

Online Sorting and Online TSP: Randomized, Stochastic, and High-Dimensional

Mikkel Abrahamsen  

University of Copenhagen, Denmark

Ioana O. Bercea  

KTH Royal Institute of Technology, Stockholm, Sweden

Lorenzo Beretta  

University of California, Santa Cruz, CA, USA

Jonas Klausen  

University of Copenhagen, Denmark

László Kozma  

Institut für Informatik, Freie Universität Berlin, Germany

Abstract

In the *online sorting problem*, n items are revealed one by one and have to be placed (immediately and irrevocably) into empty cells of a size- n array. The goal is to minimize the sum of absolute differences between items in consecutive cells. This natural problem was recently introduced by Aamand, Abrahamsen, Beretta, and Kleist (SODA 2023) as a tool in their study of online geometric packing problems. They showed that when the items are reals from the interval $[0, 1]$ a competitive ratio of $O(\sqrt{n})$ is achievable, and no deterministic algorithm can improve this ratio asymptotically.

In this paper, we extend and generalize the study of online sorting in three directions:

- *randomized*: we settle the open question of Aamand et al. by showing that the $O(\sqrt{n})$ competitive ratio for the online sorting of reals cannot be improved even with the use of randomness;
- *stochastic*: we consider inputs consisting of n samples drawn uniformly at random from an interval, and give an algorithm with an improved competitive ratio of $\tilde{O}(n^{1/4})$. The result reveals connections between online sorting and the design of efficient hash tables;
- *high-dimensional*: we show that $\tilde{O}(\sqrt{n})$ -competitive online sorting is possible even for items from \mathbb{R}^d , for arbitrary fixed d , in an adversarial model. This can be viewed as an online variant of the classical TSP problem where tasks (cities to visit) are revealed one by one and the salesperson assigns each task (immediately and irrevocably) to its timeslot. Along the way, we also show a tight $O(\log n)$ -competitiveness result for *uniform metrics*, i.e., where items are of different *types* and the goal is to order them so as to minimize the *number of switches* between consecutive items of different types.

2012 ACM Subject Classification Theory of computation → Design and analysis of algorithms

Keywords and phrases sorting, online algorithm, TSP

Digital Object Identifier 10.4230/LIPIcs.ESA.2024.5

Related Version *Full Version*: <https://arxiv.org/abs/2406.19257> [2]

Funding *Ioana O. Bercea, Lorenzo Beretta and Jonas Klausen*: Supported by grant 16582, Basic Algorithms Research Copenhagen (BARC), from the VILLUM Foundation.

Mikkel Abrahamsen: Supported by Starting Grant 1054-00032B from the Independent Research Fund Denmark under the Sapere Aude research career programme and part of Basic Algorithms Research Copenhagen (BARC), supported by the VILLUM Foundation grant 16582.

László Kozma: Supported by DFG grant KO 6140/1-2.



© Mikkel Abrahamsen, Ioana O. Bercea, Lorenzo Beretta, Jonas Klausen, and László Kozma; licensed under Creative Commons License CC-BY 4.0

32nd Annual European Symposium on Algorithms (ESA 2024).

Editors: Timothy Chan, Johannes Fischer, John Iacono, and Grzegorz Herman; Article No. 5; pp. 5:1–5:15

Leibniz International Proceedings in Informatics



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

1 Introduction

The following natural problem called *online sorting* was recently introduced by Aamand, Abrahamsen, Beretta, and Kleist [1]: Given a sequence x_1, \dots, x_n of real values, assign them bijectively to array cells $A[1], \dots, A[n]$. Crucially, after receiving x_j , for $j = 1, 2, \dots, n$, we must immediately and irrevocably set $A[i] = x_j$, for some previously unused array index $i \in [n]$. The goal is to minimize $\sum_{i=1}^{n-1} |A[i+1] - A[i]|$, the sum of absolute differences between items in consecutive cells.

Aamand et al. [1] study the problem as modelling certain geometric online packing problems. In particular, they use it to show lower bounds on the competitive ratio of such problems¹. It is easy to see that the optimal (offline) solution of online sorting is to place the entries in sorted (increasing) order, and the question is how to approximate this solution in an online setting where the items are revealed one by one.

The problem evokes familiar scenarios where we must commit *step-by-step* to an ordering of items, such as when scheduling meetings in a calendar, writing recipes in a notebook, planting trees in a garden, or writing data into memory cells. In such situations, some local coherence or sortedness is often desirable, but once the location of an item has been assigned, it is expensive to change it; for instance, it may be difficult to reschedule meetings or to migrate data items once memory locations are referenced from elsewhere. One must then carefully balance between placing similar items next to each other and leaving sufficiently large gaps amid uncertainty about future arrivals².

Moreover, online sorting can be firmly placed among familiar and well-studied online problems; we briefly mention two. As they differ from online sorting in crucial aspects, a direct transfer of techniques appears difficult.

- *list labeling* or *order maintenance* [17, 12, 7, 27, 8]: in this problem, a sequence of values are to be assigned labels consistent with their ordering (in effect placing the values into an array). The main difference from online sorting is that the sequence must be fully sorted and the goal is to minimize *recourse*, i.e., movement of previously placed items.
- *matching on the line* [16, 25, 15, 5]: here, a sequence of clients (e.g., drawn from $[n]$) are to be matched uniquely and irrevocably to servers (say, locations in $[n]$). The problem differs from online sorting mainly in its cost function; the goal here is to minimize the sum of distances between matched client-server pairs.

One of the main results of Aamand et al. [1] is an algorithm for online sorting with competitive ratio $O(\sqrt{n})$. Aamand et al. require the input entries to come from the unit interval $[0, 1]$ and to contain the endpoints 0 and 1. Conveniently, this makes the offline cost equal to 1 and the competitive ratio equal to the online cost. We show that the same competitiveness result can be obtained even if these assumptions are relaxed.

Aamand et al. [1] also show that the $O(\sqrt{n})$ bound cannot be improved by any deterministic algorithm, and leave it as an open question whether it can be improved through randomization (assuming that the adversary is oblivious, i.e., that it does not know the “coin

¹ The *competitive ratio* of an online algorithm is the worst-case ratio of its cost to the optimum (offline) cost over inputs of a certain size n . The competitive ratio of a problem is the best competitive ratio achievable by an online algorithm for the problem.

² The cost measure of online sorting is also natural as a measure of the *unsortedness* of a sequence. For this, the *number of inversions* (i.e., the number of pairs $a < b$ where b appears before a) is perhaps more widely used. We argue, however, that with this latter cost, one cannot obtain a nontrivial competitive ratio. Consider an adversary that outputs $n/2$ copies of 0.5, followed by either all 0s or all 1s until the end. It is easy to see that one of these choices results in at least $n^2/8$ inversions for any online algorithm, whereas the optimal (sorted) sequence has zero inversions.

flips” of our algorithm). Note that in several online problems such as *paging* or *k-server*, randomization can lead to asymptotic improvements in competitiveness (against an oblivious adversary); e.g., see [11]. As our first main result, we show that for online sorting, randomization (essentially) does not help.

