# Interval Selection in the Streaming Model 

Sergio Cabello* Pablo Pérez-Lantero ${ }^{\dagger}$

October 2, 2018


#### Abstract

A set of intervals is independent when the intervals are pairwise disjoint. In the interval selection problem we are given a set $\mathbb{I}$ of intervals and we want to find an independent subset of intervals of largest cardinality. Let $\alpha(\mathbb{I})$ denote the cardinality of an optimal solution. We discuss the estimation of $\alpha(\mathbb{I})$ in the streaming model, where we only have one-time, sequential access to the input intervals, the endpoints of the intervals lie in $\{1, \ldots, n\}$, and the amount of the memory is constrained.

For intervals of different sizes, we provide an algorithm in the data stream model that computes an estimate $\hat{\alpha}$ of $\alpha(\mathbb{I})$ that, with probability at least $2 / 3$, satisfies $\frac{1}{2}(1-\varepsilon) \alpha(\mathbb{I}) \leq$ $\hat{\alpha} \leq \alpha(\mathbb{I})$. For same-length intervals, we provide another algorithm in the data stream model that computes an estimate $\hat{\alpha}$ of $\alpha(\mathbb{I})$ that, with probability at least $2 / 3$, satisfies $\frac{2}{3}(1-\varepsilon) \alpha(\mathbb{I}) \leq \hat{\alpha} \leq \alpha(\mathbb{I})$. The space used by our algorithms is bounded by a polynomial in $\varepsilon^{-1}$ and $\log n$. We also show that no better estimations can be achieved using $o(n)$ bits of storage.

We also develop new, approximate solutions to the interval selection problem, where we want to report a feasible solution, that use $O(\alpha(\mathbb{I}))$ space. Our algorithms for the interval selection problem match the optimal results by Emek, Halldórsson and Rosén [Space-Constrained Interval Selection, ICALP 2012], but are much simpler.


## 1 Introduction

Several fundamental problems have been explored in the data streaming model; see [3, 16 for an overview. In this model we have bounds on the amount of available memory, the data arrives sequentially, and we cannot afford to look at input data of the past, unless it was stored in our limited memory. This is effectively equivalent to assuming that we can only make one pass over the input data.

In this paper, we consider the interval selection problem. Let us say that a set of intervals is independent when all the intervals are pairwise disjoint. In the interval selection problem, the input is a set $\mathbb{I}$ of intervals and we want to find an independent subset of largest cardinality. Let us denote by $\alpha(\mathbb{I})$ this largest cardinality. There are actually two different problems: one problem is finding (or approximating) a largest independent subset, while the other problem is estimating $\alpha(\mathbb{I})$. In this paper we consider both problems in the data streaming model.

There are many natural reasons to consider the interval selection problem in the data streaming model. Firstly, the interval selection problem appears in many different contexts and several extensions have been studied; see for example the survey 14 .

[^0]Secondly, the interval selection problem is a natural generalization of the distinct elements problem: given a data stream of numbers, identify how many distinct numbers appeared in the stream. The distinct elements problem has a long tradition in data streams; see Kane, Nelson and Woodruff 12 for an optimal algorithm and references therein for a historical perspective.

Thirdly, there has been interest in understanding graph problems in the data stream model. However, several problems cannot be solved within the memory constraints usually considered in the data stream model. This leads to the introduction by Feigenbaum et al. 77 of the semi-streaming model, where the available memory is $O\left(|V| \log ^{O(1)}|V|\right)$, being $V$ the vertex set of the corresponding graph. Another model closely related to preemptive online algorithms was considered by Halldórsson et al. [8]: there is an output buffer where a feasible solution is always maintained.

Finally, geometrically-defined graphs provide a rich family of graphs where certain graph problems may be solved within the traditional model. We advocate that graph problems should be considered for geometrically-defined graphs in the data stream model. The interval selection problem is one such case, since it is exactly finding a largest independent set in the intersection graph of the input intervals.

Previous works. Emek, Halldórsson and Rosén [6] consider the interval selection problem with $O(\alpha(\mathbb{I}))$ space. They provide a 2 -approximation algorithm for the case of arbitrary intervals and a (3/2)-approximation for the case of proper intervals, that is, when no interval contains another interval. Most importantly, they show that no better approximation factor can be achieved with sublinear space. Since any $O(1)$-approximation obviously requires $\Omega(\alpha(\mathbb{I}))$ space, their algorithms are optimal. They do not consider the problem of estimating $\alpha(\mathbb{I})$. Halldórsson et al. 8 consider maximum independent set in the aforementioned online streaming model. As mentioned before, estimating $\alpha(\mathbb{I})$ is a generalization of the distinct elements problems. See Kane, Nelson and Woodruff (12] and references therein.

Our contributions. We consider both the estimation of $\alpha(\mathbb{I})$ and the interval selection problem, where a feasible solution must be produced, in the data streaming model. We next summarize our results and put them in context.
(a) We provide a 2-approximation algorithm for the interval selection problem using $O(\alpha(\mathbb{I}))$ space. Our algorithm has the same space bounds and approximation factor than the algorithm by Emek, Halldórsson and Rosén [6], and thus is also optimal. However, our algorithm is considerably easier to explain, analyze and understand. Actually, the analysis of our algorithm is nearly trivial. This result is explained in Section 3 .
(b) We provide an algorithm to obtain a value $\hat{\alpha}(\mathbb{I})$ such that $\frac{1}{2}(1-\varepsilon) \alpha(\mathbb{I}) \leq \hat{\alpha}(\mathbb{I}) \leq \alpha(\mathbb{I})$ with probability at least $2 / 3$. The algorithm uses $O\left(\varepsilon^{-5} \log ^{6} n\right)$ space for intervals with endpoints in $\{1, \ldots, n\}$. As a black-box subroutine we use a 2 -approximation algorithm for the interval selection problem. This result is explained in Section 4.
(c) For same-length intervals we provide a (3/2)-approximation algorithm for the interval selection problem using $O(\alpha(\mathbb{I}))$ space. Again, Emek, Halldórsson and Rosén [6] provide an algorithm with the same guarantees and give a lower bound showing that the algorithm is optimal. We believe that our algorithm is simpler, but this case is more disputable. This result is explained in Section 5
(d) For same-length intervals with endpoints in $\{1, \ldots, n\}$, we show how to find in $O\left(\varepsilon^{-2} \log (1 / \varepsilon)+\right.$ $\log n)$ space an estimate $\hat{\alpha}(\mathbb{I})$ such that $\frac{2}{3}(1-\varepsilon) \alpha(\mathbb{I}) \leq \hat{\alpha}(\mathbb{I}) \leq \alpha(\mathbb{I})$ with probability at
least $2 / 3$. This algorithm is an adaptation of the new algorithm in (c). This result is explained in Section 6 .
(e) We provide lower bounds showing that the approximation ratios in (b) and (d) are essentially optimal, if we use $o(n)$ space. Note that the lower bounds of Emek, Halldórsson and Rosén [6] hold for the interval selection problem but not for the estimation of $\alpha(\mathbb{I})$. We employ a reduction from the one-way randomized communication complexity of INDEX. Details appear in Section 7 .

The results in (a) and (c) work in a comparison-based model and we assume that a unit of memory can store an interval. The results in (b) and (d) are based on hash functions and we assume that a unit of memory can store values in $\{1, \ldots, n\}$. Assuming that the input data, in our case the endpoints of the intervals, is from $\{1, \ldots, n\}$ is common in the data streaming model. The lower bounds of (e) are stated at bit level.

It is important to note that estimating $\alpha(\mathbb{I})$ requires considerably less space than computing an actual feasible solution with $\Theta(\alpha(\mathbb{I}))$ intervals. While our results in (a) and (c) are a simplification of the work of Emek et al., the results in (b) and (d) were unknown before.

As usual, the probability of success can be increased to $1-\delta$ using $O(\log (1 / \delta))$ parallel repetitions of the algorithm and choosing the median of values computed in each repetition.

## 2 Preliminaries

We assume that the input intervals are closed. Our algorithms can be easily adapted to handle inputs that contain intervals of mixed types: some open, some closed, and some half-open.

We will use the term 'interval' only for the input intervals. We will use the term 'window' for intervals constructed through the algorithm and 'segment' for intervals associated with the nodes of a segment tree. (This segment tree is explained later on.) The windows we consider may be of any type regarding the inclusion of endpoints.

For each natural number $n$, we let $[n]$ be the integer range $\{1, \ldots, n\}$. We assume that $0<\varepsilon<1 / 2$.

### 2.1 Leftmost and rightmost interval

Consider a window $W$ and a set of intervals $\mathbb{I}$. We associate to $W$ two input intervals.

- The interval Leftmost $(W)$ is, among the intervals of $\mathbb{I}$ contained in $W$, the one with smallest right endpoint. If there are many candidates with the same right endpoint, Leftmost $(W)$ is one with largest left endpoint.
- The interval Rightmost $(W)$ is, among the intervals of $\mathbb{I}$ contained in $W$, the one with largest left endpoint. If there are many candidates with the same left endpoint, $\operatorname{Rightmost}(W)$ is one with smallest right endpoint.

When $W$ does not contain any interval of $\mathbb{I}$, then $\operatorname{Leftmost}(W)$ and $\operatorname{Rightmost}(W)$ are undefined. When $W$ contains a unique interval $I \in \mathbb{I}$, we have $\operatorname{Leftmost}(W)=\operatorname{Rightmost}(W)=$ $I$. Note that the intersection of all intervals contained in $W$ is precisely $\operatorname{Leftmost}(W) \cap$ Rightmost $(W)$.

In fact, we will consider $\operatorname{Leftmost}(W)$ and $\operatorname{Rightmost}(W)$ with respect to the portion of the stream that has been treated. We relax the notation by omitting the reference to $\mathbb{I}$ or the portion of the stream we have processed. It will be clear from the context with respect to which set of intervals we are considering $\operatorname{Leftmost}(W)$ and $\operatorname{Rightmost}(W)$.

### 2.2 Sampling

We next describe a tool for sampling elements from a stream. A family of permutations $\mathcal{H}=\{h:[n] \rightarrow[n]\}$ is $\varepsilon$-min-wise independent if

$$
\forall X \subset[n] \text { and } \forall y \in X: \quad \frac{1-\varepsilon}{|X|} \leq \operatorname{Pr}_{h \in \mathcal{H}}[h(y)=\min h(X)] \leq \frac{1+\varepsilon}{|X|}
$$

Here, $h \in \mathcal{H}$ is chosen uniformly at random. The family of all permutations is 0 -min-wise independent. However, there is no compact way to specify an arbitrary permutation. As discussed by Broder, Charikar and Mitzenmacher [2], the results of Indyk [10] can be used to construct a compact, computable family of permutations that is $\varepsilon$-min-wise independent. See [1,4,5] for other uses of $\varepsilon$-min-wise independent permutations.

Lemma 1. For every $\varepsilon \in(0,1 / 2)$ and $n>0$ there exists a family of permutations $\mathcal{H}(n, \varepsilon)=\{h$ : $[n] \rightarrow[n]\}$ with the following properties: (i) $\mathcal{H}(n, \varepsilon)$ has $n^{O(\log (1 / \varepsilon))}$ permutations; (ii) $\mathcal{H}(n, \varepsilon)$ is $\varepsilon$-min-wise independent; (iii) an element of $\mathcal{H}(n, \varepsilon)$ can be chosen uniformly at random in $O(\log (1 / \varepsilon))$ time; (iv) for $h \in \mathcal{H}(n, \varepsilon)$ and $x, y \in[n]$, we can decide with $O(\log (1 / \varepsilon))$ arithmetic operations whether $h(x)<h(y)$.

