

# Memoryless Worker-Task Assignment with Polylogarithmic Switching Cost

Aaron Berger   




MIT, Cambridge, MA, USA

William Kuszmaul   

MIT, Cambridge, MA, USA

Adam Polak   

EPFL, Lausanne, Switzerland

Jonathan Tidor   

MIT, Cambridge, MA, USA

Nicole Wein   

DIMACS, Piscataway, NJ, USA

---

## Abstract

We study the basic problem of assigning memoryless workers to tasks with dynamically changing demands. Given a set of  $w$  workers and a multiset  $T \subseteq [t]$  of  $|T| = w$  tasks, a memoryless worker-task assignment function is any function  $\phi$  that assigns the workers  $[w]$  to the tasks  $T$  based only on the current value of  $T$ . The assignment function  $\phi$  is said to have switching cost at most  $k$  if, for every task multiset  $T$ , changing the contents of  $T$  by one task changes  $\phi(T)$  by at most  $k$  worker assignments. The goal of memoryless worker task assignment is to construct an assignment function with the smallest possible switching cost.

In past work, the problem of determining the optimal switching cost has been posed as an open question. There are no known sub-linear upper bounds, and after considerable effort, the best known lower bound remains 4 (ICALP 2020).

We show that it is possible to achieve polylogarithmic switching cost. We give a construction via the probabilistic method that achieves switching cost  $O(\log w \log(wt))$  and an explicit construction that achieves switching cost  $\text{polylog}(wt)$ . We also prove a super-constant lower bound on switching cost: we show that for any value of  $w$ , there exists a value of  $t$  for which the optimal switching cost is  $w$ . Thus it is not possible to achieve a switching cost that is sublinear *strictly* as a function of  $w$ .

Finally, we present an application of the worker-task assignment problem to a metric embeddings problem. In particular, we use our results to give the first low-distortion embedding from sparse binary vectors into low-dimensional Hamming space.

**2012 ACM Subject Classification** Theory of computation  $\rightarrow$  Distributed algorithms; Theory of computation  $\rightarrow$  Random projections and metric embeddings; Mathematics of computing  $\rightarrow$  Combinatorics

**Keywords and phrases** Distributed Task Allocation, Metric Embeddings, Probabilistic Method

**Digital Object Identifier** 10.4230/LIPIcs.ICALP.2022.19

**Category** Track A: Algorithms, Complexity and Games

**Related Version** *Full Version:* <https://arxiv.org/abs/2008.10709>

**Funding** *Aaron Berger:* NSF Graduate Research Fellowship Program DGE-1745302.

*William Kuszmaul:* NSF GRFP fellowship and a Fannie and John Hertz Fellowship. Research was partially sponsored by the United States Air Force Research Laboratory and was accomplished under Cooperative Agreement Number FA8750-19-2-1000. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official



© Aaron Berger, William Kuszmaul, Adam Polak, Jonathan Tidor, and Nicole Wein; licensed under Creative Commons License CC-BY 4.0

49th International Colloquium on Automata, Languages, and Programming (ICALP 2022).

Editors: Mikołaj Bojańczyk, Emanuela Merelli, and David P. Woodruff;

Article No. 19; pp. 19:1–19:19



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



policies, either expressed or implied, of the United States Air Force or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

*Adam Polak:* Swiss National Science Foundation project *Lattice Algorithms and Integer Programming* (185030). Part of this work was done at Jagiellonian University, supported by National Science Center of Poland grant 2017/27/N/ST6/01334.

*Jonathan Tidor:* NSF Graduate Research Fellowship Program DGE-1745302.

*Nicole Wein:* This work was done at MIT, supported by NSF Grant CCF-1514339.

## 1 Introduction

The general problem of **distributed task allocation**, where a group of agents must collectively allocate themselves to tasks, has been studied in a wide variety of settings where the agents have varying degrees of communication, memory, knowledge of the system, and faultiness (see Georgiou and Shvartsman’s book [15] for a survey). In this paper we are interested in the dynamic version of a distributed task allocation problem, that is, where the demands for each task are changing over time. Dynamic task allocation has been the focus of a great deal of both empirical and theoretical work in areas such as swarm robotics [20, 32, 26, 27, 21, 23, 34] and collective insect behavior [4, 31, 30, 33, 34].

Although there are many possible variations of the dynamic task allocation problem (in particular, in terms of what the capabilities of the workers are), most share a basic common structure. For positive integers  $w$  and  $t$ , there are  $w$  workers  $1, 2, \dots, w$ , and there is a multiset  $T \subseteq [t]$  of  $|T| = w$  tasks.<sup>1</sup> The multiset  $T$  changes gradually over time: in each time step, one new task is added and one old task is removed. The goal is to maintain a dynamic assignment of workers  $\{1, 2, \dots, w\}$  to tasks  $T$  such that the **switching cost** is as small as possible, i.e., the number of worker-task assignments that change each time that  $T$  changes is bounded by some small quantity.

The goal of studying dynamic distributed task allocation is to answer the following question: to what degrees are various capabilities (i.e., memory, communication, knowledge of the system, computational power, etc.) needed for the workers to guarantee a small switching cost?

Motivated in part by applications to swarm robotics and collective ant behavior, recent work [33, 34] has focused on the question of what happens when workers are *completely memoryless*. At any given moment each worker  $i \in \{1, 2, \dots, w\}$  must determine which task  $\tau \in T$  they are assigned to solely as a function of the current task multiset  $T$ ; the workers do not remember anything about the past system states or worker assignments. Note that this memoryless-ness requirement is sufficiently strong that it also implies communicationless-ness – indeed, if workers cannot remember where they were in previous steps, and all that each worker knows is the current multiset of tasks, then there is no worker-specific information to be shared. Being memoryless and communicationless is especially important to settings where a worker might suffer a fault and thus “reboot”; in this case, the worker can determine which task they are assigned to without relying on what was stored in its memory and without knowledge of which tasks other workers are assigned to at the moment [34].

---

<sup>1</sup> If a task  $j$  appears  $m_T(j)$  times in the multiset  $T$ , then one should think of the task as having a current demand of  $m_T(j)$  workers. By writing  $T \subseteq [t]$  we mean that the elements of multiset  $T$  are from  $\{1, 2, \dots, t\}$ , but the multiplicity of each element can be arbitrarily large.

It remains an open question [34, 33] whether memoryless workers can achieve even a *sub-linear* switching cost. In this paper we show that, not only is sub-linear switching cost possible, but it is even possible to achieve a polylogarithmic switching cost of  $\text{polylog}(wt)$ . We also prove a lower bound for any algorithm that wishes to parameterize by only  $w$  and not  $t$ : for any  $w$ , if  $t \gg w$  is sufficiently large, then the trivial switching cost of  $w$  becomes optimal.

**Formal problem statement.** Throughout this paper, we shall refer to the memoryless dynamic distributed task allocation problem simply as the **worker-task assignment problem**. Formally, the worker-task assignment problem is defined as follows. There are  $w$  workers  $1, 2, \dots, w$  and  $t$  tasks  $1, 2, \dots, t$ . A **worker-task assignment function**  $\phi$  is a function that takes as input a multiset  $T$  of  $w$  tasks, and produces an assignment of workers to tasks such that the number of workers assigned to a given task  $\tau \in T$  is equal to the multiplicity of  $\tau$  in  $T$ .

Two task multisets  $T_1, T_2$  of size  $w$  are said to be *adjacent* if they agree on exactly  $w - 1$  elements; that is,  $|T_1 \setminus T_2| = |T_2 \setminus T_1| = 1$ .<sup>2</sup> The switching cost between two adjacent task multisets  $T_1, T_2$  of size  $w$  is defined as the number of workers whose assignment changes between  $\phi(T_1)$  and  $\phi(T_2)$ . The **switching cost of  $\phi$**  is defined to be the maximum switching cost over *all* pairs of adjacent task multisets. The goal of the worker-task assignment problem is to design a worker-task assignment function with the minimum possible switching cost.

**Prior work in the memoryless setting.** The optimal switching cost is trivially between 1 and  $w$ . Improving either of these bounds substantially has proven difficult, however.

Su, Su, Dornhaus, and Lynch [33] initiated the study of the worker-task assignment problem and observed that assigning the workers to tasks in numerical order achieves a switching cost of  $\min(t - 1, w)$ . They also proved a lower bound of 2 on switching cost, and showed a matching upper bound in the case where  $w \leq 6$  and  $t \leq 4$ .

Subsequent work by Su and Wein [34], in ICALP 2020, pushed further on the lower-bound side of the problem. They proved that a switching cost of 2 is not always possible in general. They show that, if  $t \geq 5$  and  $w \geq 3$ , then any worker-task assignment function must have switching cost at least 3; and if  $t$  is sufficiently large in terms of  $w$  (i.e., it is a tower of height  $w - 1$ ), then the switching cost must be at least 4.

The bounds by [33] and [34] have until now remained state-of-the-art. It remains unknown whether the optimal switching cost is small (it could be as small as 4) or large (it could be as large as  $\min(t - 1, w)$ ). And even achieving a lower bound of 4 on switching cost [34] has required a quite involved argument.

## 1.1 This paper

The contributions of this paper are twofold. First, we present significant progress on the worker-task assignment problem, resulting in both a polylogarithmic upper bound and a super-constant lower bound for the optimal switching cost – our results are interesting in part because of their use of randomized techniques to construct a deterministic assignment function. Second, we explore a natural connection between the worker-task assignment problem and the metric embedding problem of densification into Hamming space, and we transform our progress on the former into new results for the latter.

<sup>2</sup> Let  $m_A(i)$  denote the number of times element  $i$  appears in multiset  $A$ . Then, for any two multisets  $A$  and  $B$ , we define multisets  $A \setminus B$ ,  $A \cup B$ , and  $A \cap B$  to be such that  $m_{A \setminus B}(i) = \max(0, m_A(i) - m_B(i))$ ,  $m_{A \cup B}(i) = \max(m_A(i), m_B(i))$ , and  $m_{A \cap B}(i) = \min(m_A(i), m_B(i))$ , for every element  $i$ .