► **Theorem 1.** *The (deterministic and randomized) competitive ratio of online sorting is $\Theta(\sqrt{n})$. The lower bound $\Omega(\sqrt{n})$ holds even when the input numbers are from $[0, 1]$.*

Online TSP. Ordering real values with the above cost (sum of differences between consecutive items) can be naturally viewed as a one-dimensional variant of the Traveling Salesperson Problem (TSP). Suppose that n cities are revealed one by one (with repetitions allowed), and a salesperson must decide, for each occurrence c of a city, on a timeslot for visiting c , i.e., the position of c in the eventual tour. The cost is then the length of the fully constructed tour³. Formally, given a sequence (x_1, \dots, x_n) of items $x_i \in S$, for a metric space S with distance function $d(\cdot, \cdot)$, the goal is to assign the items bijectively to array cells $A[1], \dots, A[n]$ in an online fashion such as to minimize $\sum_{i=1}^{n-1} d(A[i], A[i+1])$. We refer to this problem⁴ as *online TSP in S* .

Online TSP in \mathbb{R} is exactly online sorting. A natural d -dimensional generalization is *online TSP in \mathbb{R}^d* , with the Euclidean distance $d(\cdot, \cdot)$ between items.⁵ A difficulty arising in dimensions two and above is that the optimal cost is no longer constant; in \mathbb{R}^d , the (offline) optimum may reach $\Theta(n^{1-\frac{1}{d}})$ even if the input points come from a unit box. Computing the optimum exactly is NP-hard even if $d = 2$ [23]. Our second main result is an online algorithm whose cost is $O(n^{1-\frac{1}{d+1}})$ and a competitiveness guarantee close to the one-dimensional case.

► **Theorem 2.** *There is a deterministic algorithm for online TSP in \mathbb{R}^d with competitive ratio $\sqrt{d} \cdot 2^d \cdot O(\sqrt{n \log n})$.*

As this setting includes online sorting as a special case, the lower bound of $\Omega(\sqrt{n})$ applies. A key step in obtaining Theorem 2 is the study of the *uniform metric* variant of the problem, i.e., the case of distance function d with $d(x, y) = 1$ if and only if $x \neq y$. This captures the natural problem where items (or tasks) fall into a certain number of *types*, and we wish to order them such as to minimize the number of *switches*, i.e., consecutive pairs of items of different types. For this case we prove a tight (deterministic and randomized) competitive ratio, independent of the number of different types, which may be of independent interest. Our algorithm is a natural greedy strategy; we analyze it by modelling the evolution of contiguous runs of empty cells as a coin-removal game between the algorithm and adversary.

► **Theorem 3.** *The competitive ratio of online sorting of n items under the uniform metric is $\Theta(\log n)$. The upper bound $O(\log n)$ is achieved by a deterministic algorithm and the lower bound $\Omega(\log n)$ also holds for randomized algorithms.*

³ A small technicality is whether the salesperson must return to the starting point or not. The effect of this in our cost regime, however, is negligible.

⁴ This model is sometimes referred to as the *online-list model*; it has been considered mostly in the context of scheduling problems [24, 14]. To our knowledge, TSP has not been studied in this setting before. Note however, that a different online model, called the *online-time model* has been used to study TSP [4, 21, 10]. In that model new cities can be revealed while the salesperson is already executing the tour; this may require changing the tour on the fly. Results in the two models are not comparable.

⁵ One may also view this task as a form of *dimensionality reduction*: we seek to embed a d -dimensional data set in a one-dimensional space (the array), while preserving some distance information.

Aamand et al. also consider the setting where the array is only partially filled (placing n items into $m > n$ cells). In this case, the cost is understood as the sum of distances $d(x, y)$ over pairs of items x, y with no other item placed between them. We thus extend our previous result to arrays of a larger size, obtaining a tight characterization.

► **Theorem 4.** *The competitive ratio (deterministic and randomized) of online sorting of n items with an array of size $\lceil \gamma n \rceil$, with $\gamma > 1$, under the uniform metric is $\Theta(1 + \log \frac{\gamma}{\gamma-1})$.*

Stochastic input. Given the (rather large) $\Omega(\sqrt{n})$ lower bound on the competitive ratio of online sorting, it is natural to ask whether we can overcome this barrier by relaxing the worst-case assumption on the input. Such a viewpoint has become influential recently in an attempt to obtain more refined and more realistic guarantees in online settings (e.g., see [26, 15]). A natural model is to view each item as drawn independently from some distribution, e.g., uniformly at random from a fixed interval. Our next main result shows an improved competitive ratio for such stochastic inputs:

► **Theorem 5.** *There is an algorithm for online sorting of n items drawn independently and uniformly at random from $(0, 1]$ that achieves competitive ratio $O((n \log n)^{1/4})$ with probability at least $1 - 2/n$.*

The algorithm illuminates a connection between online sorting and hash-based dictionaries [20, §6.4]. In the latter, the task is to place a sequence of n keys in an array of size $O(n)$, with the goal of minimizing *search time*. Fast searches are achieved by hashing keys to locations in the array. Due to hash collisions, not all elements can be stored exactly at their hashed location, and various paradigms have been employed to ensure that they are stored “nearby” (for the search to be fast). We adapt two such paradigms to online sorting. The first is to hash elements into buckets and solve the problem separately in each bucket. Each bucket has a fixed capacity, and so an additional *backyard* is used to store keys that do not fit in their hashed bucket [3, 9]. In Theorem 5, we use the values of the entries to assign them to buckets and employ a similar backyard design. We note some critical technical differences: in our design, all buckets must be full and we solve the problem recursively in each bucket; we also operate in a much tighter balls-into-bins regime, as we are hashing n elements into exactly n locations.

When the array size is allowed to be bigger than n , we employ yet another way for resolving hash collisions: the linear probing approach of Knuth [19]. Here, the value $\lceil \alpha n \rceil$ serves as the hash location of an entry $\alpha \in (0, 1)$. Intuitively, this is a good approximation for where the entry would appear in the sorted order. If we make sure that we place the entry close enough to this intended location, we can hope for a small overall cost. That is, we use the fact that linear probing places similar values close to each other. We get the following:

► **Theorem 6.** *For any $\gamma > 1$, there is an algorithm for online sorting of n items drawn independently and uniformly from $(0, 1)$ into an array of size $\lceil \gamma n \rceil$ that achieves expected competitive ratio $O\left(1 + \frac{1}{\gamma-1}\right)$.*

Paper structure. In §2 we present our results for the standard (one-dimensional) online sorting, in particular the lower bound for the randomized competitive ratio (Theorem 1). Results for online TSP in \mathbb{R}^d and results for the uniform metric (Theorems 2, 3, 4) are in §3. Our results for online sorting with stochastic input (Theorems 5, 6) are in §4. We conclude with a list of open questions in §5. We defer some proofs and remarks to the full version of the paper [2].

2 Competitiveness for online sorting

Given a sequence $X = (x_1, \dots, x_n) \in \mathbb{R}^n$ and a bijection $f: [n] \rightarrow [n]$ assigning “array cells” $A[i] = x_{f(i)}$, let $D_f(X) = \sum_{i=1}^{n-1} |A[i+1] - A[i]|$. Let $\text{OPT}(X)$ denote the quantity $\min_f D_f(X)$, i.e., the *offline optimum*.

An *online* algorithm is one that constructs the mapping f incrementally. Upon receiving x_j , for $j = 1, \dots, n$, the algorithm immediately and irrevocably assigns $f(i) = j$, for some previously unassigned $i \in [n]$. For an online algorithm \mathcal{A} , we denote by $\mathcal{A}(X)$ the cost $D_f(X)$ for the function f constructed by algorithm \mathcal{A} on input sequence X . We are interested in the competitive ratio $\mathcal{C}_{\mathcal{A}}(n)$ of an algorithm \mathcal{A} , i.e., the supremum of $\mathcal{A}(X)/\text{OPT}(X)$ over all inputs X of a certain size n , and in the best possible competitive ratio $\mathcal{C} = \mathcal{C}(n) = \inf_{\mathcal{A}} \mathcal{C}_{\mathcal{A}}(n)$ obtainable by any online algorithm \mathcal{A} for a given problem.

