

Space-Constrained Interval Selection

Yuval Emek, Technion - Israel Institute of Technology
 Magnús M. Halldórsson, Reykjavik University, Iceland
 Adi Rosén, CNRS and Université Paris Diderot, France

We study streaming algorithms for the interval selection problem: finding a maximum cardinality subset of disjoint intervals on the line. A deterministic 2-approximation streaming algorithm for this problem is developed, together with an algorithm for the special case of proper intervals, achieving improved approximation ratio of $3/2$. We complement these upper bounds by proving that they are essentially best possible in the streaming setting: it is shown that an approximation ratio of $2 - \epsilon$ (or $3/2 - \epsilon$ for proper intervals) cannot be achieved unless the space is linear in the input size. In passing, we also answer an open question of Adler and Azar (J. Scheduling 2003) regarding the space complexity of constant-competitive randomized preemptive online algorithms for the same problem.

Categories and Subject Descriptors: F.2.0 [Analysis of Algorithms and Complexity]: General

General Terms: Algorithms, Theory

Additional Key Words and Phrases: streaming algorithms, interval selection, online algorithms, lower bounds

ACM Reference Format:

Yuval Emek, Magnús M. Halldórsson, and Adi Rosén, 2016. Space-Constrained Interval Selection. *ACM V*, N, Article A (March 2016), 8 pages.
 DOI: 1.1

1. INTRODUCTION

In this paper we consider the *interval selection* problem, namely, finding a maximum cardinality subset of disjoint intervals from a given collection of intervals on the real line. A celebrated algorithm of Gavril [17] that picks the intervals greedily in increasing order of their right endpoint yields an optimal solution in the classic setting where the complete set of intervals is known to the algorithm. Here we study this problem in the *streaming* model [19; 26], where the input is given to the algorithm as a stream of items (intervals in our case), one at a time, and the algorithm has a limited memory that precludes storing the whole input. Yet, the algorithm is still required to output a feasible solution, with a good approximation ratio.

The motivation for the streaming model stems, inter alia, from applications of managing very large data sets, such as biological data, (cf. [24]), network traffic data, (cf. [28]), and more. Although some function of the whole data set is to be computed, it is impossible to store the whole input. Depending on the setting, different variants of the streaming model have been considered in the literature, such as the classical

A preliminary version of this paper appeared in the proceedings of ICALP 2012 [11].

Research by the second author is partially supported by Iceland Research Foundation grants 120032011 and 152679-051. Research by the third author is partially supported by ANR projects QRAC, ALADDIN and RDAM.

Author's addresses: Y. Emek, Faculty of Industrial Engineering and Management, Technion - Israel Institute of Technology, Technion City, Haifa 3200003, Israel, yemek@ie.technion.ac.il; M. M. Halldórsson, ICE-TCS, School of Computer Science, Reykjavik University, Iceland, mmh@ru.is; Adi Rosén, IRIF - Case 7014, Université Paris Diderot, 75205 Paris Cedex 13, France, adiro@liafa.univ-paris-diderot.fr.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2016 Copyright held by the owner/author(s). 0000-0000/2016/03-ARTA \$15.00

DOI: 1.1

streaming model [19] or the so-called *semi-streaming* model [14]. Common to all of them is the fact that the space used by the streaming algorithm is linear in some natural upper bound on the size of the output it returns (sometimes, a multiplicative polylogarithmic overhead is allowed).

In many problems considered in the streaming literature, the size of the output is fully determined by some parameter of the input, and thus, one would typically express the space complexity as a function of this parameter (cf. [4; 15]). However, in other problems, the size of the output cannot be a priori expressed that way as it depends on the given instance; in such settings it is natural to seek a streaming algorithm whose space complexity is not much larger than the output size of the given instance (cf. [18]). Clearly, as long as the computational model of the streaming algorithm is based on a Turing machine with no distinction between the working tape and the output tape, the size of the output is an inherent lower bound on the required space.

In this paper, we consider a setting where the algorithm is given a stream of real-line intervals, each one defined by its two endpoints, and the goal is to compute a maximum cardinality subset of disjoint intervals (or an approximation thereof). This problem finds many applications, e.g., in resource allocation problems, and it has been extensively studied in the online and offline settings in many variants, see, e.g., [17; 5; 27; 9; 7; 6]. We seek algorithms with a good upper bound on the space they use for a given instance, expressed in terms of the size of the output for that specific instance. Typically, we seek algorithms that use space which is at most linear in the size of the output and yet guarantee a good approximation ratio.

Related Work. The offline interval selection problem corresponds to finding a maximum independent set in an interval graph. An optimal greedy algorithm was discovered early [17] and has since been a staple of algorithms textbooks [10; 20]. It should be noted that the input can be given in (at least) two different ways: as an intersection graph with the nodes corresponding to the intervals, or as a set of intervals given by their endpoints. This distinction makes little difference in the traditional offline setting, where switching between these representations can be done efficiently, but, it can be important in access- or resource-constrained settings. We choose to study the interval selection problem assuming the latter representation — that is, the input is given as a set of intervals — since we believe that it makes more sense in applications related to the online and streaming settings (most previous works on online interval selection make the same choice).

The study of space-constrained algorithms goes back at least to the 1980 work of Munro and Paterson on selection and sorting [25]. More recently, the streaming model was developed to capture the processing of massive data-sets that arise in practice [26]. Most streaming algorithms deal with the approximate computation of various statistics, as exemplified by the celebrated paper of Alon, Matias, and Szegedy [4].

A number of classic graph theoretic problems have been treated in the streaming setting, for example, matching problems [23; 13], diameter and shortest paths [14; 15], min-cut [3], and graph spanners [15]. These were mostly studied under the *semi-streaming* model, introduced by Feigenbaum et al. [14]; in this model, the algorithm is allowed to use $n \log^{O(1)}(n)$ space on an n -vertex graph (i.e., $\log^{O(1)}(n)$ bits per vertex). Closest to our problem, the independent set problem in general sparse graphs (and hypergraphs) was studied in the streaming setting by Halldórsson et al. [18]. Geometric streaming algorithms have also appeared in recent years, especially dealing with extent and ranges, such as [2].

There is a plethora of literature on interval selection in the online setting, where the intervals arrive one-by-one and the algorithm should maintain a valid solution at all times so that an interval can be added to the solution only at its arrival time. Some

papers capture the problem as a call admission problem on a linear network, with the objective of maximizing the number (or weight) of accepted calls. Awerbuch et al. [5] present a strongly $\lceil \log N \rceil$ -competitive algorithm for the problem, where N is the number of nodes on the line (corresponding to the number of possible interval endpoints). This yields an $O(\log \Delta)$ -competitive algorithm for the weighted case, where Δ is the ratio between the longest to the shortest interval. On the negative side, they establish a lower bound of $\Omega(\log N)$ on the competitive ratio of randomized non-preemptive online interval selection algorithms.¹ In the context of the real line, this immediately implies that such algorithms cannot have a competitive ratio that is independent of the length of the input. In fact, Bachmann et al. [7] showed that the competitive ratio of randomized non-preemptive online algorithms for interval selection on the real line must be linear in the number of intervals in the input. For the weighted case, preemptive online scheduling has a lower bound of $\Omega(\log \kappa)$, where κ is the minimum of the maximum interval length, the maximum interval benefit, and the maximum interval value-density (i.e., the interval benefit divided by its length) [6]. In comparison, much better results are possible for preemptive online algorithms in the unweighted setting: Adler and Azar [1] devise a 16-competitive algorithm. One way of easing the task of the algorithm in the weighted case is to assume arrival by time, i.e., the intervals arrive in order of left endpoints. This has been treated for different weighted problems [27; 22; 12; 16] (but the lower bound of [6] still holds).

Subsequent to the initial publication of the present results [11], Cabello and Pérez-Lanterno [8] presented streaming algorithms that estimate the size of the maximum independent set out of a set of intervals. Their algorithms give for a general instance a $(2 + \epsilon)$ approximation, and for unit intervals a $(3/2 + \epsilon)$ approximation, using space polynomial in $\frac{1}{\epsilon}$ and in $\log n$. They also presented new, simpler than ours, algorithms for finding the approximated independent set, which in some cases match our bounds as to the approximation ratio and the space used.

Our results. We give tight results for the interval selection problem in the streaming setting. Our main positive result is a deterministic 2-approximation streaming algorithm that uses space linear in the size of the output (Sec. 3). This is complemented by a matching lower bound (Sec. 5), stating that an approximation ratio of $2 - \epsilon$ cannot be obtained by any randomized streaming algorithm with space significantly smaller than the size of the input (which can be much larger than the size of the output). The special case of proper interval collections (i.e., collections of intervals with no proper containments) is also considered, for which a deterministic 3/2-approximation streaming algorithm that uses space linear in the output size is presented (Sec. 4) and a matching lower bound on the approximation ratio is established (Sec. 5) for streams of unit intervals (a special case of proper intervals). The upper bounds are extended to *multiple-pass* streaming algorithms: we show that an approximation ratio $1 + 1/(2p - 1)$ can be obtained in p passes over the input (Sec. 6).

In passing, we also answer an open question posed by Adler and Azar [1] in the context of randomized preemptive online algorithms for the interval selection problem. Adler and Azar point out that the decisions made by their online algorithm depend on the whole history (i.e., the input seen so far) and that natural attempts to remove this dependency seem to fail. Consequently, they write (using the term “active call” for an interval in the solution maintained by the online algorithm) that “*it seems very interesting to find out whether there exist constant-competitive algorithms where each*

¹A *preemptive* online interval selection algorithm is allowed to withdraw intervals that were previously picked to be included in the solution. In contrast, if a *non-preemptive* online interval selection algorithm adds an interval to the solution, then this interval will remain in the solution indefinitely.

decision depends only on the currently active calls and maybe on additional bounded information". We answer this question in the affirmative by slightly modifying our main algorithm to achieve a randomized preemptive online algorithm that admits constant competitive ratio and uses space linear in the size of the optimal solution, rather than the size of the input, as the algorithm of Adler and Azar does (Sec. 7).²

2. PRELIMINARIES

We think of the real line \mathbb{R} as stretching from left to right so that an *interval* I contains all points between its left *endpoint* $\text{left}(I)$ and its right endpoint $\text{right}(I)$, where $\text{left}(I) < \text{right}(I)$. Each endpoint can be either *open* (exclusive) or *closed* (inclusive). A *half-open* interval has a closed left endpoint and an open right endpoint. (This is, perhaps, the natural interval type to use in most resource allocation applications.) Observe that the assumption that $\text{left}(I) < \text{right}(I)$ implies that every interval contains an open set (in the topological sense) and that half-open intervals are always well defined.

The interval related notions of *intersection*, *disjointness*, and *containment* follow the standard view of an interval as a set of points. Throughout, the symbols \subset and \supset denote strict containment, whereas the symbols \subseteq and \supseteq denote non-strict containment. Two intervals I, J *properly* intersect if they intersect without containment; I properly contains J if I contains J and J does not contain I . An interval collection \mathcal{I} is said to be *proper* (and the intervals in the collection, *proper* intervals) if no two intervals in \mathcal{I} exhibit proper containment. The *load* of \mathcal{I} is defined to be $\max_{p \in \mathbb{R}} |\{I \in \mathcal{I} \mid p \in I\}|$.

The *interval selection* problem asks for a maximum cardinality subset of pairwise disjoint intervals out of a given set S of intervals. In the streaming model, the input interval set S is considered to be an ordered set (a.k.a. a *stream*) and the intervals arrive one by one according to that order. The intervals are specified by their endpoints, where each endpoint is represented by a bit string of length b (the same b for all endpoints). This may potentially provide a streaming algorithm with the edge of knowing in advance some bounds on the number of intervals that will arrive and on the number of intervals that can be placed between two already-received intervals. However, our algorithms do not take advantage of this extra information and our lower bounds show that it is essentially useless. An optimal solution to a given instance S of the interval selection problem is denoted by $\text{Opt}(S)$.