**Results on worker-task assignment.** Our first result establishes that it is possible to construct a worker-task assignment function with  $O(\log w \log(wt))$  switching cost. This resolves the open question as to whether memoryless worker can allocate themselves to tasks with strong worst-case guarantees.

► **Theorem 1.** *There exists a worker-task assignment function that achieves switching cost  $O(\log w \log(wt))$ .*

Theorem 1 is proven via the probabilistic method and is thus non-constructive. By replacing random hash functions with strong dispersers, however, we show that one can construct an explicit worker-task assignment function with polylogarithmic switching cost.

► **Theorem 2.** *There is an explicit worker-task assignment function that achieves switching cost  $O(\text{polylog}(wt))$ .*

Both Theorems 1 and 2 continue to hold in the more general setting where the size of  $T$  changes over time. That is,  $T$  is permitted to be any multiset of  $[t]$  of size  $w$  or smaller. Two task multisets  $T_1, T_2$  of different sizes are considered adjacent if they satisfy  $||T_1| - |T_2|| = 1$  and  $|(T_1 \cup T_2) \setminus (T_1 \cap T_2)| = 1$ . If  $|T| < w$ , then our worker-task assignment function assign workers  $1, \dots, |T|$  to tasks, and leaves workers  $|T| + 1, \dots, w$  unassigned.

Finally, from the lower bounds side, we prove that no algorithm can achieve sub-linear switching cost as a function of only  $w$ . Theorem 3 says that, if  $t$  is sufficiently larger than  $w$ , then for any worker-task assignment function, there must exist a pair of adjacent task multi-sets that forces *all* of the  $w$  workers to be reassigned. In the statement of the theorem, and throughout the paper,  $\text{tow}(n)$  is defined to be a tower of twos of height  $n$  (i.e., the inverse of the  $\log^*$  function).

► **Theorem 3.** *For every  $w$  and  $t \geq \text{tow}(\Omega(w))$ , every worker-task assignment function has switching cost  $w$ .*

This represents the first super-constant lower bound for the switching cost of a worker-task assignment function. Another way to think about the theorem is that for every  $t$ , there is some  $w$  for which any worker-task assignment function has switching cost at least  $\Omega(\log^*(t))$ . Therefore our bounds leave a gap between  $\log^*$  and  $\text{polylog}$  in terms of dependence on  $t$ .

**An application to metric embeddings: Densification into Hamming space.** The problem of embedding one metric space  $\mathcal{M}_1$  into another metric space  $\mathcal{M}_2$  with small distortion has been widely studied in many contexts and has found many algorithmic applications [29, 7, 9, 8, 13, 2, 6, 10, 1, 5, 19, 24, 3, 22].

Bourgain [5] initiated the study of metric embeddings (into normed spaces) by showing that  $O(\log |M|)$ -distortion embeddings into  $\ell_2$  are possible for any space  $M$ . Much of the subsequent work has focused either on embeddings between exponentially large metric spaces [29, 7, 9, 8, 13, 2, 10, 3], or on embeddings with sub-logarithmic distortion [19, 24, 3, 10].

One natural question is that of **densification**: can one embed sparse vectors from a high-dimensional  $\ell_1$ -space into a low-dimensional  $\ell_1$ -space? That is, if  $V_n^k$  is the set of  $n$ -dimensional vectors with  $k$  non-zero entries, what is the smallest  $m$  for which  $V_n^k$  can be embedded into  $m$ -dimensional  $\ell_1$ -space with low distortion? Charikar and Sahai [10] were the first to consider this problem, and showed how to achieve an output dimension of  $m = O((k/\epsilon)^2 \log n)$  with distortion  $1 + \epsilon$ . They also showed how to apply densification to the related problem of embedding arbitrary tree metrics into low-dimensional  $\ell_1$ -space [10]. Subsequently, Berinde et al. [3] used expander graphs in order to achieve  $m = O(k \log(n/k)/\epsilon^2)$  with distortion

$1 + \epsilon$ . They then used their densification embedding as a tool to perform sparse signal recovery [3, 17, 16, 18]. Both of the known densification algorithms [10, 3] rely on linear sketches, in which each vector  $\vec{x} \in V_n^k$  is mapped to a vector of the form  $\sum_i x_i \vec{b}_i$  for some set of vectors  $\vec{b}_1, \dots, \vec{b}_n$ .

The prior work on densification [10, 3] has focused on embedding into  $\ell_1$ -space. In Section 5, we consider the same problem over Hamming space, where the distance between two vectors  $\vec{x}, \vec{y}$  is given by  $\text{Ham}(\vec{x}, \vec{y}) = |\{i \mid x_i \neq y_i\}|$ . Densification over Hamming space requires new techniques due to the fact that summations of vectors (and thus linear sketches) do not behave well in Hamming space.

Let  $\mathcal{H}_n^k$  denote the set of  $n$ -dimensional binary vectors with  $k$  ones. Let  $\mathcal{H}_k(n)$  denote the set of  $k$ -dimensional vectors with entries from  $[n]$ . We show that  $\mathcal{H}_n^k$  can be embedded into  $\mathcal{H}_k(n)$  with distortion  $O(\log n \log k)$ .

► **Theorem 4.** *There exists a map  $\phi : \mathcal{H}_n^k \rightarrow \mathcal{H}_k(n)$  such that, for every  $\vec{x}, \vec{y} \in \mathcal{H}_n^k$ ,*

$$\text{Ham}(\vec{x}, \vec{y})/2 \leq \text{Ham}(\phi(\vec{x}), \phi(\vec{y})) \leq O(\log n \log k) \text{Ham}(\vec{x}, \vec{y}).$$

The densification embedding is a simple application of the worker-task assignment problem. In order to embed a vector  $\vec{x} \in \mathcal{H}_n^k$  into  $\mathcal{H}_k(n)$ , we simply assign the workers  $\{1, 2, \dots, k\}$  to the task set  $T = \{i \mid \vec{x}_i = 1\}$ , and then we construct the vector  $\vec{y}$  whose  $i$ -th coordinate denotes the task in  $T$  to which worker  $i$  is assigned. This map transforms the switching cost in the worker-task assignment problem into the distortion of the metric embedding.

The densification embedding is optimal in two senses. First, the target space of the embedding must have  $\Omega(k)$  coordinates simply in order to allow for distances of  $\Omega(k)$ . Second, when  $k \ll n$ , any embedding of  $\mathcal{H}_n^k$  to  $k$ -dimensional Hamming space must use  $\Omega(\log n)$  bits per coordinate, simply in order so that the embedding is an injection. It is not clear whether the distortion achieved by our embedding is optimal, however, and it remains open whether smaller distortion can be achieved by allowing for a larger target-space dimension.

We remark that the basic relationship between worker-task assignment and densification embeddings problem has already implicitly been observed in previous work on lower bounds, as a way to formalize what makes the worker-task assignment problem difficult [34]. In contrast, here we are using the relationship as an avenue to obtain improved upper bounds for the densification problem.

## 2 Technical overview

This section gives an overview of the main technical ideas in the paper. For simplicity, the section will treat the task multiset  $T \subseteq [t]$  as always being a set (rather than a multiset). As discussed in Section 3, one can formally reduce from the multiset case to the set case, at the cost of  $t$  being replaced with  $t' = wt$ .

### 2.1 A warmup: The random-permutation algorithm

We begin by describing a simple assignment function that we call the **random-permutation algorithm**. The random-permutation algorithm does not necessarily achieve small switching cost, but it does have the property that for any two adjacent task sets  $T_1, T_2 \subseteq [t]$ , the switching cost between  $T_1$  and  $T_2$  is  $O(\log w)$  with high probability in  $w$ .

## 19:6 Memoryless Worker-Task Assignment with Polylogarithmic Switching Cost

**The algorithm.** The random-permutation algorithm assigns to each worker  $i \in [w]$  a random permutation

$$\sigma_i = \langle \sigma_i(1), \sigma_i(2), \dots, \sigma_i(t) \rangle$$

of the numbers  $[t]$ . We think of worker  $i$  as **preferring** task  $\sigma_i(j)$  over task  $\sigma_i(j+1)$  for all  $j \in [t-1]$ .

Suppose we wish to assign workers to tasks  $T$ . The random-permutation algorithm assigns the workers  $1, 2, \dots, w$  to tasks  $\tau_1, \tau_2, \dots, \tau_w \in T$  one by one in order of worker ID, assigning worker  $i$  to the task that it most prefers out of the tasks in  $T$  that have not yet been assigned a worker.

For each  $i \in [w]$ , we define the  **$i$ -remainder tasks** to be the tasks  $T \setminus \{\tau_1, \dots, \tau_i\}$ . That is, the  $i$ -remainder tasks are the tasks that remain after the first  $i$  workers are assigned. This means that worker  $i+1$  is assigned to the  $i$ -remainder task that it most prefers.

**Analyzing expected switching cost.** Let  $T_1, T_2 \subseteq [t]$  be adjacent task sets of size  $w$ . We begin by showing that the expected switching cost from  $T_1$  to  $T_2$  is  $O(\log w)$ .

Let  $r$  and  $s$  be such that  $T_1 = (T_2 \cup \{r\}) \setminus \{s\}$ . Let  $\psi_1$  and  $\psi_2$  denote the assignments produced by the random permutation algorithm for  $T_1$  and  $T_2$ , respectively. Let  $A_i$  and  $B_i$  denote the set of  $i$ -remainder tasks during the constructions of  $\psi_1$  and  $\psi_2$ , respectively.

The key to analyzing the random-permutation algorithm is to compare the  $i$ -remainder sets  $A_i$  and  $B_i$  for each  $i \in [w]$ . We claim that  $|A_i \setminus B_i| \leq 1$  for all  $i \in [w]$ . We prove this by induction on  $i$ : suppose that  $A_{i-1} = (B_{i-1} \cup \{a\}) \setminus \{b\}$ , and suppose for contradiction that  $|A_i \setminus B_i| \geq 2$ . If either  $\psi_1$  assigns worker  $i$  to task  $a$ , or  $\psi_2$  assigns worker  $i$  to task  $b$ , then we would be guaranteed that  $|A_i \setminus B_i| \leq 1$ , a contradiction. Thus  $\psi_1$  and  $\psi_2$  must each assign worker  $i$  to a task in  $A_{i-1} \cap B_{i-1}$ . But this means that, in both assignments, worker  $i$  is assigned to the task in  $A_{i-1} \cap B_{i-1}$  that worker  $i$  most prefers. Thus  $\psi_1$  and  $\psi_2$  assign worker  $i$  to the same task, again contradicting that  $|A_i \setminus B_i| \geq 2$ .