Upper bound. Aamand et al. [1] show that $\mathcal{C} \in \Theta(\sqrt{n})$, under the additional restrictions that $x_i \in [0, 1]$ for all $i \in [n]$, and that $\{0, 1\} \subseteq \{x_1, \dots, x_n\}$. As our first result, we employ a careful doubling strategy to show the same upper bound without the two restrictions, for general sequences of n reals (the lower bound clearly continues to hold).

▷ **Claim 7 (Proof in full version [2]).** There is a deterministic online algorithm \mathcal{A} for online sorting of an arbitrary sequence of n reals, with $\mathcal{C}_{\mathcal{A}} \in O(\sqrt{n})$.

Lower bound. The competitive ratio $\mathcal{C}_{\mathcal{A}}$ of a *randomized* algorithm \mathcal{A} is the supremum of $\mathbb{E}[\mathcal{A}(X)]/\text{OPT}(X)$ over all inputs X , where the expectation is over the random choices of \mathcal{A} . Aamand et al. [1] leave open the question of whether a lower bound of $\Omega(\sqrt{n})$ on the competitive ratio also holds for randomized algorithms. We settle this question in the affirmative. It is important to emphasize that the random choices of \mathcal{A} are not known to the adversary, i.e., we assume the *oblivious* model [11].

▷ **Claim 8.** $\mathcal{C}_{\mathcal{A}} \in \Omega(\sqrt{n})$ for every (possibly randomized) online algorithm \mathcal{A} .

In the remainder of the section we prove Claim 8, which implies Theorem 1. As usual, to lower bound the performance of a randomized algorithm (a distribution over deterministic algorithms), we lower bound instead the performance of a deterministic algorithm on an (adversarially chosen) distribution over input sequences.

Input distribution. Assume for simplicity that \sqrt{n} is an integer. Repeat the following until a sequence of length n is obtained. With probability $\frac{1}{2\sqrt{n}}$ set the remainder of the sequence to 0s. With probability $\frac{1}{2\sqrt{n}}$ set the remainder of the sequence to 1s. With probability $1 - \frac{1}{\sqrt{n}}$ offer $\frac{k}{\sqrt{n}}$ for $k = 0, \dots, \sqrt{n} - 1$ as the next \sqrt{n} elements of the sequence, and flip a new coin. We refer to the \sqrt{n} elements produced by the third option as an *epoch*. Notice that $\text{OPT} \leq 1$, thus it is enough to lower bound the expected cost of the algorithm.

Notation. Let T be a partially filled array, $|T|$ be the number of stored elements and $G(T)$ be the number of “gaps”, i.e., maximally contiguous groups of empty cells.

Let $f(T)$ be the minimum cost of *filling* T by an online algorithm, when the input sequence is chosen according to the distribution described above, where the cost of two neighboring elements $T[i], T[i+1]$ are accounted for when $T[i]$ or $T[i+1]$ is filled, whichever happens the latest.

More formally, define the cost of a partially filled array T to be

$$c(T) = \sum_{\substack{i=0 \\ T[i], T[i+1] \text{ are nonempty}}}^{n-1} |T[i] - T[i+1]|,$$

and let $\mathcal{A}(T)$ be the final array produced by an online algorithm \mathcal{A} when the remaining $n - |T|$ elements of the input sequence are generated as described above. Then $f(T) = c(\mathcal{A}(T)) - c(T)$ is the difference between the final cost and the cost of the partially filled array T . Thus $f(T) = 0$ if $|T| = n$, and $\mathbb{E}[f(\emptyset)]$ is the value we wish to bound, using \emptyset for the empty array.

For mappings T_1, T_2 where $T_2[i] = T_1[i]$ for all indices where $T_1[i]$ is nonempty we say that T_2 can be obtained from T_1 . For such a pair of mappings let $c(T_1, T_2) = c(T_2) - c(T_1)$ be the cost of transforming T_1 into T_2 according to the way of accounting given above.

Let $\mathcal{T}(T)$ be the set of mappings that can be obtained by inserting an epoch into T . That is, $\mathcal{T}(T)$ contains all mappings T' which can be obtained from T where the difference between T' and T corresponds to the elements of an epoch. Note that $|T'| = |T| + \sqrt{n}$ for $T' \in \mathcal{T}(T)$.

Let $\mathcal{L} = \{T: G(T) \leq \frac{\sqrt{n}}{8}\}$ and $\mathcal{H} = \{T: G(T) > \frac{\sqrt{n}}{8}\}$ be the sets of all partially filled arrays with a low/high number of gaps.

► **Lemma 9.** *Let $T_1, T_2 \in \mathcal{L}$ with $T_2 \in \mathcal{T}(T_1)$. Then $c(T_1, T_2) \geq \frac{3}{16}$.*

Proof. If an element is placed between two empty cells, the number of gaps will increase by one. If it is placed next to one or two occupied cells, the number of gaps stays constant or is reduced by one – call such an insertion an *attachment*. As $G(T_2) \leq G(T_1) + \frac{\sqrt{n}}{8}$ at least $\frac{7\sqrt{n}}{16}$ of the \sqrt{n} insertions of the epoch must be attachments.

An attachment incurs cost at least $\frac{1}{\sqrt{n}}$ unless the value(s) in the neighboring cell(s) and the newly inserted value are identical. As all values within an epoch are distinct, an attachment can only be without cost if the neighboring cell was occupied in T_1 (that is, before the epoch). As $G(T_1) \leq \frac{\sqrt{n}}{8}$ at most $\frac{\sqrt{n}}{4}$ occupied cells border an empty cell. The remaining $\sqrt{n}(\frac{7}{16} - \frac{1}{4}) = \frac{3\sqrt{n}}{16}$ attachments will incur non-zero cost. ◀

As a shorthand, let $(T+0)$ and $(T+1)$ be the mappings obtained by filling all gaps in T with 0s/1s, respectively.

► **Lemma 10.** $c(T, (T+0)) + c(T, (T+1)) \geq G(T)$.

Proof. Each gap in T is bordered by at least one non-empty cell $A[i]$. We have $|A[i] - 0| + |A[i] - 1| = 1$. ◀

For a mapping T the expected remaining cost can be bounded from below by

$$\begin{aligned} \mathbb{E}[f(T)] &\geq \frac{1}{2\sqrt{n}} \cdot c(T, (T+0)) + \frac{1}{2\sqrt{n}} \cdot c(T, (T+1)) \\ &\quad + \left(1 - \frac{1}{\sqrt{n}}\right) \cdot \min_{T' \in \mathcal{T}(T)} \{c(T, T') + \mathbb{E}[f(T')]\} \end{aligned}$$

and by Lemma 10 we thus have

$$\mathbb{E}[f(T)] \geq \frac{G(T)}{2\sqrt{n}} + \left(1 - \frac{1}{\sqrt{n}}\right) \cdot \min_{T' \in \mathcal{T}(T)} \{c(T, T') + \mathbb{E}[f(T')]\}. \quad (1)$$

$$\text{Let } \mathcal{L}(i) = \min_{T \in \mathcal{L}} \mathbb{E}[f(T)], \quad \text{and} \quad \mathcal{H}(i) = \min_{T \in \mathcal{H}} \mathbb{E}[f(T)] \\ |T|=n-i\cdot\sqrt{n} \qquad \qquad \qquad |T|=n-i\cdot\sqrt{n}$$

be the minimum expected cost of filling *any* array with $i \cdot \sqrt{n}$ empty cells, and which contains a low/high number of gaps, for $i \in \{0, 1, \dots, \sqrt{n}\}$, respectively $i \in \{0, 1, \dots, \sqrt{n} - 1\}$. (Note that $\mathcal{H}(\sqrt{n})$ is undefined as an empty array cannot have a high number of gaps.)

