

Succinct Dynamic Ordered Sets with Random Access*

GIULIO ERMANNANO PIBIRI, ISTI-CNR, Italy

ROSSANO VENTURINI, University of Pisa, Italy

The representation of a dynamic ordered set of n integer keys drawn from a universe of size m is a fundamental data structuring problem. Many solutions to this problem achieve optimal time but take polynomial space, therefore preserving time optimality in the *compressed* space regime is the problem we address in this work. For a polynomial universe $m = n^{\Theta(1)}$, we give a solution that takes $\text{EF}(n, m) + o(n)$ bits, where $\text{EF}(n, m) \leq n \lceil \log_2(m/n) \rceil + 2n$ is the cost in bits of the *Elias-Fano* representation of the set, and supports random access to the i -th smallest element in $O(\log n / \log \log n)$ time, updates and predecessor search in $O(\log \log n)$ time. These time bounds are optimal.

1 INTRODUCTION

The *dynamic ordered set problem* with integer keys is to represent a set $\mathcal{S} \subseteq [m] = \{0, \dots, m-1\}$, with $|\mathcal{S}| = n$, such that the following operations are supported: $\text{Search}(x)$ determines whether $x \in \mathcal{S}$; $\text{Insert/Delete}(x)$ inserts/deletes x in/from \mathcal{S} ; $\text{Predecessor/Successor}(x)$ returns the next smaller/larger element from \mathcal{S} ; $\text{Minimum/Maximum}()$ returns the smallest/largest element from \mathcal{S} . This is among the most studied problems in Computer Science (see the introduction to parts III and V of the book by Cormen et al. [6]). Many solutions to this problem are known to require an optimal amount of time per operation within polynomial space. For example, under the comparison-based model that allows only two keys to be compared in $O(1)$ time, it is well-known that any self-balancing search tree data structure, such as AVL or Red-Black, solves the problem optimally in $O(\log n)$ worst-case time and $O(n)$ words of space. (Unless otherwise specified, all logarithms are binary throughout the article).

However, working with integer keys makes it possible to beat the $O(\log n)$ -time bound with a RAM model having word size $w = \Theta(\log m)$ bits [10, 19, 21, 25]. In this scenario, classical solutions include the *van Emde Boas tree* [21–23], *x/y-fast trie* [25] and the *fusion tree* [10] – historically the first data structure that broke the barrier of $\Omega(\log n)$, by exhibiting an improved running time of $O(\log_w n) = O(\log n / \log \log m)$.

In this work, we are interested in preserving the asymptotic time optimality for the operations under *compressed space*. A simple information-theoretic argument [15] shows that one needs at least $B(n, m) = \lceil \log \binom{m}{n} \rceil = n \log(em/n) - \Theta(n^2/m) - O(\log n)$ bits to represent \mathcal{S} ($e = 2.718$ is the base of the natural logarithm), because there are $\binom{m}{n}$ possible ways of selecting n integers out of m . The meaning of this bound is that *any* solution solving the problem in optimal time *but* taking polynomial space, i.e., $O(n^{\Theta(1)} \log m)$ bits, is actually $\Omega(n \log n)$ bits larger than necessary.

Interestingly, the *Elias-Fano* representation [7, 8] of the ordered set \mathcal{S} uses $\text{EF}(n, m) \leq n \lceil \log(m/n) \rceil + 2n$ bits which is at most $n \log(m/n) + 3n$ bits. For $n = o(\sqrt{m})$, we have that $B(n, m) \approx n \log(m/n) +$

*A preliminary version of this work was published in the Proceedings of the 28-th Annual Symposium on Combinatorial Pattern Matching (CPM 2017) [16], when the first author was a Ph.D. student at the University of Pisa, Italy.

Authors' addresses: Giulio Ermanno Pibiri, ISTI-CNR, Via Giuseppe Moruzzi 1, 56124, Pisa, Italy; Rossano Venturini, University of Pisa, Largo Bruno Pontecorvo 3, 56127, Pisa, Italy, giulio.ermanno.pibiri@isti.cnr.it, rossano.venturini@unipi.it.

1.44n bits, showing that Elias-Fano takes $B(n, m) + 1.56n$ bits. We conclude that Elias-Fano is at most $1.56n$ bits away from the information-theoretic minimum [11]. We describe Elias-Fano in Section 2.1.

Given the total order of \mathcal{S} , it is natural to extend the problem by also considering the operation Access that, given an index $0 \leq i < n$, returns the i -th smallest element from \mathcal{S} . (This operation is also known as Select.) It should also be noted that, for any key x , the operation Search(x) can be implemented by running Successor(x) and checking whether the returned value is equal to x or not. Furthermore, it is well-known that Predecessor and Successor have the same complexities and are solved similarly, thus we only discuss Predecessor. Lastly, returning the smallest/largest integer from \mathcal{S} can be trivially done by storing these elements explicitly in $O(\log m)$ bits (which is negligible compared to the space needed to represent \mathcal{S}) and updating them as needed upon insertions/deletions. For these reasons, the problem we consider in this article is formalized as follows.

PROBLEM 1. *Dynamic ordered set with random access* – Given a non-negative integer m , represent an ordered set $\mathcal{S} \subseteq [m]$ with $|\mathcal{S}| = n$, such that the following operations are supported for any x and $0 \leq i < n$:

- Access(i) returns the i -th smallest element from \mathcal{S} ,
- Insert(x) sets $\mathcal{S} = \mathcal{S} \cup \{x\}$,
- Delete(x) sets $\mathcal{S} = \mathcal{S} \setminus \{x\}$,
- Predecessor(x) = $\max\{y \in \mathcal{S} \mid y < x\}$.