We now analyze the probability of  $\psi_1$  and  $\psi_2$  differing in their assignment of worker  $i$ . Since  $A_i$  contains at most one element  $a$  not in  $B_i$ , the probability that worker  $i$  prefers  $a$  over all elements in  $B_i$  is at most  $1/|B_i| = 1/(w-i+1)$ . Similarly, since  $B_i$  contains at most one element  $b$  not in  $A_i$ , the probability that worker  $i$  prefers  $b$  over all elements in  $A_i$  is at most  $1/|A_i| = 1/(w-i+1)$ . By the union bound, it follows that the probability of  $\psi_1$  and  $\psi_2$  assigning worker  $i$  to different tasks is at most  $2/(w-i+1)$ .

By linearity of expectation, the expected switching cost between  $T_1$  and  $T_2$  is at most

$$\sum_{i=1}^w \frac{2}{w-i+1} = O(\log w).$$

**Why a union bound fails for worst-case switching cost.** By using Chernoff-style bounds, one can modify the above analysis of the random-permutation algorithm to show that, with high probability in  $w$  (i.e., probability  $1 - 1/\text{poly } w$ ), the switching cost between  $T_1$  and  $T_2$  is  $O(\log w)$ .

On the other hand, achieving a switching cost of  $O(\log w)$  for *all* pairs  $(T_1, T_2)$  of adjacent task sets presents a challenge because there are  $\binom{w+1}{2} \binom{t}{w+1}$  such pairs that must be considered. When  $w = t/2$ , the number of distinct pairs  $(T_1, T_2)$  of adjacent task sets exceeds  $2^t \geq 2^w$ .

Thus, the probability bounds achieved by the random-permutation algorithm are nowhere near high enough to enable a union bound over all adjacent worker-set pairs. We call this the **union-bound magnitude issue**.

## 2.2 An algorithm with small switching cost

We now describe a randomized assignment algorithm  $\mathcal{A}$  that, with high probability in  $t$ , achieves switching cost  $O(\log w \log t)$  on *all* adjacent task-sets  $T_1, T_2 \subseteq [t]$  of size  $w$ . That is, with high probability,  $\mathcal{A}$  produces an assignment function satisfying the requirements of Theorem 1. The algorithm  $\mathcal{A}$  is called the **multi-round balls-to-bins algorithm**.

The multi-round balls-to-bins algorithm essentially flips the approach taken by the random-permutation algorithm. One can think of the random-permutation algorithm as consisting of  $w$  phases in which each phase *deterministically* assigns exactly one worker to a task, and then the phases *probablistically* incur small switching cost. In contrast, the multi-round balls-to-bins algorithm consists of  $O(\log w)$  phases, where each phase *probabilistically* assigns some number of workers to tasks, and each phase *deterministically* incurs small switching cost. Whereas the failure mode of the random-permutation algorithm is that a high-switching cost may occur, the failure mode of the multi-round balls-to-bins algorithm is that some workers may be left unassigned. As we shall see later, this distinction plays an important role in solving the union-bound magnitude issue.

**Structure of the multi-round balls-to-bins algorithm.** We begin with a succinct description of the algorithm  $\mathcal{A}$ . For each  $i$  from 1 to  $\log_{1.1} w$ , repeat the following hashing procedure  $c \log t$  many times. Initialize a hash table consisting of  $w/(1.1)^i$  bins and randomly hash each unassigned worker and each unassigned task into this table. For each bin that contains at least one worker and one task, assign the minimum worker in that bin to the minimum task in that bin.

In more detail, the algorithm  $\mathcal{A}$  is the composition of  $O(\log w)$  sub-algorithms  $\mathcal{A}_1, \mathcal{A}_2, \dots$ . Each of  $\mathcal{A}_1, \mathcal{A}_2, \dots$  are **partial-assignment** algorithms, meaning that  $\mathcal{A}_i$  assigns some subset of the workers to some subset of the tasks in  $T$ , possibly leaving workers and tasks unassigned. Note that the input to algorithm  $\mathcal{A}_i$  is the set of workers/tasks that remain unassigned by  $\mathcal{A}_1, \dots, \mathcal{A}_{i-1}$ . Thus one can think of the input to  $\mathcal{A}_i$  as being a pair  $(W, T)$  where  $W \subseteq [w]$  is a set of workers,  $T \subseteq [t]$  is a set of tasks, and  $|W| = |T|$ .

The algorithm  $\mathcal{A}_1$ 's responsibility is to assign enough workers to tasks so that at most  $w/1.1$  workers remain unassigned. Algorithm  $\mathcal{A}_2$  is then executed on the remaining (i.e., not-yet-assigned) workers and tasks, and is responsible for assigning enough workers to tasks so that at most  $w/(1.1)^2$  workers remain unassigned. Continuing like this, algorithm  $\mathcal{A}_i$  is executed on the workers/tasks that remain unassigned by all of  $\mathcal{A}_1, \dots, \mathcal{A}_{i-1}$ , and is responsible for assigning enough workers to tasks that at most  $r_i = w/(1.1)^i$  workers in  $W$  remain unassigned.

Each of the  $\mathcal{A}_i$ 's are randomized algorithms, meaning that they have some probability of failure. The failure mode for  $\mathcal{A}_i$  is *not* high-switching cost, however. In fact, as we shall see later, each  $\mathcal{A}_i$  deterministically contributes at most  $O(\log w)$  to the switching cost. Instead, the way in which  $\mathcal{A}_i$  can fail is that it may leave more than  $r_i$  workers unassigned. This means that the failure mode for the full algorithm  $\mathcal{A}$  is that it may fail to assign all of the workers in  $W$  to tasks in  $T$ .

**Applying the probabilistic method to  $\mathcal{A}_1, \mathcal{A}_2, \dots$**  Before describing the partial-assignment algorithms  $\mathcal{A}_i$  in detail, we first describe how our analysis of algorithm  $\mathcal{A}$  overcomes the union-bound magnitude issue.

Recall that each algorithm  $\mathcal{A}_i$  is responsible for reducing the number of remaining workers to  $r_i = w/(1.1)^i$ . We will later see that each  $\mathcal{A}_i$  has a failure probability  $p_i$  that is a function of  $r_i$  and  $t$ , namely,

$$p_i = \frac{1}{t^{\Omega(r_i)}}.$$



As  $i$  grows, the failure probability  $p_i$  of  $\mathcal{A}_i$  becomes larger, making it impossible to union-bound over exponentially many pairs of task sets  $T_1, T_2$ .

An important insight is that, if all of  $\mathcal{A}_1, \dots, \mathcal{A}_{i-1}$  succeed (i.e., they each assign the number of workers that they are responsible for assigning) then the number of workers and tasks that  $\mathcal{A}_{i-1}$  is executed on is only  $O(r_i)$ . That is, if we think of the inputs to  $\mathcal{A}_i$  as being pairs  $(W, T)$  where  $W \subseteq [w]$  is a set of workers and  $T \subseteq [t]$  is a set of tasks, the set of inputs  $(W, T)$  that algorithm  $\mathcal{A}_{i-1}$  must succeed on is only the inputs for which  $|W| = |T| \leq O(r_i)$ . The number of such inputs is at most  $t^{O(r_i)}$ . In other words, even though the failure probability  $p_i$  of algorithm  $\mathcal{A}_i$  increases with  $i$ , the number of inputs over which we must apply a union bound decreases. By a union bound, we can deduce that  $\mathcal{A}_i$  has a high probability in  $t$  of succeeding on all relevant inputs  $(W, T)$ . Combining this analysis over all of the partial-assigning algorithms  $\mathcal{A}_1, \mathcal{A}_2, \dots$ , we get that the full assignment algorithm  $\mathcal{A}$  also succeeds with high probability in  $t$ . In particular, we have proven that there exists a deterministic assignment function with the desired switching cost, and that such a function can be obtained with high probability by the randomized algorithm  $\mathcal{A}$ .

**Designing  $\mathcal{A}_i$ .** Each algorithm  $\mathcal{A}_i$  is a composition of  $\Theta(\log t)$  algorithms  $\mathcal{A}_{i,1}, \mathcal{A}_{i,2}, \mathcal{A}_{i,3}, \dots$ , each of which individually is a partial assignment algorithm.

Each algorithm  $\mathcal{A}_{i,j}$  takes a simple balls-in-bins approach to assigning some subset of the remaining workers to some subset of the remaining tasks.

In particular,  $\mathcal{A}_{i,j}$  places the workers into bins  $1, 2, \dots, r_i$  by hashing each worker to a bin (using a random function from  $[w]$  to  $[r_i]$ ). Similarly, the tasks are placed into bins  $1, 2, \dots, r_i$  by hashing each task to a bin. If a bin  $b$  contains both at least one worker and at least one task, then the smallest-numbered worker in bin  $b$  is assigned to the smallest-number task in bin  $b$ .

Note that each of the algorithms  $\mathcal{A}_{i,1}, \mathcal{A}_{i,2}, \mathcal{A}_{i,3}, \dots$  are identical copies of one-another, except using different random bits. Also note all of the  $\mathcal{A}_i$ 's are defined in the same way as each other, except the number of bins hashed to decreases as  $i$  increases. As we shall see shortly, the reason for having  $\mathcal{A}_i$  consist of  $\Theta(\log t)$  sub-algorithms is to enable probability amplification later in the analysis.

**Bounding the switching cost.** The partial assignment algorithms  $\mathcal{A}_{i,j}$  are designed to satisfy two essential properties, which we prove formally in the full proof. These two properties can then be combined to bound the switching cost of the full algorithm  $\mathcal{A}$ .