Combining Equation (1) with Lemma 9 we obtain

$$\mathcal{L}(i) \geq \left(1 - \frac{1}{\sqrt{n}}\right) \cdot \min\{3/16 + \mathcal{L}(i-1), \mathcal{H}(i-1)\}, \\ \mathcal{H}(i) \geq 1/16 + \left(1 - \frac{1}{\sqrt{n}}\right) \cdot \min\{\mathcal{L}(i-1), \mathcal{H}(i-1)\},$$

with $\mathcal{L}(0) = \mathcal{H}(0) = 0$. Our next lemma, proved in [2], leads to the lower bound.

► **Lemma 11.** $\mathcal{L}(\sqrt{n}) \in \Omega(\sqrt{n})$.

As the empty mapping \emptyset is contained in \mathcal{L} we have $\mathbb{E}[f(\emptyset)] \geq \mathcal{L}(\sqrt{n}) \in \Omega(\sqrt{n})$. Yao's minmax principle [22, Prop. 2.6] implies the lower bound on the expected cost of randomized algorithms for a worst-case input.

3 Competitiveness for online TSP

We now consider the generalization of online sorting that we call *online TSP*. Given a sequence $X = (x_1, \dots, x_n) \in S^n$ for some metric space S , and a bijection $f: [n] \rightarrow [n]$, and $A[i] = x_{f(i)}$, let $D_f(X) = \sum_{i=1}^{n-1} d(A[i+1], A[i])$. Here, $d(\cdot, \cdot)$ is a metric over S . As before, $\text{OPT}(X) = \min_f D_f(X)$, i.e., the *offline optimum*, and $\mathcal{A}(X)$ is the cost $D_f(X)$ for a function f constructed by an online algorithm \mathcal{A} on input sequence X . We define $\mathcal{C}_{\mathcal{A}}$ and \mathcal{C} as before.

Our main interest is in the case $S = \mathbb{R}^d$, and particularly $d = 2$, with $d(\cdot, \cdot)$ the Euclidean distance. As a tool in the study of the Euclidean case, we first look at a simpler, *uniform metric* problem in §3.1, showing a tight $\Theta(\log n)$ bound on the competitive ratio. Then, in §3.2 we study the Euclidean \mathbb{R}^2 and \mathbb{R}^d cases. As the uniform metric case is natural in itself, we revisit it in §3.3 in the setting where the array size is larger than n .

3.1 Uniform metric

Let S be an arbitrary discrete set and consider the distance function $d(x, y) = 0$ if $x = y$ and $d(x, y) = 1$ otherwise. Let $K = K(X)$ denote the number of distinct entries in the input sequence X , i.e., $K = |\{x_1, \dots, x_n\}|$. We give instance-specific bounds on the cost in terms of K and n . The following claim is obvious.

▷ **Claim 12.** $\text{OPT}(X) = K - 1$.

Next, we give a bound on the online cost and show that it is asymptotically optimal.

▷ **Claim 13.** There is an online algorithm \mathcal{A} with cost $\mathcal{A}(X) \leq K \log_2 n$.

▷ **Claim 14.** For every (possibly randomized) algorithm \mathcal{A} and all $K \geq 3$, there is an input distribution X with $K = K(X)$ such that $\mathbb{E}[\mathcal{A}(X)] \in \Omega(K \log n)$.

Note that the condition $K \geq 3$ is essential; if $K = 1$, then $\mathcal{A}(X) = \text{OPT}(X) = 0$ for every algorithm \mathcal{A} , and if $K = 2$, then there is an online algorithm \mathcal{A} that achieves $\mathcal{A}(X) = \text{OPT}(X) = 1$. (Place the two types of elements at the opposite ends of the array.)

Claims 12, 13, 14 together yield the main result of this subsection.

► **Theorem 3.** *The competitive ratio of online sorting of n items under the uniform metric is $\Theta(\log n)$. The upper bound $O(\log n)$ is achieved by a deterministic algorithm and the lower bound $\Omega(\log n)$ also holds for randomized algorithms.*

Proof of Claim 13. We describe algorithm \mathcal{A} (an alternative is presented in [2]), noting that it is not required to know K in advance. Assume without loss of generality that $\{x_1, \dots, x_n\} = \{1, 2, \dots, K\}$. For each j , with $1 \leq j \leq K$, maintain a *cursor* $c_j \in [n]$ indicating the array cell where the next item x_i is placed if it equals j . More precisely, if $x_i = j$, then let $f(c_j) = i$, and move the cursor to the right: $c_j = \min\{c_j + 1, n\}$.

If $f(c_j)$ is already assigned (i.e., the cell $A[c_j]$ is already written), set c_j to the mid-point of the largest empty contiguous interval. Similarly, when j is encountered the first time, then initialize c_j at the mid-point of the largest empty contiguous interval. Thus, initially, $c_{x_1} = \lfloor n/2 \rfloor$, i.e., place the first element at the middle of the array.

Clearly, \mathcal{A} can always place x_i *somewhere*, so \mathcal{A} correctly terminates. It remains to prove the upper bound on the number of unequal neighbors at the end of the process.

We model the execution of \mathcal{A} as a **coin-game**. Consider a number of up to K *piles* of coins. The game starts with a single pile of n coins. An adversary repeatedly performs one of two possible operations: (1) remove one coin from an arbitrary pile, (2) split the *largest* pile into two equal parts. Operation (2) is only allowed when the number of piles is less than K , and only for a pile of at least two coins. The game ends when all coins have been removed.

It is easy to see that this game models the execution of \mathcal{A} in the sense that for any execution of \mathcal{A} on X there is an execution of the coin-game in which the sizes of the piles correspond at each step to the lengths of the contiguous empty intervals in the array. Moreover, the number of consecutive unequal pairs at the end of algorithm \mathcal{A} (i.e., the cost $\mathcal{A}(X)$) is at most the number of splits (i.e., operations (2)) of the coin-game execution. It is thus sufficient to upper bound the number of splits in *any* execution of the coin game.

Let n_i denote the size of a pile before its split, for the i -th split operation. As n_i is the size of the largest pile and piles can only get smaller, the sequence n_i is non-increasing. Suppose the i -th split replaces a pile of size t with two piles of size $t/2$. Then, after the split there are at most $K - 2$ piles with sizes in $[t/2, t]$ and no pile greater than t . Thus, after at most $K - 2$ further splits, we split a pile of size at most $t/2$. (Possible operations (1) can only strengthen this claim, as they make some piles smaller.) It follows that $n_{i+K} \leq n_i/2$ for all i . As $n_1 \leq n$, and $n_i \geq 1$ for all i , the number of splits is at most $K \log_2 n$. ◁

Proof of Claim 14. Let \mathcal{A} be an algorithm filling the items into an array A of size n . We present a distribution over inputs X incurring cost $\mathbb{E}[X] \geq \Omega(K \log n)$. By Yao's minmax principle [22, Prop. 2.6] we thus get that every randomized algorithm has a worst-case input of cost $\Omega(K \log n)$.