We may sometimes talk about *segments*, rather than intervals, when we want to emphasize that the entities under consideration are not necessarily part of the input. Given a set \mathcal{I} of intervals, a maximal continuous segment in $\bigcup_{I \in \mathcal{I}} I$ is called a *connected component* (or, for simplicity, *component*).

3. THE MAIN ALGORITHM

3.1. Overview

Given a stream S of intervals, our algorithm maintains a (proper interval) collection $A \subseteq S$, referred to as the *actual* intervals, from which the output $\text{Alg}(S) = \text{Opt}(A)$ is taken. It also maintains a collection V of *virtual* intervals, where each virtual interval is the intersection of two actual intervals that existed in A at some point. The role of the virtual intervals is to prevent undesired intervals from joining A : an arriving interval $I \in S$ joins A if and only if it does not contain any currently maintained virtual or actual interval.

²The technique employed in Sec. 7 is based on a “classify and randomly select” argument that guarantees that the solution produced by the online algorithm is a constant approximation of the optimal solution with constant probability. Using the technique of [21] (reformulated as Theorem 4.1 in [1]), this can be strengthened to guarantee a constant approximation with high probability.

Our algorithm is designed to guarantee that each interval $I \in S$ leaves a *trace* in either A or V , namely, there exists some $J \in A \cup V$ such that $J \subseteq I$. Moreover, if $I, I' \in A$ properly intersect, then $I \cap I' \in V$. This essentially means that an arriving interval is rejected if and only if it contains some previous interval of S or the intersection of two properly intersecting previous intervals in S that have belonged to A .

Following that, it is not too difficult to show that the load of the interval collection A is at most 2. Based on a careful analysis of the structure of the (connected) components in A and the locations of the virtual intervals within these components and between them, we can argue that $|V| \leq |A|$. This immediately yields the desired upper bound on the space of our algorithm as $|A| \leq 2 \cdot |\text{Opt}(A)|$. The bound on the approximation ratio essentially stems from the observation that $|\text{Opt}(S)| \leq |\text{Opt}(A \cup V)|$ (a direct corollary of the fact that each interval in S leaves a trace in $A \cup V$) and from the invariant that each actual interval contains at most 2 virtual intervals.

It is interesting to point out that our algorithm is in fact a deterministic preemptive online algorithm that maintains a load-2 interval collection (the collection A). Since the main result of Adler and Azar [1] also relies on such an algorithm, one may wonder if the two algorithms can be compared. Actually, the algorithm of Adler and Azar bases its rejection (and preemption) decisions on similar conditions: an arriving interval is rejected if and only if it contains some previous interval of S or the intersection of two properly intersecting intervals in A . (Adler and Azar use a different terminology, but the essence is very similar.) The difference lies in the latter condition: Whereas the algorithm of Adler and Azar considers only the properly intersecting intervals that are currently in A , our algorithm also (implicitly) considers properly intersecting intervals that belonged to A in the past and were preempted since. This seemingly small difference turns out to be crucial as it allows our algorithm to use much less memory, thus giving rise to an interesting phenomena: by remembering extra information (i.e., intersecting intervals that belonged to A in the past and are not in A anymore), we actually end up using less memory.

3.2. The algorithm

Consider a stream $S = (I_1, \dots, I_n)$ of intervals on the real line. It will be convenient to assume that all endpoints are distinct, i.e., $\{\text{left}(I), \text{right}(I)\} \cap \{\text{left}(J), \text{right}(J)\} = \emptyset$ for every two intervals $I, J \in S$. Unless stated otherwise, we will also assume that the intervals mentioned in this section are closed on both endpoints. These two assumptions are lifted in Appendix A.

Our algorithm, denoted `Alg`, maintains a collection $A \subseteq S$ of *actual* intervals and a collection V of *virtual* intervals, where each virtual interval is realized by endpoints of intervals in S . That is, any virtual interval $I \in V$ satisfies $\{\text{left}(I), \text{right}(I)\} \subseteq \{\text{left}(J), \text{right}(J) \mid J \in S\}$. The algorithm initially sets $A, V \leftarrow \emptyset$. Then, upon arrival of a new interval $I \in S$, `Alg` proceeds according to the policy³ presented in Algorithm 1.

The algorithm first verifies that the new interval I does not contain any currently stored (actual or virtual) interval; if it does, then the new interval is ignored (rejected). Therefore, if `Alg` reaches line 3, then we can assume that $I \not\supseteq J$ for any interval $J \in A \cup V$. Next, in lines 4–7 `Alg` removes all the actual and virtual intervals that contain I . Lines 8–12 form the heart of the algorithm: updating the virtual intervals that remain in V . The idea here is that a virtual interval that intersects with I is “trimmed” until it is contained in I ; if an actual interval intersects with I , then the intersection is introduced as a new virtual interval. Finally, any actual interval J that exclusively

³Note that `Alg` can be thought of as an online algorithm with preemption with respect to the set A .

Algorithm 1 The policy of Alg upon arrival of an interval $I \in S$.

```

1: if  $\exists J \in A \cup V$  s.t.  $J \subseteq I$  then
2:   reject  $I$  and halt
3:  $A \leftarrow A \cup \{I\}$ 
4: for all  $J \in A - \{I\}$  s.t.  $J \supseteq I$  do
5:    $A \leftarrow A - \{J\}$ 
6: for all  $J \in V$  s.t.  $J \supseteq I$  do
7:    $V \leftarrow V - \{J\}$ 
8: for  $p \in \{\text{left}(I), \text{right}(I)\}$  do
9:   if  $\exists J \in V$  s.t.  $p \in J$  then
10:     $V \leftarrow V - \{J\} \cup \{I \cap J\}$ 
11:   else if  $\exists J \in A$  s.t.  $p \in J$  then
12:     $V \leftarrow V \cup \{I \cap J\}$ 
13: for all  $J \in A$  and  $K \in V$  do
14:   if  $\text{left}(J) < \text{left}(K) < \text{right}(K) < \text{right}(J)$  then
15:     $A \leftarrow A - \{J\}$ 

```

contains some virtual interval K (that is, J contains K even if we remove J 's endpoints) is removed from the actual interval collection A in lines 13–15.

After the last interval I_n is processed, Alg outputs $\text{Alg}(S) = \text{Opt}(A)$, that is, an optimal subset of the interval collection A (computed, say, by the greedy left-to-right algorithm). In the remainder of this section we prove that: (a) at all times, $|V| \leq |A| \leq 2 \cdot |\text{Alg}(S)|$; and (b) $|\text{Alg}(S)| \geq |\text{Opt}(S)|/2$. Together, we obtain the desired approximation, using space at most constant times larger than the size of the optimal output. The figures in Appendix C may provide helpful illustrations for the effect that the new interval I has on the sets A and V .

3.3. Analysis

Throughout the analysis, we let $1 \leq t \leq n$ denote the time at which Alg completed processing interval $I_t \in S$; time $t = 0$ denotes the beginning of the execution. We refer to the period between time $t - 1$ and time t as *round* t . The stream prefix (I_1, \dots, I_t) is denoted by S_t . The collections A and V at time t are denoted by A_t and V_t , respectively, although, when t is clear from the context, we may omit the subscript. We begin by showing that each virtual interval is indeed realized by (at most) two actual intervals and that the new interval I is not removed immediately after joining A .

PROPOSITION 3.1. *At any time t , if $\rho \in V_t$, then there exist some $\sigma, \sigma' \in S_t$ such that $\rho = \sigma \cap \sigma'$.*

PROOF. By induction on t . The case $t = 0$ is trivial as $V_0 = \emptyset$. For time $t > 0$, we observe that any new virtual interval ρ added to V in round t is either the intersection of two actual intervals (line 12) or the intersection of an actual interval and a virtual interval in V_{t-1} (line 10). In the former case, the assertion follows immediately; in the latter case, the assertion follows by the inductive hypothesis. \square

PROPOSITION 3.2. *For every $1 \leq t \leq n$, if Alg reaches line 3 when processing $I_t = I$, then $I \in A_t$.*

PROOF. In line 3, I is added to A and subsequently, it can only be removed from A if a virtual interval ρ that is contained in I but does not have an endpoint in common with V is found (line 15). Such an interval ρ cannot be in V_{t-1} since otherwise, I would have been rejected in line 2. The assertion follows since every virtual interval added to V in round t has a common endpoint with I . \square

Lemma 3.3 lies at the core of our analysis: it states that each interval in S leaves some trace in either A or V . This will be employed later on to argue that $\text{Alg}(S)$ is not much smaller than $\text{Opt}(S)$.

LEMMA 3.3. *For every interval $I_t \in S$ and for every time $t' \geq t$, there exists some interval $\rho \in A_{t'} \cup V_{t'}$ such that $\rho \subseteq I_t$.*

PROOF. A new coming interval I is added to A in line 3 unless some interval $\rho \subseteq I$ is found in $A \cup V$. An actual interval $\rho \in A$ is removed from A only if another actual interval $I \subseteq \rho$ has just joined A (line 5) or if a virtual interval $\sigma \subset \rho$ is found in V (line 15). A virtual interval $\rho \in V$ is removed from V only if an actual interval $I \subseteq \rho$ has just joined A (line 7) or if it is replaced in V by another virtual interval $\sigma \subseteq \rho$ (line 10). The assertion follows. \square

3.3.1. The structural lemma. We now turn to establish our main lemma regarding the updating phase in lines 8–12 and the resulting structure of the interval collections A and V . Lemma 3.4 states seven invariants maintained by our algorithm; these invariants are then proved simultaneously by induction on t , essentially by straightforward analysis of the policy presented in Algorithm 1.

LEMMA 3.4. *For any round $1 \leq t \leq n$, the updating phase satisfies the following two properties:*

(P1) *If ρ is added to V in round t , then $\rho \in V_t$.*

(P2) *If ρ and σ are added to V in round t , then $\rho \cap \sigma = \emptyset$.*

Moreover, for any time $0 \leq t \leq n$, the interval collections A and V satisfy the following five properties:

(P3) *For every $\rho \in A$ and $\sigma \in V$, if $\rho \cap \sigma \neq \emptyset$, then $\sigma \subset \rho$ with a common endpoint.*

(P4) *For every $\rho, \sigma \in A$, if $\rho \cap \sigma \neq \emptyset$, then $\rho \cap \sigma \in V$.*

(P5) *Every point $p \in \mathbb{R}$ is contained in at most 1 virtual interval.*

(P6) *Every point $p \in \mathbb{R}$ is contained in at most 2 actual intervals.*

(P7) *There do not exist two actual intervals $\rho, \sigma \in A$ such that $\rho \subseteq \sigma$.*

PROOF. We first establish (P1) regardless of the other six properties.

Establishing (P1). It is sufficient to show that if ρ is added to V in line 10 or line 12 of the execution for $p = \text{left}(I)$, then it is not removed from V in line 10 of the execution for $p = \text{right}(I)$. Indeed, if ρ is added to V in the execution for $p = \text{left}(I)$, then $\rho = I \cap \sigma$ for some interval $\sigma \in A_{t-1} \cup V_{t-1}$ such that $\text{left}(I) \in \sigma$. Since σ cannot contain I (as otherwise, it would have been removed in line 5 or line 7), it follows that $\text{left}(\sigma) < \text{left}(I) < \text{right}(\sigma) < \text{right}(I)$, so $\rho = [\text{left}(I), \text{right}(\sigma)]$. Therefore, $\text{right}(I) \notin \rho$ and ρ is not removed from V in line 10 of the execution for $p = \text{right}(I)$.

Next, we establish (P2), (P3), (P4), and (P5) simultaneously by induction on t . The case $t = 0$ is trivial: (P2) holds vacuously, while (P3), (P4), and (P5) hold as $A_0 = V_0 = \emptyset$. Assume that the four properties hold for $t - 1$ and consider the execution of Alg upon arrival of interval $I = I_t$ for some $1 \leq t \leq n$.