Our contribution. In this article we describe a solution to Problem 1 whose space in bits is expressed in terms of $\text{EF}(n, m)$ – the cost of representing \mathcal{S} with Elias-Fano – and achieves optimal running times. We consider a unit-cost RAM model with word size $w = \Theta(\log m)$ bit, allowing multiplication. We study the asymptotic behaviour of the data structures, therefore we also assume, without loss of generality, that n is larger than a sufficiently big constant [15].

For the important and practical case where the integers come from a polynomial universe of size $m = n^{\Theta(1)}$, we give a solution that uses $\text{EF}(n, m) + o(n)$ bits, thus introducing a *sublinear redundancy* with respect to $\text{EF}(n, m)$, and supports: Access in $O(\log n / \log \log n)$ time, Insert, Delete and Predecessor in $O(\log \log n)$ time. The time bound for random access under updates matches a lower bound given by Fredman and Saks [9] for *dynamic selection*. Dynamic predecessor search, instead, matches a lower bound given by Pătraşcu and Thorup [17]. Our result significantly improves the space of the best known solution by Pătraşcu and Thorup [19] which takes optimal time but polynomial space, i.e., $O(n \log m)$ bits.

In Section 2 we discuss related work and preliminaries. The main result is described in Section 3. In Section 4 we develop a solution that achieves a better update time under the assumption that we can only add a key larger than the maximum in the set (and delete the maximum).

2 PRELIMINARIES

In this section we illustrate the context of our work, whose discussion is articulated in three parts. We first describe the static Elias-Fano representation because it is a key ingredient of our solutions. Then we discuss the results concerning the static predecessor and dynamic ordered set (and related) problems, stressing what lower bounds apply to these problems. Recall that we use a RAM model with word size $w = \Theta(\log m)$ bits.

Table 1. An example of Elias-Fano encoding.

S	3	4	7	13	14	15	21	25	36	38		54	62
<i>high</i>	0	0	0	0	0	0	0	0	1	1	1	1	1
	0	0	0	0	0	0	1	1	0	0	0	1	1
	0	0	0	1	1	1	0	1	0	0	1	0	1
<i>low</i>	0	1	1	1	1	1	1	0	1	1		1	1
	1	0	1	0	1	1	0	0	0	1		1	1
	1	0	1	1	0	1	1	1	0	0		0	0
H	1110			1110			10	10	110		0	10	10
L	011.100.111			101.110.111			101	001	100.110		110		110

2.1 Static Elias-Fano representation

LEMMA 1. Elias-Fano [7, 8]. *An ordered set $S \subseteq [m]$, with $|S| = n$, can be represented in $\text{EF}(n, m) + o(n)$ bits such that Access is supported in $O(1)$ and Predecessor in $O(1 + \log(m/n))$, where $\text{EF}(n, m) \leq n \lceil \log(m/n) \rceil + 2n$.*

Space complexity. Let $S[i]$ indicate the i -th smallest of S . We write each $S[i]$ in binary using $\lceil \log m \rceil$ bits. The binary representation of each integer is then split into two parts: a *low* part consisting in the right-most $\ell = \lceil \log(m/n) \rceil$ bits that we call *low bits* and a *high* part consisting in the remaining $\lceil \log m \rceil - \ell$ bits that we similarly call *high bits*. Let us call ℓ_i and h_i the values of low and high bits of $S[i]$ respectively. The integers $L = [\ell_0, \dots, \ell_{n-1}]$ are written explicitly in $n \lceil \log(m/n) \rceil$ bits and they represent the encoding of the low parts. Concerning the high bits, we represent them in *negated unary* using a bitmap of $n + 2^{\lceil \log n \rceil} \leq 2n$ bits as follows. We start from a 0-valued bitmap H and set the bit in position $h_i + i$, for $i = 0, \dots, n - 1$. It is easy to see that the k -th unary value of H , say n_k , indicates that n_k integers of S have high bits equal to k . For example, if H is $\{1110, 1110, 10, 10, 110, 0, 10, 10\}$ (as in Table 1), we have that $H[1] = 1110$, so we know that there are 3 integers in S having high bits equal to 1.

Summing up the costs of high and low parts, we derive that Elias-Fano takes $\text{EF}(n, m) \leq n \lceil \log(m/n) \rceil + 2n$ bits. Although we can opt for an arbitrary split into high and low parts, ranging from 0 to $\lceil \log m \rceil$, it can be shown that $\ell = \lceil \log(m/n) \rceil$ minimizes the overall space of the encoding [7]. As explained in Section 1, the space of Elias-Fano is related to the information-theoretic minimum: it is at most $1.56n$ bits redundant.

Example. Table 1 shows a graphical example for the sorted set $S = [3, 4, 7, 13, 14, 15, 21, 25, 36, 38, 54, 62]$. The *missing high bits* embody the representation of the fact that using $\lceil \log_2 n \rceil$ bits to represent the high part of an integer, we have *at most* $2^{\lceil \log_2 n \rceil}$ distinct high parts because not all of them could be present. In Table 1, we have $\lceil \log_2 12 \rceil = 3$ and we can form up to 8 distinct high parts. Notice that, for example, no integer has high part equal to 101 which are, therefore, “missing” high bits.

Random access. A remarkable property of Elias-Fano is that it can be *indexed* to support Access in $O(1)$ worst-case. The operation is implemented by building an auxiliary data structure on top of H that answers Select_1 queries. The answer to a $\text{Select}_b(i)$ query over a bitmap is the position of the i -th bit set to b . This auxiliary data structure is *succinct* in the sense that it is negligibly small in asymptotic terms, compared to $\text{EF}(n, m)$, requiring only $o(n)$ additional bits [13, 24], hence bringing the total space of the encoding to $\text{EF}(n, m) + o(n)$ bits. For a given $i \in [0, n]$, we proceed as follows. The low bits ℓ_i are trivially retrieved as $L[i\ell, (i + 1)\ell]$. The retrieval of the high bits is,

instead, more complicated. Since we write in negated unary how many integers share the same high part, we have a 1 bit for every integer in \mathcal{S} and a 0 for every distinct high part. Therefore, to retrieve h_i , we need to know how many 0s are present in $H[0, \text{Select}_1(i))$. This quantity is evaluated on H in $O(1)$ as $\text{Select}_1(i) - i$. Lastly, re-linking the high and low bits together is as simple as: $\text{Access}(i) = ((\text{Select}_1(i) - i) \ll \ell) \mid \ell_i$, where \ll indicates the left shift operator and \mid is the bitwise OR.