For each free cell $A[i]$ we say that $A[i]$ is *friendly* for the two (possibly identical) elements y_k and y_l that are placed closest to the left and right of $A[i]$ in A . When asked to place an element y into A , \mathcal{A} will increase the cost of the partial solution *unless* y is placed in a cell that is friendly for y . When placing y in a cell that is friendly for y , the number of friendly cells for y will decrease by one. The number of friendly cells for other elements may or may not decrease, but no element will have more friendly cells than before the insertion of y .

Let $z_1 \leq z_2 \leq \dots \leq z_K$ be the number of friendly cells for each of the K values, sorted nondecreasingly. With n' free cells in A , $\sum_{i=1}^K z_i \leq 2n'$, and hence $z_1 \leq 2n'/K$. Consider the median $z_{\lfloor K/2 \rfloor}$. As $\sum_{i=\lfloor K/2 \rfloor}^K z_i \leq 2n'$, we have $z_1 \leq \dots \leq z_{\lfloor K/2 \rfloor} \leq \frac{2n'}{\frac{1}{2}K} = 4n'/K$.

We now describe our input distribution, proceeding in *epochs*: At the start of each epoch a value y is chosen uniformly at random from $\{1, \dots, K\}$, and this element is presented all through the epoch. If $K \in \{3, 4\}$, the epoch will consist of $2n'/K$ copies of y , where n' is

the number of unoccupied cells at the start of the epoch. By the first observation above, $\Pr[z_y \leq 2n'/K] \geq 1/K \geq 1/4$, in which case the input forces \mathcal{A} to increase the cost of the partial solution.

If $K > 4$, we instead let the epoch be of length $4n'/K$, and by the second observation above we obtain that $\Pr[z_y \leq 4n'/K] \geq 1/2$, again increasing the cost of the partial solution with constant probability.

To establish a lower bound for the expected cost produced by this input, it remains to lower bound the number of epochs processed. For small K , the epoch leaves $n' \cdot \frac{K-2}{K}$ free cells for the coming epochs. For $K > 4$ the epoch leaves $n' \cdot \frac{K-4}{K}$ free cells. Thus, in both cases, the number of epochs is at least

$$\log_{\frac{K}{K-4}}(n) = \frac{\log_2(n)}{\log_2\left(1 + \frac{4}{K-4}\right)} \geq \frac{K-4}{4} \log_2(n).$$

Hence, the input will force \mathcal{A} to produce a solution of expected cost $\Omega(K \log n)$. \triangleleft

3.2 Online TSP in \mathbb{R}^d

We now proceed to the case where $S = \mathbb{R}^d$ and $d(\cdot, \cdot)$ is the Euclidean distance. We start with the first new case, $d = 2$. For ease of presentation, we omit floors and ceilings in the analysis. We assume that the input points are from the unit box $[0, 1]^2$ and that the optimum length is at least 1. These assumptions can be relaxed by a similar doubling-approach as in the proof of Claim 7.

The following result is well known [6, 18, 13] (consider, e.g., a $\sqrt{n} \times \sqrt{n}$ uniform grid).

\triangleright **Claim 15.** For all sequences X of n points in $[0, 1]^2$, we have $\text{OPT}(X) \in O(\sqrt{n})$. Moreover, there exists a sequence X of n points in $[0, 1]^2$ such that $\text{OPT}(X) \in \Omega(\sqrt{n})$.

As a warm-up before our competitiveness result, we give an upper bound on the cost of an online algorithm and we show the tightness of this bound. The arguments extend the one-dimensional ones in a straightforward way. The proofs of the following can be found in the full version of the paper [2].

\triangleright **Claim 16.** There is an online algorithm \mathcal{A} such that $\mathcal{A}(X) \in O(n^{2/3})$ for all $X \in [0, 1]^2$.

\triangleright **Claim 17.** For every deterministic \mathcal{A} there is an input X such that $\mathcal{A}(X) \in \Omega(n^{2/3})$.

Now we move to the study of the competitive ratio. Note that the lower bound of Claim 8 immediately applies to our setting. Combined with Claim 16, it follows that $\mathcal{C} \in O(n^{2/3}) \cap \Omega(n^{1/2})$. In the following we (almost) close this gap, proving the following.

\triangleright **Claim 18.** The competitive ratio of online TSP in \mathbb{R}^2 is $O(\sqrt{n \log n})$.

Partition the box $[0, 1]^2$ into $t \times t$ boxes of sizes $1/t \times 1/t$, for t to be set later. Let K be the number of boxes that are *touched*, i.e., that contain some input point x_i . We first need to lower bound the optimum.

\blacktriangleright **Lemma 19.** $\text{OPT}(X) \geq K/4t$.

Proof. If $K/4t < 1$, we are done, as $\text{OPT}(X) \geq 1$ by assumption. Assume therefore $K \geq 4t$. Choose an arbitrary representative point of X from each touched box. Observe that the optimum cannot be shorter than the optimal tour of the representatives (by triangle inequality). Among any five consecutive vertices of the representatives-tour, two must be from non-neighboring boxes, thus the total length of the four edges connecting these points is at least $1/t$. It follows that at least $K/4$ edges have this length, yielding the claim. \blacktriangleleft

Now we need an algorithm that does well in terms of K . For any two neighboring entries in the array we consider only whether they come from the same box. This allows us to reduce our problem to the uniform metric case (Claim 13).

▷ **Claim 20.** There is an online algorithm \mathcal{A} such that $\mathcal{A}(X) \in O(K \log n + n/t)$.

Proof. We treat points from the same box as having the *same value*. We run the algorithm for the uniform metric given in Claim 13, incurring a total number of $O(K \log n)$ differing pairs of neighbors. For these pairs we account for a maximum possible cost of $\sqrt{2}$. All other neighboring pairs have the same value (= come from the same box), incurring a cost of at most $O(1/t)$ each, for a total of $O(n/t)$. This completes the proof. ◀

Together with Lemma 19, this yields Claim 18. Indeed, set $t = \sqrt{\frac{n}{\log n}}$. If $K \leq t$, then using $\text{OPT} \geq 1$, we have $\mathcal{A}/\text{OPT} \leq \frac{K \log n + n/t}{1} \in O(\sqrt{n \log n})$. If $K > t$, then using Lemma 19, $\mathcal{A}/\text{OPT} \leq \frac{K \log n + n/t}{K/4t} = 4t \log n + 4n/K \in O(\sqrt{n \log n})$.

Online TSP in \mathbb{R}^d . We now extend the bound on the competitive ratio (Claim 18) to the higher dimensional case, also considering the dependence on d , leading to the claimed result.

► **Theorem 2.** *There is a deterministic algorithm for online TSP in \mathbb{R}^d with competitive ratio $\sqrt{d} \cdot 2^d \cdot O(\sqrt{n \log n})$.*

Analogously to Claim 15, we first state an absolute bound on the optimal TSP cost in d dimensions, treating d as a constant [6, 18, 13].

▷ **Claim 21.** For all sets X of n points in $[0, 1]^d$, we have $\text{OPT}(X) \in O(n^{1-1/d})$. Moreover, there exists a set X of n points in $[0, 1]^d$ such that $\text{OPT}(X) \in \Omega(n^{1-1/d})$.