Establishing (P2). As each iteration of the for loop in lines 8–12 adds at most one virtual interval to V , we may assume that ρ is added in the execution for $p = \text{left}(I)$ and σ is added in the execution for $p = \text{right}(I)$. This means that $\rho = I \cap \tau_\ell$ and $\sigma = I \cap \tau_r$ for some intervals $\tau_\ell, \tau_r \in A_{t-1} \cup V_{t-1}$ such that $\text{left}(I) \in \tau_\ell$ and $\text{right}(I) \in \tau_r$. We argue that τ_ℓ and τ_r do not intersect, which implies that ρ and σ do not intersect.

To that end, assume by way of contradiction that they do and let $\tau_\cap = \tau_\ell \cap \tau_r$. If both τ_ℓ and τ_r are virtual intervals, then we immediately reach a contradiction due the inductive hypothesis on (P5). If both τ_ℓ and τ_r are actual intervals, which means that ρ and σ are added to V in line 12, then by the inductive hypothesis on (P4), $\tau_\cap \in V_{t-1}$. By definition, τ_\cap must intersect with I . On the other hand, neither $\text{left}(I)$ nor $\text{right}(I)$

can belong to τ_\cap as otherwise, the else condition in line 11 would not have passed, thus $\tau_\cap \subset I$. But this means that Alg should not have reached line 3 and in particular, ρ and σ would not have been added to V .

So, assume that τ_ℓ is actual and τ_r is virtual (the proof of the converse possibility is identical). By the inductive hypothesis on (P3), we know that $\tau_r \subset \tau_\ell$. But this implies that both endpoints of I belong to τ_ℓ , namely, $I \subseteq \tau_\ell$, and τ_ℓ should have been removed from A in line 5.

Establishing (P3). Consider some $\rho \in A_t$ and $\sigma \in V_t$ such that $\rho \cap \sigma \neq \emptyset$. If $\rho \in A_{t-1}$ and $\sigma \in V_{t-1}$, then the property holds by the inductive hypothesis. Assume first that ρ is added to A in round t , so ρ is the last arriving interval I . Notice that σ cannot be in V_{t-1} as this implies that either (i) $\sigma \subseteq I$, in which case I would have been rejected in line 2; (ii) $\sigma \supseteq I$, in which case σ would have been removed from V in line 7; or (iii) σ and I properly intersect, in which case σ is removed from V in line 10. Thus, σ is added to V in round t either in line 10 or in line 12. In both cases, σ is contained in I with a common endpoint.

It remains to consider the case in which $\rho \in A_{t-1}$ and σ is added to V in round t . If σ is added to V in line 10, then it replaces in V some interval $\tau \in V_{t-1}$ such that $\sigma \subseteq \tau$. Hence, τ must also intersect with ρ and by the inductive hypothesis, $\tau \subset \rho$, so σ must be contained in ρ . Since ρ is not removed in line 15, ρ and σ must have a common endpoint. If σ is added to V in line 12, then $\sigma = I \cap \tau$ for some interval $\tau \in A_{t-1}$ such that an endpoint p of I is contained in τ . The property is established by arguing that τ and ρ must be the same interval.

To that end, suppose toward a contradiction that $\tau \neq \rho$. Assume without loss of generality that $p = \text{left}(I)$, so $\text{left}(\tau) < \text{left}(I) < \text{right}(\tau) < \text{right}(I)$. Since $\sigma = I \cap \tau = [\text{left}(I), \text{right}(\tau)]$ intersects with ρ , both I and τ must also intersect with ρ . By the inductive hypothesis on (P4), we know that $\sigma_\cap = \rho \cap \tau \in V_{t-1}$. We also know that σ_\cap intersects with I as both ρ and τ intersect with I . Since $\tau \not\subseteq I$, it follows that $\sigma_\cap \not\subseteq I$, hence σ_\cap must still be in V when Alg reaches line 8. If $\text{left}(I) \in \sigma_\cap$, then the else condition in line 11 would not have passed and σ would not have been added to V in line 12, so $\text{left}(I) \notin \sigma_\cap$. But $\text{right}(I) \notin \sigma_\cap$ as $\text{right}(I) \notin \tau$, hence $\sigma_\cap \subseteq I$ and I should have been rejected in line 2. In any case, we conclude that ρ and τ are indeed the same interval.

Establishing (P4). Consider two intersecting intervals $\rho, \sigma \in A_t$. If both ρ and σ are also in A_{t-1} , then by the inductive hypothesis, $\tau = \rho \cap \sigma \in V_{t-1}$. If $\tau \notin V_t$, then it must have been removed from V either in line 7 because $I \subseteq \tau$, in which case I is also contained in both ρ and σ and they would have been removed from A in line 5, or in line 10, where it is replaced in V by some other virtual interval $\tau' \subset \tau$ (the strict containment follows from the distinct endpoints assumption), in which case at least one of the intervals ρ and σ should have been removed in line 15. Therefore, $\tau \in V_t$ and the property holds in that case.

So, suppose that $\rho \in A_{t-1}$, while $\sigma = I$ is added to A in round t . Since $\rho, I \in A_t$, both ρ and I are in A when Alg reaches line 8, thus they cannot contain each other. Assume without loss of generality that $\text{left}(\rho) < \text{left}(I) < \text{right}(\rho) < \text{right}(I)$. If $\text{left}(I)$ does not belong to any virtual interval in V_{t-1} , then in line 12 the virtual interval $\tau = \rho \cap I$ is added to V and it must still be there at time t due to (P1). So, assume that $\text{left}(I)$ belongs to some virtual interval $\tau \in V_{t-1}$. Since τ intersects with ρ , the inductive hypothesis on (P3) implies that $\tau \subset \rho$ with a common endpoint. In line 10, τ is replaced in V by the new virtual interval $\tau' = \tau \cap I$, which, by (P1) remains in V at time t . The interval τ' intersects with both ρ and I , hence, by (P3) (applied to time t), it is contained in both of them, having a common endpoint with each, thus $\tau' = \rho \cap \sigma$ and the property holds.




Fig. 1. A component C of A . The solid lines depict the actual intervals ρ_i , $i = 1, \dots, 5$; the dashed lines depict the virtual intervals that intersect with C (all of them must be contained in C).

Establishing (P5). Suppose toward a contradiction that there exists two distinct intervals $\rho, \sigma \in V_t$ such that $\rho \cap \sigma \neq \emptyset$. Assume without loss of generality that σ was added to V after ρ . By the inductive hypothesis, σ is added to V in round t , while (P2) guarantees that $\rho \in V_{t-1}$. If σ is added to V in line 10, then $\sigma = I \cap \tau$ for some virtual interval τ which is guaranteed to be in V_{t-1} by (P2). But then the inductive hypothesis implies that $\rho \cap \tau = \emptyset$, thus $\rho \cap \sigma = \emptyset$.

So, assume that σ is added to V in line 12. In that case $\sigma = I \cap \tau$ for some $\tau \in A_{t-1}$. Assume without loss of generality that $\text{left}(\tau) < \text{left}(I) < \text{right}(\tau) < \text{right}(I)$ so that $\sigma = [\text{left}(I), \text{right}(\tau)]$ is added to V for $p = \text{left}(I)$. Since ρ intersects with σ , it must also intersect with both I and τ . We know that p cannot belong to ρ as otherwise, the else condition in line 11 would not have passed. But, by the inductive hypothesis on (P3), $\rho \subset \tau$, thus $\rho \subseteq I$ and I should have been rejected in line 2.

Properties (P6) and (P7) can now be established based on the other properties.

Establishing (P6). Consider some point $p \in \mathbb{R}$ and suppose toward a contradiction that there exist three distinct intervals $\rho_1, \rho_2, \rho_3 \in A_t$ such that $p \in \rho_i$ for every $1 \leq i \leq 3$. By (P4), the intersections $\sigma_{1,2} = \rho_1 \cap \rho_2$, $\sigma_{1,3} = \rho_1 \cap \rho_3$, and $\sigma_{2,3} = \rho_2 \cap \rho_3$ are all in V_t . But (P5) implies that $\sigma_{1,2}$, $\sigma_{1,3}$, and $\sigma_{2,3}$ are pairwise disjoint, in contradiction to their definition.

Establishing (P7). Consider any two intervals $\rho, \sigma \in A_t$. If $\rho \cap \sigma \neq \emptyset$, then (P4) implies that $\rho \cap \sigma \in V_t$. By (P3), $\rho \cap \sigma$ is strictly contained in both ρ and σ , hence ρ cannot be a subset of σ (nor can σ be a subset of ρ). \square

3.3.2. The components. We employ Lemma 3.4 in order to understand the structure of the components of A and their relations with the intervals in V . To that end, fix some time t and consider an arbitrary component C formed as the union of the actual intervals $\rho_1, \dots, \rho_k \in A_t$. We denote the leftmost and rightmost points in (the segment) C by $\text{left}(C)$ and $\text{right}(C)$, respectively.

We argue that the actual and virtual intervals that intersect with component C are restricted to a certain structure (refer to Figure 1 for an illustration). To that end, assume without loss of generality that $\text{left}(\rho_i) < \text{left}(\rho_{i+1})$ for every $1 \leq i \leq k-1$. Lemma 3.4(P6) and (P7) then guarantee that

$$\text{left}(\rho_{i-1}) < \text{left}(\rho_i) < \text{right}(\rho_{i-1}) < \text{left}(\rho_{i+1}) < \text{right}(\rho_i) < \text{right}(\rho_{i+1})$$

for every $2 \leq i \leq k-1$. By Lemma 3.4(P4), we conclude that $\rho_i \cap \rho_{i+1} \in V_t$ for every $1 \leq i \leq k-1$, while Lemma 3.4(P3) implies that the segment $[\text{left}(\rho_2), \text{right}(\rho_{k-1})]$ does not intersect with any other virtual interval in V_t . The segment C possibly contains two more virtual intervals at time t : an interval $\sigma_\ell \subseteq [\text{left}(\rho_1), \text{left}(\rho_2))$ and an interval $\sigma_r \subseteq (\text{right}(\rho_{k-1}), \text{right}(\rho_k)]$, but then Lemma 3.4(P3) guarantees that $\text{left}(\sigma_\ell) = \text{left}(\rho_1) = \text{left}(C)$ and $\text{right}(\sigma_r) = \text{right}(\rho_k) = \text{right}(C)$. There may also exist virtual intervals in between the components of A , but we will soon show that their number and structure are fairly limited. This requires the definition of the following notions.

Partition into portions. Point $p \in \mathbb{R}$ is said to be an $\langle x, y \rangle$ -point at time t if there exist exactly x actual intervals in A_t and y virtual intervals in V_t that

contain p . By Lemma 3.4, it suffices to consider $\langle x, y \rangle$ -points for $\langle x, y \rangle$ pairs in the set $\mathcal{F} = \{\langle 0, 0 \rangle, \langle 1, 0 \rangle, \langle 0, 1 \rangle, \langle 1, 1 \rangle, \langle 2, 1 \rangle\}$, referred to as *feasible pairs*. Fixing some feasible pair $\langle x, y \rangle \in \mathcal{F}$, a maximal connected set of $\langle x, y \rangle$ -points is referred to as an $\langle x, y \rangle$ -*portion*. Given some segment $G \subseteq \mathbb{R}$, let $\varphi(G)$ be the string of feasible pairs that encodes the types of portions encountered in a left-to-right scan of G ; e.g., if C is the component illustrated in Figure 1, then $\varphi(C) = \langle 1, 0 \rangle \langle 2, 1 \rangle \langle 1, 0 \rangle \langle 2, 1 \rangle \langle 1, 0 \rangle \langle 2, 1 \rangle \langle 1, 0 \rangle \langle 2, 1 \rangle \langle 1, 0 \rangle \langle 1, 1 \rangle$.

More generally, it follows from Lemma 3.4 and the discussion above that for every component C of A it holds that

$$\varphi(C) \in \langle 1, 1 \rangle^? \langle 1, 0 \rangle (\langle 2, 1 \rangle \langle 1, 0 \rangle)^* \langle 1, 1 \rangle^?,$$

where following the common notation in regular expressions, we use a superscript question mark to denote 0 or 1 occurrences and superscript asterisk to denote zero or more occurrences (cf. the Kleene star). Notice that by the distinct endpoints assumption, given $\varphi(C)$, it is easy to determine (uniquely) the numbers of actual and virtual intervals that C contains and the relative (total) order of their endpoints. Thus the string $\varphi(C)$ gives the so-called *topological structure* of component C .