Predecessor search. The query $\text{Predecessor}(x)$ is supported in $O(1 + \log(m/n))$ time as follows. Let h_x be the high bits of x . Then for $h_x > 0$, $i = \text{Select}_0(h_x) - h_x + 1$ indicates that there are i integers in \mathcal{S} whose high bits are less than h_x . On the other hand, $j = \text{Select}_0(h_x + 1) - h_x$ gives us the position at which the elements having high bits larger than h_x start. The corner case $h_x = 0$ is handled by setting $i = 0$. These two preliminary operations take $O(1)$. Now we can conclude the search in the range $\mathcal{S}[i, j]$, having *skipped* a potentially large range of elements that, otherwise, would have required to be compared with x . The range may contain up to u/n integers that we search with binary search. The time bound follows. In particular, it could be that $\text{Predecessor}(x) < \mathcal{S}[i]$: in this case $\mathcal{S}[i - 1]$ is the element to return if $i > 0$.

Partitioning the representation. In this article we will use extensively the following property of Elias-Fano.

PROPERTY 1. *Given an ordered set $\mathcal{S} \subseteq \{0, \dots, m\}$, with $|\mathcal{S}| = n$, let $\text{EF}(\mathcal{S}[i, j])$ indicate the Elias-Fano representation of $\mathcal{S}[i, j]$, for any $0 \leq i < j \leq n$. Then given an index $k \in [1, n]$, we have that $\text{EF}(\mathcal{S}[0, k]) + \text{EF}(\mathcal{S}'[k, n]) \leq \text{EF}(\mathcal{S}[0, n])$, where $\mathcal{S}'[l] = \mathcal{S}[l] - \mathcal{S}[k - 1] + 1$, for $l = k, \dots, n - 1$.*

The property tells us that splitting the Elias-Fano encoding of \mathcal{S} does not increase its space of representation. This is possible because each segment can be encoded with a *reduced universe*, by subtracting to each integer the last value of the preceding segment (the first segment is left as it is). Informally, we say that a segment is “re-mapped” relatively to its own universe. The property can be easily extended to work with an arbitrary number of splits. Let us now prove it.

PROOF. We know that $\text{EF}(\mathcal{S}[0, n])$ takes $n\phi + n + \lceil m/2^\phi \rceil$ bits, where $\phi = \lceil \log(m/n) \rceil$. Similarly, $\text{EF}(\mathcal{S}[0, k]) = k\phi_1 + k + \lceil m_1/2^{\phi_1} \rceil$ and $\text{EF}(\mathcal{S}'[k, n]) = (n-k)\phi_2 + (n-k) + \lceil m_2/2^{\phi_2} \rceil$, where $m_1 = \mathcal{S}[k-1]$ and $m_2 = \mathcal{S}[n-1] - m_1 + 1$, are minimized by choosing $\phi_1 = \lceil \log(m_1/k) \rceil$ and $\phi_2 = \lceil \log(m_2/(n-k)) \rceil$. Any other choice of ϕ_1 and ϕ_2 yields a larger cost, therefore: $\text{EF}(\mathcal{S}[0, k]) + \text{EF}(\mathcal{S}[k, n]) \leq k\phi + (n-k)\phi + k + (n-k) + \lceil m_1/2^\phi \rceil + \lceil m_2/2^\phi \rceil \leq n\phi + n + \lceil m/2^\phi \rceil = \text{EF}(\mathcal{S}[0, n])$. \square

An important consideration to make is that Property 1 needs the knowledge of the value $\mathcal{S}[k - 1]$ to work — the pivoting element — which can be stored in $O(\log m)$ bits. This means that for small values of n it can happen that the space reduction does *not* exceed $O(\log m)$ bits. Since we do not deal with such values of n , we always assume that this is not the case.

2.2 The static predecessor problem

Simple solutions. There are two simple solutions to the static predecessor problem. The first uses an array $P[0..m]$ where we store the answers to all possible queries. In this case $\text{Predecessor}(x) = P[x]$ for any $x < m$ ($P[0] = -\infty$), thus the problem is solved in $O(1)$ worst-case time and $m \lceil \log m \rceil$ bits. The second solution stores \mathcal{S} as a sorted array and answers the queries using binary search, therefore taking $n \lceil \log m \rceil$ bits and $O(\log n)$ worst-case time. Both solutions are unsatisfactory: the first one because of its space; the second one because of its time.

Lower bounds. Ajtai [1] proved the first $\omega(1)$ -time lower bound for polynomial space, i.e., $O(n^{\Theta(1)})$ memory words, claiming that $\forall w, \exists n$ that gives $\Omega(\sqrt{\log w})$ query time. Miltersen [14] elaborated on Ajtai's result and also showed that $\forall n, \exists w$ that gives $\Omega(\sqrt[3]{\log n})$ query time.

For the *dense* case of $m = n(\log n)^{O(1)}$, Pagh [15] gave a static data structure taking $\mathcal{B} + o(n)$ bits and answering membership and predecessor queries in $O(1)$ worst-case time. (We consider larger universes in this article.)

Beame and Fich [3, 4] proved two strong bounds for any cell-probe data structure. They proved that $\forall w, \exists n$ that requires $\Omega(\log w / \log \log w)$ query time and that $\forall n, \exists w$ that requires $\Omega(\sqrt{\log n / \log \log n})$ query time. They also gave a static data structure achieving

$$O(\min\{\log w / \log \log w, \sqrt{\log n / \log \log n}\})$$

which is, therefore, optimal.