**Compatibility:** Let  $I_1 = (W_1, T_1)$  and  $I_2 = (W_2, T_2)$  be inputs to  $\mathcal{A}_{i,j}$ . Suppose  $I_1$  and  $I_2$  are **unit distance**, meaning that

$$|W_1 \setminus W_2| + |W_2 \setminus W_1| + |T_1 \setminus T_2| + |T_2 \setminus T_1| \leq 2.$$

Let  $I'_1 = (W'_1, T'_1)$  and  $I'_2 = (W'_2, T'_2)$  be the workers and tasks that remain unassigned when  $\mathcal{A}_{i,j}$  is executed on each of  $I_1$  and  $I_2$ , respectively. Then  $I'_1$  and  $I'_2$  are guaranteed to also be unit-distance.

**Low Switching Cost:** The switching cost of  $\mathcal{A}_{i,j}$  is  $O(1)$ . That is, if  $I_1 = (W_1, T_2)$  and  $I_2 = (W_2, T_2)$  are inputs to  $\mathcal{A}_{i,j}$ , and  $I_1$  and  $I_2$  are unit-distance, then the worker-task assignments made by  $\mathcal{A}_{i,j}$  on each of  $I_1$  and  $I_2$  differ by at most  $O(1)$  assignments.

Consider two adjacent task sets  $T_1$  and  $T_2$ . When we execute  $\mathcal{A}$  on  $T_1$  and  $T_2$ , respectively, we use  $I_1^{i,j}$  and  $I_2^{i,j}$ , respectively, to denote the worker/task input that are given to partial-assignment algorithm  $\mathcal{A}_{i,j}$ .



The Compatibility property of the  $\mathcal{A}_{i,j}$ 's guarantees by induction that, for each  $\mathcal{A}_{i,j}$  the worker/task inputs  $I_1^{i,j}$  and  $I_2^{i,j}$  are unit-distance (or zero-distance). The Low-Switching-Cost property then guarantees that each  $\mathcal{A}_{i,j}$  contributes at most  $O(1)$  to the switching cost of  $\mathcal{A}$ . Since there are  $O(\log t \log w)$   $\mathcal{A}_{i,j}$ 's, this bounds the total switching cost of  $\mathcal{A}$  by  $O(\log t \log w)$ .

**Deriving the success probabilities.** Next we analyze the probability of  $\mathcal{A}_i$  failing on a given worker/task input  $(W, T)$ . Recall that the only way in which  $\mathcal{A}_i$  might fail is if more than  $r_i$  workers remain unassigned after  $\mathcal{A}_i$  finishes. Additionally, since we need only consider cases where  $\mathcal{A}_{i-1}$  succeeds, we can assume that  $r_i \leq |W|, |T| \leq 1.1r_i$ .

Let  $q$  denote the number of workers that  $\mathcal{A}_{i,1}$  assigns to tasks. Given that  $r_i \leq |W|, |T| \leq 1.1r_i$ , a simple analysis of  $\mathcal{A}_{i,1}$  shows that  $\mathbb{E}[q] \geq r_i/5$ . On the other hand, using McDiarmid's inequality, one can perform a balls-in-bins style analysis in order to show that  $\Pr[\mathbb{E}[q] - q > r_i/10] \leq 2^{-\Omega(r_i)}$ . This means that  $\mathcal{A}_{i,1}$  has probability at most  $2^{-\Omega(r_i)}$  of leaving more than  $r_i$  workers unassigned.

In order for  $\mathcal{A}_i$  to fail (i.e.,  $\mathcal{A}$  leaves more than  $r_i$  workers unassigned), *all* of sub-algorithms  $\mathcal{A}_{i,1}, \mathcal{A}_{i,2}, \dots$  would have to individually fail. Since there are  $\Theta(\log t)$  sub-algorithms, the probability of them all failing is

$$p_i = 2^{-\Omega(r_i \log t)} = t^{-\Omega(r_i)}.$$

This allows us to apply the probabilistic method to the  $\mathcal{A}_i$ 's in order to bound the probability of any  $\mathcal{A}_i$  failing on any input, as desired.

**An explicit construction with polylogarithmic switching cost.** The multi-round balls-to-bins algorithm gives a non-explicit approach to constructing an assignment function with low switching cost. The approach is non-explicit because it relies on the probabilistic method.

In the full version of the paper<sup>3</sup> we show how to obtain an explicit algorithm with switching cost  $\text{polylog } wt$ . The basic idea is to replace random hash functions, used to place workers and tasks into bins, with functions obtained from pseudorandom objects called strong dispersers. Instead of trying a number of random hash functions within the  $\mathcal{A}_{i,j}$ 's, we instead iterate over all of the hash functions from a small family given by a strong disperser [28].

### 2.3 A lower bound on switching cost

Define  $s_{w,t}$  to be the optimal switching cost for assignment functions that assign workers  $1, 2, \dots, w$  to multisets of  $w$  tasks from the universe  $[t]$ . The upper bounds in this paper establish that  $s_{w,t} \leq O(\log w \log(wt))$ . It is natural to wonder whether smaller bounds can be achieved, and in particular, whether a small switching cost that depends only on  $w$  can be achieved.

It trivially holds that  $s_{w,t} \leq w$ . We show that when  $t$  is sufficiently large relative to  $w$ , there is a matching lower bound of  $s_{w,t} \geq w$ . In fact, our lower bound only uses the evaluation of the assignment function on sets (as opposed to multisets).

Consider an assignment function  $\phi$  that, given a multiset  $T$  of tasks with elements from  $[t]$  of  $w$  tasks, produces an assignment of workers  $[w]$  to tasks  $T$ . Our goal will be to find tasks  $\tau_1 < \tau_2 < \dots < \tau_{w+1}$  such that if  $\phi(\{\tau_1, \dots, \tau_w\})$  assigns worker  $i$  to task  $\tau_{\pi(i)}$  for some permutation  $\pi$  of  $[w]$ , then  $\phi(\{\tau_2, \dots, \tau_{w+1}\})$  assigns worker  $i$  to task  $\tau_{\pi(i)+1}$ . The existence of such a configuration immediately implies that  $\phi$  has switching cost  $w$ .

<sup>3</sup> <https://arxiv.org/abs/2008.10709>

## 19:10 Memoryless Worker-Task Assignment with Polylogarithmic Switching Cost

We use an application of the hypergraph Ramsey theorem to show that, when  $t$  is large enough, a configuration of the type described in the above paragraph must exist. Let  $K_t^{(w)}$  denote the complete  $w$ -uniform hypergraph on  $t$  vertices. This is just the set of  $w$ -element subsets of  $[t]$ , which correspond to sets of tasks. For each hyperedge  $T = \{\tau_1, \dots, \tau_w\}$ , where  $1 \leq \tau_1 < \dots < \tau_w \leq t$ , we color the hyperedge  $T$  by a color  $\pi$  where  $\tau_{\pi(i)}$  is the task assigned to worker  $i$ .

This gives a coloring of the hyperedges of  $K_t^{(w)}$  by  $w!$  colors, each color being a permutation of  $[w]$ . By the hypergraph Ramsey theorem, if  $t$  is large enough in terms of  $w$ , there must exist  $w+1$  vertices  $\tau_1, \dots, \tau_{w+1}$  so all the hyperedges formed by the vertices have the same color  $\pi$ . By examining the hyperedges  $\{\tau_1, \dots, \tau_w\}$  and  $\{\tau_2, \dots, \tau_{w+1}\}$ , it follows that  $\phi(\{\tau_1, \dots, \tau_w\})$  assigns each worker  $i$  to task  $\tau_{\pi(i)}$  and that  $\phi(\{\tau_2, \dots, \tau_{w+1}\})$  assigns each worker  $i$  to task  $\tau_{\pi(i)+1}$ , as desired.

### 3 Achieving switching cost $O(\log w \log(wt))$

In this section, we prove the following theorem.

► **Theorem 1.** *There exists a worker-task assignment function that achieves switching cost  $O(\log w \log(wt))$ .*

We demonstrate the existence of such a function via the probabilistic method, showing that there is a randomized construction that produces a low-switching cost worker-task assignment function with nonzero probability. In the full version of the paper, we also show how to derandomize the construction at the cost of a few extra log factors.

**From multisets to sets.** We begin by showing that, without loss of generality, we can restrict our attention to task multisets  $T$  that are sets (rather than multisets). We reduce from the multiset version of the problem with  $w$  workers and  $t$  tasks to the set version of the problem with  $w$  workers and  $wt$  tasks.

► **Lemma 5.** *Define  $n = wt$ . Let  $\phi$  be a worker-task assignment function that assigns workers  $[w]$  to task sets  $T \subseteq [n]$  (note that  $\phi$  is defined only on task sets  $T$ , and not on multisets). Let  $s$  be the switching cost of  $\phi$  (considering only pairs of adjacent subsets of  $[n]$ , rather than adjacent sub-multisets). Then there exists a worker-task assignment function  $\phi'$  assigning workers  $[w]$  to task multisets  $T \subseteq [t]$ , such that  $\phi'$  also has switching cost  $s$ .*

**Proof.** When discussing the assignment function  $\phi$ , we think of its input task-set  $T$  as consisting of elements from  $[t] \times [w]$  rather than elements of  $[tw]$ .

With this in mind, we construct  $\phi'$  as follows. Given a task multiset  $T \subseteq [t]$ , define the set  $\mathbf{S}(T) \subseteq [t] \times [w]$  to be  $\bigcup_{i=1}^t \{(i, 1), \dots, (i, m_T(i))\}$ , where  $m_T(i)$  is the multiplicity of  $i$  in  $T$ . The worker-task assignment  $\phi$  produces some bijection  $\psi_{\mathbf{S}(T)} : [w] \rightarrow \mathbf{S}(T)$ . Similarly,  $\phi'$  should produce some bijection  $\psi'_T : [w] \rightarrow T$ . This bijection is defined naturally by projection: if  $\psi_{\mathbf{S}(T)}$  assigns worker  $j$  to task  $(i, x)$ , let  $\psi'_T$  assign worker  $j$  to task  $i$ .