Proof. Indyk [10] showed that there exist constants $c_{1}, c_{2}>1$ such that, for any $\varepsilon>0$ and any family $\mathcal{H}^{\prime}$ of $c_{2} \log (1 / \varepsilon)$-wise independent hash functions $[m] \rightarrow[m]$, it holds the following:

$$
\begin{aligned}
\forall X \subset[m] \text { with }|X| \leq \frac{\varepsilon m}{c_{1}} \text { and } \forall y & \in[m] \backslash X: \\
& \frac{1-\varepsilon}{|X|+1} \leq \operatorname{Pr}_{h^{\prime} \in \mathcal{H}^{\prime}}\left[h^{\prime}(y)<\min h^{\prime}(X)\right] \leq \frac{1+\varepsilon}{|X|+1} .
\end{aligned}
$$

Set $m=c_{1} n / \varepsilon>n$ and let $\mathcal{H}^{\prime}=\left\{h^{\prime}:[m] \rightarrow[m]\right\}$ be a family of $c_{2} \log (1 / \varepsilon)$-wise independent hash functions. Since $n=\varepsilon m / c_{1}$, the result of Indyk implies that

$$
\forall X \subset[n] \text { and } \forall y \in[n] \backslash X: \quad \frac{1-\varepsilon}{|X|+1} \leq \operatorname{Pr}_{h^{\prime} \in \mathcal{H}^{\prime}}\left[h^{\prime}(y)<\min h^{\prime}(X)\right] \leq \frac{1+\varepsilon}{|X|+1}
$$

Each hash function $h^{\prime} \in \mathcal{H}^{\prime}$ can be used to create a permutation $\widehat{h^{\prime}}:[n] \rightarrow[n]$ : define $\widehat{h^{\prime}}(i)$ as the position of $\left(h^{\prime}(i), i\right)$ in the lexicographic order of $\left\{\left(h^{\prime}(i), i\right) \mid i \in[n]\right\}$. Consider the set of permutations $\widehat{\mathcal{H}^{\prime}}=\left\{\widehat{h^{\prime}}:[n] \rightarrow[n] \mid h^{\prime} \in \mathcal{H}^{\prime}\right\}$. For each $X \subset[n]$ and $y \in[n] \backslash X$ we have

$$
\frac{1-\varepsilon}{|X|+1} \leq \operatorname{Pr}_{h^{\prime} \in \mathcal{H}^{\prime}}\left[h^{\prime}(y)<\min h^{\prime}(X)\right] \leq \underset{\widehat{h^{\prime}} \in \widehat{\mathcal{H}^{\prime}}}{\operatorname{Pr}}\left[\widehat{h^{\prime}}(y)<\min \widehat{h^{\prime}}(X)\right]
$$

and

$$
\begin{aligned}
\widehat{h^{\prime} \in \widehat{\mathcal{H}^{\prime}}} \operatorname{Pr}\left[\widehat{h^{\prime}}(y)<\min \widehat{h^{\prime}}(X)\right] & \leq \operatorname{Pr}_{h^{\prime} \in \mathcal{H}^{\prime}}\left[h^{\prime}(y) \leq \min h^{\prime}(X)\right] \\
& \leq \operatorname{Pr}_{h^{\prime} \in \mathcal{H}^{\prime}}\left[h^{\prime}(y)<\min h^{\prime}(X)\right]+\operatorname{Pr}_{h^{\prime} \in \mathcal{H}^{\prime}}\left[h^{\prime}(y)=\min h^{\prime}(X)\right] \\
& \leq \frac{1+\varepsilon}{|X|+1}+\frac{1}{m} \leq \frac{1+2 \varepsilon}{|X|+1},
\end{aligned}
$$

where we have used that $h^{\prime}(y)=\min h^{\prime}(X)$ corresponds to a collision and $m>n / \varepsilon$. We can rewrite this as

$$
\forall X \subset[n] \text { and } \forall y \in X: \quad \frac{1-\varepsilon}{|X|} \leq \widehat{\operatorname{h}^{\prime} \in \widehat{\mathcal{H}^{\prime}}} \underset{ }{ }\left[\widehat{h^{\prime}}(y)=\min \widehat{h^{\prime}}(X)\right] \leq \frac{1+2 \varepsilon}{|X|}
$$

Using $\varepsilon / 2$ instead of $\varepsilon$ in the discussion, the lower bound becomes $(1-\varepsilon / 2) /|X| \geq(1-\varepsilon) /|X|$ and the upper bound becomes $(1+\varepsilon) /|X|$, as desired. Standard constructions using polynomials over finite fields can be used to construct a family $\mathcal{H}^{\prime}=\left\{h^{\prime}:[m] \rightarrow[m]\right\}$ of $c_{2} \log (1 / \varepsilon)$-wise independent hash functions such that: $\mathcal{H}^{\prime}$ has $n^{O(\log (1 / \varepsilon))}$ hash functions; an element of $\mathcal{H}^{\prime}$ can be chosen uniformly at random in $O(\log (1 / \varepsilon))$ time; for $h^{\prime} \in \mathcal{H}^{\prime}$ and $x \in[n]$ we can compute $h^{\prime}(x)$ using $O(\log (1 / \varepsilon))$ arithmetic operations.

This gives an implicit description of our desired set of permutations $\widehat{\mathcal{H}^{\prime}}$ satisfying (i)-(iii). Moreover, while computing $\widehat{h^{\prime}}(x)$ for $h^{\prime} \in \mathcal{H}^{\prime}$ is demanding, we can easily decide whether $\widehat{h^{\prime}}(x)<\widehat{h^{\prime}}(y)$ by computing and comparing $\left(h^{\prime}(x), x\right)$ and $\left(h^{\prime}(y), y\right)$.

Let us explain now how to use Lemma 1 to make a (nearly-uniform) random sample. We learned this idea from Datar and Muthukrishnan [5]. Consider any fixed subset $X \subset[n]$ and let $\mathcal{H}=\mathcal{H}(n, \varepsilon)$ be the family of permutations given in Lemma 1 . An $\mathcal{H}$-random element $s$ of $X$ is obtained by choosing a hash function $h \in \mathcal{H}$ uniformly at random, and setting $s=\arg \min \{h(t) \mid t \in X\}$. It is important to note that $s$ is not chosen uniformly at random from $X$. However, from the definition of $\varepsilon$-min-wise independence we have

$$
\forall x \in X: \quad \frac{1-\varepsilon}{|X|} \leq \operatorname{Pr}[s=x] \leq \frac{1+\varepsilon}{|X|}
$$

In particular, we obtain the following

$$
\forall Y \subset X: \quad \frac{(1-\varepsilon)|Y|}{|X|} \leq \operatorname{Pr}[s \in Y] \leq \frac{(1+\varepsilon)|Y|}{|X|}
$$

This means that, for a fixed $Y$, we can estimate the ratio $|Y| /|X|$ using $\mathcal{H}$-random samples from $X$ repeatedly, and counting how many belong to $Y$.

Using $\mathcal{H}$-random samples has two advantages for data streams with elements from $[n]$. Through the stream, we can maintain an $\mathcal{H}$-random sample $s$ of the elements seen so far. For this, we select $h \in \mathcal{H}$ uniformly at random, and, for each new element $a$ of the stream, we check whether $h(a)<h(s)$ and update $s$, if needed. An important feature of sampling in such way is that $s$ is almost uniformly at random among those appearing in the stream, without counting multiplicities. The other important feature is that we select $s$ at its first appearance in the stream. Thus, we can carry out any computation that depends on $s$ and on the portion of the stream after its first appearance. For example, we can count how many times the $\mathcal{H}$-random element $s$ appears in the whole data stream.

We will also use $\mathcal{H}$ to make conditional sampling: we select $\mathcal{H}$-random samples until we get one satisfying a certain property. To analyze such technique, the following result will be useful.
Lemma 2. Let $Y \subset X \subset[n]$ and assume that $0<\varepsilon<1 / 2$. Consider the family of permutations $\mathcal{H}=\mathcal{H}(n, \varepsilon)$ from Lemma 1, and take a $\mathcal{H}$-random sample s from $X$. Then

$$
\forall y \in Y: \quad \frac{1-4 \varepsilon}{|Y|} \leq \operatorname{Pr}[s=y \mid s \in Y] \leq \frac{1+4 \varepsilon}{|Y|}
$$

Proof. Consider any $y \in Y$. Since $s=y$ implies $s \in Y$, we have

$$
\operatorname{Pr}[s=y \mid s \in Y]=\frac{\operatorname{Pr}[s=y]}{\operatorname{Pr}[s \in Y]} \leq \frac{\frac{1+\varepsilon}{|X|}}{\frac{(1-\varepsilon)|Y|}{|X|}}=\frac{1+\varepsilon}{1-\varepsilon} \cdot \frac{1}{|Y|} \leq(1+4 \varepsilon) \frac{1}{|Y|},
$$

where in the last inequality we used that $\varepsilon<1 / 2$. Similarly, we have

$$
\operatorname{Pr}[s=y \mid s \in Y] \geq \frac{\frac{1-\varepsilon}{|X|}}{\frac{(1+\varepsilon)|Y|}{|X|}}=\frac{1-\varepsilon}{1+\varepsilon} \cdot \frac{1}{|Y|} \geq(1-4 \varepsilon) \frac{1}{|Y|}
$$

which completes the proof


Figure 1: At the bottom there is a partition of the real line. Filled-in disks indicate that the endpoint is part of the interval; empty disks indicate that the endpoint is not part of the interval. At the top, we show the split of some optimal solution $\mathbb{J}^{*}$ (dotted blue) into $\mathbb{J}_{\subset}^{*}$ and $\mathbb{J}_{\cap}^{*}$.

## 3 Largest independent subset of intervals

In this section we show how to obtain a 2 -approximation to the largest independent subset of $\mathbb{I}$ using $O(\alpha(\mathbb{I}))$ space.

A set $\mathbb{W}$ of windows is a partition of the real line if the windows in $\mathbb{W}$ are pairwise disjoint and their union is the whole $\mathbb{R}$. The windows in $\mathbb{W}$ may be of different types regarding the inclusion of endpoints. See Figure 1 for an example.

Lemma 3. Let $\mathbb{I}$ be a set of intervals and let $\mathbb{W}$ be a partition of the real line with the following properties:

- Each window of $\mathbb{W}$ contains at least one interval from $\mathbb{I}$.
- For each window $W \in \mathbb{W}$, the intervals of $\mathbb{I}$ contained in $W$ pairwise intersect.

Let $\mathbb{d}$ be any set of intervals constructed by selecting for each window $W$ of $\mathbb{W}$ an interval of $\mathbb{I}$ contained in $W$. Then $|\mathbb{J}|>\alpha(\mathbb{I}) / 2$.

Proof. Let us set $k=|\mathbb{W}|$. Consider a largest independent set of intervals $\mathbb{J}^{*} \subseteq \mathbb{I}$. We have $\left|\mathbb{J}^{*}\right|=\alpha(\mathbb{I})$. Split the set $\mathbb{J}^{*}$ into two sets $\mathbb{J}_{\subset}^{*}$ and $\mathbb{J}_{\cap}^{*}$ as follows: $\mathbb{J}_{\subset}^{*}$ contains the intervals contained in some window of $\mathbb{W}$ and $\mathbb{J}_{\cap}^{*}$ contains the other intervals. See Figure 1 for an example. The intervals in $\mathbb{J}_{\cap}^{*}$ are pairwise disjoint and each of them intersects at least two consecutive windows from $\mathbb{W}$, thus $\left|\mathbb{J}_{\cap}^{*}\right| \leq k-1$. Since all intervals contained in a window of $\mathbb{W}$ pairwise intersect, $\mathbb{J}_{\subset}^{*}$ has at most one interval per window. Thus $\left|\mathbb{J}_{\subset}^{*}\right| \leq k$. Putting things together we have

$$
\alpha(\mathbb{I})=\left|\mathbb{J}^{*}\right|=\left|\mathbb{J}_{\cap}^{*}\right|+\left|\mathbb{J}_{\subset}^{*}\right| \leq k-1+k=2 k-1 .
$$

Since $\mathbb{J}$ contains exactly $|\mathbb{W}|=k$ intervals, we obtain

$$
2 \cdot|\mathbb{J}|=2 k>2 k-1 \geq \alpha(\mathbb{I})
$$

which implies the result.
We now discuss the algorithm. Through the processing of the stream, we maintain a partition $\mathbb{W}$ of the line so that $\mathbb{W}$ satisfies the hypothesis of Lemma 3. To carry this out, for each window $W$ of $\mathbb{W}$ we store the intervals $\operatorname{Leftmost}(W)$ and $\operatorname{Rightmost}(W)$. See Figure 2 for an example. To initialize the structures, we start with a unique window $\mathbb{W}=\{\mathbb{R}\}$ and set $\operatorname{Leftmost}(W)=\operatorname{Rightmost}(W)=I_{0}$, where $I_{0}$ is the first interval of the stream. With such initialization, the hypothesis of Lemma 3 hold and $\operatorname{Leftmost}()$ and $\operatorname{Rightmost}()$ have the correct values.


Figure 2: Data maintained by the algorithm.
With a few local operations, we can handle the insertion of a new interval in $\mathbb{I}$. Consider a new interval $I$ of the stream. If $I$ is not contained in any window of $\mathbb{W}$, we do not need to do anything. If $I$ is contained in a window $W$, we check whether it intersects all intervals contained $W$. If $I$ intersects all intervals contained in $W$, we may have to update $\operatorname{Leftmost}(W)$ and $\operatorname{Rightmost}(W)$. If $I$ does not intersect all the intervals contained in $W$, then $I$ is disjoint from $\operatorname{Leftmost}(W) \cap \operatorname{Rightmost}(W)$. In such a case we can use one endpoint of $\operatorname{Leftmost}(W) \cap$ $\operatorname{Rightmost}(W)$ to split the window $W$ into two windows $W_{1}$ and $W_{2}$, one containing $I$ and the other containing either $\operatorname{Leftmost}(W)$ or $\operatorname{Rightmost}(W)$, so that the assumptions of Lemma 3 are restored. We also have enough information to obtain $\operatorname{Leftmost}()$ and $\operatorname{Rightmost}()$ for the new windows $W_{1}$ and $W_{2}$. Figures 4 and 5 show two possible scenarios. See the pseudocode in Figure 3 for a more detailed description.

Lemma 4. The policy described in Figure 3 maintains a set of windows $\mathbb{W}$ and a set of intervals $\mathbb{J}$ that satisfy the assumptions of Lemma 3.