Building on a long line of research, Pătraşcu and Thorup [17, 18] finally proved the following optimal (up to constant factors) space/time trade-off.

THEOREM 1. Pătraşcu and Thorup [17, 18]. *A static data structure representing n integer keys in z bits, takes time*

$$\Theta\left(\min\left\{\frac{\log n}{\log w}, \log \frac{w - \log n}{a}, \frac{\log(w/a)}{\log(\frac{a}{\log n} \log(w/a))}, \frac{\log(w/a)}{\log(\log(w/a)/\log \frac{\log n}{a})}\right\}\right)$$

to answer a Predecessor query, where $a = \log(z/n)$.

This lower bound holds for cell-probe, RAM, trans-dichotomous RAM, external memory and communication game models. The first branch of the trade-off indicates that, whenever one integer fits in one memory word, fusion trees [10] are optimal as they have $O(\log n / \log w)$ query time. The second branch holds for *polynomial universes*, i.e., when $m = n^\gamma$, for any $\gamma = \Theta(1)$. In such important case we have that $w = \Theta(\log m) = \Theta(\log n)$, therefore y -fast tries [25] and van Emde Boas trees [21–23] are optimal with query time $O(\log \log n)$. The last two bounds of the trade-off, instead, treat the case for *super-polynomial* universes and are out of scope for this work.

For example, given a space budget of $O(n \text{ polylog } n)$ words we have $a = O(\log \log n)$, thus implying that y -fast tries and van Emde Boas trees are optimal if $w = O(\text{polylog } n)$ and fusion trees are optimal if $\log w = \Omega(\sqrt{\log n} \cdot \log \log n)$.

Predecessor queries in succinct space. We are now interested in determining the optimal running time of Predecessor given the Elias-Fano space bound of $\text{EF}(n, m) + o(n)$ bits from Lemma 1, knowing that the time for dynamic predecessor with logarithmic update time can not be better than that of static predecessor (allowing polynomial space) [19].

We make the following observation.

OBSERVATION 1. *Given any linear-space data structure supporting Predecessor in $O(t)$ worst-case time, an ordered set $\mathcal{S} \subseteq [m]$ with $|\mathcal{S}| = n$ can be represented in $\text{EF}(n, m) + O(n/2^{ct} \cdot \log m) + o(n)$ bits such that Access is supported in $O(1)$ and Predecessor in $O(t)$ worst-case time, for any constant $c > 1$.*

We represent \mathcal{S} with Elias-Fano and (logically) divide it into $\lceil n/2^{ct} \rceil$ blocks of 2^{ct} integers each (the last block may contain less integers). We can solve Predecessor queries in a block in $O(t)$ time by applying binary search, given that each access is performed in $O(1)$ time. The first element of each block (and its position in \mathcal{S}) is also stored in the linear-space data structure solving Predecessor in $O(t)$ time. The space of such data structure is $O(n/2^{ct} \cdot \log m)$ bits.

COROLLARY 1. *An ordered set $\mathcal{S} \subseteq [m]$, with $|\mathcal{S}| = n$ and $m = n^{\Theta(1)}$, can be represented in $\text{EF}(n, m) + o(n)$ bits such that Access is supported in $O(1)$ and Predecessor in optimal $O(\min\{1 + \log(m/n), \log \log n\})$ worst-case time.*

The linear-space data structure in Observation 1 is chosen to be an y -fast trie, whose $t = \log \log n$ query time is optimal for polynomial universes (second branch of Theorem 1). The space of the y -fast trie is $O(n/(\log n)^{c-1}) = o(n)$ bits.

Let $m = n^\gamma$, for any $\gamma = \Theta(1)$. The bound $O(\log \log n)$ only depends on n , whereas the plain Elias-Fano bound of $O(1 + \log(m/n))$ depends on both n and m , thus varying γ only one of the two bounds is optimal. In fact, we have that $1 + \log(m/n) \leq \log \log n$ whenever $m \leq \frac{n}{2} \log n$, i.e., when $n^{\gamma-1} \leq \frac{1}{2} \log n$. From this last condition we derive that the plain Elias-Fano bound is less than $\log \log n$ when $1 \leq \gamma \leq 1 + \log \log n / \log n$. When, instead, $\gamma > 1 + \log \log n / \log n$, the query time $O(\log \log n)$ is optimal and exponentially better than Elias-Fano. Therefore, $O(\min\{1 + \log(m/n), \log \log n\})$ is an accurate characterization of the Predecessor time bound with $\text{EF}(n, m) + o(n)$ bits.

However for the rest of the discussion, we assume that m is sufficiently large so that $\log \log n < 1 + \log(m/n)$, that is $m > \frac{n}{2} \log n$.

2.3 Dynamic problems

Ordered set problem. As far as the Access operation is *not* supported, the following results hold. The van Emde Boas tree [21–23] is a recursive data structure that maintains \mathcal{S} in $O(m \log m)$ bits and $O(\log w)$ worst-case time. Willard [25] improved the space bound to $O(n \log m)$ bits with the y -fast trie. (The bound for Insert/Delete is amortized rather than worst-case). When polynomial universes are considered, Pătraşcu and Thorup [17] proved that van Emde Boas trees and y -fast tries have an optimal query time for the dynamic predecessor problem too, that is $O(\log \log n)$ worst-case.

Fredman and Willard [10] solved that dynamic predecessor problem in $O(\log n / \log \log m)$ time and $O(n)$ space with the fusion tree. This data structure is a B -tree with branching factor $B = w^{O(1)}$ that stores in each internal node a *fusion node* a small data structure able of answering predecessor queries in $O(1)$ for sets up to B integers.