Likewise, if P is a segment between two adjacent components of A , or the segment to the left of the leftmost component of A or the segment to the right of the rightmost component of A , then it holds that

$$\varphi(P) \in \langle 0, 0 \rangle (\langle 0, 1 \rangle \langle 0, 0 \rangle)^*.$$

We refer to the $\langle 0, 1 \rangle$ -portions as the *isolated* virtual intervals. The following lemma imposes a crucial restriction on the topological structure of $\varphi(P)$ (and also on $\varphi(C)$).

LEMMA 3.5. *Let $\varphi_t = \varphi(\mathbb{R})$ be the string of feasible pairs that encodes the types of portions encountered in a left-to-right scan of the real line at time $1 \leq t \leq n$. Then, every $\langle 0, 0 \rangle$ entry in φ_t is immediately preceded or immediately followed by a $\langle 1, 0 \rangle$ entry.*

PROOF. By induction on t . The assertion holds trivially for $t = 1$ as $\varphi_1 = \langle 0, 0 \rangle \langle 1, 0 \rangle \langle 0, 0 \rangle$. Suppose that the assertion holds at time $t - 1$ and consider time t . If Alg does not reach line 3 when processing $I = I_t$, then $\varphi_t = \varphi_{t-1}$ and the assertion holds by the inductive hypothesis, so assume hereafter that Alg does reach line 3 when processing I which means that $I \not\subseteq J$ for any $J \in A_{t-1} \cup V_{t-1}$.

The proof continues by case analysis that considers the different types of possible intersections between I and the portions corresponding to φ_{t-1} . Specifically, in Table I we identify a total of 27 types of intersections (up to symmetry) and address each one of them by depicting the string φ_{t-1} together with a solid line stretched above the entries corresponding to the portions with which interval I intersects. The resulting string φ_t is given on the third column. The strings $\alpha, \beta \in \mathcal{F}^*$ denote the substrings of φ_{t-1} encoding the portions in segments of \mathbb{R} which are either not affected by the newcoming interval I or affected in a limited way. We also use the notation

$$\alpha' = \begin{cases} \gamma, & \alpha = \gamma \langle 0, 0 \rangle \\ \gamma \langle x - 1, y \rangle, & \alpha = \gamma \langle x, y \rangle, x > 0 \end{cases} \quad \beta' = \begin{cases} \gamma, & \beta = \langle 0, 0 \rangle \gamma \\ \langle x - 1, y \rangle \gamma, & \beta = \langle x, y \rangle \gamma, x > 0 \end{cases}$$

which facilitates a reduction in the number of cases we have to consider. Figures illustrating each of these 27 cases are provided in Appendix C (the corresponding figure numbers are indicated in the fourth column of Table I).

Observe that I can intersect at most 2 components of $A \cup V$, since otherwise there exists some $J \in A \cup V$ such that $J \subseteq I$ and I would not reach line 3 of the algorithm. Therefore, to cover all cases of the different types of possible intersections between I and the portions corresponding to φ_{t-1} . Table I is divided into the following 5 main categories:

- C1: I intersects with 0 components of A and with 0 isolated virtual intervals.
- C2: I intersects with 0 components of A and with 1 isolated virtual interval.
- C3: I intersects with 1 component of A and with 0 isolated virtual intervals.
- C4: I intersects with 1 component of A and with 1 isolated virtual interval.
- C5: I intersects with 2 components of A and with 0 isolated virtual intervals.

For convenience, category C3 is further divided into two parts according to the relation of I with the component C of A with which it intersects: the top part includes the cases where $I - C \neq \emptyset$; the bottom part includes the cases where $I - C = \emptyset$.

The assertion follows since in all cases, every $\langle 0, 0 \rangle$ entry in φ_t is immediately preceded or immediately followed by a $\langle 1, 0 \rangle$ entry (taking into consideration the constraints imposed on α and β due to the inductive hypothesis). \square

COROLLARY 3.6. *Let C_1, \dots, C_m be the components of A_t , $0 \leq t \leq n$, indexed from left to right. For every $1 \leq i \leq m - 1$, at most one of the following three events occur:*

- (a) $\varphi(C_i) \in \langle 1, 1 \rangle^? \langle 1, 0 \rangle \langle 2, 1 \rangle \langle 1, 0 \rangle^* \langle 1, 1 \rangle$;
- (b) $\varphi(C_{i+1}) \in \langle 1, 1 \rangle \langle 1, 0 \rangle \langle 2, 1 \rangle \langle 1, 0 \rangle^* \langle 1, 1 \rangle^?$; or
- (c) *the segment P between C_i and C_{i+1} satisfies $\varphi(P) = \langle 0, 0 \rangle \langle 0, 1 \rangle \langle 0, 0 \rangle$ (i.e., there is a single isolated virtual interval between C_i and C_{i+1}).*

Moreover, there is no isolated virtual interval to the left of C_1 or to the right of C_m .

3.3.3. Accounting. We are now ready to establish the following lemma.

LEMMA 3.7. $|\text{Alg}(S_t)| \geq |\text{Opt}(S_t)|/2$ at every time $0 \leq t \leq n$.

PROOF. Lemma 3.3 guarantees that $|\text{Opt}(S_t)| \leq |\text{Opt}(A_t \cup V_t)|$. As $|\text{Alg}(S_t)| = |\text{Opt}(A_t)|$, it is sufficient to bound the ratio $R = \frac{|\text{Opt}(A_t \cup V_t)|}{|\text{Opt}(A_t)|}$, showing that it is at most 2.

Let C_1, \dots, C_m be the components of A_t . Examining the three cases of Corollary 3.6, we observe that the ratio R is maximized if event (c) always occurs, that is, $C_i \in \langle 1, 0 \rangle \langle 2, 1 \rangle \langle 1, 0 \rangle^*$ for every $1 \leq i \leq m$ and there is an isolated virtual interval between C_i and C_{i+1} for every $1 \leq i \leq m - 1$. We will increase R even further by assuming that there exists an isolated virtual interval to the right of C_m .

Consider some component $C = C_i$, $1 \leq i \leq m$, and let $A(C) = \{\rho \in A_t \mid \rho \subseteq C\}$ and $V(C) = \{\sigma \in V_t \mid \sigma \subseteq C\}$. It is easy to verify that $|V(C)| = |A(C)| - 1$ and that $|\text{Opt}(A(C))| = \lceil |A(C)|/2 \rceil$, whereas

$$|\text{Opt}(A(C) \cup V(C))| = \begin{cases} |V(C)| & \text{if } V(C) \neq \emptyset \\ 1 & \text{otherwise.} \end{cases}$$

Accounting for the isolated virtual interval to the right of C_i , we conclude that each component C_i , $1 \leq i \leq m$, contributes:

- (i) 1 to the denominator of R and 2 to the numerator of R if $V(C_i) = \emptyset$; and
- (ii) $\lceil |A(C_i)|/2 \rceil$ to the denominator of R and $|A(C_i)| - 1 + 1 = |A(C_i)|$ to the numerator of R if $V(C_i) \neq \emptyset$.

The assertion follows. \square

COROLLARY 3.8. $|\text{Alg}(S)| \geq |\text{Opt}(S)|/2$.

It remains to bound from above the amount of space used by our algorithm, showing that it is linear in the length of the bit string representing $\text{Alg}(S)$. At each time t , the space of Alg is linear in the length of the bit strings representing A_t and V_t . As $|\text{Opt}(S_t)|/2 \leq |\text{Alg}(S_t)| \leq |\text{Opt}(S_t)|$ for every $0 \leq t \leq n$, and since $|\text{Opt}(S_t)|$ is non-decreasing with t , it is sufficient to show that $|A_t| + |V_t| = O(|\text{Alg}(S_t)|) = O(|\text{Opt}(A_t)|)$.

By Lemma 3.4(P6), we know that the actual intervals in A_t can be colored in two colors such that if two intervals belong to the same color class, then they do not intersect. Thus, $|A_t| \leq 2 \cdot |\text{Opt}(A_t)|$ at every time t . On the other hand, Corollary 3.6 implies that if we scan the real line from left to right counting the actual and virtual intervals, then the number of virtual intervals never exceeds that of the actual intervals. Therefore, $|V_t| \leq |A_t|$ which establishes the following corollary.

COROLLARY 3.9. *At every time t , the space of Alg is linear in the length of the bit string representing Alg(S).*

4. PROPER INTERVALS

In this section we consider the interval selection problem for proper intervals. There is an easy deterministic 2-approximate streaming (and online) algorithm that uses no extra space in addition to storing the output: simply greedily add an interval whenever possible. We give here a streaming algorithm with an improved approximation ratio of $3/2$ that uses output-linear space. As we show in Sec. 5, that is optimal in the semi-streaming setting.

We first give an informal overview of the operation of the algorithm. It maintains a collection of disjoint segments, called *zones*, which partitions the subset of the real line covered by the intervals seen so far. For each zone, the algorithm keeps track of two intervals from the input stream: the one with the leftmost left endpoint (reps., rightmost right endpoint) among those with a left (resp., right) endpoint in the zone. The output of the algorithm is the maximum interval selection from this set of recorded intervals. The performance guarantee follows from showing that this solution contains at least two intervals within any span of three intervals of the optimal solution.

We now define the zones. Define the *support* of a set X of intervals, denoted $\text{supp}(X)$, to be the subset of the real line covered by intervals in X , or $\text{supp}(X) = \cup_{I \in X} I$. *Connected components* of the support of X are defined in a natural manner. The maximal segments on the real line that are outside connected components of the support of X are called *out-regions* (with respect to X). The zones are segments on the real line and the zone collection Z_t at time t is a partition of $\text{supp}(S_t)$. A zone may initially be *flexible* in that its endpoints might change, and in that it may be merged with other regions or zones into a single zone. Zones at the extreme ends of connected components are flexible, with the sole exception of the initial zone of a connected component; other zones are *fixed* and unchangeable, with permanent endpoints.

We now specify how the algorithm creates and maintains the zones. Initially, there are no zones and the whole line is considered one out-region. When an interval I is received at time t , we consider the following cases depending on the positions of its endpoints:

- (1) *Both endpoints of I are in an out-region:* Create a new fixed zone equal to I , in a connected component of its own.
- (2) *Both endpoints of I fall within the same connected component:* Do nothing.
- (3) *The endpoints of I belong to zones in different connected components:* First, for each of the two zones in which the endpoints fall, fix it (if it is not already fixed) without changing its endpoints. Then, create a new fixed zone, denoted by z , which includes the out-region between the respective components. If I properly includes one or two flexible zones, then these zones are merged into z ; the resulting zone z is then a fixed zone with right (reps., left) endpoint changed to be the right (resp., left) endpoint of the flexible zone to the right (reps., left) of z .
- (4) *One endpoint of I falls in a zone and the other in an out-region:* Let k be the zone, and C be the connected component, in which one endpoint of I falls. Without loss of generality assume that the right endpoint of I falls in the out-region, i.e., the

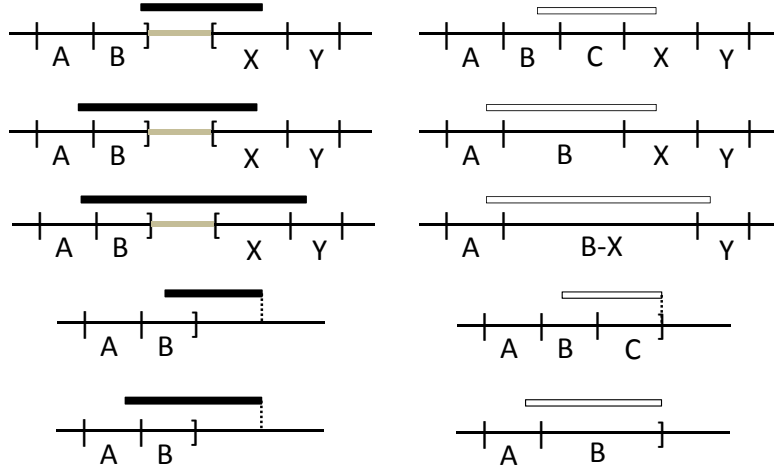


Fig. 2. Illustration of the different cases in zone management, with three variations of Case 3 and two for Case 4. On the left, an incoming interval is shown in bold with the prior settings of the zones; the resulting zone configuration is shown on the right. Flexible ends of zones are shown with brackets and out-regions in grey.

left endpoint of I falls in k . Fix zone k (if it is not already fixed) without changing its endpoints. Create a new flexible zone, denoted by z , which covers $I \setminus \text{supp}(C)$. If I properly includes a flexible zone k' , then k' is merged into z , where the resulting zone z has as its left endpoint the left endpoint of k' .

