing*, *Algorithmica*, 61 (2011), pp. 606–637.
- [17] A. K. CHANDRA AND C. K. WONG, *Worst-case analysis of a placement algorithm related to storage allocation*, *SIAM J. Comput.*, 4 (1975), pp. 249–263, <https://doi.org/10.1137/0204021>.
- [18] C. CHEKURI AND S. KHANNA, *On multidimensional packing problems*, *SIAM J. Comput.*, 33 (2004), pp. 837–851, <https://doi.org/10.1137/S0097539799356265>.
- [19] G. CHRISTODOULOU, V. S. MIRROKNI, AND A. SIDIROPOULOS, *Convergence and approximation in potential games*, *Theoret. Comput. Sci.*, 438 (2012), pp. 13–27.
- [20] R. A. CODY AND E. G. COFFMAN, JR., *Record allocation for minimizing expected retrieval costs on drum-like storage devices*, *J. ACM*, 23 (1976), pp. 103–115.
- [21] U. FAIGLE, W. KERN, AND G. TURÁN, *On the performance of on-line algorithms for partition problems*, *Acta Cybernet.*, 9 (1989), pp. 107–119.
- [22] R. FLEISCHER AND M. WAHL, *Online scheduling revisited*, in European Symposium on Algorithms, Lecture Notes in Comput. Sci. 1879, Springer, Berlin, pp. 202–210.
- [23] A. GHODSI, M. ZAHARIA, B. HINDMAN, A. KONWINSKI, S. SHENKER, AND I. STOICA, *Dominant resource fairness: Fair allocation of multiple resource types*, in Proceedings of the 8th USENIX Symposium on Networked Systems Design and Implementation, USENIX Association Berkeley, CA, 2011, pp. 323–336.
- [24] A. GHODSI, M. ZAHARIA, S. SHENKER, AND I. STOICA, *Choosy: Max-min fair sharing for datacenter jobs with constraints*, in EuroSys '13 Proceedings of the 8th ACM European Conference on Computer Systems, ACM, New York, 2013, pp. 365–378.
- [25] T. GORMLEY, N. REINGOLD, E. TORNG, AND J. WESTBROOK, *Generating adversaries for request-answer games*, in Proceedings of the Eleventh Annual ACM-SIAM Symposium on Discrete Algorithms, SIAM, Philadelphia, 2000, pp. 564–565.
- [26] R. L. GRAHAM, *Bounds for certain multiprocessing anomalies*, *Bell Syst. Tech. J.*, 45 (1966), pp. 1563–1581.
- [27] R. GRANDL, G. ANANTHANARAYANAN, S. KANDULA, S. RAO, AND A. AKELLA, *Multi-resource packing for cluster schedulers*, in ACM Special Interest Group on Data Communication, 2014, pp. 455–466.
- [28] A. GULATI, G. SHANMUGANATHAN, A. HOLLER, C. WALDSPURGER, M. JI, AND X. ZHU, *Vmware distributed resource management: Design, implementation, and lessons learned*, available online at <https://labs.vmware.com/vmtj/vmware-distributed-resource-management-design-implementation-and-lessons-learned>.
- [29] M. M. HALLDÓRSSON AND M. SZEGEDY, *Lower bounds for on-line graph coloring*, *Theoret. Comput. Sci.*, 130 (1994), pp. 163–174.
- [30] D. G. HARRIS AND A. SRINIVASAN, *The Moser-Tardos framework with partial resampling*, in Proceedings of the IEEE Symposium on Foundations of Computer Science, IEEE, Los Alamitos, CA, 2013, pp. 469–478.
- [31] J. F. RUDIN III, *Improved Bound for the On-Line Scheduling Problem*, Ph.D. thesis, The University of Texas at Dallas, 2001.
- [32] S. IM, N. KELL, J. KULKARNI, AND D. PANIGRAHI, *Tight bounds for online vector scheduling*, in Proceedings of the IEEE Symposium on Foundations of Computer Science, IEEE, Los Alamitos, CA, 2015, pp. 525–544.
- [33] G. JUNG, K. R. JOSHI, M. A. HILTUNEN, R. D. SCHLICHTING, AND C. PU, *Generating adaptation policies for multi-tier applications in consolidated server environments*, in International Conference on Autonomic Computing, Washington, DC, 2008, pp. 23–32.
- [34] D. R. KARGER, S. J. PHILLIPS, AND E. TORNG, *A better algorithm for an ancient scheduling problem*, *J. Algorithms*, 20 (1996), pp. 400–430.
- [35] V. S. A. KUMAR, M. V. MARATHE, S. PARTHASARATHY, AND A. SRINIVASAN, *A unified approach to scheduling on unrelated parallel machines*, *J. ACM*, 56 (2009), 28.
- [36] G. LEE, B.-G. CHUN, AND R. H. KATZ, *Heterogeneity-aware resource allocation and scheduling in the cloud*, in USENIX Workshop on Hot Topics in Cloud Computing, USENIX Association Berkeley, CA, 2011, 4.

- [37] A. MEYERSON, A. ROYTMAN, AND B. TAGIKU, *Online multidimensional load balancing*, in International Workshop on Approximation Algorithms for Combinatorial Optimization Problems, Lecture Notes in Comput. Sci. 8096, Springer, Heidelberg, 2013, pp. 287–302.
- [38] R. MOTWANI AND P. RAGHAVAN, *Randomized Algorithms*, Cambridge University Press, Cambridge, 1995.
- [39] T. MÜTZE, T. RAST, AND R. SPÖHEL, *Coloring random graphs online without creating monochromatic subgraphs*, Random Structures Algorithms, 44 (2014), pp. 419–464.
- [40] S. PELLELY, D. MEISNER, T. F. WENISCH, AND J. W. VANGILDER, *Understanding and abstracting total data center power*, in Workshop on Energy-Efficient Design, 2009.
- [41] K. PRUHS, J. SGALL, AND E. TORNG, *Online Scheduling*, Handbook of Scheduling: Algorithms, Models, and Performance Analysis, 2004.
- [42] J. SGALL, *On-line scheduling*, in Online Algorithms, Lecture Notes in Comput. Sci. 1442, Springer, Berlin, 1998, pp. 196–231.
- [43] S. SURI, C. D. TÓTH, AND Y. ZHOU, *Selfish load balancing and atomic congestion games*, Algorithmica, 47 (2007), pp. 79–96.
- [44] WIKIPEDIA, *Minkowski inequality*, [https://en.wikipedia.org/wiki/Minkowski\\_inequality](https://en.wikipedia.org/wiki/Minkowski_inequality).
- [45] F. YAO, A. DEMERS, AND S. SHENKER, *A scheduling model for reduced CPU energy*, in IEEE Symposium on Foundations of Computer Science, 1995, pp. 374–382.
- [46] S. ZHANG, B. WANG, B. ZHAO, AND J. TAO, *An energy-aware task scheduling algorithm for a heterogeneous data center*, in Proceedings of the 12th IEEE International Conference on Trust, Security and Privacy in Computing and Communications, 2013, pp. 1471–1477.