Extending their result to the dynamic predecessor problem, Beame and Fich [3, 4] proved that any cell-probe data structure using $(\log m)^{O(1)}$ bits per memory cell and $n^{O(1)}$ worst-case time for insertions, requires $\Omega(\sqrt{\log n / \log \log n})$ worst-case query time. They also proved that, under a RAM model, the dynamic predecessor problem can be solved in $O(\min\{\log \log n \cdot \log w / \log \log w, \sqrt{\log n / \log \log n}\})$, using linear space. This bound was matched by Andersson and Thorup [2] with the so-called *exponential search tree*. This data structure has an optimal bound of $O(\sqrt{\log n / \log \log n})$ worst-case time for searching and updating \mathcal{S} , using polynomial space.

Set problems with random access. The lower bound for the problem changes by considering the Access operation because this operation is related to the *partial sums problem* that is, given an integer array $A[0..n)$, support $\text{Sum}(i)$ returning the sum of the first $i + 1$ integers, $\text{Update}(i, \Delta)$ which sets $A[i]$ to $A[i] + \Delta$ and $\text{Search}(x)$ which returns the index $i \in [0, n)$ such that $\text{Sum}(i) < x \leq \text{Sum}(i + 1)$. Fredman and Saks [9] proved a bound of $\Omega(\log n / \log \log n)$ amortized time for this problem (see also the extended version of the article by Pătraşcu and Thorup [19] – Section 5). Therefore, this is the lower bound that applies to our problem as well. Bille et al. [5] extended the problem as to also support dynamic changes to the array.

Fredman and Saks [9] also proved that $\Omega(\log n / \log \log n)$ amortized is necessary for the *list representation problem*, that is to support Access, Insert and Delete. However, this problem is slightly different than the one tackled here, because one can specify the *position* of insertion of a

key. Likewise, the Delete operation specify the position of the key, rather than the key itself. Raman, Raman, and Rao [20] also addressed the list representation problem (referred to as the *dynamic array problem*) and provide two solutions. The first solution is given by the following lemma.

LEMMA 2. Raman, Raman, and Rao [20]. *A dynamic array containing n elements can be implemented to support Access in $O(1)$, Insert and Delete in $O(n^\epsilon)$ time using $O(n^{1-\epsilon})$ pointers, where ϵ is any fixed positive constant.*

The second solution supports all the three operations in $O(\log n / \log \log n)$ amortized time. Both solutions take $o(n)$ bits of redundancy (besides the space needed to store the array) and the time bounds are optimal.

Since it takes, $O(B^4)$ time to construct and update a fusion node with B keys, Pătraşcu and Thorup [19] showed that it is possible to “dynamize” the fusion node and obtained the following result.

LEMMA 3. Pătraşcu and Thorup [19]. *An ordered set $\mathcal{S} \subseteq [m]$, with $|\mathcal{S}| = n$, can be represented in $O(n \log m)$ bits and supporting Insert, Delete, Rank, Select and Predecessor in $O(\log n / \log \log m)$ per operation.*

The time bound of $O(\log n / \log \log m)$ is optimal, matching a lower bound by Fredman and Saks [9] for dynamic ranking and selection, and that of predecessor queries for non-polynomial universes (first branch of the trade-off from Theorem 1).

3 SUCCINCT DYNAMIC ORDERED SETS WITH RANDOM ACCESS

In this section we illustrate our main result for polynomial universes: a solution to Problem 1 that uses $EF(n, m) + o(n)$ bits and supports all operations in optimal time. From Section 2.3, we recall that a bound of $\Omega(\log n / \log \log n)$ applies to the Access operation under updates; Predecessor search needs, instead, $\Omega(\log \log n)$ time as explained in Section 2.2.

THEOREM 2. *An ordered set $\mathcal{S} \subseteq [m]$, with $|\mathcal{S}| = n$ and $m = n^{\Theta(1)}$, can be represented in $EF(n, m) + o(n)$ bits such that Access is supported in $O(\log n / \log \log n)$, Insert, Delete and Predecessor in $O(\log \log n)$ time.*

We first show how to handle small sets of integers efficiently in Section 3.1. Then we use this solution to give the final construction in Section 3.2.

3.1 Handling small sets

In this section, we give a solution to Problem 1 working for a small set of integers.

The following lemma is useful.

LEMMA 4. Jansson et al. [12] *Given a collection of k blocks, each of size $O(b)$ bits, we can store it using $O(k \log k + b^2)$ bits of redundancy to support Address in $O(1)$ time and Realloc in $O(b/w)$ time.*

We say that the data structure of Lemma 4 has parameters (k, b) . The operation $\text{Address}(i)$ returns a pointer to where the i -th block is stored in memory; the operation $\text{Realloc}(i, b')$ changes the length of the i -th block to b' bits.

Now we show the following theorem.

THEOREM 3. *Let $\mathcal{S} \subseteq [m]$ be an ordered set with $|\mathcal{S}| = n$ and $m = n^{\Theta(1)}$. Then a subset \mathcal{S}' of \mathcal{S} , with $|\mathcal{S}'| = n' = \Theta((\log n \cdot \log \log n)^2)$ and $\mathcal{S}' \subseteq [m']$, can be represented with $EF(n', m') + O((\log n)^2 \cdot \log \log n) + o(n')$ bits and supporting Access, Insert, Delete and Predecessor in $O(\log \log n)$ time.*

Memory management. We divide the ordered elements of \mathcal{S}' into blocks of size $\Theta((\log \log n)^2)$ and represent each block with Elias-Fano. We have $O((\log n)^2)$ blocks. Physically, the high and low parts of the Elias-Fano representations are stored using two different data structures.