This completes the specification of the zones. See Fig. 2 for an illustration of the different cases.

It is clear from the definition of the zones and their endpoints that the zone collection partitions the support of the intervals seen so far. We shall assume a total order on the zones, from left to right. The following lemma is crucial for the correctness and performance guarantee of the algorithm.

LEMMA 4.1. *The zone collection satisfies the following invariants:*

- (#). *Each zone is contained in an input interval, and*
- (*). *Each input interval contains at most two zones.*

PROOF. The former invariant follows directly from how zones are created (in Cases 1, 3, and 4).

To prove the latter invariant by *reductio ad absurdum*, assume that there is an input sequence for which the invariant fails. Let I_t be the first interval presented that (at some time) contains three zones, denoted k_1, k_2, k_3 in order (from left to right). Since zones never shrink, time t is then also the first time that the invariant fails. Let $t' \leq t$ be the time when zone k_2 achieves the extent that it has at time t . Consider the interval $I_{t'}$ introduced at time t' . Since no zone is created or changed in Case 2, it does not apply to $I_{t'}$.

Let $K = [\text{left}(k_1), \text{right}(k_3)]$ denote the minimal segment of the real line containing all three zones. We observe that $I_{t'} \subsetneq K \subseteq I_t$, implying that the input intervals are not proper, which is a contradiction. Namely, in Case 1, the zone created corresponds to the input interval, while, as is clear from Fig. 2, in Cases 3 and 4, the input interval has its endpoints in the zones (k_1, k_3) on either side of the center zone (k_2) that is created or changed. \square

We next define the intervals maintained by the algorithm. For an interval I and time t , let $bz(I) = bz_t(I)$ (resp. $fz(I) = fz_t(I)$) be the zone in which $\text{left}(I)$ (resp. $\text{right}(I)$) falls at time t . For each zone k , and time t , the algorithm stores the endpoints of the intervals $L_k = \arg \min_{I \in S_t, bz(I)=k} \text{left}(I)$, i.e., the one with the leftmost left endpoint in the zone, and $R_k = \arg \max_{I \in S_t, fz(I)=k} \text{right}(I)$, i.e., the one with the rightmost right endpoint in the zone. Note that L_k (resp., R_k) will be undefined if no interval has its left (resp., right) endpoint in zone k .

The output of the algorithm is the maximum interval selection of the set $A = \cup_k \{L_k, R_k\}$, obtained by, e.g., the classic left-right greedy algorithm.

We next detail how the algorithm maintains the intervals stored. When interval I arrives at time t , we consider the four cases defined above.

If Case 1 applies, then I defines a new zone k , and both L_k and R_k are permanently set to equal I .

If Case 2 applies, no zones change; the L - and R -intervals are updated for the zones in which the endpoints of I fall. E.g., if $b(I)$ falls in zone k , and $b(I) < b(L_k)$, then L_k is updated as I .

If Case 3 applies, the L - and R -intervals are updated, for the zones in which the endpoints of I fall. Then, a new zone is created, but all the endpoints, if any, that fall in that zone have already appeared, and thus we need only merge the information from the at most two previous flexible zones that are merged into the new zone.

If Case 4 applies, then one endpoint of I falls in a zone k , which becomes fixed (if it was not fixed before). Without loss of generality, assume it is the left endpoint of I that falls in k . Then, L_k is updated appropriately. Also, a new zone z is created. If I does not properly include any flexible zone (Subcase 1), then z covers $I \setminus \text{supp}(C)$, R_z is defined to be I , and L_z is undefined. If I does properly include a flexible zone k' (Subcase 2), then k' is merged into z . R_z is defined to be I , and L_z is defined to be $L_{k'}$.

We conclude that the algorithm can indeed maintain the intervals L_k and R_k for each zone k as defined.

An immediate corollary of Lemma 4.1 is that the union of any five adjacent zones must properly contain an input interval. By induction, any $5t$ adjacent zones must include at least t disjoint intervals. Hence, we obtain an upper bound on the space used by the algorithm by noting that for each zone k , only six values need to be recorded: the endpoints of L_k , of R_k and of the zone k itself. We thus have the following.

LEMMA 4.2. *The number of zones is at most $5|OPT| + 4$. The space used by the algorithm is at most $O(|OPT|)$.*

Finally, we prove the performance guarantee of the algorithm. The following lemma captures the core of the argument. For two proper intervals I and J , we write $I \leq J$ to denote that $\text{left}(I) \leq \text{left}(J)$ (and thus also $\text{right}(I) \leq \text{right}(J)$), and similarly $I < J$ if $\text{left}(I) < \text{left}(J)$. The *span* of a set X of intervals is the segment defined by the leftmost left endpoint and the rightmost right endpoint of intervals in X .

LEMMA 4.3. *Let $R \subseteq S_t$ be a collection of three disjoint input intervals. Then, at the end of round t , A contains a pair of intervals of S_t that are contained in the span of R .*

PROOF. Let the three disjoint intervals be $O_1 < O_2 < O_3$. Our claim is that A contains a pair of disjoint intervals $I, I' \in S_t$, $I, I' \subset [\text{left}(O_1), \text{right}(O_3))$.

Let $b_i = bz(O_i)$ (respectively, $f_i = fz(O_i)$) be the zone in which the left (resp., right) endpoint of O_i falls, for $i = 1, 2, 3$, respectively. Clearly, $b_1 \leq f_1 \leq b_2 \leq f_2 \leq b_3 \leq f_3$. We observe that $b_1 < b_2 < b_3$. Suppose, e.g., that $b_1 = f_1 = b_2$. Then, O_1 is contained within zone b_1 , and since zone b_1 contains the point $\text{left}(O_2)$ that is not contained in O_1 , it

follows that O_1 is properly contained in zone b_1 . So, by Lemma 4.1(Invariant (#)), there is an interval that properly contains O_1 , which is a contradiction. Therefore, $f_1 < b_3$.

Consider R_{f_1} and L_{b_3} , which are well-defined intervals since the right endpoint of O_1 falls in zone f_1 and the left endpoint of O_3 falls in zone b_3 . By definition, $O_1 \leq R_{f_1}$ and $L_{b_3} \leq O_3$, so $R_{f_1}, L_{b_3} \subset [\text{left}(O_1), \text{right}(O_3))$. Since $f_1 < b_3$, it follows that R_{f_1} and L_{b_3} are disjoint, establishing the claim. \square

We conclude with the following theorem.

THEOREM 4.4. $|ALG| \geq 2|OPT|/3$.

PROOF. Let $|S| = n$, and let O_0, O_1, \dots, O_{p-1} be the intervals in the optimal solution OPT in order of left endpoints, where $p = |OPT|$. Let $s = \lfloor \frac{p}{3} \rfloor$, and let OPT' consist of the first (in that order) $3s$ intervals from OPT . By Lemma 4.3, A contains a pair of intervals within each of the segments $[\text{left}(O_{3r}), \text{right}(O_{3r+2}))$, for $0 \leq r < s$, and thus at least $2s = 2\lfloor \frac{p}{3} \rfloor$ intervals from $[\text{left}(O_0), \text{right}(O_{3s-1}))$, the span of OPT' . If $3 \mid p$, then this establishes the theorem.

Suppose that $3 \nmid p$. Then, the last interval in OPT , O_{p-1} , is not contained in OPT' . Observe that $O_{p-1} \leq R_{k_{max}}$, where k_{max} is the last zone at the end of the last round. Thus, A contains $2s$ disjoint intervals within in the span of OPT' and one that doesn't intersect the span of OPT' . Combined, we get that $|ALG| \geq 2s+1 = 2\lfloor \frac{p}{3} \rfloor + 1 \geq 2p/3$. \square

5. LOWER BOUNDS

In this section we establish lower bounds on the approximation ratio of randomized streaming algorithms for the interval selection problem, establishing the following two theorems.

THEOREM 5.1 (LOWER BOUND FOR GENERAL INTERVALS). *For every real $\epsilon > 0$, integer $q_0 > 0$, and sublinear function $s : \mathbb{N} \rightarrow \mathbb{N}$, there exist $q_0 \leq q \leq c \cdot q_0$, where c is a universal constant, an arbitrarily large integer z , and an interval stream S such that*

- (1) $|S| = z$;
- (2) $|Opt(S)| = q$; and
- (3) $Alg(S) < q(1/2 + \epsilon)$ for any randomized interval selection streaming algorithm Alg with space $s(z) \cdot b$, where $b = \Theta(\log z \cdot \log q)$ is the length of the bit strings representing the endpoints.

THEOREM 5.2 (LOWER BOUND FOR UNIT INTERVALS). *For every real $\epsilon > 0$, integer $q_0 > 0$, and sublinear function $s : \mathbb{N} \rightarrow \mathbb{N}$, there exist $q_0 \leq q \leq c \cdot q_0$, where c is a universal constant, an arbitrarily large integer z , and a unit interval stream S such that*

- (1) $|S| = z$;
- (2) $|Opt(S)| = q$; and
- (3) $Alg(S) < q(2/3 + \epsilon)$ for any randomized proper interval selection streaming algorithm Alg with space $s(z) \cdot b$, where $b = \Theta(\log z)$ is the length of the bit strings representing the endpoints.

Notice that in Theorems 5.1 and 5.2, the stream length z can be made arbitrarily large with respect to the output size q . This means, in particular, that the lower bounds given by these two theorems hold for any randomized streaming algorithm that can store up to $s'(q)$ intervals (rather than $s(z)$ intervals) for any function $s' : \mathbb{N} \rightarrow \mathbb{N}$.

Our lower bounds are proved by designing a random interval stream S for which every deterministic algorithm performs badly on expectation; the assertion then follows by Yao's principle. (Our construction uses half-open intervals, but this can be easily altered.) We remark that under the setting used by our lower bounds, the algorithm is

required to output a collection \mathcal{C} of disjoint intervals, and the quality of the solution is then determined to be the cardinality of $\mathcal{C} \cap S$. In other words, the algorithm is allowed to output non-existing intervals (that is, intervals that never arrived in the input), but it will not be credited for them. This, obviously, can only increase the power of the algorithm, and makes our lower bound stronger.

The (k, n) -gadget. Fix some positive integer m whose role is to bound the space of the algorithm. Our lower bounds rely on the following framework, characterized by the parameters $k, n \in \mathbb{Z}_{>0}$, denoted a (k, n) -gadget. Consider an extensive form two-player zero-sum game played between the algorithm (MAX) and the adversary (MIN), depicted by a sequence of k phases. Informally, in each phase t , the adversary chooses a permutation $\pi_t \in P_n$, where P_n is the collection of all permutations on $[n]$, and an index $i_t \in [n]$.⁴ The algorithm observes π_t (but not i_t) and produces a *memory image* M_t , i.e., a bit string of length m . The index i_t is handed to the algorithm after the memory image is produced. At the end of the last phase the algorithm tries to *recover* $\pi_t(i_t)$ for $t = 1, \dots, k$: it outputs some $i_t^* \in [n]$ based on the memory image M_t , index i_t , and all other memory images and indices. For each t such that $i_t^* = \pi_t(i_t)$, the algorithm scores a (positive) point.