Proof. A simple case analysis shows that the policy maintains the assumptions of Lemma 3 and the properties of Leftmost () and Rightmost ().

Consider, for example, the case when the new interval $[x, y]$ is contained in a window $W \in \mathbb{W}$ and $[\ell, r]=\operatorname{Leftmost}(W) \cap \operatorname{Rightmost}(W)$ is to the left of $[x, y]$. In this case, the algorithm will update the structures in lines 11-16 and lines 24-27. See Figure 5 for an example. By inductive hypothesis, all the intervals in $\mathbb{I} \backslash\{[x, y]\}$ contained in $W$ intersect $[\ell, r]$. Note that $W_{1}=W \cap(-\infty, r]$, and thus only the intervals contained in $W$ with right endpoint $r$ are contained in $W_{1}$. By the inductive hypothesis, $\operatorname{Leftmost}(W)$ has right endpoint $r$ and has largest left endpoint among all intervals contained in $W_{1}$. Thus, when we set $\operatorname{Rightmost}\left(W_{1}\right)=\operatorname{Leftmost}\left(W_{1}\right)=\operatorname{Leftmost}(W)$, the correct values for $W_{1}$ are set. As for $W_{2}=W \cap(r,+\infty)$, no interval of $\mathbb{I} \backslash\{[x, y]\}$ is contained in $W_{2}$, thus $[x, y]$ is the only interval contained in $W_{2}$ and setting $\operatorname{Rightmost}\left(W_{2}\right)=\operatorname{Leftmost}\left(W_{2}\right)=[x, y]$ we get the correct values for $W_{2}$. Lines 24-27 take care to replace $W$ in $\mathbb{W}$ by $W_{1}$ and $W_{2}$. For $W_{1}$ and $W_{2}$ we set the correct values of $\operatorname{Leftmost}()$ and $\operatorname{Rightmost}()$ and the assumptions of Lemma 3 hold. For the other windows of $\mathbb{W} \backslash\{W\}$ nothing is changed.

We can store the partition of the real line $\mathbb{W}$ using a dynamic binary search tree. With this, line 1 and lines $24-25$ take $O(\log |\mathbb{W}|)=O(\log \alpha(\mathbb{I}))$ time. The remaining steps take constant time. The space required by the data structure is $O(|\mathbb{W}|)=O(\alpha(\mathbb{I}))$. This shows the following result.

Theorem 5. Let $\mathbb{I}$ be a set of intervals in the real line that arrive in a data stream. There is a data stream algorithm to compute a 2 -approximation to the largest independent subset of $\mathbb{I}$ that uses $O(\alpha(\mathbb{I}))$ space and handles each interval of the stream in $O(\log \alpha(\mathbb{I}))$ time.

Process interval $I=[x, y]$

1. find the window $W$ of $\mathbb{W}$ that contains $x$
2. $\quad[\ell, r] \leftarrow \operatorname{Leftmost}(W) \cap \operatorname{Rightmost}(W)$
3. if $y \in W$ then
4. if $[\ell, r] \cap[x, y] \neq \emptyset$ then
5. if $\ell<x$ or $(\ell=x$ and $[x, y] \subset \operatorname{Rightmost}(W))$ then
6. $\quad \operatorname{Rightmost}(W) \leftarrow[x, y]$
7. $\quad$ if $y<r$ or $(y=r$ and $[x, y] \subset \operatorname{Leftmost}(W))$ then
$\operatorname{Leftmost}(W) \leftarrow[x, y]$
else $\quad\left(*[\ell, r]\right.$ and $[x, y]$ are disjoint; split $\left.W^{*}\right)$ if $x>r$ then $\quad\left(*[\ell, r]\right.$ to the left of $\left.[x, y]^{*}\right)$
make new windows $W_{1}=W \cap(-\infty, r]$ and $W_{2}=W \cap(r,+\infty)$
$\operatorname{Leftmost}\left(W_{1}\right) \leftarrow \operatorname{Leftmost}(W)$
$\operatorname{Rightmost}\left(W_{1}\right) \leftarrow \operatorname{Leftmost}(W)$
$I^{\prime} \leftarrow \operatorname{Leftmost}(W)$
$\operatorname{Leftmost}\left(W_{2}\right) \leftarrow[x, y]$
$\operatorname{Rightmost}\left(W_{2}\right) \leftarrow[x, y]$
else $\quad\left(* y<\ell,[\ell, r]\right.$ to the right of $\left.[x, y]^{*}\right)$
make new windows $W_{1}=W \cap(-\infty, \ell)$ and $W_{2}=W \cap[\ell,+\infty)$
$\operatorname{Leftmost}\left(W_{1}\right) \leftarrow[x, y]$
$\operatorname{Rightmost}\left(W_{1}\right) \leftarrow[x, y]$
$\operatorname{Leftmost}\left(W_{2}\right) \leftarrow \operatorname{Rightmost}(W)$
$\operatorname{Rightmost}\left(W_{2}\right) \leftarrow \operatorname{Rightmost}(W)$
$I^{\prime} \leftarrow \operatorname{Rightmost}(W)$
remove $W$ from $\mathbb{W}$
add $W_{1}$ and $W_{2}$ to $\mathbb{W}$
remove from $\mathbb{J}$ the interval that is contained in $W$
add to $\mathbb{J}$ the intervals $[x, y]$ and $I^{\prime}$
8. (* If $y \notin W$ then $[x, y]$ is not contained in any window ${ }^{*}$ )

Figure 3: Policy to process a new interval $[x, y] . \mathbb{W}$ maintains a partition of the real line and $\mathbb{J}$ maintains a 2 -approximation to $\alpha(\mathbb{I})$.


Figure 4: Example handled by lines 5-6 of the algorithm. The endpoints represented by crosses may be in the window or not.


Figure 5: Example handled by lines 11-16 of the algorithm. The endpoints represented by crosses may be in the window or not.

## 4 Size of largest independent set of intervals

In this section we show how to obtain a randomized estimate of the value $\alpha(\mathbb{I})$. We will assume that the endpoints of the intervals are in $[n]$.

Using the approach of Knuth [13], the algorithm presented in Section 3 can be used to define an estimator whose expected value lies between $\alpha(\mathbb{I}) / 2$ and $\alpha(\mathbb{I})$. However, it has large variance and we cannot use it to obtain an estimate of $\alpha(\mathbb{I})$ with good guarantees. The precise idea is as follows.

The windows appearing through the algorithm of Section 3 naturally define a rooted binary tree $T$, where each node represents a window. At the root of $T$ we have the whole real line. Whenever a window $W$ is split into two windows $W^{\prime}$ and $W^{\prime \prime}$, in $T$ we have nodes for $W^{\prime}$ and $W^{\prime \prime}$ with parent $W$. The size of the output is the number of windows in the final partition, which is exactly the number of leaves in $T$. Knuth [13] shows how to obtain an unbiased estimator of the number of leaves of a tree. This estimator is obtained by choosing random root-to-leaf paths. (At each node, one can use different rules to select how the random path continues.) Unfortunately, the estimator has very large variance and cannot be used to obtain good guarantees. Easy modifications of the method do not seem to work, so we develop a different method.

Our idea is to carefully split the window $[1, n]$ into segments, and compute for each segment a 2-approximation. If each segment contains enough disjoint intervals from the input, then we do not do much error combining the results of the segments. We then have to estimate the number of segments in the partition of $[1, n]$ and the number of independent intervals in each segment. First we describe the ingredients, independent of the streaming model, and discuss their properties. Then we discuss how those ingredients can be computed in the data streaming model.

### 4.1 Segments and their associated information

Let $T$ be a balanced segment tree on the $n$ segments $[i, i+1), i \in[n]$. Each leaf of $T$ corresponds to a segment $[i, i+1)$ and the order of the leaves in $T$ agrees with the order of their corresponding intervals along the real line. Each node $v$ of $T$ has an associated segment, denoted by $S(v)$, that is the union of all segments stored at its descendants. It is easy to see that, for any interval node $v$ with children $v_{\ell}$ and $v_{r}$, the segment $S(v)$ is the disjoint union of $S\left(v_{\ell}\right)$ and $S\left(v_{r}\right)$. See Figure 6 for an example. We denote the root of $T$ by $r$. We have $S(r)=[1, n+1)$.


Figure 6: Segment tree for $n=16$.
Let $\mathbb{S}$ be the set of segments associated with all nodes of $T$. Note that $\mathbb{S}$ has $2 n-1$ elements. Each segment $S \in \mathbb{S}$ contains the left endpoint and does not contain the right endpoint.

For any segment $S \in \mathbb{S}$, where $S \neq S(r)$, let $\pi(S)$ be the "parent" segment of $S$ : this is the segment stored at the parent of $v$, where $S(v)=S$.

For any $S \in \mathbb{S}$, let $\beta(S)$ be the size of the largest independent subset of $\{I \in \mathbb{I} \mid I \subset S\}$. That is, we consider the restriction of the problem to intervals of $\mathbb{I}$ contained in $S$. Similarly, let $\hat{\beta}(S)$ be the size of a feasible solution computed for $\{I \in \mathbb{I} \mid I \subset S\}$ by the 2-approximation algorithm described in Section 3 or by the algorithm of Emek, Halldórsson and Rosén [6]. We thus have $\beta(S) \geq \hat{\beta}(S) \geq \beta(S) / 2$ for all $S \in \mathbb{S}$.

Lemma 6. Let $\mathbb{S}^{\prime} \subset \mathbb{S}$ be a set of segments with the following properties:
(i) $S(r)$ is the disjoint union of the segments in $\mathbb{S}^{\prime}$, and,
(ii) for each $S \in \mathbb{S}^{\prime}$, we have $\beta(\pi(S)) \geq 2 \varepsilon^{-1}\lceil\log n\rceil$.

Then,

$$
\alpha(\mathbb{I}) \geq \sum_{S \in \mathbb{S}^{\prime}} \hat{\beta}(S) \geq\left(\frac{1}{2}-\varepsilon\right) \alpha(\mathbb{I}) .
$$

Proof. Since the segments in $\mathbb{S}^{\prime}$ are disjoint because of hypothesis (i), we can merge the solutions giving $\beta(S)$ independent intervals, for all $S \in \mathbb{S}^{\prime}$, to obtain a feasible solution for the whole $\mathbb{I}$. We conclude that

$$
\alpha(\mathbb{I}) \geq \sum_{S \in \mathbb{S}^{\prime}} \beta(S) \geq \sum_{S \in \mathbb{S}^{\prime}} \hat{\beta}(S) .
$$

This shows the first inequality.
Let $\tilde{\mathbb{S}}$ be the set of leafmost elements in the set of parents $\left\{\pi(S) \mid S \in \mathbb{S}^{\prime}\right\}$. Thus, each $\tilde{S} \in \tilde{\mathbb{S}}$ has some child in $\mathbb{S}^{\prime}$ and no descendant in $\tilde{\mathbb{S}}$. For each $\tilde{S} \in \tilde{\mathbb{S}}$, let $\Pi_{T}(\tilde{S})$ be the path in $T$ from the root to $\tilde{S}$. By construction, for each $S \in \mathbb{S}^{\prime}$ there exists some $\tilde{S} \in \tilde{\mathbb{S}}$ such that the parent of $S$ is on $\Pi_{T}(\tilde{S})$. By assumption (ii), for each $\tilde{S} \in \tilde{\mathbb{S}}$, we have $\beta(\tilde{S}) \geq 2 \varepsilon^{-1}\lceil\log n\rceil$. Each $\tilde{S} \in \tilde{\mathbb{S}}$ is going to "pay" for the error we make in the sum at the segments whose parents belong to $\Pi_{T}(\tilde{S})$.