A straightforward generalization of Claims 16 and 17 yields (for all constant d):

▷ **Claim 22.** There is an online algorithm \mathcal{A} for online TSP in \mathbb{R}^d such that $\mathcal{A}(X) \in O(n^{1-\frac{1}{d+1}})$ for all X . For every deterministic algorithm \mathcal{A} there is an input X such that $\mathcal{A}(X) \in \Omega(n^{1-\frac{1}{d+1}})$.

Proof of Theorem 2. We follow the proof of Claim 18, with minor changes. We assume the input to come from the d -dimensional unit box $[0, 1]^d$. We partition this box into t^d boxes of sizes $(1/t)^d$. When adapting Lemma 19, we have to consider 2^d (instead of 4) consecutive tour edges, thus obtaining $\text{OPT}(X) \geq \frac{K}{2^d t}$. When adapting Claim 20, our upper bound increases by a factor of \sqrt{d} , the distance between antipodal vertices of a unit d -cube, replacing the implicit $\sqrt{2}$ in the earlier bound. Thus, when bounding the competitive ratio, we incur an overall factor of $\sqrt{d} \cdot 2^d$, yielding the result. ◀

3.3 Uniform metric with a larger array

We study the problem of placing n elements $\{1, \dots, K\}$ into an array of size $\lceil \gamma n \rceil$ for a fixed $\gamma > 1$. The algorithm and distribution from Claims 13 and 14 can be reused in this setting, stopping either process after the insertion of the first n elements.

▷ **Claim 23.** For every algorithm \mathcal{A} and any number $K \geq 3$ of input values, there is an input distribution X such that $\mathbb{E}[\mathcal{A}(X)] \in \Omega(K \cdot (1 + \log(\gamma/(\gamma - 1))))$.

▷ **Claim 24.** There is an online algorithm \mathcal{A} with cost $\mathcal{A}(X) \in O(K \cdot (1 + \log(\gamma/(\gamma - 1))))$.

The two claims together imply Theorem 4. We defer their proofs to the full version of the paper [2]. They rely on similar arguments as the proofs of the claims mentioned above, with a refined argument for counting the number of epochs/moves performed.

4 Competitiveness for stochastic online sorting

In this section, we prove Theorem 5 and Theorem 6 (§4.1).

► **Theorem 5.** *There is an algorithm for online sorting of n items drawn independently and uniformly at random from $(0, 1]$ that achieves competitive ratio $O((n \log n)^{1/4})$ with probability at least $1 - 2/n$.*

We start by describing the general design of the algorithm from Theorem 5 and showing some fundamental properties. We then describe a recursive algorithm and give the parameterizations that achieve the desired competitive cost.

The general design. We let α and β denote two parameters in $(0, 1)$ which we will set later. We decompose the array A of n cells as such: the first $N = n - n^\beta$ cells are divided into $M = n^\alpha$ consecutive sub-arrays A_1, \dots, A_{n^α} , and the remaining n^β cells (at the end of the array) form one single subarray denoted by B . We refer to each of the sub-arrays A_i as a *bucket* and to B as the *backyard*. We note that each bucket A_i can hold $C = N/M$ elements, which we refer to as the *capacity* of the bucket.

We use the values of the elements to hash them into the array. Namely, for an element $x \in (0, 1)$, we define $h: (0, 1] \rightarrow \{1, \dots, M\}$ by setting $h(x) = \lceil x \cdot M \rceil$. In other words, the elements in the interval $(0, 1/M]$ will all hash to bucket 1, elements in the interval $(1/M, 2/M]$ will hash to bucket 2, etc. Since the elements are chosen independently and uniformly at random from $(0, 1]$, we get that h assigns the elements independently and uniformly at random into the M buckets.⁶

SortUnif₁(A, n): the first algorithm. Let $\text{SortDet}(A, n)$ denote the deterministic algorithm from [1]. We now define the algorithm $\text{SortUnif}_1(A, n)$ as such: upon receiving x , it checks if the bucket $A_{h(x)}$ has any empty cells. If so, it forwards x to $\text{SortDet}(A_{h(x)}, C)$. Otherwise, it forwards x to $\text{SortDet}(B, n^\beta)$ (and we say that the corresponding bucket is *full*).

We first prove that $\text{SortUnif}_1(A, n)$ successfully places all items with high probability. Note that this is not always guaranteed: it could happen that both the bucket and the backyard are full (before we have managed to place all the n elements). If this happens, then among all n elements, strictly less than C hash into some bucket. We call this event a *failure* and show the following by employing a Chernoff bound (proof is given in [2]):

▷ **Claim 25.** Given any $c > 0$, if $\beta \geq \frac{1}{2} \cdot \left(1 + \alpha + \frac{\ln \ln n + \ln(2(c+1))}{\ln n}\right)$, then $\text{SortUnif}_1(A, n)$ fails with probability at most $1/n^c$.

We now bound the cost of $\text{SortUnif}_1(A, n)$:

▷ **Claim 26.** If $\text{SortUnif}(A, n)$ does not fail, then its cost is at most $O(\sqrt{C} + n^{\beta/2})$.

Proof. Since each bucket receives at least C elements, the cost of placing elements inside bucket A_1 is given by the cost of $\text{SortDet}(A_1, C)$ on elements from $(0, 1/M)$. We bound this by $O(\sqrt{C} \cdot 1/M]$. For the remaining buckets, the elements are drawn from $(i/M, (i+1)/M]$. This is equivalent to sorting elements from $(0, 1/M]$, and so their cost is the same as that of A_1 . Therefore, in total, the cost from each individual bucket is $O(\sqrt{C})$. In addition, we also have the cost from crossing from one bucket to the next. This is at most $2/M$, since

⁶ We assume that there is no element of value 0. The probability of this happening is 0.

the maximum difference between elements from consecutive buckets is at most $2/M$. Since there are M buckets, this amounts to a cost of at most 2. The cost of crossing from bucket A_M to the backyard is 1. Finally, the cost in the backyard is $O(n^{\beta/2})$, since it employs $\text{SortDet}(B, n^\beta)$ on elements from $(0, 1)$. \triangleleft

We instantiate α and β such that we get the following (proof in [2]):

► **Lemma 27.** *Given any $c > 0$, $\text{SortUnif}_1(A, n)$ has cost at most $O(n^{1/3} \cdot \ln^{1/6} n \cdot (2(c+1))^{1/6})$ with probability at least $1 - 1/n^c$.*

SortUnif_k: recursing on the buckets. We take advantage of the fact that within each bucket, the elements are chosen uniformly at random. That is, we can apply the same strategy recursively inside the bucket. We get a series of algorithms SortUnif_k for $k \geq 2$. In SortUnif_k , we let α and β be defined as in SortUnif_1 . When we see an element x with $h(x) = i$, if its bucket A_i is not full, we place $(xM - i + 1)$ using $\text{SortUnif}_{k-1}(A_{h(x)}, C)$. Note that we have already conditioned on the fact that $\lceil x \cdot M \rceil = i$. In this case, $(xM - i + 1)$ becomes uniformly distributed in $(0, 1)$. If A_i is full, we place x in the backyard according to $\text{SortDet}(B, n^\beta)$.