The high parts of all blocks are stored using the data structure of Lemma 4, with parameters $(O((\log n)^2), \Theta((\log \log n)^2))$. For this choice of parameters, we support both Address and Realloc in $O(1)$ time and pay a redundancy of $O((\log n)^2 \cdot \log \log n + (\log \log n)^4) = O((\log n)^2 \cdot \log \log n)$ bits. This allows to manipulate the high part of a block in $O(1)$ time upon Access, Insert and Delete.

The low parts are stored in a collection of $\Theta((\log \log n)^2)$ dynamic arrays, each being an instance of the data structure of Lemma 2. We maintain an array A of $\Theta((\log \log n)^2)$ pointers to such data structures, taking $O((\log \log n)^2 \cdot \log n) = O((\log n)^2)$ bits. Each array stores $O((\log n)^2)$ integers and supports Access in $O(1)$, Insert and Delete in $O(\log \log n)$ as soon as $\epsilon < 1/6$ in Lemma 2. The redundancy to maintain the arrays is $o(n')$ bits.

Indexing. The blocks are indexed with a τ -ary tree \mathcal{T} , with $\tau = (\log n)^\sigma$ and $0 < \sigma < 1$. It follows that the height of the tree is constant and equal to $h = O(\log_\tau(\log n)^2) = O(1/\sigma)$. The tree operates as a B-tree where internal nodes store $\Theta(\tau)$ children. In particular, each node stores $\Theta(\tau)$ counters, telling how many integers are present in the leaves that descend from each child. These counters are kept in prefix-sum fashion to enable binary search. Such counters takes $O(\tau \log \log n) = o(\log n)$ bits which fit in (less than) a machine word. This allows us to update such counters in $O(1)$ time upon insertions/deletions.

Each leaf node also stores two offsets per block, each taking $O(\log \log n)$ bits. The first offset is the position in A of the pointer to the dynamic array storing the low parts of the Elias-Fano representation of the block. The second offset tells where the low parts of the block are stored inside the dynamic array. Thus the overhead per block is $O(\log \log n)$ bits. As usual, each internal node also stores a pointer per child, thus maintaining the tree topology imposes an overhead per block equal to $O(\log n/\tau) = O((\log n)^{1-\sigma}) = O(\log \log n)$ bits as soon as $\sigma \geq 2/3$. Since the overhead per block is $O(\log \log n)$ bits, it follows that the total space of \mathcal{T} is $O((\log n)^2 \cdot \log \log n)$ bits.

Operations. To support Access, we navigate the tree spending $O(\log \tau) = O(\sigma \log \log n)$ per level by binary searching the counters. The proper block is therefore identified in $O(h \times \sigma \log \log n) = O(\log \log n)$ and the wanted integer is returned in $O(1)$ time from it knowing the local offset of the integer inside the block calculated during the traversal.

To support Insert, we need to identify the proper block where to insert the new integer. (The Delete operation is symmetric.) Again, we use binary search on each level of the tree but searching among the last values of the indexed blocks. We can retrieve the last value of a block in $O(1)$, having the pointer to the block and its size information from the counters. This is trivial at the leaves. In the internal nodes, instead, if the upper bound of the i -th child is needed for comparison for some $1 \leq i \leq \Theta(\tau)$, we access the block storing such value by following the pointer to the *right-most* block indexed in the sub-tree rooted in the i -th child. Accessing the right-most block takes $O(1)$ time. Having located the proper block, we insert the new integer in $O(\log \log n)$ time, as explained before. Updating the counters in each node of the tree along the root-to-leaf path takes $O(1)$ time as they fit in $o(\log n)$ bits. If a split or merge of a block happens, it is handled as in a B-tree and solved in a constant number of $O(1)$ -time operations.

During a Predecessor search we identify the proper block in $O(\log \log n)$ time as explained for Insert and return the predecessor by binary searching the block's values. The total time of the search is $O(\log \log n)$.

Space complexity. We now analyze the space taken by the Elias-Fano representations of the blocks. Our goal is to show that such space can be bounded by $EF(n', m')$, that is the space of

encoding the set \mathcal{S}' with Elias-Fano. Since the universe of representation of a block could be as large as m' , storing the lower bounds of the blocks in order to use reduced universes — as for Property 1 — would require $O((\log n)^3)$ bits of redundancy. This is excessive because if the data structure is replicated every n' integers to represent a larger dynamic set \mathcal{S} with $|\mathcal{S}| = n$, then these lower bounds would cost $O(n/(\log \log n)^2 \cdot \log n)$ bits, which is *not* sub-linear in n . We show that this extra space can be avoided, observing that the number of bits used to represent the low part of Elias-Fano remains the same for a sufficiently long sequence of p updates.

From Section 2.1 recall that Elias-Fano represents each low part with $\lceil \phi \rceil = \lceil \log(m'/n') \rceil$ bits. Now, suppose that the low parts of the blocks are encoded using a sub-optimal value $\lceil \mu \rceil$ instead of $\lceil \phi \rceil$. After we perform p updates, $\lceil \mu \rceil = \lceil \log(m'/(n' \pm p)) \rceil$ is set to $\lceil \phi \rceil$ by rebuilding the blocks. It is easy to see that n' updates are required to let $\lceil \mu \rceil$ become $\lceil \phi \rceil \pm 1$, because $\lceil \log(\cdot) \rceil$ changes by $+1$ (-1) whenever its argument doubles (halves). Therefore we have $\lceil \mu \rceil = \lceil \phi \rceil$ for any $p < n'$. In our case $n' = \Theta((\log n \cdot \log \log n)^2)$. In order to guarantee an amortized cost for update equal to $O(\log \log n)$, we set $p = O((\log n)^2 \cdot \log \log n)$. Storing the current value of $\lceil \mu \rceil$ adds a global redundancy of $\Theta(\log n)$ bits which is negligible.

3.2 Final construction

Now we prove the final result – Theorem 2 – whose key ingredient is the data structure given in Theorem 3.