Let $\mathbb{J}^{*} \subset \mathbb{I}$ be an optimal solution to the interval selection problem. For each segment $S \in \mathbb{S}$, $\mathbb{J}^{*}$ has at most 2 intervals that intersect $S$ but are not contained in $S$. Therefore, for all $S \in \mathbb{S}$ we have that

$$
\begin{equation*}
\left|\left\{J \in \mathbb{J}^{*} \mid J \cap S \neq \emptyset\right\}\right| \leq\left|\left\{J \in \mathbb{J}^{*} \mid J \subset S\right\}\right|+2 \leq \beta(S)+2 . \tag{1}
\end{equation*}
$$

The segments in $\tilde{\mathbb{S}}$ are pairwise disjoint because in $T$ none is a descendant of the other. This means that we can join solutions obtained inside the segments of $\tilde{\mathbb{S}}$ into a feasible solution. Combining this with hypothesis (ii) we get

$$
\begin{equation*}
\left|\mathbb{J}^{*}\right| \geq \sum_{\tilde{S} \in \tilde{\mathbb{S}}} \beta(\tilde{S}) \geq|\tilde{\mathbb{S}}| \cdot 2 \varepsilon^{-1}\lceil\log n\rceil \tag{2}
\end{equation*}
$$

For each $\tilde{S} \in \tilde{\mathbb{S}}$, the path $\Pi_{T}(\tilde{S})$ has at most $\lceil\log n\rceil$ vertices. Since each $S \in \mathbb{S}^{\prime}$ has a parent in $\Pi_{T}(\tilde{S})$, for some $\tilde{S} \in \tilde{\mathbb{S}}$, we obtain from equation (2) that

$$
\begin{equation*}
\left|\mathbb{S}^{\prime}\right| \leq 2\lceil\log n\rceil \cdot|\tilde{\mathbb{S}}| \leq 2\lceil\log n\rceil \cdot \frac{\left|\mathbb{J}^{*}\right|}{2 \varepsilon^{-1}\lceil\log n\rceil}=\varepsilon \cdot\left|\mathbb{J}^{*}\right| \tag{3}
\end{equation*}
$$

Using that $S(r)$ is the union of the segments in $\mathbb{S}^{\prime}$ and equation (1) we obtain

$$
\begin{aligned}
\left|\mathbb{J}^{*}\right| & \leq \sum_{S \in \mathbb{S}^{\prime}}\left|\left\{J \in \mathbb{J}^{*} \mid J \cap S \neq \emptyset\right\}\right| \\
& \leq \sum_{S \in \mathbb{S}^{\prime}}(\beta(S)+2) \\
& =2 \cdot\left|\mathbb{S}^{\prime}\right|+\sum_{S \in \mathbb{S}^{\prime}} \beta(S) \\
& \leq 2 \varepsilon \cdot\left|\mathbb{J}^{*}\right|+\sum_{S \in \mathbb{S}^{\prime}} \beta(S),
\end{aligned}
$$

where in the last inequality we used equation (3). Now we use that

$$
\forall S \in \mathbb{S}: \quad 2 \cdot \hat{\beta}(S) \geq \beta(S)
$$

to conclude that

$$
\left|\mathbb{J}^{*}\right| \leq 2 \varepsilon \cdot\left|\mathbb{J}^{*}\right|+\sum_{S \in \mathbb{S}^{\prime}} \beta(S) \leq 2 \varepsilon \cdot\left|\mathbb{J}^{*}\right|+\sum_{S \in \mathbb{S}^{\prime}} 2 \cdot \hat{\beta}(S)
$$

The second inequality that we want to show follows because $\left|\mathbb{J}^{*}\right|=\alpha(\mathbb{I})$.
We would like to find a set $\mathbb{S}^{\prime}$ satisfying the hypothesis of Lemma 6. However, the definition should be local: to know whether a segment $S$ belongs to $\mathbb{S}^{\prime}$ we should use only local information around $S$. The estimator $\hat{\beta}(S)$ is not suitable. For example, it may happen that, for some segment $S \in \mathbb{S} \backslash\{S(r)\}$, we have $\hat{\beta}(\pi(S))<\hat{\beta}(S)$, which is counterintuitive and problematic. We introduce another estimate that is an $O(\log n)$-approximation but is monotone nondecreasing along paths to the root.

For each segment $S \in \mathbb{S}$ we define

$$
\gamma(S)=\mid\left\{S^{\prime} \in \mathbb{S} \mid S^{\prime} \subset S \text { and } \exists I \in \mathbb{I} \text { s.t. } I \subset S^{\prime}\right\} \mid
$$

Thus, $\gamma(S)$ is the number of segments of $\mathbb{S}$ that are contained in $S$ and contain some input interval.

Lemma 7. For all $S \in \mathbb{S}$, we have the following properties:
(i) $\gamma(S) \leq \gamma(\pi(S))$, if $S \neq S(r)$,
(ii) $\gamma(S) \leq \beta(S) \cdot\lceil\log n\rceil$,
(iii) $\gamma(S) \geq \beta(S)$, and
(iv) $\gamma(S)$ can be computed in $O(\gamma(S)$ ) space using the portion of the stream after the first interval contained in $S$.

Proof. Property (i) is obvious from the definition because any $S^{\prime}$ contained in $S$ is also contained in the parent $\pi(S)$.

For the rest of the proof, fix some $S \in \mathbb{S}$ and define

$$
\mathbb{S}^{\prime}=\left\{S^{\prime} \in \mathbb{S} \mid S^{\prime} \subset S \text { and } \exists I \in \mathbb{I} \text { s.t. } I \subset S^{\prime}\right\}
$$

Note that $\gamma(S)$ is the size of $\mathbb{S}^{\prime}$. Let $T_{S}$ the subtree of $T$ rooted at $S$.
For property (ii), note that $T_{S}$ has at most $\lceil\log n\rceil$ levels. By the pigeonhole principle, there is some level $L$ of $T_{S}$ that contains at least $\gamma(S) /\lceil\log n\rceil$ different intervals of $\mathbb{S}^{\prime}$. The segments of $\mathbb{S}^{\prime}$ contained in level $L$ are disjoint, and each of them contains some intervals of $\mathbb{I}$. Picking an interval from each $S^{\prime} \in L$, we get a subset of intervals from $\mathbb{I}$ that are pairwise disjoint, and thus $\beta(S) \geq \gamma(S) /\lceil\log n\rceil$.

For property (iii), consider an optimal solution $\mathbb{J}^{*}$ for the interval selection problem in $\{I \in \mathbb{I} \mid I \subset S\}$. Thus $\left|\mathbb{J}^{*}\right|=\beta(S)$. For each interval $J \in \mathbb{J}^{*}$, let $S(J)$ be the smallest $S \in \mathbb{S}$ that contains $J$. Then $S(J) \in \mathbb{S}^{\prime}$. Note that $J$ contains the middle point of $S(J)$, as otherwise there would be a smaller segment in $\mathbb{S}$ containing $J$. This implies that the segments $S(J), J \in \mathbb{J}^{*}$, are all distinct. (However, they are not necessarily disjoint.) We then have

$$
\gamma(S)=\left|\mathbb{S}^{\prime}\right| \geq\left|\left\{S(J) \mid J \in \mathbb{J}^{*}\right\}\right|=\left|\mathbb{J}^{*}\right|=\beta(S)
$$

For property (iv), we store the elements of $\mathbb{S}^{\prime}$ in a binary search tree. Whenever we obtain an interval $I$, we check whether the segments contained in $S$ and containing $I$ are already in the search tree and, if needed, update the structure. The space needed in a binary search tree is proportional to the number of elements stored and thus we need $O(\gamma(S))$ space.

A segment $S$ of $\mathbb{S}, S \neq S(r)$, is relevant if

$$
\gamma(\pi(S)) \geq 2 \varepsilon^{-1}\lceil\log n\rceil^{2} \text { and } 1 \leq \gamma(S)<2 \varepsilon^{-1}\lceil\log n\rceil^{2}
$$

Let $\mathbb{S}_{r e l} \subset \mathbb{S}$ be the set of relevant segments. If $\mathbb{S}_{r e l}$ is empty, then we take $\mathbb{S}_{r e l}=\{S(r)\}$.
Because of Lemma $7(\mathrm{i}), \gamma(\cdot)$ is nondecreasing along a root-to-leaf path in $T$. Using Lemmas 6 and 7, we obtain the following.

Lemma 8. We have

$$
\alpha(\mathbb{I}) \geq \sum_{S \in \mathbb{S}_{\text {rel }}} \hat{\beta}(S) \geq\left(\frac{1}{2}-\varepsilon\right) \alpha(\mathbb{I})
$$

Proof. If $\gamma(S(r))<2 \varepsilon^{-1}\lceil\log n\rceil^{2}$, then $\mathbb{S}_{r e l}=\{S(r)\}$ and the result is clear. Thus we can assume that $\gamma(S(r)) \geq 2 \varepsilon^{-1}\lceil\log n\rceil^{2}$, which implies that $S(r) \notin \mathbb{S}_{r e l}$.

Define

$$
\mathbb{S}_{0}=\left\{S \in \mathbb{S} \backslash\{S(r)\} \mid \gamma(S)=0 \text { and } \gamma(\pi(S)) \geq 2 \varepsilon^{-1}\lceil\log n\rceil^{2}\right\}
$$

First note that the segments of $\mathbb{S}_{r e l} \cup \mathbb{S}_{0}$ form a disjoint union of $S(r)$. Indeed, for each elementary segment $[i, i+1) \in \mathbb{S}$, there exists exactly one ancestor that is either relevant or in $\mathbb{S}_{0}$. Lemma 7 (ii), the definition of relevant segment, and the fact $\gamma(S(r)) \geq 2 \varepsilon^{-1}\lceil\log n\rceil^{2}$ imply that

$$
\forall S \in \mathbb{S}_{r e l} \cup \mathbb{S}_{0}: \quad \beta(\pi(S)) \geq \gamma(\pi(S)) /\lceil\log n\rceil \geq 2 \varepsilon^{-1}\lceil\log n\rceil
$$

Therefore, the set $\mathbb{S}^{\prime}=\mathbb{S}_{r e l} \cup \mathbb{S}_{0}$ satisfies the conditions of Lemma 6sing that for all $S \in \mathbb{S}_{0}$ we have $\gamma(S)=\hat{\beta}(S)=0$, we obtain the claimed inequalities.


Figure 7: Active segments because of an interval $I$.

Let $N_{\text {rel }}$ be the number of relevant segments. A segment $S \in \mathbb{S}$ is active if $S=S(r)$ or its parent contains some input interval. See Figure 7 for an example. Let $N_{\text {act }}$ be the number of active segments in $\mathbb{S}$. We are going to estimate $N_{a c t}$, the ratio $N_{r e l} / N_{a c t}$, and the average value of $\hat{\beta}(S)$ over the relevant segments $S \in \mathbb{S}_{\text {rel }}$. With this, we will be able to estimate the sum considered in Lemma 8. The next section describes how the estimations are obtained in the data streaming model.

### 4.2 Algorithms in the streaming model

For each interval $I$, we use $\sigma_{\mathbb{S}}(I)$ for the sequence of segments from $\mathbb{S}$ that are active because of interval $I$, ordered non-increasingly by size. Thus, $\sigma_{\mathbb{S}}(I)$ contains $S(r)$ followed by the segments whose parents contain $I$. The selected ordering implies that a parent $\pi(S)$ appears before $S$, for all $S$ in the sequence $\sigma_{\mathbb{S}}(I)$. Note that $\sigma_{\mathbb{S}}(I)$ has at most $2\lceil\log n\rceil$ elements because $T$ is balanced.

Lemma 9. There is an algorithm in the data stream model that uses $O\left(\varepsilon^{-2}+\log n\right)$ space and computes a value $\hat{N}_{\text {act }}$ such that

$$
\operatorname{Pr}\left[\left|N_{a c t}-\hat{N}_{a c t}\right| \leq \varepsilon \cdot N_{a c t}\right] \geq \frac{11}{12}
$$

Proof. We estimate $N_{\text {act }}$ using, as a black box, known results to estimate the number of distinct elements in a data stream. The stream of intervals $\mathbb{I}=I_{1}, I_{2}, \ldots$ defines a stream of segments $\sigma=\sigma_{\mathbb{S}}\left(I_{1}\right), \sigma_{\mathbb{S}}\left(I_{2}\right), \ldots$ that is $O(\log n)$ times longer. The segments appearing in the stream $\sigma$ are precisely the active segments.

We have reduced the problem to the problem of how many distinct elements appear in a stream of segments from $\mathbb{S}$. The result of Kane, Nelson and Woodruff 12 for distinct elements uses $O\left(\varepsilon^{-2}+\log |\mathbb{S}|\right)=O\left(\varepsilon^{-2}+\log n\right)$ space and computes a value $\hat{N}_{\text {act }}$ such that

$$
\operatorname{Pr}\left[(1-\varepsilon) N_{a c t} \leq \hat{N}_{a c t} \leq(1+\varepsilon) N_{a c t}\right] \geq \frac{11}{12}
$$

Note that, to process an interval of the stream $\mathbb{I}$, we have to process $O(\log n)$ segments of $\mathbb{S}$.
Lemma 10. There is an algorithm in the data stream model that uses $O\left(\varepsilon^{-4} \log ^{4} n\right)$ space and computes a value $\hat{N}_{\text {rel }}$ such that

$$
\operatorname{Pr}\left[\left|N_{r e l}-\hat{N}_{r e l}\right| \leq \varepsilon \cdot N_{r e l}\right] \geq \frac{10}{12}
$$

Proof. The idea is the following. We estimate $N_{\text {act }}$ by $\hat{N}_{\text {act }}$ using Lemma 9 . We take a sample of active segments, and count how many of them are relevant. To get a representative sample, it will be important to use a lower bound on $N_{\text {rel }} / N_{\text {act }}$. With this we can estimate $N_{\text {rel }}=\left(N_{\text {rel }} / N_{\text {act }}\right) \cdot N_{\text {act }}$ accurately. We next provide the details.