Note that there are now several causes for failure: either some bucket A_i is not full, or one of the algorithms inside the bucket fails. We bound the probability of either of these happening as follows, proving the claim by induction over k (proof in [2]):

► **Lemma 28.** *Given any $c > 0$, $\text{SortUnif}_k(A, n)$ has cost at most*

$$O\left(n^{1/f_k} \cdot \ln^{1/4} n \cdot (2(c+1))^{B_k}\right)$$

with probability at least $1 - 2/n^c$, where $f_k = 4 - 1/2^{k-1}$ and $B_k = k/4$.

Setting $c = 1$ and $k = \log_2\left(2/(4 + \ln n - \sqrt{\ln^2 n + 8 \ln n})\right)$ yields the bounds claimed in Theorem 5; the details are given in the full paper [2].

4.1 Stochastic online sorting in larger arrays

► **Theorem 6.** *For any $\gamma > 1$, there is an algorithm for online sorting of n items drawn independently and uniformly from $(0, 1)$ into an array of size $\lceil \gamma n \rceil$ that achieves expected competitive ratio $O\left(1 + \frac{1}{\gamma-1}\right)$.*

The algorithm. Define $\alpha = (\gamma - 1)/10$ and $\beta = \gamma - \alpha$ such that array A has size $(\beta + \alpha)n$. We will refer to the final αn cells of A as the *buffer*. The first βn cells we consider to correspond to the interval $[0, 1)$. More specifically, cell $A[i]$ represents the interval $\left[\frac{i}{\beta n}, \frac{i+1}{\beta n}\right)$.

Let $h: [0, 1) \rightarrow \{0, \beta n - 1\}$ be the function which maps values $x \in [0, 1)$ to the index of the corresponding cell: $h(x) = \left\lfloor \frac{x}{\beta n} \right\rfloor$. Our algorithm \mathcal{A} inserts each value x into $A[h(x)]$ if the cell is available, and otherwise into the *first* available cell following $h(x)$, possibly wrapping around from $A[\gamma n - 1]$ to $A[0]$ (although, as we argue in the full version [2], this is unlikely). In other words, x is inserted into the cell $A[h(x) + i \bmod \gamma n]$ where i is the smallest non-negative integer such that the specified cell is available.

Counting steps. As the values are drawn independently and uniformly from $[0, 1)$, the indices $h(x)$ are also independently and uniformly distributed in $\{0, \dots, \beta n - 1\}$. The algorithm thus mirrors the scheme of linear probing commonly used for implementing hash tables. In linear probing each key x that is to be added to the table (and which is generally not assumed to come from a known distribution) is hashed to an index of the array, and x is inserted into the first following free cell.

When implementing a hash table, one is interested in analyzing the *speed* of insertions, which for linear probing corresponds to the number of cells that are probed before a free cell is found. Although speed is not a concern for our online algorithm, we will nonetheless show that the number of steps performed is an important measure for bounding the cost of our solution. Formally, let $s(k)$ be the number of steps performed when inserting the k -th value x_k of the input. If $A[h(x_k)]$ is free, then $s(k) = 1$; otherwise $s(k) = i + 1$, when x_k is inserted into cell $A[h(x_k) + i \bmod \gamma n]$.

Now consider the linear probing process \mathcal{A}_{LP} inserting n elements into a table T of size βn (that is, without the buffer space). Analogously to $s(i)$ we let $s_{LP}(i)$ be the number of steps performed by \mathcal{A}_{LP} when inserting the i -th element. We couple the processes \mathcal{A} and \mathcal{A}_{LP} so that they encounter the same stream of indices $h(x_1), h(x_2), \dots, h(x_n)$ during execution. We then have $s(i) \leq s_{LP}(i)$. Indeed, the only difference between the processes is that \mathcal{A} has an extra αn cells of buffer space to prevent wraparound at the end of A . More generally, we thus have $\mathbb{E}[s(i)] \leq \mathbb{E}[s_{LP}(i)]$. Combining the pieces we get that:

$$\triangleright \text{Claim 29. } \mathbb{E}[\mathcal{A}(X)] \leq O\left(1 + \frac{1}{\beta n} \sum_{i=1}^n \mathbb{E}[s(i)]\right).$$

The proof can be found in the full version [2]. We then invoke the following classic result by Knuth [19] to bound $\sum_i \mathbb{E}[s_{LP}(i)]$:

▶ **Theorem 30** ([19]). *Consider the process of inserting $(1 - \varepsilon)m$ elements into an array of size m by linear probing. When hash values are assigned uniformly and independently,*

$$\sum_{i=1}^{(1-\varepsilon)m} \mathbb{E}[s_{LP}(i)] \in O\left((1 - \varepsilon)m \cdot \frac{1}{\varepsilon}\right).$$

To apply Theorem 30, set $m = \beta n$ and $\varepsilon = (\beta - 1)/\beta$ such that $(1 - \varepsilon)m = n$, and thus $\sum_{i=1}^n \mathbb{E}[s(i)] \leq \sum_{i=1}^n \mathbb{E}[s_{LP}(i)] \in O\left(n \cdot \left(\frac{\beta}{\beta-1}\right)\right)$. As $\beta - 1 = \frac{\beta}{10} \cdot (\gamma - 1)$, we have $\mathbb{E}[\mathcal{A}(X)] \in O\left(1 + \frac{1}{\beta} \cdot \frac{\beta}{\beta-1}\right) \subseteq O\left(1 + \frac{1}{\gamma-1}\right)$, proving Theorem 6. ◀

5 Conclusion and open questions

In this paper, we studied the online sorting problem in some of its variants. Several questions and directions to further study remain, we mention those that we find most interesting.

Online sorting of reals. For large arrays ($m = \lceil \gamma n \rceil$, for $\gamma > 1$) significant gaps remain between the upper and lower bounds on the competitive ratio of online sorting (see [1]).

In a different direction, consider online sorting and *online list labeling*, mentioned in §1. These can be seen as two extremes of an “error vs. recourse” trade-off: online sorting allows no recourse, while list labeling allows no error. Achieving a smooth trade-off between the two optimization problems by some hybrid approach is an intriguing possibility.

Stochastic and other models. In the stochastic setting of online sorting, we could improve upon the worst-case bound, but the obtained ratio (Theorem 5) is likely not optimal; we are not aware of nontrivial lower bounds. For stochastic online sorting with large arrays (Theorem 6), a matching lower bound is likewise missing. We restricted our study to uniform distribution over an interval; extending the results to other distributions remains open.

Other models that go beyond the worst-case assumption could yield further insight; we mention two possible models: (1) online sorting *with advice* (e.g., from a machine learning model): suppose that each input item comes with a possibly unreliable estimate of its rank in the optimal (sorted) sequence; what is the best way to make use of this advice? (2) online sorting *with partially sorted input*: for instance, suppose that the *input cost* $\sum_{i=1}^{n-1} |x_{i+1} - x_i|$ is small. What is the best guarantee for the online sorting cost in this case?

Online TSP in various metrics. For online TSP in \mathbb{R}^d , for fixed d , a small gap remains between the lower bound $\Omega(\sqrt{n})$ and upper bound $O(\sqrt{n} \log n)$ on the competitive ratio (Theorem 2). The dependence on d (when $d \in \omega(1)$) is not known to be optimal. Online TSP in \mathbb{R}^d with large array, and/or with stochastic input are also interesting directions.

Some non-Euclidean metrics pose natural questions, e.g., L_0 or L_∞ in \mathbb{R}^d , or tree- or doubling metrics. More broadly, is the optimal competitive ratio $O(\sqrt{n})$ for arbitrary metrics?