More formally, the adversarial strategy is defined by the choices of the permutations π_t and the indices i_t for $t = 1, \dots, k$. We commit the adversary to make those choices uniformly at random (so, the adversary reveals its mixed strategy), namely, $\pi_t \in_r P_n$ and $i_t \in_r [n]$ for every t , where all the random choices are independent. The strategy of the algorithm is defined by the function sequences $\{f_t\}_{t=1}^k$ and $\{g_t\}_{t=1}^k$, where

$$f_t : P_n \times (\{0, 1\}^m \times [n])^{t-1} \rightarrow \{0, 1\}^m \quad \text{and} \quad g_t : \{0, 1\}^m \times [n] \times (\{0, 1\}^m \times [n])^{k-1} \rightarrow [n].$$

Let Γ_0 be the empty string and recursively define $\Gamma_t = \Gamma_{t-1} \circ f_t(\pi_t, \Gamma_{t-1}) \circ i_t$.⁵ The payoff of the algorithm is the number of t s, $1 \leq t \leq k$, such that

$$g_t \left(f_t(\pi_t, \Gamma_{t-1}), i_t, \{f_{t'}(\pi_{t'}, \Gamma_{t'-1}), i_{t'}\}_{t' \neq t} \right) = \pi_t(i_t).$$

In the language of the aforementioned informal description, the role of the function f_t is to produce the memory image M_t based on the permutation π_t and all previous memory images and indices (whose concatenation is given by Γ_{t-1}). The role of the function g_t is to recover $\pi_t(i_t)$ based on the memory image M_t , index i_t , and all other memory images and indices.

Note that the memory images $M_{t'}$ and indices $i_{t'}$, $t' \neq t$, do not contain any information on the permutation π_t on top of that contained in M_t . In particular, the entropy in $\pi_t(i_t)$ given M_t , i_t , and $\{M_{t'}, i_{t'}\}_{t' \neq t}$ is equal to the entropy in $\pi_t(i_t)$ given M_t and i_t . Therefore, it will be convenient to decompose the domain of the function $g_t : \{0, 1\}^m \times [n] \times (\{0, 1\}^m \times [n])^{k-1} \rightarrow [n]$ so that the $(\{0, 1\}^m \times [n])^{k-1}$ -part determines which function $\hat{g}_t : \{0, 1\}^m \times [n] \rightarrow [n]$ is chosen, and then this function \hat{g}_t is used to produce i_t^* based on M_t and i_t . Similarly, we decompose the domain of the function $f_t : P_n \times (\{0, 1\}^m \times [n])^{t-1} \rightarrow \{0, 1\}^m$ so that the $(\{0, 1\}^m \times [n])^{t-1}$ -part determines which function $\hat{f}_t : P_n \rightarrow \{0, 1\}^m$ is chosen, and then this function \hat{f}_t is used to produce M_t based on π_t .

We now turn to bound from above the expected payoff of the algorithm as a function of k , m , and n . The key ingredient in this context is the following lemma, which is essentially a well known fact in slightly different settings; a proof is provided in Appendix B for completeness.

⁴Throughout, we use the notation $[n]$, where n is a positive integer, to denote the set $\{1, \dots, n\}$.

⁵We use the notation $u \circ v$ to denote the concatenation of the string u to string v .




Fig. 3. The relative locations of the intervals in an (n, π) -stack for $n = 4$. The left and right endpoints of interval J_i are located in the segments depicted by the bidirectional arrows whose length is $\lambda/2$. The exact location within this segment is determined by $\pi(i)$. In the construction of the 2-lower bound for general intervals, the bold rectangles correspond to the segments σ_ℓ and σ_r in which the stacks (or auxiliary intervals) identified with the left and right children, respectively, of the current node are deployed assuming that the good interval is interval J_2 (these segments do not intersect with the segments corresponding to the bidirectional arrows).

LEMMA 5.3. *For every real $\alpha > 0$ and integer $n_0 > 0$, there exists an integer $n > n_0$ such that for every two functions $\hat{f} : P_n \rightarrow \{0, 1\}^m$ and $\hat{g} : \{0, 1\}^m \times [n] \rightarrow [n]$, where $m = \lfloor \alpha n \log n \rfloor$, we have $\mathbb{P}_{\pi \in_r P_n, i \in_r [n]} (\hat{g}(\hat{f}(\pi), i) = \pi(i)) < 2\alpha$.*

COROLLARY 5.4. *For every real $\alpha > 0$ and integers $k, n_0 > 0$, there exists an integer $n > n_0$ such that if $m \leq \alpha n \log n$, then the expected payoff of the algorithm player in a (k, n) -gadget with space m is smaller than $2\alpha k$.*

Construction's punchline. We shall establish Theorems 5.1 and 5.2 by a reduction from the (k, n) -gadget framework. Specifically, given the parameters k and n , we construct a random interval stream (resp., a unit interval stream) S that satisfies

- (i) the length of the stream is $z = |S| = \Theta(nk)$;
 - (ii) the size of the optimal solution is $q = \text{Opt}(S) = 2k + 1$ (resp., $q = \text{Opt}(S) = 3k$);
 - (iii) the endpoints of the intervals in S can be represented by bit strings of length $b = \Theta(\log n \cdot \log k)$ (resp., $b = \Theta(\log n + \log k)$); and
 - (iv) a deterministic streaming algorithm with space s that outputs in expectation $k + 1 + \ell$ (resp., $2k + \ell$) disjoint intervals of S implies an algorithm player for the (k, n) -gadget with memory images of size s that achieves an expected payoff at least ℓ .
- Theorems 5.1 and 5.2 follow by Corollary 5.4 since we can pick $\alpha > 0$ and a sufficiently large n so that $2\alpha k \leq \epsilon q$ and $s(z) \cdot b \leq \alpha n \log n$.

The (n, π) -stack. We now turn to implement a (k, n) -gadget via a carefully designed interval stream. As a first step, we introduce the (n, π) -stack construction. Given an integer $n > 0$ and a permutation $\pi \in P_n$, an (n, π) -stack deployed in the segment $[x, y)$, $x < y$, is a collection of n intervals J_1, \dots, J_n satisfying:

- (1) all intervals J_i are half open;
- (2) all intervals J_i have the same length, i.e., $\text{right}(J_i) - \text{left}(J_i) = \lambda n$, where $\lambda = \frac{y-x}{2n-1/2}$; and
- (3) $\text{left}(J_i) = x + \lambda(i-1) + \epsilon\pi(i)$ for every $i \in [n]$, where $\epsilon = \lambda/(2n)$.

Notice that this deployment ensures that

$$\text{left}(J_1) < \text{left}(J_2) < \dots < \text{left}(J_n) < \text{right}(J_1) < \text{right}(J_2) < \dots < \text{right}(J_n).$$

This implies in particular that the half open segment $[\text{left}(J_n), \text{right}(J_1))$ is contained in J_i for every $i \in [n]$. Moreover, the union of the intervals in the stack does not necessarily cover the whole segment $[x, y)$; it is always contained in $[x, y)$, though. The structure of an (n, π) -stack is illustrated in Figure 3.

The (k, n) -gadget is implemented by introducing k stacks, each corresponding to one phase, and some *auxiliary* intervals; the stack corresponding to phase t is referred to as stack t . The permutation π used in the construction of stack t is π_t . The index i_t will dictate the choice of one *good* interval out of the n intervals in that stack. What exactly makes this interval good will be clarified soon; informally, the algorithm has no incentive to output an interval in a stack unless this interval is good.

The k stacks are used both by the construction of the 2-lower bound for general interval streams and by that of the $(3/2)$ -lower bound for unit intervals. The difference between the two constructions lies in the manner in which these stacks are deployed in the real line, and in the addition of the auxiliary intervals.

A $(3/2)$ -lower bound for unit intervals. The interval stream that realizes the (k, n) -gadget for the $(3/2)$ -lower bound for unit intervals is constructed as follows. It contains k sufficiently spaced apart stacks, where the intervals in each stack are scaled to a unit length (so $\lambda = 1/n$). Consider stack t and suppose that it is deployed in the segment $[x, y]$, where $y = x + 2 - 1/(2n)$. Recall that the permutation that determines the exact location of the intervals in the stack is π_t and that the good interval is J_{i_t} .

After the arrival of the n intervals in the stack, two more half open unit auxiliary intervals are presented:

$$L_t = \left[x + \frac{i_t - 1}{n} - 1, x + \frac{i_t - 1}{n} \right) \quad \text{and} \quad R_t = \left[x + \frac{i_t - 1/2}{n} + 1, x + \frac{i_t - 1/2}{n} + 2 \right).$$

In other words, the interval L_t (respectively, R_t) is located to the left (resp., right) of the leftmost (resp., rightmost) point in which $\text{left}(J_{i_t})$ (resp., $\text{right}(J_{i_t})$) may be deployed. It is easy to verify that except for the good interval J_{i_t} that does not intersect with L_t and R_t , every interval in the stack intersects with exactly one of these two auxiliary intervals.

The best response of the algorithm would be to output the two auxiliary intervals and to try to recover the good interval J_{i_t} . (Note that the payoff guaranteed by this strategy is at least 2 per stack, whereas any other strategy yields a payoff of at most 2 per stack.) For that purpose, the algorithm has to recover the exact locations of the endpoints of J_{i_t} that implicitly encode $\pi_t(i_t)$. Observing that the endpoints in this construction can be represented by bit strings of length $\log(n) + \log(k)$, Theorem 5.2 follows by Corollary 5.4.

A 2-lower bound for general intervals. The interval stream that realizes the (k, n) -gadget for the 2-lower bound for general intervals is constructed as follows. Assume that $k = 2^\kappa - 1$ for some positive integer κ and consider a perfect binary tree T of depth κ . The k stacks are identified with the internal nodes of T so that stack t precedes stack $t + 1$ in a pre-order traversal of T . (In other words, if stack t is identified with node u and stack t' is identified with a child of u , then $t < t'$.) In addition to the intervals in the stacks, we also introduce $2^\kappa = k + 1$ auxiliary intervals which are identified with the leaves of T ; these auxiliary intervals arrive last in the stream. We say that an interval J is *assigned* to node $u \in T$ if J belongs to the stack identified with u or if u is a leaf and J is the auxiliary interval identified with it.

The deployment of the stacks and the auxiliary intervals in \mathbb{R} is performed as follows. Stack 1 (identified with T 's root) is deployed in $[0, 1)$. Given the deployment of stack t identified with internal node $u \in T$ in the segment $[x, y)$, we deploy the stacks identified with the left and right children of u in the segments

$$\sigma_\ell = [x + \lambda(i_t - 3/2), x + \lambda(i_t - 1)) \quad \text{and} \quad \sigma_r = [x + \lambda(i_t + n - 1/2), x + \lambda(i_t + n)) ,$$

respectively, where recall that $\lambda = \frac{y-x}{2n-1/2}$. If the children of u are leaves in T , then we deploy auxiliary intervals in those two segments instead of stacks, that is, one auxiliary interval in σ_ℓ and one in σ_r . Refer to Figure 3 for an illustration.

The key observation regarding the choice of σ_ℓ and σ_r is that

$$\begin{aligned} \text{left}(J_{i_t-1}) \leq \text{left}(\sigma_\ell) < \text{right}(\sigma_\ell) \leq \text{left}(J_{i_t}) \quad \text{and} \\ \text{right}(J_{i_t}) \leq \text{left}(\sigma_r) < \text{right}(\sigma_r) \leq \text{right}(J_{i_t+1}). \end{aligned}$$

In particular, this implies that: (1) the good interval in the stack identified with node $u \in T$ does not intersect with any interval assigned to a descendant of u in T ; and (2) a non-good interval in the stack identified with node $u \in T$ contains every interval assigned to a descendant of either the left child of u or the right child of u in T .

Since the $k+1$ auxiliary intervals are non-intersecting and arrive last in the stream, they can be included in the output of the algorithm without requiring any additional space (on top of that dedicated to their representation). Moreover, an auxiliary interval intersects with at most one interval in any valid solution, hence it is a dominant strategy for of the algorithm to output all the auxiliary intervals. Therefore, the best response of the algorithm can include an interval J_i of stack t , $1 \leq t \leq k$, in the output only if it is the good interval of that stack, namely, $i = i_t$.