In $T$, each relevant segment $S^{\prime} \in \mathbb{S}_{\text {rel }}$ has $2 \gamma\left(S^{\prime}\right)<4 \varepsilon^{-1}\lceil\log n\rceil^{2}$ active segments below it and at most $2\lceil\log n\rceil$ active segments whose parent is an ancestor of $S^{\prime}$. This means that for each relevant segment there are at most

$$
4 \varepsilon^{-1}\lceil\log n\rceil^{2}+2\lceil\log n\rceil \leq 6 \varepsilon^{-1}\lceil\log n\rceil^{2}
$$

active segments. We obtain that

$$
\begin{equation*}
\frac{N_{r e l}}{N_{a c t}} \geq \frac{1}{6 \varepsilon^{-1}\lceil\log n\rceil^{2}}=\frac{\varepsilon}{6\lceil\log n\rceil^{2}} \tag{4}
\end{equation*}
$$

Fix any injective mapping $b$ between $\mathbb{S}$ and $\left[n^{2}\right]$ that can be easily computed. For example, for each segment $S=[x, y)$ we may take $b(S)=n(x-1)+(y-1)$. Consider a family $\mathcal{H}=\mathcal{H}\left(n^{2}, \varepsilon\right)$ of permutations $\left[n^{2}\right] \rightarrow\left[n^{2}\right]$ guaranteed by Lemma 1. For each $h \in \mathcal{H}$, the function $h \circ b$ gives an order among the elements of $\mathbb{S}$. We use them to compute $\mathcal{H}$-random samples among the active segments.

Set $k=\left\lceil 72\lceil\log n\rceil^{2} /\left(\varepsilon^{3}(1-\varepsilon)\right]\right\rceil=\Theta\left(\varepsilon^{-3} \log ^{2} n\right)$, and choose permutations $h_{1}, \ldots, h_{k} \in \mathcal{H}$ uniformly and independently at random. For each permutation $h_{j}$, where $j=1, \ldots, k$, let $S_{j}$ be the active segment of $\mathbb{S}$ that minimizes $\left(h_{j} \circ b\right)(\cdot)$. Thus

$$
S_{j}=\arg \min \left\{h_{j}(S) \mid S \in \mathbb{S} \text { is active }\right\} .
$$

The idea is that $S_{j}$ is nearly a random active segment of $\mathbb{S}$. Therefore, if we define the random variable

$$
X=\mid\left\{j \in\{1, \ldots, k\} \mid S_{j} \text { is relevant }\right\} \mid
$$

then $N_{\text {rel }} / N_{\text {act }}$ is roughly $X / k$. Below we discuss the computation of $X$. To analyze the random variable $X$ more precisely, let us define

$$
p=\operatorname{Pr}_{h_{j} \in \mathcal{H}}\left[S_{j} \text { is relevant }\right] .
$$

Since $S_{j}$ is selected among the active segments, the discussion after Lemma 1 implies

$$
\begin{equation*}
p \in\left[\frac{(1-\varepsilon) N_{\text {rel }}}{N_{\text {act }}}, \frac{(1+\varepsilon) N_{\text {rel }}}{N_{\text {act }}}\right] . \tag{5}
\end{equation*}
$$

In particular, using the estimate (4) and the definition of $k$ we get

$$
\begin{equation*}
k p \geq \frac{72\lceil\log n\rceil^{2}}{\varepsilon^{3}(1-\varepsilon)} \cdot \frac{(1-\varepsilon) N_{r e l}}{N_{\text {act }}} \geq \frac{72\lceil\log n\rceil^{2}}{\varepsilon^{3}} \cdot \frac{\varepsilon}{6\lceil\log n\rceil^{2}}=\frac{12}{\varepsilon^{2}} . \tag{6}
\end{equation*}
$$

Note that $X$ is the sum of $k$ independent random variables taking values in $\{0,1\}$ and $\mathbb{E}[X]=k p$. It follows from Chebyshev's inequality and the lower bound in (6) that

$$
\begin{aligned}
\operatorname{Pr}\left[\left|\frac{X}{k}-p\right| \geq \varepsilon p\right] & =\operatorname{Pr}[|X-k p| \geq \varepsilon k p] \leq \frac{\operatorname{Var}[X]}{(\varepsilon k p)^{2}} \\
& =\frac{k p(1-p)}{(\varepsilon k p)^{2}}<\frac{1}{k p \varepsilon^{2}} \leq \frac{1}{12} .
\end{aligned}
$$

To finalize, let us define the estimator $\hat{N}_{\text {rel }}=\hat{N}_{\text {act }} \cdot\left(\frac{X}{k}\right)$ of $N_{\text {rel }}$, where $\hat{N}_{\text {act }}$ is the estimator of $N_{\text {act }}$ given in Lemma 9. When the events

$$
\left[\left|N_{a c t}-\hat{N}_{a c t}\right| \leq \varepsilon N_{a c t}\right] \quad \text { and } \quad\left[\left|\frac{X}{k}-p\right| \leq \varepsilon p\right]
$$

occur, then we can use equation (5) and $\varepsilon<1 / 2$ to see that

$$
\begin{aligned}
\hat{N}_{r e l} & \leq(1+\varepsilon) N_{a c t} \cdot(1+\varepsilon) p \leq(1+\varepsilon)^{2} N_{\text {act }} \cdot \frac{(1+\varepsilon) N_{\text {rel }}}{N_{\text {act }}}=(1+\varepsilon)^{3} N_{r e l} \\
& \leq(1+7 \varepsilon) N_{\text {rel }},
\end{aligned}
$$

and also

$$
\begin{aligned}
\hat{N}_{r e l} & \geq(1-\varepsilon) N_{a c t} \cdot(1-\varepsilon) p \geq(1-\varepsilon)^{2} N_{a c t} \cdot \frac{(1-\varepsilon) N_{r e l}}{N_{\text {act }}}=(1-\varepsilon)^{3} N_{r e l} \\
& \geq(1-7 \varepsilon) N_{r e l} .
\end{aligned}
$$

We conclude that

$$
\begin{aligned}
\operatorname{Pr}\left[(1-7 \varepsilon) N_{\text {rel }} \leq \hat{N}_{\text {rel }} \leq(1+7 \varepsilon) N_{\text {rel }}\right] & \geq 1-\operatorname{Pr}\left[\left|N_{\text {act }}-\hat{N}_{\text {act }}\right| \geq \varepsilon N_{a c t}\right]-\operatorname{Pr}\left[\left|\frac{X}{k}-p\right| \geq \varepsilon p\right] \\
& \geq 1-\frac{1}{12}-\frac{1}{12} \geq \frac{10}{12}
\end{aligned}
$$

Replacing in the argument $\varepsilon$ by $\varepsilon / 7$, we obtain the desired bound.
It remains to discuss how $X$ can be computed. For each $j$, where $j=1, \ldots, k$, we keep a variable that stores the current segment $S_{j}$ for all the segments that are active so far, keep information about the choice of $h_{j}$, and keep information about $\gamma\left(S_{j}\right)$ and $\gamma\left(\pi\left(S_{j}\right)\right)$, so that we can decide whether $S_{j}$ is relevant.

Let $I_{1}, I_{2}, \ldots$ be the data stream of input intervals. We consider the stream of segments $\sigma=\sigma_{\mathbb{S}}\left(I_{1}\right), \sigma_{\mathbb{S}}\left(I_{2}\right), \ldots$. When handling a segment $S$ of the stream $\sigma$, we may have to update $S_{j}$; this happens when $h_{j}(S)<h_{j}\left(S_{j}\right)$. Note that we can indeed maintain $\gamma\left(\pi\left(S_{j}\right)\right)$ because $S_{j}$ becomes active the first time that its parent contains some input interval. This is also the first time when $\gamma\left(\pi\left(S_{j}\right)\right)$ becomes nonzero, and thus the forthcoming part of the stream has enough information to compute $\gamma\left(S_{j}\right)$ and $\gamma\left(\pi\left(S_{j}\right)\right.$ ). (Here it is convenient that $\sigma_{\mathbb{S}}(I)$ gives segments in decreasing size.) To maintain $\gamma\left(S_{j}\right)$ and $\gamma\left(\pi\left(S_{j}\right)\right)$, we use Lemma 7 (iv).

To reduce the space used by each index $j$, we use the following simple trick. If at some point we detect that $\gamma\left(S_{j}\right)$ is larger than $2 \varepsilon^{-1}\lceil\log n\rceil^{2}$, we just store that $S_{j}$ is not relevant. If at some point we detect that $\gamma\left(\pi\left(S_{j}\right)\right)$ is larger than $2 \varepsilon^{-1}\lceil\log n\rceil^{2}$, we just store that $\pi\left(S_{j}\right)$ is large enough that $S_{j}$ could be relevant. We conclude that, for each $j$, we need at most $O\left(\log (1 / \varepsilon)+\varepsilon^{-1} \log ^{2} n\right)$ space. Therefore, we need in total $O\left(k \varepsilon^{-1} \log ^{2} n\right)=O\left(\varepsilon^{-4} \log ^{4} n\right)$ space.

Let

$$
\rho=\left(\sum_{S \in \mathbb{S}_{\text {rel }}} \hat{\beta}(S)\right) /\left|\mathbb{S}_{\text {rel }}\right| .
$$

The next result shows how to estimate $\rho$.
Lemma 11. There is an algorithm in the data stream model that uses $O\left(\varepsilon^{-5} \log ^{6} n\right)$ space and computes a value $\hat{\rho}$ such that

$$
\operatorname{Pr}[|\rho-\hat{\rho}| \leq \varepsilon \rho] \geq \frac{10}{12} .
$$

Proof. Fix any injective mapping $b$ between $\mathbb{S}$ and $\left[n^{2}\right]$ that can be easily computed, and consider a family $\mathcal{H}=\mathcal{H}\left(n^{2}, \varepsilon\right)$ of permutations $\left[n^{2}\right] \rightarrow\left[n^{2}\right]$ guaranteed by Lemma 11. For each $h \in \mathcal{H}$, the function $h \circ b$ gives an order among the elements of $\mathbb{S}$. We use them to compute $\mathcal{H}$-random samples among the active segments.

Let $\mathbb{S}_{\text {act }}$ be the set of active segments.
Consider a random variable $Y_{1}$ defined as follows. We repeatedly sample $h_{1} \in \mathcal{H}$ uniformly at random, until we get that $\arg \min _{S \in \mathbb{S}_{\text {act }}} \hat{\beta}(S)$ is a relevant segment. Let $S_{1}$ be the resulting relevant segment $\arg \min _{S \in \mathbb{S}_{a c t}} \hat{\beta}(S)$, and set $Y_{1}=\hat{\beta}\left(S_{1}\right)$. Because of Lemma 2, where $X=\mathbb{S}_{\text {act }}$ and $Y=\mathbb{S}_{\text {rel }}$, we have

$$
\forall S \in \mathbb{S}_{\text {rel }}: \quad \frac{1-4 \varepsilon}{\left|\mathbb{S}_{\text {rel }}\right|} \leq \operatorname{Pr}\left[S_{1}=S\right] \leq \frac{1+4 \varepsilon}{\left|\mathbb{S}_{\text {rel }}\right|}
$$

We thus have

$$
\mathbb{E}\left[Y_{1}\right]=\sum_{S \in \mathbb{S}_{\text {rel }}} \operatorname{Pr}\left[S_{1}=S\right] \cdot \hat{\beta}(S) \leq \sum_{S \in \mathbb{S}_{\text {rel }}} \frac{1+4 \varepsilon}{\left|\mathbb{S}_{\text {rel }}\right|} \cdot \hat{\beta}(S)=(1+4 \varepsilon) \cdot \rho,
$$

and similarly

$$
\mathbb{E}\left[Y_{1}\right] \geq \sum_{S \in \mathbb{S}_{\text {rel }}} \frac{1-4 \varepsilon}{\left|\mathbb{S}_{\text {rel }}\right|} \cdot \hat{\beta}(S)=(1-4 \varepsilon) \cdot \rho
$$

For the variance we can use $\hat{\beta}(S) \leq \gamma(S)$ and the definition of relevant segments to get

$$
\begin{aligned}
\operatorname{Var}\left[Y_{1}\right] & \leq \mathbb{E}\left[Y_{1}^{2}\right] \\
& =\sum_{S \in \mathbb{S}_{\text {rel }}} \operatorname{Pr}\left[S_{1}=S\right] \cdot(\hat{\beta}(S))^{2} \\
& \leq \sum_{S \in \mathbb{S}_{\text {rel }}} \frac{1+4 \varepsilon}{\left|\mathbb{S}_{\text {rel }}\right|} \cdot \hat{\beta}(S) \cdot \frac{2\lceil\log n\rceil^{2}}{\varepsilon} \\
& \leq \frac{2(1+4 \varepsilon)\lceil\log n\rceil^{2}}{\varepsilon} \cdot \rho \\
& \leq \frac{6\lceil\log n\rceil^{2}}{\varepsilon} \cdot \rho
\end{aligned}
$$