We now compute the switching cost of  $\phi'$ . Let  $T$  and  $T'$  be two adjacent task multisets, so  $T' = T \cup \{a\} \setminus \{b\}$  for some  $a, b \in [t]$ . Then  $\mathbf{S}(T') = \mathbf{S}(T) \cup \{(a, m_T(a) + 1)\} \setminus \{(b, m_T(b))\}$ , and so  $\mathbf{S}(T')$  is adjacent to  $\mathbf{S}(T)$ . Since  $\phi$  has switching cost  $s$ ,  $\psi_{\mathbf{S}(T)}$  and  $\psi_{\mathbf{S}(T')}$  agree on  $w - s$  workers. By construction,  $\psi'_T$  and  $\psi'_{T'}$  must agree on these  $w - s$  workers as well, and so it too has switching cost at most  $s$ . ◀

In the remainder of the section, we will make the assumption that  $T$  is a subset of  $[n]$ , and we will show how to design an assignment function with switching cost  $O(\log w \log n)$  on all pairs of adjacent subsets of  $[n]$ . By Lemma 5, setting  $n = wt$  then implies Theorem 1.

**Designing an assignment function as an algorithm.** It will be helpful to think of the function we construct for assigning workers to tasks as an algorithm  $\mathcal{A}$ , which we call the **multi-round balls-to-bins algorithm**. The algorithm  $\mathcal{A}$  takes as input a set  $T \subseteq [n]$  of tasks with  $|T| = w$  and must produce a bijection from the workers  $[w]$  to  $T$ .

The algorithm constructs this bijection in stages. Each stage is what we call a **partial assignment algorithm**, which takes as input the current sets of workers and tasks that have yet to be matched and assigns some subset of these workers to some subset of the tasks. Formally, we define a partial assignment algorithm to be any function  $\psi$  which accepts as input any pair of sets  $T \subseteq [n], W \subseteq [w]$  with  $|T| = |W|$  and produces a matching between some subset of  $T$  and some subset of  $W$ . After applying  $\psi$  to  $(T, W)$ , there may remain some unmatched elements  $T' \subseteq T, W' \subseteq W$ . We call  $(T, W)$  the **worker-task input** to  $\psi$  and  $(T', W')$  the **worker-task output**. Since a matching must remove exactly as many elements from  $T$  as it does from  $W$ , we must also have  $|W'| = |T'|$ . Consequently, there is a natural notion of the **composition** of two partial assignment algorithms: the composition  $\psi' \circ \psi$  applies  $\psi$  and then  $\psi'$ , letting the worker-task output of  $\psi$  be the worker-task input to  $\psi'$ .

**The algorithm.** We recall the description of the algorithm  $\mathcal{A}$ . For each  $i$  from 1 to  $c \log w$ , repeat the following hashing procedure  $c \log n$  many times. Initialize a hash table consisting of  $w/(1.1)^i$  bins and randomly hash each unassigned worker and each unassigned task into this table. For each bin that contains at least one worker and one task, assign the minimum worker in that bin to the minimum task in that bin.

In more detail, our algorithm  $\mathcal{A}$  is the composition of  $\log_{1.1} w$  partial-assignment algorithms,

$$\mathcal{A} = \mathcal{A}_1 \circ \mathcal{A}_2 \circ \cdots \circ \mathcal{A}_{\log_{1.1} w}.$$

Let  $c$  be a large positive constant. Each  $\mathcal{A}_i$  is itself the composition of  $c \log n$  partial-assignment algorithms,

$$\mathcal{A}_i = \mathcal{A}_{i,1} \circ \mathcal{A}_{i,2} \circ \cdots \circ \mathcal{A}_{i,c \log n}.$$

**Designing the parts.** Each  $\mathcal{A}_{i,j}$  assigns workers to tasks using what we call a  $w/(1.1)^i$ -**bin hash**, which we define as follows.

For a given parameter  $k$ , a  **$k$ -bin hash** selects functions  $h_1 : [w] \rightarrow [k]$  and  $h_2 : [n] \rightarrow [k]$  independently and uniformly at random. For each worker  $\omega \in [w]$ , we say that  $\omega$  is **assigned** to bin  $h_1(\omega)$ . Similarly, for each  $\tau \in [n]$  we say  $\tau$  is assigned to  $h_2(\tau)$ . These functions are then used to construct a partial assignment. Given a worker-task input  $(W, T)$ , we restrict our attention to only the assignments of workers in  $W$  and tasks in  $T$ . In each bin  $\kappa \in [k]$  with at least one worker and one task assigned, match the smallest such worker to the smallest such task. Importantly, once  $h_1$  and  $h_2$  are fixed, the algorithm  $\mathcal{A}_{i,j}$  uses this same pair of hash functions for every worker-task input, which (as we will see later) is what allows it to make very similar assignments for similar inputs and achieve low switching cost.

We set each  $\mathcal{A}_{i,j}$  to be an independent random instance of the  $k$ -bin hash, where  $k = w/(1.1)^i$ . Formally, this means that the algorithm  $\mathcal{A} = \mathcal{A}_{1,1} \circ \cdots \circ \mathcal{A}_{\log_{1.1} w, c \log n}$  is a random variable whose value is a partial-assignment function. Our task is thus to prove that, with non-zero probability,  $\mathcal{A}$  fully assigns all workers to tasks and has small switching cost.

**Analyzing the algorithm.** In Section 3.2, we show that  $\mathcal{A}$  *deterministically* has switching cost  $O(\log w \log n)$ .

Although  $\mathcal{A}$  always has small switching cost, the algorithm is *not* always a legal worker-task assignment function. This is because the algorithm may sometimes act as a *partial* worker-task assignment function, leaving some workers and tasks unassigned.

In Section 3.1, we show that with probability greater than 0 (and, in fact, with probability  $1 - 1/\text{poly } n$ ), the algorithm  $\mathcal{A}$  succeeds at fully assigning workers to tasks for *all* worker-task inputs  $(W, T)$ . Theorem 1 follows by the probabilistic method.

### 3.1 Bounding the probability of failure

Call a partial-assignment algorithm  $\psi$  **fully-assigning** if for every worker/task input  $(W, T)$ ,  $\psi$  assigns all of the workers in  $W$  to tasks in  $T$ . That is,  $\psi$  never leaves workers unassigned.

► **Proposition 6.** *The multi-round balls-to-bins algorithm  $\mathcal{A}$  is fully-assigning with high probability in  $n$ . That is, for any polynomial  $p(n)$ , if the constant  $c$  used to define  $\mathcal{A}$  is sufficiently large, then  $\mathcal{A}$  is fully-assigning with probability at least  $1 - O(1/p(n))$ .*

Proposition 6 tells us that with high probability in  $n$ ,  $\mathcal{A}$  succeeds at assigning all workers on *all* inputs. We remark that this is a much stronger statement than saying that  $\mathcal{A}$  succeeds with high probability in  $n$  on a *given* input  $(W, T)$ .

The key to proving Proposition 6 is to show that each  $\mathcal{A}_i$  performs what we call  $(w/(1.1)^i)$ -**halving**. A partial-assignment function  $\psi$  is said to perform  $k$ -**halving** if for every worker/task input  $(W, T)$  of size at most  $1.1k$ , the worker-task output  $(W', T')$  for  $\psi(W, T)$  has size at most  $k$ .

If every  $\mathcal{A}_i$  performs  $w/(1.1)^i$ -halving, then it follows that

$$\mathcal{A}_1 \circ \cdots \circ \mathcal{A}_{\log_{1.1} w}$$

is a fully-assigning algorithm. Thus our task is to show that each  $\mathcal{A}_i$  performs  $w/(1.1)^i$ -halving with high probability in  $n$ .

We begin by analyzing the  $k$ -bin hash on a given worker/task input  $(W, T)$ .

► **Lemma 7.** *Let  $\psi$  a randomly selected  $k$ -bin hash. Let  $(W, T)$  be a worker/task input satisfying  $|W| = |T| \leq 1.1k$ , and let  $(W', T')$  be the worker/task output of  $\psi(W, T)$ . The probability that  $(W', T')$  has size  $k$  or larger is  $2^{-\Omega(k)}$ .*

**Proof.** We may assume that  $|W| = |T| \geq k$ , else the conclusion is trivially true. Let  $X$  be the random variable denoting the number of worker/task assignments made by  $\psi(W, T)$ . Equivalently,  $X$  counts the number of bins to which at least one worker is assigned and at least one task is assigned – call these the **active bins**. We will show that  $\Pr[X < \frac{k}{8}] \leq 2^{-\Omega(k)}$ . Since  $|W| = |T| \leq 1.1k$ , this immediately implies that  $|W'| = |T'| \leq 1.1k - 0.125k \leq k$  with probability  $1 - 2^{-\Omega(k)}$ , as desired.

We begin by computing  $\mathbb{E}[X]$ . For each bin  $j \in [k]$ , the probability no workers are assigned to bin  $j$  is  $(1 - 1/k)^{|W|} \leq (1 - 1/k)^k \leq 1/e$ . Similarly, the probability that no tasks are assigned to bin  $j$  is at most  $(1 - 1/k)^{|T|} \leq 1/e$ . The probability of bin  $j$  being active is therefore at least  $1 - 2/e \geq 1/4$ . By linearity of expectation,  $\mathbb{E}[X] \geq k/4$ .