By definition, in order to include interval J_{i_t} of stack t in the output, the algorithm must hold the exact locations of its endpoints. The construction of stack t based on permutation π_t implies that the exact locations of interval J_{i_t} 's endpoints encode the value of $\pi_t(i_t)$. Observing that the endpoints in this construction can be represented by bit strings of length $\log(n) \cdot \log(k)$, Theorem 5.1 follows by Corollary 5.4.

6. MULTIPLE-PASS ALGORITHMS

We extend now the streaming algorithms to use multiple passes through the data. First, some notations. For an interval I , let $\text{next}(I)$ be the interval in the input that has the leftmost right endpoint among those having a left endpoint to the right of the right endpoint of I , and let $\text{prev}(I)$ be the interval that has the rightmost left endpoint among those having a right endpoint to the left of the left endpoint of I . We use the notation $\text{next}^i(I)$ defined recursively by $\text{next}^0(I) = I$ and $\text{next}^i(I) = \text{next}(\text{next}^{i-1}(I))$, for $i > 0$. Define $\text{prev}^i(I)$ similarly. Observe that if I is available before a pass, then a streaming algorithm can easily compute $\text{next}(I)$ and $\text{prev}(I)$ by the end of the pass, while maintaining $O(1)$ intervals in the memory at all times.

The multi-pass algorithm runs as follows. The first pass consists of our one-pass algorithm, either as the algorithm of Sec. 3 for general intervals, or the algorithm of Sec. 4 for proper intervals. The result of this pass is the set A , whichever base algorithm is used. Let $N_0 = P_0 = A$. In round $p > 1$, the algorithm inductively computes $N_{p-1} = \{\text{next}(I) : I \in N_{p-2}\}$ and $P_{p-1} = \{\text{prev}(I) : I \in P_{p-2}\}$. Let $A_p = \cup_{i \geq 0}^{p-1} (N_i \cup P_i) = \{\text{next}^i(I), \text{prev}^i(I) : I \in A, 0 \leq i \leq p-1\}$ denote the combined set of intervals stored after pass p . When requested, the algorithm produces as output the maximum interval selection in A_p . This completes the specification of the algorithm.

We first observe that $|A_p| \leq (2p-1)|A|$, hence the space used in phase p is at most $2p-1$ times larger than the length of the bit string representing A .

Define the *span* of a set R of intervals to be the segment with endpoints being the leftmost left endpoint of the intervals in R , and the rightmost right endpoint of the intervals in R . The main idea behind the performance argument is to show that each group of disjoint intervals (in particular, the optimal solution) must contain within its span nearly as many disjoint intervals in A_p . We first handle the base case for general intervals.

LEMMA 6.1. *Given an input of general intervals, the set A computed by the algorithm Alg of Sec. 3 satisfies the following property: for any pair of disjoint input intervals I_1 and I_2 , A contains an interval within the span of $\{I_1, I_2\}$ (given by $[\text{left}(I_1), \text{right}(I_2))$), assuming $I_1 < I_2$.*

PROOF. Let I_1 and I_2 be disjoint input intervals. By Lemma 3.3, there are intervals ρ_1, ρ_2 in $A \cup V$ such that $\rho_1 \subseteq I_1$ and $\rho_2 \subseteq I_2$. If $\rho_i \in A$ for either $i = 1$ or $i = 2$, then the proof is completed. Otherwise ($\rho_1, \rho_2 \in V$), Lemma 3.4(P3) and Corollary 3.6 ensure that the span of $\{\rho_1, \rho_2\}$ contain some interval in A . The proof is completed since the span of $\{\rho_1, \rho_2\}$ is a subset of the span of $\{I_1, I_2\}$. \square

The following lemmas apply both to general or proper intervals.

LEMMA 6.2. *Let I be an interval in A_p and let $s \leq p - 1$. Let $R \subseteq S$, such that $I \in R$, be a set of $s + 1$ disjoint intervals. Then, A_p contains a set of $s + 1$ intervals within the span of R .*

PROOF. Suppose R contains intervals I, I_1, I_2, \dots, I_s , s.t. $I_1 < I_2 \dots I_j < I < I_{j+1} < \dots < I_s$. By definition, the intervals $\text{next}^i(I)$, $0 \leq i \leq s - j$, are disjoint and contained in A_{s+1} and thus also in A_p . Also, by induction, $\text{next}^i(I) \leq I_i$, for $i = 1, \dots, s - j$, and thus they fall within the span of R . A similar claim holds for the intervals $\text{prev}^i(I)$, $0 \leq i \leq j$. \square

LEMMA 6.3. *Consider any set R of m disjoint intervals in S , where $m = 2p$ for general intervals and $m = 2p + 1$ for proper intervals. Then, A_p contains $m - 1$ intervals within the span of R .*

PROOF. Consider first general intervals. Partition R into R_l, R', R_r in order, where R_l and R_r each contain $p - 1$ intervals and R' contains two. By Lemma 6.1, $A = A_1$ contains an interval I within the span of R' . By Lemma 6.2, A_p contains p intervals within the span of $R_r \cup \{I\}$ and p intervals within the span of $R_l \cup \{I\}$, for a total of $2p - 1 = m - 1$ intervals.

For proper intervals, partition R into R_l, R', R_r in order, where R_l and R_r each contain $p - 2$ intervals and R' contains three. By Lemma 4.3, $A = A_1$ contains two disjoint intervals, I_l and I_r , within the span of R' . By Lemma 6.2, A_p contains $p - 1$ intervals within the span of $R_r \cup \{I_r\}$ and $p - 1$ intervals within the span of $R_l \cup \{I_l\}$, for a total of $2(p - 1) = m - 1$ intervals. \square

An interval is said to be *end-simplicial* with respect to a set X of intervals, if it contains either the leftmost right endpoint or the rightmost left endpoint of its connected component with respect to X .

LEMMA 6.4. *The set A contains all the end-simplicial intervals with respect to S .*

PROOF. End-simpliciality means that the interval cannot properly contain other input intervals. Regarding general intervals, recall from Proposition 3.1 that virtual intervals in $\text{Alg}(S)$ are formed by the intersection of two intervals in the input. Thus, if I is end-simplicial with respect to S , it contains no virtual interval. Hence, I is admitted to A and never rejected.

For proper intervals, observe that by the proper intervals algorithm (Section 4), an end-simplicial interval I on the left (reps., right) of its component in S , will always be the left representative, L_k , (resp., the right representative, R_k ,) of the left (resp., right) zone contained in I . \square

THEOREM 6.5. *The multi-pass algorithm run for p passes is a $1 + \frac{1}{2^{p-1}}$ -approximation (respectively, $1 + \frac{1}{2^p}$ -approximation) algorithm for general (reps., proper) intervals. The space used by the algorithm run for p passes is $O(p)$ times the size of the output.*

PROOF. Define $m = 2p$ for general intervals and $m = 2p + 1$ for proper intervals. Consider an optimal interval selection with intervals $I_1, \dots, I_{|OPT|}$, where $\alpha = |OPT|$. Let $r = \alpha \bmod m$, and $q = \lfloor \alpha/m \rfloor$. Also let $t = \lceil r/2 \rceil$ and $t' = \lfloor r/2 \rfloor$. For each $R_i = \{I_{t+1+im}, \dots, I_{t+(i+1)m}\}$, where $i = 0 \dots, q-1$, it holds by Lemma 6.3 that A_p contains $m-1$ intervals within the span of R_i . By Lemmas 6.4 and 6.2, A_p also contains t disjoint intervals within the span of $[\text{left}(S), \text{right}(I_t)]$ and t' disjoint intervals within the span of $[\text{left}(I_{m-t'+1}), \text{right}(S)]$. Hence, A_p contains at least $q(m-1) + t + t' = \alpha - q \geq \alpha(m-1)/m$ disjoint intervals. \square

7. ONLINE ALGORITHM

In this section we briefly show how to use the streaming algorithm presented in Sec. 3 to derive a randomized preemptive online interval selection algorithm. Our algorithm is 6-competitive and on top of maintaining at any time the set of currently accepted intervals A^* , its only additional memory is an interval set of cardinality linear in the size of the current optimum. We thus answer an open question of Adler and Azar [1] about the space complexity of randomized preemptive online algorithms for our problem.

Recall that our streaming algorithm maintains a set A of intervals. With respect to that set, our algorithm is a deterministic preemptive online algorithm, adding an interval to A only when that interval arrives, and possibly preempting it later. By Corollary 3.8, the cardinality of the set A is at least half the cardinality of the optimal solution of the input seen so far. Moreover, combining Lemma 3.4(P6) and Lemma 3.4(P7), we conclude that every interval added to A intersects with at most 2 previous intervals in A . Therefore, A is *online 3-colorable*: upon addition into A , each interval can be assigned one of three colors, such that any two intersecting intervals always have two different colors.

Our preemptive algorithm is now simple. We initially pick a random color c in $\{1, 2, 3\}$. We then run the streaming algorithm on each received interval I , adding I to A , and preempting intervals from A as does the streaming algorithm. If I is added to A we assign it a valid color from $\{1, 2, 3\}$ in a first-fit manner. Our solution ALG consists of every interval J in A whose color is c . Clearly, $E[|ALG|] = |A|/3 \geq |OPT|/6$, that is, the algorithm is 6-competitive.

Acknowledgments

We thank Jaikumar Radhakrishnan and Oded Regev for helpful discussions.

REFERENCES

- R. Adler and Y. Azar. Beating the logarithmic lower bound: Randomized preemptive disjoint paths and call control algorithms. *J. Scheduling*, 6(2):113–129, 2003.
- P. K. Agarwal and R. Sharathkumar. Streaming algorithms for extent problems in high dimensions. In *SODA*, pages 1481–1489, 2010.
- K. Ahn and S. Guha. Graph sparsification in the semi-streaming model. In *ICALP*, pages 328–338, 2009.
- N. Alon, Y. Matias, and M. Szegedy. The space complexity of approximating the frequency moments. *J. Comput. Syst. Sci.*, 58(1):137–147, 1999.
- B. Awerbuch, Y. Bartal, A. Fiat, and A. Rosén. Competitive non-preemptive call control. In *SODA*, pages 312–320, 1994.
- Y. Azar and O. Gilon. Buffer management for packets with processing times. In N. Bansal and I. Finocchi, editors, *Algorithms - ESA 2015 - 23rd Annual European Symposium, Patras, Greece, September 14-16, 2015, Proceedings*, volume 9294 of *Lecture Notes in Computer Science*, pages 47–58. Springer, 2015.
- U. T. Bachmann, M. M. Halldórsson, and H. Shachnai. Online scheduling intervals and t -intervals. In *SWAT*, 2010.
- S. Cabello and P. Pérez-Lantero. Interval selection in the streaming model. In F. Dehne, J. Sack, and U. Stege, editors, *Algorithms and Data Structures - 14th International Symposium, WADS 2015, Victoria, BC, Canada, August 5-7, 2015. Proceedings*, volume 9214 of *Lecture Notes in Computer Science*, pages 127–139. Springer, 2015.
- R. Canetti and S. Irani. Bounding the power of preemption in randomized scheduling. *SIAM J. Comput.*, 27(4):993–1015, 1998.
- T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms*. MIT Press and McGraw-Hill, 3rd edition, 2009.
- Y. Emek, M. Halldórsson, and A. Rosén. Space-constrained interval selection. In *ICALP*, pages 302–313, 2012.
- L. Epstein and A. Levin. Improved randomized results for the interval selection problem. *Theor. Comput. Sci.*, 411(34-36):3129–3135, 2010.
- L. Epstein, A. Levin, J. Mestre, and D. Segev. Improved approximation guarantees for weighted matching in the semi-streaming model. In *STACS*, pages 347–358, 2010.
- J. Feigenbaum, S. Kannan, A. McGregor, S. Suri, and J. Zhang. On graph problems in a semi-streaming model. *Theor. Comput. Sci.*, 348:207–216, December 2005.
- J. Feigenbaum, S. Kannan, A. McGregor, S. Suri, and J. Zhang. Graph distances in the data-stream model. *SIAM J. Comput.*, 38(5):1709–1727, 2008.
- S. Fung, C. Poon, and F. Zheng. Improved randomized online scheduling of unit length intervals and jobs. In *WAOA*, 2008.
- F. Gavril. Algorithms for minimum coloring, maximum clique, minimum covering by cliques, and maximum independent set of a chordal graph. *SIAM J. Comput.*, 1(2):180–187, 1972.
- B. V. Halldórsson, M. M. Halldórsson, E. Losievskaja, and M. Szegedy. Streaming algorithms for independent sets. In *ICALP*, pages 641–652, 2010.
- M. R. Henzinger, P. Raghavan, and S. Rajagopalan. Computing on data streams. In J. M. Abello and J. S. Vitter, editors, *External Memory Algorithms*, pages 107–118. American Mathematical Society, Boston, MA, USA, 1999.
- J. Kleinberg and E. Tardos. *Algorithm Design*. Addison Wesley, 2005.
- S. Leonardi, A. Marchetti-Spaccamela, A. Presciutti, and A. Rosén. On-line randomized call control revisited. *SIAM J. Comput.*, 31(1):86–112, 2001.
- R. J. Lipton and A. Tomkins. Online interval scheduling. In *SODA*, pages 302–311, 1994.
- A. McGregor. Finding graph matchings in data streams. In *APPROX-RANDOM*, pages 170–181, 2005.
- P. Melsted and B. V. Halldórsson. Kmerstream: streaming algorithms for k -mer abundance estimation. *Bioinformatics*, 30(24):3541–3547, 2014.
- J. I. Munro and M. Paterson. Selection and sorting with limited storage. *Theor. Comput. Sci.*, 12:315–323, 1980.
- S. Muthukrishnan. Data streams: Algorithms and applications. *Foundations and Trends in Theoretical Computer Science*, 1(2), 2005.
- G. J. Woeginger. On-line scheduling of jobs with fixed start and end times. *Theor. Comput. Sci.*, 130(1):5–16, 1994.