Note also that $\gamma(S) \geq 1$ implies $\hat{\beta}(S) \geq 1$. Therefore, $\rho \geq 1$.
Consider an integer $k$ to be chosen later. Let $Y_{2}, \ldots, Y_{k}$ be independent random variables with the same distribution that $Y_{1}$, and define $\hat{\rho}=\left(\sum_{i=1}^{k} Y_{i}\right) / k$. Using Chebyshev's inequality and $\rho \geq 1$ we obtain

$$
\begin{aligned}
\operatorname{Pr}\left[\left|\hat{\rho}-\mathbb{E}\left[Y_{1}\right]\right| \geq \varepsilon \rho\right] & =\operatorname{Pr}\left[\left|\hat{\rho} k-\mathbb{E}\left[Y_{1}\right] k\right| \geq \varepsilon k \rho\right] \leq \frac{\operatorname{Var}[\hat{\rho} k]}{(\varepsilon k \rho)^{2}} \\
& =\frac{k \operatorname{Var}\left[Y_{1}\right]}{(\varepsilon k \rho)^{2}}<\frac{\frac{6\lceil\log n]^{2}}{\varepsilon} \cdot \rho}{k \varepsilon^{2} \rho^{2}} \\
& \leq \frac{6\lceil\log n\rceil^{2}}{k \varepsilon^{3}}
\end{aligned}
$$

Setting $k=6 \cdot 12 \cdot\lceil\log n\rceil^{2} / \varepsilon^{3}$, we have

$$
\operatorname{Pr}\left[\left|\hat{\rho}-\mathbb{E}\left[Y_{1}\right]\right| \geq \varepsilon \rho\right] \leq \frac{1}{12} .
$$

We then proceed similar to the proof of Lemma 10 . Set $k_{0}=12\lceil\log n\rceil^{2} k / \varepsilon(1-\varepsilon)=$ $\Theta\left(\varepsilon^{-4} \log ^{4} n\right)$. For each $j \in\left[k_{0}\right]$, take a function $h_{j} \in \mathcal{H}$ uniformly at random and select $S_{j}=\arg \min \{h(b(S)) \mid S$ is active $\}$. Let $X$ be the number of relevant segments in $S_{1}, \ldots, S_{k_{0}}$ and let $p=\operatorname{Pr}\left[S_{1} \in \mathbb{S}_{\text {rel }}\right]$. Using the analysis of Lemma 10 we have

$$
k_{0} p \geq \frac{\left(12\lceil\log n\rceil^{2}\right) k}{\varepsilon(1-\varepsilon)} \cdot \frac{(1-\varepsilon) N_{r e l}}{N_{a c t}} \geq k \cdot \frac{12\lceil\log n\rceil^{2}}{\varepsilon(1-\varepsilon)} \cdot \frac{(1-\varepsilon) \varepsilon}{6\lceil\log n\rceil^{2}}=2 k
$$

and

$$
\operatorname{Pr}\left[\left|X-k_{0} p\right| \geq k_{0} p / 2\right] \leq \frac{\operatorname{Var}[X]}{\left(k_{0} p / 2\right)^{2}}=\frac{4 k_{0} p(1-p)}{k_{0}^{2} p^{2}}<\frac{4}{k_{0} p} \leq \frac{4}{2 k} \leq \frac{1}{12}
$$

This means that, with probability at least $11 / 12$, the sample $S_{1}, \ldots, S_{k_{0}}$ contains at least $(1 / 2) k_{0} p \geq k$ relevant segments. We can then use the first $k$ of those relevant segments to compute the estimate $\hat{\rho}$.

With probability at least $1-1 / 12-1 / 12=10 / 12$ we have the events

$$
\left[\left|X-k_{0} p\right| \geq k_{0} p / 2\right] \quad \text { and } \quad\left[\left|\hat{\rho}-\mathbb{E}\left[Y_{1}\right]\right| \geq \varepsilon \rho\right]
$$

In such a case

$$
\hat{\rho} \leq \varepsilon \rho+\mathbb{E}\left[Y_{1}\right] \leq \varepsilon \rho+(1+4 \varepsilon) \rho=(1+5 \varepsilon) \rho
$$

and similarly

$$
\hat{\rho} \geq \mathbb{E}\left[Y_{1}\right]-\varepsilon \rho \geq(1-4 \varepsilon) \rho-\varepsilon \rho=(1-5 \varepsilon) \rho
$$

Therefore,

$$
\operatorname{Pr}[|\rho-\hat{\rho}| \leq 5 \varepsilon \rho] \geq \frac{10}{12}
$$

Changing the role of $\varepsilon$ and $\varepsilon / 5$, the claimed probability is obtained.
It remains to show that we can compute $\hat{\rho}$ in the data stream model. Like before, for each $j \in\left[k_{0}\right]$, we have to maintain the segment $S_{j}$, information about the choice of the permutation $h_{j}$, information about $\gamma\left(S_{j}\right)$ and $\gamma\left(\pi\left(S_{j}\right)\right.$ ), and the value $\hat{\beta}\left(S_{j}\right)$. Since $\hat{\beta}\left(S_{j}\right) \leq \beta\left(S_{j}\right) \leq \gamma\left(S_{j}\right)$ because of Lemma 7 (iii), we need $O\left(\varepsilon^{-1} \log ^{2} n\right)$ space per index $j$. In total we need $O\left(k_{0} \varepsilon^{-1} \log ^{2} n\right)=$ $O\left(\varepsilon^{-5} \log ^{6} n\right)$ space.

Theorem 12. Assume that $\varepsilon \in(0,1 / 2)$ and let $\mathbb{I}$ be a set of intervals with endpoints in $\{1, \ldots, n\}$ that arrive in a data stream. There is a data stream algorithm that uses $O\left(\varepsilon^{-5} \log ^{6} n\right)$ space and computes a value $\hat{\alpha}$ such that

$$
\operatorname{Pr}\left[\frac{1}{2}(1-\varepsilon) \cdot \alpha(\mathbb{I}) \leq \hat{\alpha} \leq \alpha(\mathbb{I})\right] \geq \frac{2}{3}
$$

Proof. We compute the estimate $\hat{N}_{\text {rel }}$ of Lemma 10 and the estimate $\hat{\rho}$ of Lemma 11. Define the estimate $\hat{\alpha}_{0}=\hat{N}_{\text {rel }} \cdot \hat{\rho}$. With probability at least $1-\frac{2}{12}-\frac{2}{12}=\frac{2}{3}$ we have the events

$$
\left[\left|N_{r e l}-\hat{N}_{r e l}\right| \leq \varepsilon \cdot N_{r e l}\right] \quad \text { and } \quad[|\rho-\hat{\rho}| \leq \varepsilon \rho] .
$$

When such events hold, we can use the definitions of $N_{r e l}$ and $\rho$, together with Lemma 8, to obtain

$$
\hat{\alpha}_{0} \leq(1+\varepsilon) N_{r e l} \cdot(1+\varepsilon) \rho=(1+\varepsilon)^{2} \sum_{S \in \mathbb{S}_{r e l}} \hat{\beta}(S) \leq(1+\varepsilon)^{2} \alpha(\mathbb{I})
$$

and

$$
\hat{\alpha}_{0} \geq(1-\varepsilon) N_{r e l} \cdot(1-\varepsilon) \rho=(1-\varepsilon)^{2} \sum_{S \in \mathbb{S}_{r e l}} \hat{\beta}(S) \geq(1-\varepsilon)^{2}\left(\frac{1}{2}-\varepsilon\right) \alpha(\mathbb{I})
$$

Therefore,

$$
\operatorname{Pr}\left[(1-\varepsilon)^{2}\left(\frac{1}{2}-\varepsilon\right) \cdot \alpha(\mathbb{I}) \leq \hat{\alpha}_{0} \leq(1+\varepsilon)^{2} \cdot \alpha(\mathbb{I})\right] \geq \frac{2}{3} .
$$

Using that $(1-\varepsilon)^{2}(1 / 2-\varepsilon) /(1+\varepsilon)^{2} \geq 1 / 2-3 \varepsilon$ for all $\varepsilon \in(0,1 / 2)$, rescaling $\varepsilon$ by $1 / 6$, and setting $\hat{\alpha}=\hat{\alpha}_{0} /(1+\varepsilon)^{2}$, the claimed approximation is obtained. The space bounds are those from Lemmas 10 and 11

## 5 Largest independent set of same-size intervals

In this section we show how to obtain a $(3 / 2)$-approximation to the largest independent set using $O(\alpha(\mathbb{I}))$ space in the special case when all the intervals have the same length $\lambda>0$.

Our approach is based on using the shifting technique of Hochbaum and Mass [9] with a grid of length $3 \lambda$ and shifts of length $\lambda$. We observe that we can maintain an optimal solution restricted to a window of length $3 \lambda$ because at most two disjoint intervals of length $\lambda$ can fit in.

For any real value $\ell$, let $W_{\ell}$ denote the window $[\ell, \ell+3 \lambda)$. Note that $W_{\ell}$ includes the left endpoint but excludes the right endpoint. For $a \in\{0,1,2\}$, we define the partition of the real line

$$
\mathbb{W}_{a}=\left\{W_{(a+3 j) \lambda} \mid j \in \mathbb{Z}\right\} .
$$

For $a \in\{0,1,2\}$, let $\mathbb{I}_{a}$ be the set of input intervals contained in some window of $\mathbb{W}_{a}$. Thus,

$$
\mathbb{I}_{a}=\{I \in \mathbb{I} \mid \nexists j \in \mathbb{Z} \text { s.t. }(a+3 j) \lambda \in I\} .
$$

Lemma 13. If all the intervals of $\mathbb{I}$ have length $\lambda>0$, then

$$
\max \left\{\alpha\left(\mathbb{I}_{0}\right), \alpha\left(\mathbb{I}_{1}\right), \alpha\left(\mathbb{I}_{2}\right)\right\} \geq \frac{2}{3} \alpha(\mathbb{I}) .
$$

Proof. Each interval of length $\lambda$ is contained in exactly two windows of $\mathbb{W}_{0} \cup \mathbb{W}_{1} \cup \mathbb{W}_{2}$. Let $\mathbb{J}^{*} \subseteq \mathbb{I}$ be a largest independent set of intervals, so that $\left|\mathbb{J}^{*}\right|=\alpha(\mathbb{I})$. We then have

$$
3 \cdot \max \left\{\alpha\left(\mathbb{I}_{0}\right), \alpha\left(\mathbb{I}_{1}\right), \alpha\left(\mathbb{I}_{2}\right)\right\} \geq \sum_{0 \leq a \leq 2} \alpha\left(\mathbb{I}_{a}\right) \geq \sum_{0 \leq a \leq 2}\left|\mathbb{J}^{*} \cap \mathbb{I}_{a}\right| \geq 2\left|\mathbb{J}^{*}\right|=2 \alpha(\mathbb{I})
$$

and the result follows.
For each $a \in\{0,1,2\}$ we store an optimal solution $\mathbb{J}_{a}$ restricted to $\mathbb{I}_{a}$. We obtain a (3/2)approximation by returning the largest among $\mathbb{J}_{0}, \mathbb{J}_{1}, \mathbb{J}_{2}$.

For each window $W$ considered through the algorithm, we store Leftmost $(W)$ and $\operatorname{Rightmost}(W)$. We also store a boolean value active $(W)$ telling whether some previous interval was contained in $W$. When $\operatorname{active}(W)$ is false, $\operatorname{Leftmost}(W)$ and $\operatorname{Rightmost}(W)$ are undefined.

With a few local operations, we can handle the insertion of new intervals in $\mathbb{I}$. For a window $W \in \mathbb{W}_{a}$, there are two relevant moments when $\mathbb{J}_{a}$ may be changed. First, when $W$ gets the first interval, the interval has to be added to $\mathbb{J}_{a}$ and $\operatorname{active}(W)$ is set to true. Second, when $W$ can first fit two disjoint intervals, then those two intervals have to be added to $\mathbb{J}_{a}$. See the pseudocode in Figure 8 for a more detailed description.

Lemma 14. For $a=0,1,2$, the policy described in Figure 8 maintains an optimal solution $\mathbb{J}_{a}$ for the intervals $\mathbb{I}_{a}$.

Process interval $[x, y]$ of length $\lambda$

1. for $a=0,1,2$ do
2. $\quad W \leftarrow$ window of $\mathbb{W}_{a}$ that contains $x$
3. if $y \in W$ then $\quad\left(*[x, y]\right.$ is contained in the window $\left.W \in \mathbb{W}_{a}{ }^{*}\right)$
4. if $\operatorname{active}(W)$ false then
5. $\quad \operatorname{active}(W) \leftarrow$ true
6. $\quad \operatorname{Rightmost}(W) \leftarrow[x, y]$
7. $\operatorname{Leftmost}(W) \leftarrow[x, y]$
8. add $[x, y]$ to $\mathbb{J}_{a}$
9. 

else if Rightmost $(W) \cap \operatorname{Leftmost}(W) \neq \emptyset$ then
$[\ell, r] \leftarrow \operatorname{Rightmost}(W) \cap \operatorname{Leftmost}(W)$
if $\ell<x$ then $\operatorname{Rightmost}(W) \leftarrow[x, y]$
if $y<r$ then Leftmost $(W) \leftarrow[x, y]$
if $\operatorname{Rightmost}(W) \cap \operatorname{Leftmost}(W)=\emptyset$ then remove from $\mathbb{J}_{a}$ the interval contained in $W$ add to $\mathbb{J}_{a}$ intervals Rightmost $(W)$ and $\operatorname{Leftmost}(W)$

Figure 8: Policy to process a new interval $[x, x+\lambda] . \mathbb{J}_{a}$ maintains an optimal solution for $\alpha\left(\mathbb{I}_{a}\right)$.