Lower level. We divide the ordered elements of \mathcal{S} into blocks of size $\Theta((\log n \cdot \log \log n)^2)$ and represents them using the tree data structure of Theorem 3. Therefore, we have a forest $\{\mathcal{T}_i\}$ of $k = \Theta(n/(\log n \cdot \log \log n)^2)$ such data structures.

Upper level. The first element of each block is (also) stored in the data structure of Lemma 3 that is a dynamic fusion tree with out-degree $\Theta(\log n)$, and in a y -fast trie. Let call these data structures \mathcal{F} and \mathcal{Y} respectively. The i -th leaf of both \mathcal{F} and \mathcal{Y} holds a pointer to the data structure \mathcal{T}_i .

Space and time complexity. The lower level costs $O(k \cdot (\log n)^2 \cdot \log \log n) + o(n) = O(n/\log \log n) + o(n) = o(n)$ bits. The total cost of the upper level is $O(k \cdot \log n) = O(n/(\log n \cdot (\log \log n)^2)) = o(n)$ bits. Since each block is re-mapped relatively to its universe, Property 1 guarantees that the space of representation is at most $\text{EF}(n, m)$ bits. The space bound claimed in Theorem 2 follows.

A total running time of $O(\log n/\log \log n)$ for Access follows because the \mathcal{F} data structure operates in this time. For Insert, Delete and Predecessor, we use the \mathcal{Y} data structure, thus attaining to $O(\log \log n)$ time. (The bound for Insert and Delete is amortized rather than worst-case).

4 APPEND-ONLY

In this section we extend the result given in Corollary 1 to the case where the integers are inserted in sorted order using an Append operation. In this case, we obtain an append-only representation.

THEOREM 4. *An ordered set $\mathcal{S} \subseteq [m]$, with $|\mathcal{S}| = n$ and $m = n^{\Theta(1)}$, can be represented in $\text{EF}(n, m) + o(n)$ bits such that Append and Access are supported in $O(1)$ time, Predecessor in $O(\log \log n)$ time.*

Data structure and space analysis. We maintain an array $A[0..k]$ of size $k = O((\log n)^c)$ where integers are appended uncompressed, for any $c > 1$. The array is periodically encoded with Elias-Fano in $\Theta(k)$ time and overwritten. Each compressed representation of the buffer is appended to another array of blocks encoded with Elias-Fano. More precisely, when A is full we encode with Elias-Fano its corresponding *differential* buffer, i.e., the buffer whose values are $A[i] - A[0]$, for $0 \leq i < k$. Each time the array is compressed, we append in another array A' the pair (*base*, *low*) = $(A[0], \lceil \log(A[k-1]/k) \rceil)$, i.e., the buffer lower bound value (*base*) and the number of bits (*low*) needed to encode the average gap of the Elias-Fano representation of the block.

As discussed for Corollary 1, we store the buffer lower bounds in a y -fast trie. More precisely, it stores a buffer lower bound and the index of the Elias-Fano-encoded block to which the lower bound belongs to. The space of this data structure is $o(n)$ bits. Besides the space of the y -fast trie, which is $o(n)$ bits, and that of the Elias-Fano-encoded blocks, the redundancy of the data structure is due to (1) $O((k + 1) \log n)$ bits for the array A and its (current) size; (2) $O(n/k \cdot \log n)$ bits for pointers to the Elias-Fano-encoded blocks; (3) $O(n/k \cdot \log n)$ bits for the array A' ; and it sums up to $o(n)$ bits.

Lastly, Property 1 guarantees that the space taken by the blocks encoded with Elias-Fano can be safely upper bounded by $\text{EF}(n, m)$ so that the overall space of the data structure is at most $\text{EF}(n, m) + o(n)$ bits.

Operations. The operations are supported as follows. Since we compress the array A each time it fills up (by taking $\Theta(k)$ time), Append is performed in $O(1)$ amortized time. Appending new integers in the buffer accumulates a credit of $\Theta(k)$ that (largely) pays the cost $O(\log \log n)$ of appending a value to the y -fast trie. To Access the i -th integer, we retrieve the element x in position $i - p \times k$ from the compressed block of index $p = \lfloor i/k \rfloor$. This is done in $O(1)$ worst-case time, since we know how many low bits are required to perform Access by reading $C[p].\text{low}$. We finally return the integer $x + C[p].\text{base}$. To solve Predecessor(x), we first resolve a partial Predecessor(x) query in the y -fast trie to identify the index k of the compressed block where the predecessor is located. This takes $O(\log \log n)$ worst-case time. We return $C[p].\text{base} + \text{Predecessor}(x - C[p].\text{base})$ by binary searching the block of index p in $O(\log \log n)$ worst-case time.

5 CONCLUSIONS

In this paper we have shown that Elias-Fano can be used to obtain a succinct dynamic data structure with optimal update and query time, solving the *dynamic ordered set with random access* problem. Our main result holds for polynomial universes and is a data structure using the same asymptotic space of Elias-Fano — $\text{EF}(n, m) + o(n)$ bits, where $\text{EF}(n, m) \leq n \lceil \log(m/n) \rceil + 2n$ — and supporting Access in $O(\log n / \log \log n)$ time, Insert, Delete and Predecessor in $O(\log \log n)$ time. All time bounds are optimal. Note that the space of the solution can be rewritten in terms of information-theoretic minimum $B(n, m) = \lceil \log \binom{m}{n} \rceil$ since $\text{EF}(n, m) = B(n, m) + 1.56n$ bits.

An interesting open problem is: *Can the space be improved to $B(n, m) + o(n)$ bits and preserving the operational bounds?*

Another question is: *Can the result be extended to non-polynomial universes?*

In this case, the lower bound for dynamic predecessor search is $O(\log_w n) = O(\log n / \log \log m)$ that corresponds to the *first* branch of the time/space trade-off in Theorem 1, as well as the one for Access, Insert and Delete [19]. It seems that a different solution than the one described here has to be found since the data structure of Theorem 2 allows us to support all operations in time $O(\log \log m)$ when non-polynomial universes are considered. Therefore, we give the following corollary that matches the asymptotic time bounds of y -fast tries and van Emde Boas trees (albeit sub-optimal) but in almost optimally compressed space.