- H. C. Zhao, A. Lall, M. Ogihara, O. Spatscheck, J. Wang, and J. J. Xu. A data streaming algorithm for estimating entropies of od flows. In C. Dovrolis and M. Roughan, editors, *Proceedings of the 7th ACM SIGCOMM Internet Measurement Conference, IMC 2007, San Diego, California, USA, October 24-26, 2007*, pages 279–290. ACM, 2007.

APPENDIX

A. LIFTING THE DISTINCT ENDPOINTS ASSUMPTION

Recall that our analysis assumes that all the intervals in the stream S are closed and that their endpoints are distinct. In this section we show that these assumptions can be lifted. A quick glance at our algorithm reveals that it is essentially *comparison-based*, namely, it can be implemented via a *comparison oracle* $\mathcal{C} : \mathbb{R}^2 \rightarrow \{-1, 0, +1\}$ without accessing the interval's endpoints in any other way; given two endpoints p, q of intervals in S , the comparison oracle returns

$$\mathcal{C}(p, q) = \begin{cases} -1 & \text{if } p < q \\ 0 & \text{if } p = q \\ +1 & \text{if } p > q. \end{cases}$$

The assumption that all endpoints are distinct means that the algorithm and its analysis rely on a comparison oracle $\mathcal{C}' : \mathbb{R}^2 \rightarrow \{-1, 0, +1\}$ with the additional guarantee that $\mathcal{C}'(p, q) \neq 0$ whenever $p \neq q$. We shall refer to such a comparison oracle \mathcal{C}' as a *distinct-endpoints comparison oracle*.

We show that for every stream S of intervals (the endpoints of these intervals may be arbitrarily open or closed) associated with a comparison oracle \mathcal{C} , there exists a distinct-endpoints comparison oracle \mathcal{C}' such that for every two intervals $I, J \in S$, the closure of I and the closure of J intersect under \mathcal{C}' if and only if I and J intersect under \mathcal{C} . Moreover, given an access to the comparison oracle \mathcal{C} , the distinct-endpoints comparison oracle \mathcal{C}' can be implemented under our streaming model's space requirements.

The distinct-endpoints comparison oracle \mathcal{C}' is designed as follows. Consider an endpoint p of an interval $I \in S$ and an endpoint q of an interval $J \in S$, $I \neq J$. If $\mathcal{C}(p, q) \neq 0$, then we set $\mathcal{C}'(p, q) = \mathcal{C}(p, q)$, so assume hereafter that $\mathcal{C}(p, q) = 0$. Consider first the case in which p is a right endpoint and q is a left endpoint (the converse case is analogous). If at least one of the endpoints is open, then set $\mathcal{C}'(p, q) = -1$; otherwise (both endpoints are closed), set $\mathcal{C}'(p, q) = +1$.

Now, consider the case in which both p and q are left endpoints (the converse case is analogous). If p is open and q is closed, then set $\mathcal{C}'(p, q) = +1$; if p is closed and q is open, then set $\mathcal{C}'(p, q) = -1$; if both p and q are open or both are closed, then we set

$$\mathcal{C}'(p, q) = \begin{cases} +1 & \text{if } I \text{ (the interval of } p) \text{ arrived before } J \text{ (the interval of } J) \\ -1 & \text{if } I \text{ (the interval of } p) \text{ arrived after } J \text{ (the interval of } J). \end{cases}$$

It is easy to verify that the closures of every two intervals intersect under \mathcal{C}' if and only if the intervals themselves intersect under \mathcal{C} . Therefore, it remains to show that \mathcal{C}' can be implemented in the streaming model. Apart from an access to the original comparison oracle \mathcal{C} , the implementation of $\mathcal{C}'(p, q)$ is based on: (1) knowing for each endpoint whether it is a left endpoint or a right endpoint; (2) knowing for each endpoint whether it is open or closed; and (3) knowing the order of arrival of intervals that share a left (respectively, right) endpoint. The first two requirements are clearly satisfied by the information provided in the input. For the third requirement, we note that if two intervals share a left (resp., right) endpoint p , then they must intersect. Thus, Lemma 3.4(P5) and Lemma 3.4(P6) guarantee that at any given time, our algorithm maintains $O(1)$ intervals that have p as their left (resp., right) endpoint. A data structure that tracks the arrival order of these intervals can therefore be implemented with $O(1)$ additional bits per interval.

B. PROOF OF LEMMA 5.3

Let n be sufficiently large so that $n(1 + \log(e)) \leq \alpha n \log(n)$. Suppose toward a contradiction that there exist two functions $\hat{f} : P_n \rightarrow \{0, 1\}^m$ and $\hat{g} : \{0, 1\}^m \times [n] \rightarrow [n]$ such that $\mathbb{P}(\hat{g}(\hat{f}(\pi), i) = \pi(i)) \geq 2\alpha$. We shall use these functions to construct a uniquely decodable coding scheme $s : P_n \rightarrow \{0, 1\}^*$ so that $\mathbb{E}_{\pi \in_r P_n} [|s(\pi)|] < \log(n!)$. This contradicts Shannon's source coding theorem as the entropy of choosing π uniformly at random from P_n is $\log(n!)$.

In order to construct the coding scheme, we first define the vector $v_\pi \in \{0, 1\}^n$ for every $\pi \in P_n$ by setting $v_\pi(i) = 1$ if $\hat{g}(\hat{f}(\pi), i) = \pi(i)$; and $v_\pi(i) = 0$ otherwise. Let $W_\pi = \{i \in [n] \mid v_\pi(i) = 0\}$. The coding scheme s is now defined by setting the codeword of each $\pi \in P_n$ to be

$$s(\pi) = v_\pi \circ \hat{f}(\pi) \circ \langle \pi(i) \rangle_{i \in W_\pi},$$

where $\bigcirc_{i \in W_\pi} \pi(i)$ denotes a concatenation of the standard binary representations of $\pi(i)$ for all $i \in W_\pi$ listed in increasing order of the index i .

We first argue that s is indeed a uniquely decodable code. To that end, notice that for every $\pi \in P_n$ and for every $i \in [n]$, we can extract the value of $\pi(i)$ from $s(\pi)$ as follows:

- (1) Check in v_π if the correct value of $\pi(i)$ can be extracted from $\hat{f}(\pi)$, that is, if $v_\pi(i) = 1$.
- (2) If it can ($v_\pi(i) = 1$), then $\pi(i)$ is extracted by computing $\hat{g}(\hat{f}(\pi), i)$ (recall that $\hat{f}(\pi)$ is found in the second segment of $s(\pi)$).
- (3) Otherwise ($v_\pi(i) = 0$), $\pi(i)$ is extracted from the third segment of $s(\pi)$.

Moreover, the coding scheme s is prefix-free (and hence uniquely decodable) since $v_\pi = v_{\pi'}$ implies that $|s(\pi)| = |s(\pi')|$ for every two permutations $\pi, \pi' \in P_n$. Thus, if the codewords $s(\pi)$ and $s(\pi')$ agree on the first n bits, then they must have the same length, which means that $s(\pi)$ cannot be a proper prefix of $s(\pi')$.

It remains to show that $\mathbb{E}_{\pi \in_r P_n} [|s(\pi)|] < \log(n!)$. By definition, $|s(\pi)| = n + m + \log(n) \cdot |W_\pi|$ for every $\pi \in P_n$, so

$$\mathbb{E}_{\pi \in_r P_n} [|s(\pi)|] = n + m + \log(n) \cdot \mathbb{E}_{\pi \in_r P_n} [|W_\pi|].$$

The assumption that $\mathbb{P}_{\pi \in_r P_n, i \in_r [n]} (\hat{g}(\hat{f}(\pi), i) = \pi(i)) \geq 2\alpha$ implies that $\mathbb{P}_{\pi \in_r P_n, i \in_r [n]} (i \in W_\pi) \leq 1 - 2\alpha$, hence $\mathbb{E}_{\pi \in_r P_n} [|W_\pi|] \leq (1 - 2\alpha)n$. Plugging $m = \alpha n \log(n)$, we conclude that


$$\mathbb{E}_{\pi \in_r P_n} [|s(\pi)|] \leq n + (1 - \alpha)n \log(n).$$

By the choice of n (satisfying $n(1 + \log(e)) \leq \alpha n \log(n)$), we derive the desired inequality since $\log(n!) > n \log(n) - n \log(e)$. The assertion follows.



C. FIGURES ILLUSTRATING TABLE I

In the following figures, the solid lines depict actual intervals, the dashed two-sided arrows depict the virtual intervals, and the gray rectangle depicts the newcoming interval I . The top part illustrates the situation in φ_{t-1} and the bottom part illustrates the situation in the resulting φ_t ; these two parts are separated by a dotted line.





Category C1









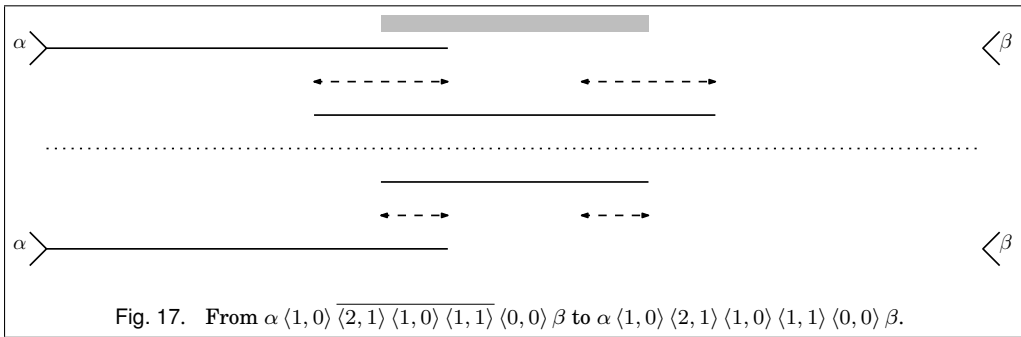

Category C2






Category C3







Category C4



Category C5