Proof. Since a window $W$ of length 3 can contain at most 2 disjoint intervals of length $\lambda$, the intervals Rightmost $(W)$ and $\operatorname{Leftmost}(W)$ suffice to obtain an optimal solution restricted to intervals contained in $W$. By the definition of $\mathbb{I}_{a}$, an optimal solution for

$$
\bigcup_{W \in \mathbb{W}_{a}}\{\operatorname{Rightmost}(W), \operatorname{Leftmost}(W)\}
$$

is an optimal solution for $\mathbb{I}_{a}$. Since the algorithm maintains such an optimal solution, the claim follows.

Since each window can have at most two disjoint intervals and each interval is contained in at most two windows of $\mathbb{W}_{0} \cup \mathbb{W}_{1} \cup \mathbb{W}_{2}$, we have at most $O(\alpha(\mathbb{I}))$ active intervals through the entire stream. Using a dynamic binary search tree for the active windows, we can perform the operations in $O(\log \alpha(\mathbb{I}))$ time. We summarize.

Theorem 15. Let $\mathbb{I}$ be a set of intervals of length $\lambda$ in the real line that arrive in a data stream. There is a data stream algorithm to compute a (3/2)-approximation to the largest independent subset of $\mathbb{I}$ that uses $O(\alpha(\mathbb{I}))$ space and handles each interval of the stream in $O(\log \alpha(\mathbb{I}))$ time.

## 6 Size of largest independent set for same-size intervals

In this section we show how to obtain a randomized estimate of the value $\alpha(\mathbb{I})$ in the special case when all the intervals have the same length $\lambda>0$. We assume that the endpoints are in $[n]$.

The idea is an extension of the idea used in Section 5. For $a=0,1,2$, let $\mathbb{W}_{a}$ and $\mathbb{I}_{a}$ be as defined in Section 5. For $a=0,1,2$, we will compute a value $\hat{\alpha}_{a}$ that $(1+\varepsilon)$-approximates $\alpha\left(\mathbb{I}_{a}\right)$ with reasonable probability. We then return $\hat{\alpha}=\max \left\{\hat{\alpha}_{0}, \hat{\alpha}_{1}, \hat{\alpha}_{2}\right\}$, which catches a fraction at least $\frac{2}{3}(1-\varepsilon)$ of $\alpha(\mathbb{I})$, with reasonable probability.

To obtain the $(1+\varepsilon)$-approximation to $\alpha\left(\mathbb{I}_{a}\right)$, we want to estimate how many windows of $\mathbb{W}_{a}$ contain some input interval and how many contain two disjoint input intervals. For this we combine known results for the distinct elements as a black box and use sampling over the windows that contain some input interval.

Lemma 16. Let a be 0,1 or 2 and let $\varepsilon \in(0,1)$. There is an algorithm in the data stream model that uses $O\left(\varepsilon^{-2} \log (1 / \varepsilon)+\log n\right)$ space and computes a value $\hat{\alpha}_{a}$ such that

$$
\operatorname{Pr}\left[\left|\alpha\left(\mathbb{I}_{a}\right)-\hat{\alpha}_{a}\right| \leq \varepsilon \cdot \alpha\left(\mathbb{I}_{a}\right)\right] \geq \frac{8}{9} .
$$

Proof. Let us fix some $a \in\{0,1,2\}$. We say that a window $W$ of $\mathbb{W}_{a}$ is of type $i$ if $W$ contains at least $i$ disjoint input intervals. Since the windows of $\mathbb{W}_{a}$ have length $3 \lambda$, they can be of type 0,1 or 2 . For $i=0,1,2$, let $\gamma_{i}$ be the number of windows of type $i$ in $\mathbb{W}_{a}$. Then $\alpha\left(\mathbb{I}_{a}\right)=\gamma_{1}+\gamma_{2}$.

We compute an estimate $\hat{\gamma}_{1}$ to $\gamma_{1}$ as follows. We have to estimate $\mid\left\{W \in \mathbb{W}_{a} \mid \exists I\right.$ s.t. $\left.I \subset W\right\} \mid$. The stream of intervals $\mathbb{I}=I_{1}, I_{2}, \ldots$ defines the sequence of windows $W(\mathbb{I})=W\left(I_{1}\right), W\left(I_{2}\right), \ldots$, where $W\left(I_{i}\right)$ denotes the window of $\mathbb{W}_{a}$ that contains $I_{i}$; if $I_{i}$ is not contained in any window of $\mathbb{W}_{a}$, we then skip $I_{i}$. Then $\gamma_{1}$ is the number of distinct elements in the sequence $W(\mathbb{I})$. The results of Kane, Nelson and Woodruff 12 imply that using $O\left(\varepsilon^{-2}+\log n\right)$ space we can compute a value $\hat{\gamma}_{1}$ such that

$$
\operatorname{Pr}\left[(1-\varepsilon) \gamma_{1} \leq \hat{\gamma}_{1} \leq(1+\varepsilon) \gamma_{1}\right] \geq \frac{17}{18} .
$$

We next explain how to estimate the ratio $\gamma_{2} / \gamma_{1} \leq 1$. Consider a family $\mathcal{H}=\mathcal{H}(n, \varepsilon)$ of permutations $[n] \rightarrow[n]$ guaranteed by Lemma 11, set $k=\left\lceil 18 \varepsilon^{-2}\right\rceil$, and choose permutations $h_{1}, \ldots, h_{k} \in \mathcal{H}$ uniformly and independently at random. For each permutation $h_{j}$, where $j=1, \ldots, k$, let $W_{j}$ be the window $[\ell, \ell+3 \lambda)$ of $\mathbb{W}_{a}$ that contains some input interval and minimizes $h_{j}(\ell)$. Thus

$$
W_{j}=\arg \min \left\{h_{j}(\ell) \mid[\ell, \ell+3 \lambda) \in \mathbb{W}_{a} \text {, some } I \in \mathbb{I} \text { is contained in }[\ell, \ell+3 \lambda)\right\} .
$$

The idea is that $W_{j}$ is a nearly-uniform random window of $\mathbb{W}_{a}$, among those that contain some input interval. Therefore, if we define the random variable

$$
M=\mid\left\{j \in\{1, \ldots, k\} \mid W_{j} \text { is of type } 2\right\} \mid
$$

then $\gamma_{2} / \gamma_{1}$ is roughly $M / k$. Below we make a precise analysis.
Let us first discuss that $M$ can be computed in space within $O(k \log (1 / \varepsilon))=O\left(\varepsilon^{-2} \log (1 / \varepsilon)\right)$. For each $j$, where $j=1, \ldots, k$, we keep information about the choice of $h_{j}$, keep a variable that stores the current window $W_{j}$ for all the intervals that have been seen so far, and store the intervals $\operatorname{Rightmost}\left(W_{j}\right)$ and $\operatorname{Leftmost}\left(W_{j}\right)$. Those two intervals tell us whether $W_{j}$ is of type 1 or 2 . When handling an interval $I$ of the stream, we may have to update $W_{j}$; this happens when $h_{j}(s)<h_{j}\left(s_{j}\right)$, where $s$ is the left endpoint of the window of $\mathbb{W}_{a}$ that contains $I$ and $s_{j}$ is the left endpoint of $W_{j}$. When $W_{j}$ is changed, we also have to reset the intervals $\operatorname{Rightmost}\left(W_{j}\right)$ and $\operatorname{Leftmost}\left(W_{j}\right)$ to the new interval $I$.

To analyze the random variable $M$ more precisely, let us define

$$
p=\operatorname{Pr}_{h_{j} \in \mathcal{H}}\left[W_{j} \text { is of type 2] } \in\left[\frac{(1-\varepsilon) \gamma_{2}}{\gamma_{1}}, \frac{(1+\varepsilon) \gamma_{2}}{\gamma_{1}}\right] .\right.
$$

Note that $M$ is the sum of $k$ independent random variables taking values in $\{0,1\}$ and $\mathbb{E}[M]=k p$. It follows from Chebyshev's inequality that

$$
\operatorname{Pr}\left[\left|\frac{M}{k}-p\right| \geq \varepsilon\right]=\operatorname{Pr}[|M-k p| \geq \varepsilon k] \leq \frac{\operatorname{Var}[M]}{(\varepsilon k)^{2}} \leq \frac{k p}{\varepsilon^{2} k^{2}} \leq \frac{1}{\varepsilon^{2} k} \leq \frac{1}{18}
$$

To finalize, let us define the estimator $\hat{\alpha}_{a}=\hat{\gamma}_{1}\left(1+\frac{M}{k}\right)$. When the events

$$
\left[(1-\varepsilon) \gamma_{1} \leq \hat{\gamma}_{1} \leq(1+\varepsilon) \gamma_{1}\right] \quad \text { and } \quad\left[\left|\frac{M}{k}-p\right| \leq \varepsilon\right]
$$

occur, then we have

$$
\begin{aligned}
\hat{\alpha}_{a} & \leq(1+\varepsilon) \gamma_{1}(1+p) \leq(1+\varepsilon) \gamma_{1}\left(1+\varepsilon+\frac{(1+\varepsilon) \gamma_{2}}{\gamma_{1}}\right) \\
& =(1+\varepsilon)^{2}\left(\gamma_{1}+\gamma_{2}\right) \leq(1+3 \varepsilon) \alpha\left(\mathbb{I}_{a}\right)
\end{aligned}
$$

and also

$$
\begin{aligned}
\hat{\alpha}_{a} & \geq(1-\varepsilon) \gamma_{1}(1-\varepsilon+p) \geq(1-\varepsilon) \gamma_{1}\left(1-\varepsilon+\frac{(1-\varepsilon) \gamma_{2}}{\gamma_{1}}\right) \\
& =(1-\varepsilon)^{2}\left(\gamma_{1}+\gamma_{2}\right) \geq(1-3 \varepsilon) \alpha\left(\mathbb{I}_{a}\right)
\end{aligned}
$$

We conclude that

$$
\operatorname{Pr}\left[(1-3 \varepsilon) \alpha\left(\mathbb{I}_{a}\right) \leq \hat{\alpha}_{a} \leq(1+3 \varepsilon) \alpha\left(\mathbb{I}_{a}\right)\right] \geq 1-\frac{1}{18}-\frac{1}{18} \geq \frac{8}{9}
$$

Replacing in the argument $\varepsilon$ by $\varepsilon / 3$, the result follows.
Theorem 17. Assume that $\varepsilon \in(0,1 / 2)$ and let $\mathbb{I}$ be a set of intervals of length $\lambda$ with endpoints in $\{1, \ldots, n\}$ that arrive in a data stream. There is a data stream algorithm that uses $O\left(\varepsilon^{-2} \log (1 / \varepsilon)+\log n\right)$ space and computes a value $\hat{\alpha}$ such that

$$
\operatorname{Pr}\left[\frac{2}{3}(1-\varepsilon) \cdot \alpha(\mathbb{I}) \leq \hat{\alpha} \leq \alpha(\mathbb{I})\right] \geq \frac{2}{3}
$$

Proof. For each $a=0,1,2$ we compute the estimate $\hat{\alpha}_{a}$ to $\alpha\left(\mathbb{I}_{a}\right)$ with the algorithm described in Lemma 16. We then have

$$
\operatorname{Pr}\left[\bigwedge_{a=0,1,2}\left[\left|\alpha\left(\mathbb{I}_{a}\right)-\hat{\alpha}_{a}\right| \leq \varepsilon \cdot \alpha\left(\mathbb{I}_{a}\right)\right]\right] \geq \frac{6}{9}
$$

When the event occurs, it follows by Lemma 13 that

$$
\frac{2}{3}(1-\varepsilon) \cdot \alpha(\mathbb{I}) \leq \max \left\{\hat{\alpha}_{0}, \hat{\alpha}_{1}, \hat{\alpha}_{2}\right\} \leq(1+\varepsilon) \alpha(\mathbb{I})
$$

Therefore, using that $(1-\varepsilon) /(1+\varepsilon) \geq 1-2 \varepsilon$ for all $\varepsilon \in(0,1 / 2)$, rescaling $\varepsilon$ by $1 / 2$, and returning $\hat{\alpha}=\max \left\{\hat{\alpha}_{0}, \hat{\alpha}_{1}, \hat{\alpha}_{2}\right\} /(1+\varepsilon)$, the result is achieved.

## 7 Lower bounds

Emek, Halldórsson and Rosén [6] showed that any streaming algorithm for the interval selection problem cannot achieve an approximation ratio of $r$, for any constant $r<2$. They also show that, for same-size intervals, one cannot obtain an approximation ratio below $3 / 2$. We are going to show that similar inapproximability results hold for estimating $\alpha(\mathbb{I})$.

The lower bound of Emek et al. uses the coordinates of the endpoints to recover information about a permutation. That is, given a solution to the interval selection problem, they can use the endpoints of the intervals to recover information. Thus, such reduction cannot be adapted to the estimation of $\alpha(\mathbb{I})$, since we do not require to return intervals. Nevertheless, there are certain similarities between their construction and ours, especially for same-length intervals.