Next we show that the random variable  $X$  is tightly concentrated around its mean. Because the bins that are active are not independent of one-another, we cannot apply a Chernoff bound. Instead, we employ McDiarmid's inequality:

► **Theorem 8** (McDiarmid '89 [25]). *Let  $A_1, \dots, A_m$  be independent random variables over an arbitrary probability space. Let  $F$  be a function mapping  $(A_1, \dots, A_m)$  to  $\mathbb{R}$ , and suppose  $F$  satisfies,*

$$\sup_{a_1, a_2, \dots, a_m, \bar{a}_i} |F(a_1, a_2, \dots, a_{i-1}, a_i, a_{i+1}, \dots, a_m) - F(a_1, a_2, \dots, a_{i-1}, \bar{a}_i, a_{i+1}, \dots, a_m)| \leq R,$$

for all  $1 \leq i \leq m$ . That is, if  $A_1, A_2, \dots, A_{i-1}, A_{i+1}, \dots, A_m$  are fixed, then the value of  $A_i$  can affect the value of  $F(A_1, \dots, A_m)$  by at most  $R$ . Then for all  $S > 0$ ,

$$\Pr[|F(A_1, \dots, A_m) - \mathbb{E}[F(A_1, \dots, A_m)]| \geq R \cdot S] \leq 2e^{-2S^2/m}.$$

The number of active bins  $X$  is a function of at most  $2.2 \cdot k$  independent random variables (namely, the hashes  $h_1(\omega)$  for each  $\omega \in W$  and the hashes  $h_2(\tau)$  for each  $\tau \in T$ ). Each of these random variables can individually change the number of active bins by at most one. It follows that we can apply McDiarmid's inequality with  $R = 1$  and  $m = 2.2k$ . Taking  $S = k/8$ , we obtain

$$\Pr[|X - \mathbb{E}[X]| \geq k/8] \leq e^{-\Omega(k)}.$$

Since  $\mathbb{E}[X] \geq k/4$ , we have that  $\Pr[X < k/8] \leq e^{-\Omega(k)}$ , which completes the proof of the lemma. ◀

Our next lemma shows that each  $\mathcal{A}_i$  is  $k$ -halving with high probability in  $n$ , where  $k = w/(1.1)^i$ .

► **Lemma 9.** *Let  $\psi_1, \dots, \psi_{c \log n}$  be independent random  $k$ -bin hashes, and let  $\psi = \psi_1 \circ \dots \circ \psi_{c \log n}$ . With high probability in  $n$ ,  $\psi$  is  $k$ -halving. That is, every worker-task input  $(W, T)$  with  $|W| = |T| \leq 1.1k$  has a worker task output  $(W', T')$  with  $|W'| = |T'| \leq k$ .*

**Proof.** Fix an arbitrary worker-task input  $(W, T)$  with  $|W| = |T| \leq 1.1k$ . Let  $(W_i, T_i)$  denote the worker-task output after applying the first  $i$  rounds,  $\psi_1 \circ \dots \circ \psi_i$ . Let  $p_i$  denote the probability that  $|W_i| = |T_i| > k$ .

First, we observe that  $p_i \leq e^{-\Omega(k)} p_{i-1}$  for all  $i > 1$ . Indeed, for  $|W_i| = |T_i| > k$ , we must necessarily have  $|W_{i-1}| = |T_{i-1}| > k$ , which occurs with probability  $p_{i-1}$ , but in this situation, the probability that  $\psi_i$  produces a worker-task output of size greater than  $k$  is a further  $e^{-\Omega(k)}$  by Lemma 7.

The probability that  $\psi$  fails to reduce the size of  $(W, T)$  to  $k$  or smaller is thus at most

$$p_{c \log n} \leq e^{-\Omega(ck \log n)} \leq n^{-\Omega(ck)}, \quad (1)$$

where  $c$  is treated as a parameter.

On the other hand, the number of possibilities for input pairs  $(W, T)$  satisfying  $|W| = |T| \leq 1.1k$  is

$$\sum_{j=0}^{1.1k} \binom{w}{j} \binom{n}{j} \leq 1.1k \cdot w^{1.1k} n^{1.1k} \leq n^{O(k)}. \quad (2)$$

Combining (1) and (2), the probability that there exists *any* pair  $(W, T)$  of size  $1.1k$  or smaller which fails to have its size reduced to  $k$  or smaller is at most  $n^{O(k) - c\Omega(k)}$ . If  $c$  is selected to be a sufficiently large constant, then it follows that  $\psi$  performs  $k$ -halving with probability at least  $1 - n^{-\Omega(k)}$ . ◀

We now prove Proposition 6.

**Proof of Proposition 6.** By Lemma 9, each algorithm  $\mathcal{A}_i$  is  $(w/(1.1)^i)$ -halving with high probability in  $n$ . By a union bound, it follows that all of  $\mathcal{A}_i \in \{\mathcal{A}_1, \dots, \mathcal{A}_{\log_{1.1} w}\}$  are  $(w/(1.1)^i)$ -halving with high probability in  $n$ . If this occurs, then

$$\mathcal{A} = \mathcal{A}_1 \circ \dots \circ \mathcal{A}_{\log_{1.1} w}$$

is fully-assigning, as desired.  $\blacktriangleleft$

### 3.2 Bounding the switching cost

Recall that two worker/task inputs  $(W_1, T_1)$  and  $(W_2, T_2)$  are said to be **unit distance** if

$$|W_1 \setminus W_2| + |W_2 \setminus W_1| + |T_1 \setminus T_2| + |T_2 \setminus T_1| \leq 2.$$

A partial-assignment algorithm  $\psi$  is  **$s$ -switching-cost bounded** if for all unit-distance pairs of worker/task inputs  $(W_1, T_1)$  and  $(W_2, T_2)$ , the set of assignments made by  $\psi(W_1, T_1)$  deterministically differs from the set of assignments made by  $\psi(W_2, T_2)$  by at most  $s$ .

In this section, we prove the following proposition.

► **Proposition 10.** *The multi-round balls-to-bins algorithm is  $O(\log w \log n)$ -switching-cost bounded.*

We begin by showing that each of the algorithms  $\mathcal{A}_{i,j}$  are  $O(1)$ -switching-cost bounded.

► **Lemma 11.** *For any  $k$ , the  $k$ -bin hash algorithm is  $O(1)$ -switching-cost bounded.*

**Proof.** Let  $\psi$  denote the  $k$ -bin hash algorithm. Consider unit-distance pairs of worker/task inputs  $(W_1, T_1)$  and  $(W_2, T_2)$ . Changing  $W_1$  to  $W_2$  can change the assignments made by  $\psi$  for at most a constant number of bins. Similarly changing  $T_1$  to  $T_2$  can change the assignments made by  $\psi$  for at most a constant number of bins. Thus  $\psi(W_1, T_1)$  differs from  $\psi(W_2, T_2)$  by at most  $O(1)$  assignments.  $\blacktriangleleft$

Recall that  $\mathcal{A}$  is the composition of the  $O(\log w \log n)$  partial-assignment algorithms  $\mathcal{A}_{i,j}$ 's. The fact that each  $\mathcal{A}_{i,j}$  is  $O(1)$ -switching-cost bounded does not directly imply that  $\mathcal{A}$  is  $O(\log w \log n)$ -switching-cost bounded, however, because switching cost does not necessarily interact well with composition. In order to analyze  $\mathcal{A}$ , we show that each  $\mathcal{A}_{i,j}$  satisfies an additional property that we call being composition-friendly.

A partial-assignment algorithm  $\psi$  is **composition-friendly**, if for all unit-distance pairs of worker/task inputs  $(W_1, T_1)$  and  $(W_2, T_2)$ , the corresponding worker/task outputs  $(W'_1, T'_1)$  and  $(W'_2, T'_2)$  are also unit-distance.

Lemma 12 shows that each  $\mathcal{A}_{i,j}$  is composition-friendly.

► **Lemma 12.** *For any  $k$ , the  $k$ -bin hash is composition-friendly.*

**Proof.** Although the algorithm  $\psi$  is formally only defined on input  $(W, T)$  for which  $|W| = |T|$ , we will abuse notation here and consider  $\psi$  even on worker/task input  $(W, T)$  satisfying  $|W| \neq |T|$ .<sup>4</sup> Define the **difference-score** of a pair of worker/task inputs  $I_1 = (W_1, T_1), I_2 = (W_2, T_2)$  to be the quantity

$$d(I_1, I_2) = |W_1 \setminus W_2| + |W_2 \setminus W_1| + |T_1 \setminus T_2| + |T_2 \setminus T_1|.$$

<sup>4</sup> Indeed, the definition of the  $k$ -bin hash does not require a worker-task input with  $|W| = |T|$ . The only reason we require this equality in general is to simplify calculations, as in practice the algorithm will only be run on worker-task inputs of equal size.

We will show the stronger statement that the difference-score  $d(O_1, O_2)$  of the corresponding worker/task outputs  $O_1 = (W'_1, T'_1), O_2 = (W'_2, T'_2)$  satisfies

$$d(O_1, O_2) \leq d(I_1, I_2). \quad (3)$$

It suffices to consider only two special cases: the case in which  $W_2 = W_1 \cup \{\omega\}$  for some worker  $\omega$  and  $T_2 = T_1$ ; and the case in which  $T_2 = T_1 \cup \{\tau\}$  for some task  $\tau$  and  $W_2 = W_1$ . Iteratively applying these two cases to transform  $I_1$  into  $I_2$  implies inequality 3.

For this purpose, the roles of  $W$  and  $T$  are identical, so suppose without loss of generality that  $W_2 = W_1 \cup \{\omega\}$  for some worker  $\omega$  and  $T_2 = T_1$ . Recall that the assignment of workers and tasks to buckets is determined by some hash functions  $h_1, h_2$  and in particular is the same whether we input  $W_1$  or  $W_2$ . We first assign (only) the elements of  $W_1$  and  $T_1$  to their respective buckets, and then look at how including the assignment of  $\omega$  changes the worker-task output. If  $h_1$  assigns  $\omega$  to either a bin with no tasks or a bin which already has some lexicographically smaller worker, then we will have  $W'_2 = W'_1 \cup \{\omega\}$  and  $T'_2 = T'_1$ . If  $h_1$  assigns worker  $\omega$  to a bin with no other workers and at least one task, we let the smallest such task be  $\tau$  and see  $W'_2 = W'_1$  and  $T'_2 = T'_1 \setminus \{\tau\}$ . Finally, if  $h_1$  assigns  $\omega$  to a bin with only larger workers and at least one task, we let the minimal such worker be  $\gamma$ , and we see  $W'_2 = W'_1 \cup \{\gamma\}$  and  $T'_2 = T'_1$ . In all three cases,  $d(O_1, O_2) = 1$ , as desired. ◀

Next, we will show that composing composition-friendly algorithms has the effect of summing switching costs.