COROLLARY 2. *An ordered set $S \subseteq \{0, \dots, m\}$, with $|S| = n$, can be represented in $\text{EF}(n, m) + o(n)$ bits such that Access, Insert, Delete and Predecessor are all supported in $O(\log \log m)$ time.*

ACKNOWLEDGMENTS

This work was partially supported by the BIGDATAGRAPES project (grant agreement #780751, European Union’s Horizon 2020 research and innovation programme) and MIUR-PRIN 2017 “Algorithms, Data Structures and Combinatorics for Machine Learning”.

REFERENCES

- [1] Miklós Ajtai. A lower bound for finding predecessors in Yao’s cell probe model. *Combinatorica*, 8(3):235–247, 1988.
- [2] Arne Andersson and Mikkel Thorup. Dynamic ordered sets with exponential search trees. *Journal of the ACM (JACM)*, 54(3):13, 2007.
- [3] Paul Beame and Faith E. Fich. Optimal bounds for the predecessor problem. In *Proceedings of Annual Symposium on Theory of Computing (STOC)*, pages 295–304, 1999.
- [4] Paul Beame and Faith E. Fich. Optimal bounds for the predecessor problem and related problems. *Journal of Computer and System Sciences (JCSS)*, 65(1):38–72, 2002.
- [5] Philip Bille, Anders Roy Christiansen, Patrick Hagge Cording, Inge Li Gørtz, Frederik Rye Skjoldjensen, Hjalte Wedel Vildhøj, and Søren Vind. Dynamic relative compression, dynamic partial sums, and substring concatenation. *Algorithmica*, 80(11):3207–3224, 2018.
- [6] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to Algorithms (3-rd Edition)*. MIT Press, 2009.
- [7] Peter Elias. Efficient storage and retrieval by content and address of static files. *Journal of the ACM (JACM)*, 21(2):246–260, 1974.
- [8] Robert Mario Fano. On the number of bits required to implement an associative memory. *Memorandum 61, Computer Structures Group, MIT, Cambridge, MA*, 1971.
- [9] Michael L. Fredman and Michael E. Saks. The cell probe complexity of dynamic data structures. In *Proceedings of the 21-st Annual Symposium on Theory of Computing (STOC)*, pages 345–354, 1989.
- [10] Michael L. Fredman and Dan E. Willard. Surpassing the information theoretic bound with fusion trees. *Journal of Computer and System Sciences (JCSS)*, 47(3):424–436, 1993.
- [11] Roberto Grossi, Alessio Orlandi, Rajeev Raman, and S Srinivasa Rao. More haste, less waste: Lowering the redundancy in fully indexable dictionaries. In *26th International Symposium on Theoretical Aspects of Computer Science STACS 2009*, pages 517–528. IBFI Schloss Dagstuhl, 2009.
- [12] Jesper Jansson, Kunihiko Sadakane, and Wing-Kin Sung. CRAM: Compressed random access memory. In *Proceedings of 39-th International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 510–521, 2012.
- [13] Veli Mäkinen and Gonzalo Navarro. Rank and select revisited and extended. *Theoretical Computer Science (TCS)*, 387(3):332–347, 2007.
- [14] Peter Bro Miltersen. Lower bounds for union-split-find related problems on random access machines. In *Proceedings of Annual Symposium on Theory of Computing (STOC)*, volume 94, pages 625–634, 1994.
- [15] Rasmus Pagh. Low redundancy in static dictionaries with constant query time. *SIAM Journal on Computing*, 31(2):353–363, 2001.
- [16] Giulio Ermanno Pibiri and Rossano Venturini. Dynamic Elias-Fano representation. In *Proceedings of the 28-th Annual Symposium on Combinatorial Pattern Matching (CPM)*, pages 30:1–30:14, 2017.
- [17] Mihai Pătraşcu and Mikkel Thorup. Time-space trade-offs for predecessor search. In *Proceedings of the 38-th Annual Symposium on Theory of Computing (STOC)*, pages 232–240, 2006.
- [18] Mihai Pătraşcu and Mikkel Thorup. Randomization does not help searching predecessors. In *Proceedings of the 18-th Annual Symposium on Discrete Algorithms (SODA)*, pages 555–564, 2007.
- [19] Mihai Pătraşcu and Mikkel Thorup. Dynamic integer sets with optimal rank, select, and predecessor search. In *Proceedings of the 55-th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 166–175, 2014.
- [20] Rajeev Raman, Venkatesh Raman, and S. Srinivasa Rao. Succinct dynamic data structures. In *Proceedings of the 7-th International Workshop on Algorithms and Data Structures (WADS)*, pages 426–437, 2001.
- [21] Peter van Emde Boas. Preserving order in a forest in less than logarithmic time. In *Proceedings of the 16-th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 75–84, 1975.

- [22] Peter van Emde Boas. Preserving order in a forest in less than logarithmic time and linear space. *Information Processing Letters (IPL)*, 6(3):80–82, 1977.
- [23] Peter van Emde Boas, Robert Kaas, and Erik Zijlstra. Design and implementation of an efficient priority queue. *Mathematical Systems Theory (MST)*, 10:99–127, 1977.
- [24] Sebastiano Vigna. Quasi-succinct indices. In *Proceedings of the 6-th International Conference on Web Search and Data Mining (WSDM)*, pages 83–92, 2013.
- [25] Dan E. Willard. Log-logarithmic worst-case range queries are possible in space $\theta(n)$. *Information Processing Letters (IPL)*, 17(2):81–84, 1983.