Consider the problem Index. The input to Index is a pair $(S, i) \in\{0,1\}^{n} \times[n]$ and the output, denoted by $\operatorname{Index}(S, i)$, is the $i$-th bit of $S$. One can think of $S$ as a subset of $[n]$, and then $\operatorname{Index}(S, i)$ is asking whether element $i$ is in the subset or not.


Figure 9: Example showing $\sigma(S, i)$ for $n=7, S=\{1,3,4,6\}, L=9$, and $i=2$ in the proof of Theorem 18. The intervals are sorted from bottom to top in the order they appear in the data stream. The empty dots represent endpoints that are not included in the interval, while the full dots represent endpoints included in the interval.

The one-way communication complexity of Index is well understood. In this scenario, Alice has $S$ and Bob has $i$. Alice sends a message to Bob and then Bob has to compute $\operatorname{Index}(S, i)$. The key question is how long should be the message in the worst case so that Bob can compute $\operatorname{InDEX}(S, i)$ correctly with probability greater than $1 / 2$, say, at least $2 / 3$. (Attaining probability $1 / 2$ is of course trivial.) It is known that, to achieve this, the message of Alice has $\Omega(n)$ bits in the worst case. See [11] for a short proof and [15] for a comprehensive treatment.

Given an input $(S, i)$ for Index, we can build a data stream of same-length intervals $\mathbb{I}$ with the property that $\alpha(\mathbb{I}) \in\{2,3\}$ and $\operatorname{Index}(S, i)=1$ if and only if $\alpha(\mathbb{I})=3$. Moreover, the first part of the stream depends only on $S$ and the second part on $i$. Thus, the state of the memory at the end of the first part can be interpreted as a message that Alice sends to Bob. This implies the following lower bound.

Theorem 18. Let $c>0$ be an arbitrary constant. Consider the problem of estimating $\alpha(\mathbb{I})$ for sets of same-length intervals $\mathbb{I}$ with endpoints in $[n]$. In the data streaming model, there is no algorithm that uses o(n) bits of memory and computes an estimate $\hat{\alpha}$ such that

$$
\begin{equation*}
\operatorname{Pr}\left[\left(\frac{2}{3}+c\right) \alpha(\mathbb{I}) \leq \hat{\alpha} \leq \alpha(\mathbb{I})\right] \geq \frac{2}{3} \tag{7}
\end{equation*}
$$

Proof. For simplicity, we use intervals with endpoints in [3n] and mix closed and open intervals in the proof.

Given an input $(S, i)$ for Index, consider the following stream of intervals. We set $L$ to some large enough value; for example $L=n+2$ will be enough. Let $\sigma_{1}(S)$ be a stream that, for each $j \in S$, contains the closed interval $[L+j, 2 L+j]$. Let $\sigma_{2}(i)$ be the length-two stream with open intervals $(i, L+i)$ and $(2 L+i, 3 L+i)$. Finally, let $\sigma(S, i)$ be the concatenation of $\sigma_{1}(S)$ and $\sigma_{2}(i)$. See Figure 9 for an example. Let $\mathbb{I}$ be the intervals in $\sigma(S, i)$. It is straightforward to see that $\alpha(\mathbb{I})$ is 2 or 3 . Moreover, $\alpha(\mathbb{I})=3$ if and only if $\operatorname{Index}(S, i)=1$.

Assume, for the sake of contradiction, that we have an algorithm in the data streaming model that uses $o(n)$ bits of space and computes a value $\hat{\alpha}$ that satisfies equation (7). Then, Alice and Bob can solve $\operatorname{Index}(S, i)$ using $o(n)$ bits, as follows.

Alice simulates the data stream algorithm on $\sigma_{1}(S)$ and sends to Bob a message encoding the state of the memory at the end of processing $\sigma_{1}(S)$. The message of Alice has $o(n)$ bits. Then, Bob continues the simulation on the last two items of $\sigma(S, i)$, that is, $\sigma_{2}(i)$. Bob has correctly computed the output of the algorithm on $\sigma(S, i)$, and therefore obtains $\hat{\alpha}$ so that equation (7) is satisfied. If $\hat{\alpha}>2$, then Bob returns the bit $\hat{\beta}=1$. If $\hat{\alpha} \leq 2$, then Bob returns $\hat{\beta}=0$. This finishes the description of the protocol.

Consider the case when $\operatorname{Index}(S, i)=1$. In that case, $\alpha(\mathbb{I})=3$. With probability at least $2 / 3$, the value $\hat{\alpha}$ computed satisfies $\hat{\alpha} \geq\left(\frac{2}{3}+c\right) \alpha(\mathbb{I})=2+3 c>2$, and therefore

$$
\begin{aligned}
\operatorname{Pr}[\hat{\beta}=1 \mid \operatorname{Index}(S, i)=1] & =\operatorname{Pr}[\hat{\alpha}>2 \mid \operatorname{Index}(S, i)=1] \\
& \geq \operatorname{Pr}\left[\left.\hat{\alpha} \geq\left(\frac{2}{3}+c\right) \alpha(\mathbb{I}) \right\rvert\, \operatorname{Index}(S, i)=1\right] \\
& \geq 2 / 3
\end{aligned}
$$



Figure 10: Example showing $\sigma(S, i)$ for $n=7, S=\{1,3,4,6\}, L=9, k=3$, and $i=2$ in the proof of Theorem 19. The intervals are sorted from bottom to top in the order they appear in the data stream. The empty dots represent endpoints that are not included in the interval, while the full dots represent endpoints included in the interval.

When $\operatorname{Index}(S, i)=0$, then $\alpha(\mathbb{I})=2$, and, with probability at least $2 / 3$, the value $\hat{\alpha}$ computed satisfies $\hat{\alpha} \leq \alpha(\mathbb{I})=2$. Therefore,

$$
\begin{aligned}
\operatorname{Pr}[\hat{\beta}=0 \mid \operatorname{Index}(S, i)=0] & =\operatorname{Pr}[\hat{\alpha} \leq 2 \mid \operatorname{Index}(S, i)=0] \\
& \geq \operatorname{Pr}[\hat{\alpha} \leq \alpha(\mathbb{I}) \mid \operatorname{Index}(S, i)=0] \\
& \geq 2 / 3 .
\end{aligned}
$$

We conclude that

$$
\operatorname{Pr}[\hat{\beta}=\operatorname{Index}(S, i)] \geq 2 / 3 .
$$

Since Bob computes $\hat{\beta}$ after a message from Alice with $o(n)$ bits, this contradicts the lower bound for Index.

For intervals of different sizes, we can use an alternative construction with the property that $\alpha(\mathbb{I})$ is either $k+1$ or $2 k+1$. This means that we cannot get an approximation ratio arbitrarily close to 2 .

Theorem 19. Let $c>0$ be an arbitrary constant. Consider the problem of estimating $\alpha(\mathbb{I})$ for sets of intervals $\mathbb{I}$ with endpoints in $[n]$. In the data streaming model, there is no algorithm that uses o(n) bits of memory and computes an estimate $\hat{\alpha}$ such that

$$
\operatorname{Pr}\left[\left(\frac{1}{2}+c\right) \alpha(\mathbb{I}) \leq \hat{\alpha} \leq \alpha(\mathbb{I})\right] \geq \frac{2}{3} .
$$

Proof. Let $k$ be a constant larger than $1 / c$. For simplicity, we will use intervals with endpoints in [2kn].

Given an input ( $S, i$ ) for Index, consider the following stream of intervals. We set $L$ to some large enough value; for example $L=n+2$ will be enough. Let $\sigma_{1}(S)$ be a stream that, for each $j \in S$, contains the $k$ open intervals $(j, L+j),(j+L, j+2 L), \ldots,(j+(k-1) L, j+k L)$. Thus $\sigma_{1}(S)$ has exactly $k|S|$ intervals. Let $\sigma_{2}(i)$ be the stream with $k+1$ zero-length closed intervals $[i, i],[i+L, i+L], \ldots,[i+k L, i+k L]$. Finally, let $\sigma(S, i)$ be the concatenation of $\sigma_{1}(S)$ and $\sigma_{2}(i)$. See Figure 10 for an example. Let $\mathbb{I}$ be the intervals in $\sigma(S, i)$. It is straightforward to see that $\alpha(\mathbb{I})$ is $k+1$ or $2 k+1$ : The greedy left-to-right optimum contains either all the intervals of $\sigma_{2}(i)$ or those together with $k$ intervals from $\sigma_{1}(S)$. This means that $\alpha(\mathbb{I})=2 k+1$ if and only if $\operatorname{Index}(S, i)=1$.

We use a protocol similar to that of the proof of Theorem 18. Alice simulates the data stream algorithm on $\sigma_{1}(S)$, Bob receives the data message from Alice and continues the algorithm on $\sigma_{2}(S)$, and Bob returns bit $\hat{\beta}=1$ if an only if $\hat{\alpha}>k+1$. With the same argument, and using the fact that $\left(\frac{1}{2}+c\right)(2 k+1)>k+1$ by the choice of $k$, we can prove that using $o(k n)=o(n)$ bits of memory one cannot distinguish, with probability at least $2 / 3$, whether $\alpha(\mathbb{I})=k+1$ or $\alpha(\mathbb{I})=2 k+1$. The result follows.

## References

[1] A. Z. Broder, M. Charikar, A. M. Frieze, and M. Mitzenmacher. Min-wise independent permutations. J. Comput. Syst. Sci. 60(3):630-659, 2000.
[2] A. Z. Broder, M. Charikar, and M. Mitzenmacher. A derandomization using min-wise independent permutations. J. Discrete Algorithms 1(1):11-20, 2003.
[3] A. Chakrabarti et al. CS49: Data stream algorithms lecture notes, fall 2011. Available at http: //www.cs.dartmouth.edu/~ac/Teach/CS49-Fall11/.
[4] E. Cohen, M. Datar, S. Fujiwara, A. Gionis, P. Indyk, R. Motwani, J. D. Ullman, and C. Yang. Finding interesting associations without support pruning. IEEE Transactions on Knowledge and Data Engineering 13(1):64-78, 2001, doi:http://doi.ieeecomputersociety.org/10.1109/69.908981.
[5] M. Datar and S. Muthukrishnan. Estimating rarity and similarity over data stream windows. ESA 2002, pp. 323-334. Springer, Lecture Notes in Computer Science 2461, 2002.
[6] Y. Emek, M. M. Halldórsson, and A. Rosén. Space-constrained interval selection. ICALP 2012 (1), pp. 302-313. Springer, Lecture Notes in Computer Science 7391, 2012. Full version available at http://arxiv.org/abs/1202.4326
[7] J. Feigenbaum, S. Kannan, A. McGregor, S. Suri, and J. Zhang. On graph problems in a semistreaming model. Theor. Comput. Sci. 348(2-3):207-216, 2005.
[8] B. V. Halldórsson, M. M. Halldórsson, E. Losievskaja, and M. Szegedy. Streaming algorithms for independent sets. ICALP 2010 (1), pp. 641-652. Springer, Lecture Notes in Computer Science 6198, 2010.
[9] D. S. Hochbaum and W. Maass. Approximation schemes for covering and packing problems in image processing and vlsi. J. ACM 32(1):130-136, 1985.
[10] P. Indyk. A small approximately min-wise independent family of hash functions. J. Algorithms 38(1):84-90, 2001.
[11] T. S. Jayram, R. Kumar, and D. Sivakumar. The one-way communication complexity of hamming distance. Theory of Computing 4(6):129-135, 2008, doi:10.4086/toc.2008.v004a006, http://www theoryof computing.org/articles/v004a006.
[12] D. M. Kane, J. Nelson, and D. P. Woodruff. An optimal algorithm for the distinct elements problem. PODS 2010, pp. 41-52, 2010.
[13] D. E. Knuth. Estimating the efficiency of backtrack programs. Mathematics of Computation 29(129):122-136, 1975.
[14] A. W. Kolen, J. K. Lenstra, C. H. Papadimitriou, and F. C. Spieksma. Interval scheduling: A survey. Naval Research Logistics (NRL) 54(5):530-543, 2007.
[15] E. Kushilevitz and N. Nisan. Communication Complexity. Cambridge University Press, New York, NY, USA, 1997.
[16] S. Muthukrishnan. Data Streams: Algorithms and Applications. Foundations and trends in theoretical computer science. Now Publishers, 2005.


[^0]:    *Department of Mathematics, IMFM, and Department of Mathematics, FMF, University of Ljubljana, Slovenia. sergio.cabello@fmf.uni-lj.si Supported by the Slovenian Research Agency, program P1-0297, projects J1-4106 and L7-5459; by the ESF EuroGIGA project (project GReGAS) of the European Science Foundation. Part of the work was done while visiting Universidad de Valparaíso, Chile.
    ${ }^{\dagger}$ Escuela de Ingeniería Civil en Informática, Universidad de Valparaíso, Chile. pablo.perez@uv. cl Supported by project Millennium Nucleus Information and Coordination in Networks ICM/FIC RC130003 (Chile).