► **Lemma 13.** *Suppose that partial-assignment algorithms  $\psi_1, \psi_2, \dots, \psi_k$  are all composition-friendly, and that each  $\psi_i$  is  $s_i$ -switching-cost bounded. Then  $\psi_1 \circ \psi_2 \circ \dots \circ \psi_k$  is composition-friendly and is  $(\sum_i s_i)$ -switching-cost-bounded.*

**Proof.** By induction, it suffices to prove the lemma for  $k = 2$ . Let  $I_1 = (W_1, T_1)$  and  $I_2 = (W_2, T_2)$  be unit-distance worker/task inputs.

For  $i \in \{1, 2\}$ , let  $I'_i = (W'_i, T'_i)$  be the worker/task output for  $\psi_1(W_i, T_i)$ , and let  $I''_i = (W''_i, T''_i)$  be the worker/task output for  $\psi_2(W'_i, T'_i)$ .

Since  $\psi_1$  is composition friendly, its outputs  $I'_1$  and  $I'_2$  are unit distance. Since  $I'_1$  and  $I'_2$  are unit distance, and since  $\psi_2$  is composition friendly, the outputs  $I''_1$  and  $I''_2$  of  $\psi_2$  are also unit distance. Thus  $\psi_1 \circ \psi_2$  is composition friendly.

Since the inputs  $I_1$  and  $I_2$  to  $\psi_1$  are unit-distance,  $\psi_1(I_1)$  and  $\psi_1(I_2)$  differ in at most  $s_1$  worker-task assignments. Since the inputs  $I'_1$  and  $I'_2$  to  $\psi_2$  are also unit distance,  $\psi_2(I'_1)$  and  $\psi_2(I'_2)$  differ in at most  $s_2$  worker-task assignments. Thus the composition  $\psi_1 \circ \psi_2$  is  $(s_1 + s_2)$ -switching-cost bounded, as desired. ◀

We can now prove Proposition 10.

**Proof of Proposition 10.** By Lemma 11, each  $\mathcal{A}_{i,j}$  is  $O(1)$ -switching-cost bounded. By Lemma 12, each  $\mathcal{A}_{i,j}$  is composition friendly. Since  $\mathcal{A}$  is the composition of the  $O(\log w \log n)$  different  $\mathcal{A}_{i,j}$ 's, it follows by Lemma 13 that  $\mathcal{A}$  is  $O(\log w \log n)$ -switching-cost bounded. ◀

## 4 Lower bounds on switching cost

Define  $s_{w,t}$  to be the optimal switching cost for assignment functions that assign workers  $1, 2, \dots, w$  to multisets of  $w$  tasks from the universe  $[t]$ . The upper bounds in this paper establish that  $s_{w,t} \leq O(\log w \log(wt))$ . It is natural to wonder whether smaller bounds can be achieved, and in particular, whether a small switching cost that depends only on  $w$  can be achieved.



## 19:16 Memoryless Worker-Task Assignment with Polylogarithmic Switching Cost

It trivially holds that  $s_{w,t} \leq w$ . We show that when  $t$  is sufficiently large relative to  $w$ , there is a matching lower bound of  $s_{w,t} \geq w$ .

► **Theorem 3.** *For every  $w$  and  $t \geq \text{tow}(\Omega(w))$ , every worker-task assignment function has switching cost  $w$ .*

**Proof.** Given any worker-task assignment function  $\phi$ , we can actually find high switching cost between a pair of task subsets, in which all demands are 0 or 1. For each  $T \subseteq [t]$  of  $w$  tasks,  $\phi$  produces a bijection of workers  $[w]$  to tasks  $T$ . In order to lower-bound the switching cost, we produce a coloring of the complete  $w$ -uniform hypergraph with  $t$  vertices. The coloring will be designed so that, if it contains a monochromatic clique on  $w + 1$  vertices, then the assignment function  $\phi$  must have worst-possible switching cost  $w$ . By applying the hypergraph Ramsey theorem, we deduce that, if  $t$  is large enough, then the coloring must contain a monochromatic  $(w + 1)$ -clique, completing the lower bound.

**Coloring the complete  $w$ -uniform hypergraph on  $t$  vertices.** Let  $K_t^{(w)}$  denote the complete  $w$ -uniform hypergraph on  $t$  vertices. Note that the hyperedges of  $K_t^{(w)}$  are just the  $w$ -element subsets of  $[t]$ , which correspond to sets of tasks.

For a task set  $T = \{\tau_1, \dots, \tau_w\}$ , where  $1 \leq \tau_1 < \dots < \tau_w \leq t$ , we color the hyperedge  $T$  with the tuple  $\pi = \langle \pi(1), \pi(2), \dots, \pi(w) \rangle$ , where  $\tau_{\pi(i)}$  is the task assigned to worker  $i$ . One can think of  $\pi$  as a permutation of numbers  $\{1, 2, \dots, w\}$ , and thus the coloring consists of at most  $w!$  colors.

**Monochromatic  $(w + 1)$ -cliques imply high switching cost.** The key property of the coloring  $C$  is that, if  $K_t^{(w)}$  contains a monochromatic  $(w + 1)$ -vertex clique (i.e.,  $K_{w+1}^{(w)}$ ), then  $\phi$  must have switching cost  $w$ .

Namely, if  $K_t^{(w)}$  contains a monochromatic  $(w + 1)$ -clique, then we can find  $w + 1$  vertices,  $\tau_1 < \tau_2 < \dots < \tau_{w+1}$ , such that every  $w$ -element subset  $T$  of these tasks is assigned the same permutation  $\pi$  as its color. In particular, this means that for the task-set  $T_1 = \{\tau_1, \dots, \tau_w\}$  each worker  $i$  is assigned to task  $\tau_{\pi(i)}$ , but for the task-set  $T_2 = \{\tau_2, \dots, \tau_{w+1}\}$  that same worker  $i$  is assigned to a different task  $\tau_{\pi(i)+1}$ . Thus there is a pair of adjacent task sets  $T_1, T_2$  that exhibit switching cost  $w$ .

**Finding a monochromatic clique.** In order to complete the lower bound, we wish to show that, if  $t$  is sufficiently large, then the coloring contains a monochromatic  $K_{w+1}^{(w)}$ . To do this, we employ the hypergraph Ramsey theorem.

► **Theorem 14** (Theorem 1 in [12]). *Let  $k \geq 2$  and  $N \geq n \geq 2$  be positive integers. The hypergraph Ramsey number  $R(k, n, N)$  is defined to be the least positive integer  $M$  such that for every  $k$ -coloring of the hyperedges of  $K_M^{(n)}$ , the complete  $n$ -uniform hypergraph on  $M$  vertices, contains a monochromatic copy of  $K_M^{(n)}$ . This quantity satisfies*

$$R(k, n, N) \leq k^{(k^{n-1})(k^{n-2}) \dots (k^2)^{k(N-n)+1}}.$$

Applying Theorem 14, we see that if  $t \geq R(w!, w, w + 1)$ , then the  $(w!)$ -coloring of  $K_t^{(w)}$  contains a monochromatic  $(w + 1)$ -clique, and the assignment function  $\phi$  must have switching cost  $w$ , as desired. By Theorem 14,  $R(w!, w, w + 1) \leq \text{tow}(O(w))$ , which implies that every worker-task assignment function has switching cost  $w$  when  $t \geq \text{tow}(\Omega(w))$ . This completes the proof of Theorem 3. ◀

Another way of viewing this argument is that a worker-task assignment function with switching cost less than  $w$  gives rise to a proper  $(w!)$ -coloring of a certain graph, with vertex set  $\binom{[t]}{w}$  and edges of the form  $(\{\tau_1, \dots, \tau_w\}, \{\tau_2, \dots, \tau_{w+1}\})$  for  $\tau_1 < \tau_2 < \dots < \tau_{w+1}$ . Such graphs are studied under the name of **shift-graphs**, see, e.g., [14, Section 3.4], where the definition and proofs of basic properties are attributed to [11]. In particular, the chromatic number of shift-graphs is known to be  $(1 + o(1)) \cdot \log^{(w-1)} t$  (with the superscript denoting iteration). This gives an alternative way to complete the proof of Theorem 3 and it gives the same asymptotic bound on  $t$  in terms of  $w$ . While the chromatic number lower bound suffices to prove the switching cost bound, the nearly matching upper bound (on chromatic number) suggests that an entirely different technique would be needed in order to asymptotically improve the switching cost bound.

## 5 Densification into Hamming space

In this section, we apply our results on worker-task assignment to the problem of densification. In particular, we show how to embed sparse high-dimensional binary vectors into dense low-dimensional Hamming space.

Let  $\mathcal{H}_n^k$  denote the set of  $n$ -dimensional binary vectors with  $k$  ones. Let  $\mathcal{H}_k(n)$  denote the set of  $k$ -dimensional vectors with entries from  $[n]$ . We show that  $\mathcal{H}_n^k$  can be embedded into  $\mathcal{H}_k(n)$  with distortion  $O(\log n \log k)$ .

► **Theorem 4.** *There exists a map  $\phi : \mathcal{H}_n^k \rightarrow \mathcal{H}_k(n)$  such that, for every  $\vec{x}, \vec{y} \in \mathcal{H}_n^k$ ,*

$$\text{Ham}(\vec{x}, \vec{y})/2 \leq \text{Ham}(\phi(\vec{x}), \phi(\vec{y})) \leq O(\log n \log k) \text{Ham}(\vec{x}, \vec{y}).$$

**Proof.** Using Theorem 1, let  $\psi$  be a worker-task assignment function mapping workers  $1, 2, \dots, k$  to a task set  $T \subseteq [n]$  with switching-cost  $O(\log n \log k)$ .