References

- 1 Anders Aamand, Mikkel Abrahamsen, Lorenzo Beretta, and Linda Kleist. Online sorting and translational packing of convex polygons. In Nikhil Bansal and Viswanath Nagarajan, editors, *Proceedings of the 2023 ACM-SIAM Symposium on Discrete Algorithms, SODA 2023, Florence, Italy, January 22-25, 2023*, pages 1806–1833. SIAM, 2023. doi:10.1137/1.9781611977554.ch69.
- 2 Mikkel Abrahamsen, Ioana O. Bercea, Lorenzo Beretta, Jonas Klausen, and László Kozma. Online sorting and online tsp: randomized, stochastic, and high-dimensional, 2024. doi:10.48550/arXiv.2406.19257.
- 3 Yuriy Arbitman, Moni Naor, and Gil Segev. Backyard cuckoo hashing: Constant worst-case operations with a succinct representation. In *2010 IEEE 51st Annual symposium on foundations of computer science*, pages 787–796. IEEE, 2010. doi:10.1109/FOCS.2010.80.
- 4 Giorgio Ausiello, Esteban Feuerstein, Stefano Leonardi, Leen Stougie, and Maurizio Talamo. Algorithms for the on-line travelling salesman 1. *Algorithmica*, 29:560–581, 2001. doi:10.1007/s004530010071.
- 5 Eric Balkanski, Yuri Faenza, and Noémie Pérvier. The power of greedy for online minimum cost matching on the line. In *Proceedings of the 24th ACM Conference on Economics and Computation*, pages 185–205, 2023. doi:10.1145/3580507.3597794.
- 6 Jillian Beardwood, J. H. Halton, and J. M. Hammersley. The shortest path through many points. *Mathematical Proceedings of the Cambridge Philosophical Society*, 55(4):299–327, 1959. doi:10.1017/S0305004100034095.
- 7 Michael A Bender, Richard Cole, Erik D Demaine, Martin Farach-Colton, and Jack Zito. Two simplified algorithms for maintaining order in a list. In *European Symposium on Algorithms*, pages 152–164. Springer, 2002. doi:10.1007/3-540-45749-6_17.
- 8 Michael A Bender, Alex Conway, Martín Farach-Colton, Hanna Komlós, William Kuszmaul, and Nicole Wein. Online list labeling: Breaking the $\log 2 n$ barrier. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 980–990. IEEE, 2022. doi:10.1109/FOCS54457.2022.00096.
- 9 Michael A. Bender, Martín Farach-Colton, John Kuszmaul, William Kuszmaul, and Mingmou Liu. On the optimal time/space tradeoff for hash tables. In Stefano Leonardi and Anupam Gupta, editors, *STOC '22: 54th Annual ACM SIGACT Symposium on Theory of Computing, Rome, Italy, June 20 - 24, 2022*, pages 1284–1297. ACM, 2022. doi:10.1145/3519935.3519969.

- 10 Antje Bjelde, Jan Hackfeld, Yann Disser, Christoph Hansknecht, Maarten Lipmann, Julie Meißner, Miriam Schlöter, Kevin Schewior, and Leen Stougie. Tight bounds for online tsp on the line. *ACM Transactions on Algorithms (TALG)*, 17(1):1–58, 2020. doi:10.1145/3422362.
- 11 Allan Borodin and Ran El-Yaniv. *Online computation and competitive analysis*. Cambridge University Press, 1998.
- 12 Paul F Dietz. Maintaining order in a linked list. In *Proceedings of the fourteenth annual ACM symposium on Theory of computing*, pages 122–127, 1982. doi:10.1145/800070.802184.
- 13 L. Few. The shortest path and the shortest road through n points. *Mathematika*, 2(2):141–144, 1955. doi:10.1112/S0025579300000784.
- 14 Amos Fiat and Gerhard J. Woeginger. On-line scheduling on a single machine: Minimizing the total completion time. *Acta Informatica*, 36(4):287–293, 1999. doi:10.1007/s002360050162.
- 15 Anupam Gupta, Guru Guruganesh, Binghui Peng, and David Wajc. Stochastic online metric matching. In Christel Baier, Ioannis Chatzigiannakis, Paola Flocchini, and Stefano Leonardi, editors, *46th International Colloquium on Automata, Languages, and Programming, ICALP 2019, July 9-12, 2019, Patras, Greece*, volume 132 of *LIPICs*, pages 67:1–67:14. Schloss-Dagstuhl-Leibniz Zentrum für Informatik, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPICs.ICALP.2019.67.
- 16 Anupam Gupta and Kevin Lewi. The online metric matching problem for doubling metrics. In *Automata, Languages, and Programming: 39th International Colloquium, ICALP 2012, Warwick, UK, July 9-13, 2012, Proceedings, Part I 39*, pages 424–435. Springer, 2012. doi:10.1007/978-3-642-31594-7_36.
- 17 Alon Itai, Alan G Konheim, and Michael Rodeh. A sparse table implementation of priority queues. In *International Colloquium on Automata, Languages, and Programming*, pages 417–431. Springer, 1981. doi:10.1007/3-540-10843-2_34.
- 18 Howard J. Karloff. How long can a euclidean traveling salesman tour be? *SIAM J. Discret. Math.*, 2(1):91–99, 1989. doi:10.1137/0402010.
- 19 Donald E. Knuth. Notes on open addressing. Unpublished memorandum. See <http://citeseer.ist.psu.edu/knuth63notes.html>, 1963.
- 20 Donald E. Knuth. *The Art of Computer Programming, Volume 3: (2nd Ed.) Sorting and Searching*. Addison Wesley Longman Publishing Co., Inc., Redwood City, CA, USA, 1998. URL: <https://www.worldcat.org/oclc/312994415>.
- 21 Nicole Megow, Martin Skutella, José Verschae, and Andreas Wiese. The power of recourse for online mst and tsp. In *Automata, Languages, and Programming: 39th International Colloquium, ICALP 2012, Warwick, UK, July 9-13, 2012, Proceedings, Part I 39*, pages 689–700. Springer, 2012. doi:10.1007/978-3-642-31594-7_58.
- 22 Rajeev Motwani and Prabhakar Raghavan. Randomized algorithms. In Mikhail J. Atallah, editor, *Algorithms and Theory of Computation Handbook*, Chapman & Hall/CRC Applied Algorithms and Data Structures series. CRC Press, 1999. doi:10.1201/9781420049503-c16.
- 23 Christos H Papadimitriou. The euclidean travelling salesman problem is np-complete. *Theoretical computer science*, 4(3):237–244, 1977. doi:10.1016/0304-3975(77)90012-3.
- 24 Kirk Pruhs, Jirí Sgall, and Eric Torng. Online scheduling. In Joseph Y.-T. Leung, editor, *Handbook of Scheduling - Algorithms, Models, and Performance Analysis*. Chapman and Hall/CRC, 2004. doi:10.1201/9780203489802.CH15.
- 25 Sharath Raghvendra. Optimal analysis of an online algorithm for the bipartite matching problem on a line. In *34th International Symposium on Computational Geometry (SoCG 2018)*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2018. doi:10.4230/LIPICs.SoCG.2018.67.
- 26 Tim Roughgarden, editor. *Beyond the Worst-Case Analysis of Algorithms*. Cambridge University Press, 2020. doi:10.1017/9781108637435.
- 27 Michael Saks. Online labeling: Algorithms, lower bounds and open questions. In *International Computer Science Symposium in Russia*, pages 23–28. Springer, 2018. doi:10.1007/978-3-319-90530-3_3.