For  $\vec{x} \in \mathcal{H}_n^k$ , define  $T(\vec{x}) = \{i \mid \vec{x}_i = 1\}$  to be the task set consisting of the positions in  $\vec{x}$  that are 1. Define  $\phi(\vec{x})$  to be the  $k$ -dimensional vector whose  $i$ -th coordinate denotes the task  $t \in T(\vec{x})$  to which  $\psi(T(\vec{x}))$  assigns worker  $i$ . For example, if  $k = 3$ ,  $\vec{x} = \langle 0, 1, 0, 1, 1, 0 \rangle$ , and  $\psi(T(\vec{x}))$  assigns workers 1, 2, 3 to tasks 4, 2, 5, respectively, then  $\phi(\vec{x}) = \langle 4, 2, 5 \rangle$ .

Since the coordinates of  $\phi(\vec{x})$  are a permutation of the positions  $T(\vec{x})$  in which  $\vec{x}$  is non-zero, it is necessarily the case that

$$\text{Ham}(\phi(\vec{x}), \phi(\vec{y})) \geq |T(\vec{x}) \setminus T(\vec{y})| \geq \text{Ham}(\vec{x}, \vec{y})/2.$$

On the other hand, since  $\psi$  has switching cost  $O(\log n \log k)$ , it is also the case that  $\psi(\vec{x})$  and  $\psi(\vec{y})$  differ by at most  $O(\log n \log k) \text{Ham}(\vec{x}, \vec{y})$  assignments, meaning that,

$$\text{Ham}(\phi(\vec{x}), \phi(\vec{y})) \leq O(\log n \log k) \text{Ham}(\vec{x}, \vec{y}).$$

This completes the proof of the theorem. ◀

We remark that Theorem 4 can be generalized to allow for the the domain space  $\mathcal{H}_n^k$  to have non-binary entries. In particular, if  $\mathcal{L}_n^k$  is the set of vectors with non-negative integer entries that sum to  $k$ , then there is an embedding  $\phi : \mathcal{L}_n^k \rightarrow \mathcal{H}_k(n)$  such that, for  $x, y \in \mathcal{L}_n^k$ ,

$$\ell_1(\vec{x}, \vec{y})/2 \leq \text{Ham}(\phi(\vec{x}), \phi(\vec{y})) \leq O(\log n \log k) \ell_1(\vec{x}, \vec{y}).$$

This follows from the same argument as Theorem 4, except that now  $T(x)$  is the *multiset* for which each element  $i \in [n]$  has multiplicity  $\vec{x}_i$ , and now  $\psi$  is the worker-task assignment mapping workers  $1, 2, \dots, k$  to a task *multiset*  $T \subseteq [n]$ .

## 6 Open problems

We leave open the question of closing the gap between upper and lower bounds for the worker-task assignment problem: the upper bound is  $\text{polylog}(wt)$  and the lower bound is  $\log^*(t)$ .

One interesting parameter regime is when  $w$  and  $t$  are comparable in size (say within a polynomial factor of each other). In this regime, no super-constant lower bound is known.

Another interesting direction is the problem of densification into Hamming space. Our upper bound for the worker-task assignment problem implies an upper bound for this problem, but our lower bound does not carry over. We leave open the problem of whether there is a better upper bound or a super-constant lower bound for this problem.

---

### References

- 1 Alexandr Andoni, Moses S Charikar, Ofer Neiman, and Huy L Nguyen. Near linear lower bound for dimension reduction in  $\ell_1$ . In *2011 IEEE 52nd Annual Symposium on Foundations of Computer Science*, pages 315–323. IEEE, 2011.
- 2 Nikhil Bansal, Niv Buchbinder, Aleksander Madry, and Joseph Naor. A polylogarithmic-competitive algorithm for the  $k$ -server problem. In *2011 IEEE 52nd Annual Symposium on Foundations of Computer Science*, pages 267–276. IEEE, 2011.
- 3 Radu Berinde, Anna C Gilbert, Piotr Indyk, Howard Karloff, and Martin J Strauss. Combining geometry and combinatorics: A unified approach to sparse signal recovery. In *2008 46th Annual Allerton Conference on Communication, Control, and Computing*, pages 798–805. IEEE, 2008.
- 4 Samuel N Beshers and Jennifer H Fewell. Models of division of labor in social insects. *Annual review of entomology*, 46(1):413–440, 2001.
- 5 Jean Bourgain. On Lipschitz embedding of finite metric spaces in Hilbert space. *Israel Journal of Mathematics*, 52(1-2):46–52, 1985.
- 6 Bo Brinkman and Moses Charikar. On the impossibility of dimension reduction in  $\ell_1$ . In *Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science*, page 514, 2003.
- 7 Diptarka Chakraborty, Elazar Goldenberg, and Michal Koucký. Streaming algorithms for embedding and computing edit distance in the low distance regime. In *Proceedings of the 48th annual ACM Symposium on Theory of Computing*, pages 712–725, 2016.
- 8 Moses Charikar, Ofir Geri, Michael P Kim, and William Kuszmaul. On estimating edit distance: Alignment, dimension reduction, and embeddings. In *45th International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 107, page 34, 2018.
- 9 Moses Charikar and Robert Krauthgamer. Embedding the Ulam metric into  $\ell_1$ . *Theory of Computing*, 2(1):207–224, 2006.
- 10 Moses Charikar and Amit Sahai. Dimension reduction in the  $\ell_1$  norm. In *The 43rd Annual IEEE Symposium on Foundations of Computer Science, 2002. Proceedings.*, pages 551–560. IEEE, 2002.
- 11 P Erdős and A Hajnal. On chromatic number of infinite graphs. In *Theory of Graphs (Proc. Colloq., Tihany, 1966)*, pages 83–98. Academic Press, 1968.
- 12 Paul Erdős and Richard Rado. Combinatorial theorems on classifications of subsets of a given set. *Proceedings of the London mathematical Society*, 3(1):417–439, 1952.
- 13 Jittat Fakcharoenphol, Satish Rao, and Kunal Talwar. A tight bound on approximating arbitrary metrics by tree metrics. In *Proceedings of the thirty-fifth annual ACM Symposium on Theory of Computing*, pages 448–455, 2003.
- 14 Stefan Felsner. *Interval orders: combinatorial structure and algorithms*. PhD thesis, Technische Universität Berlin, 1992. URL: <http://page.math.tu-berlin.de/~felsner/Paper/diss.pdf>.

- 15 Chryssis Georgiou and Alexander A Shvartsman. Cooperative task-oriented computing: Algorithms and complexity. *Synthesis Lectures on Distributed Computing Theory*, 2(2):1–167, 2011.
- 16 A. C. Gilbert, M. J. Strauss, J. A. Tropp, and R. Vershynin. Algorithmic linear dimension reduction in the  $\ell_1$  norm for sparse vectors. In *Allerton 2006 (44th Annual Allerton Conference on Communication, Control, and Computing)*, 2006.
- 17 Anna Gilbert and Piotr Indyk. Sparse recovery using sparse matrices. *Proceedings of the IEEE*, 98(6):937–947, 2010.
- 18 Piotr Indyk. Explicit constructions for compressed sensing of sparse signals. In *Proceedings of the nineteenth annual ACM-SIAM Symposium on Discrete Algorithms*, pages 30–33, 2008.
- 19 William B Johnson and Joram Lindenstrauss. Extensions of Lipschitz mappings into a Hilbert space. *Contemporary mathematics*, 26(189-206):1, 1984.
- 20 Michael JB Krieger, Jean-Bernard Billeter, and Laurent Keller. Ant-like task allocation and recruitment in cooperative robots. *Nature*, 406(6799):992, 2000.
- 21 Kristina Lerman, Chris Jones, Aram Galstyan, and Maja J Matarić. Analysis of dynamic task allocation in multi-robot systems. *The International Journal of Robotics Research*, 25(3):225–241, 2006.
- 22 Nathan Linial, Eran London, and Yuri Rabinovich. The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15(2):215–245, 1995.
- 23 Kathryn Sarah Macarthur, Ruben Stranders, Sarvapali D Ramchurn, and Nicholas R Jennings. A distributed anytime algorithm for dynamic task allocation in multi-agent systems. In *AAAI*, pages 701–706, 2011.
- 24 Jiří Matoušek. On variants of the Johnson–Lindenstrauss lemma. *Random Structures & Algorithms*, 33(2):142–156, 2008.
- 25 Colin McDiarmid. On the method of bounded differences. *Surveys in combinatorics*, 141(1):148–188, 1989.
- 26 James McLurkin and Daniel Yamins. Dynamic task assignment in robot swarms. In *Robotics: Science and Systems*, volume 8. Citeseer, 2005.
- 27 James Dwight McLurkin. *Stupid robot tricks: A behavior-based distributed algorithm library for programming swarms of robots*. PhD thesis, Massachusetts Institute of Technology, 2004.
- 28 Raghu Meka, Omer Reingold, and Yuan Zhou. Deterministic Coupon Collection and Better Strong Dispersers. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM)*, volume 28 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 872–884, 2014. doi:10.4230/LIPIcs.APPROX-RANDOM.2014.872.
- 29 Rafail Ostrovsky and Yuval Rabani. Low distortion embeddings for edit distance. *Journal of the ACM (JACM)*, 54(5):23–es, 2007.
- 30 Tsvetomira Radeva, Anna Dornhaus, Nancy Lynch, Radhika Nagpal, and Hsin-Hao Su. Costs of task allocation with local feedback: Effects of colony size and extra workers in social insects and other multi-agent systems. *PLoS computational biology*, 13(12):e1005904, 2017.
- 31 Gene E Robinson. Regulation of division of labor in insect societies. *Annual review of entomology*, 37(1):637–665, 1992.
- 32 Erol Şahin. Swarm robotics: From sources of inspiration to domains of application. In *International workshop on swarm robotics*, pages 10–20. Springer, 2004.
- 33 Hsin-Hao Su, Lili Su, Anna Dornhaus, and Nancy Lynch. Ant-inspired dynamic task allocation via gossiping. In *International Symposium on Stabilization, Safety, and Security of Distributed Systems*, pages 157–171. Springer, 2017.
- 34 Hsin-Hao Su and Nicole Wein. Lower Bounds for Dynamic Distributed Task Allocation. In *47th International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 168, pages 99:1–99:14, 2020. doi:10.4230/LIPIcs.ICALP.2020.99.