

Pattern-avoiding access in binary search trees

Parinya Chalermsook¹, Mayank Goswami¹, László Kozma², Kurt Mehlhorn¹, and Thatchaphol Saranurak^{*3}

¹Max-Planck Institute for Informatics, Germany. {parinya,gmayank,mehlhorn}@mpi-inf.mpg.de

²Saarland University, Germany. kozma@cs.uni-saarland.de

³KTH Royal Institute of Technology, Sweden. thasar@kth.se

Abstract

The *dynamic optimality conjecture* is perhaps the most fundamental open question about binary search trees (BST). It postulates the existence of an asymptotically optimal online BST, i.e. one that is *constant factor competitive* with any BST on any input access sequence. The two main candidates for dynamic optimality in the literature are *splay trees* [Sleator and Tarjan, 1985], and GREEDY [Lucas, 1988; Munro, 2000; Demaine et al. 2009]. Despite BSTs being among the simplest data structures in computer science, and despite extensive effort over the past three decades, the conjecture remains elusive. Dynamic optimality is trivial for *almost all* sequences: the optimum access cost of most length- n sequences is $\Theta(n \log n)$, achievable by any balanced BST. Thus, the obvious missing step towards the conjecture is an understanding of the “easy” access sequences, and indeed the most fruitful research direction so far has been the study of specific sequences, whose “easiness” is captured by a parameter of interest. For instance, splay provably achieves the bound of $O(nd)$ when d roughly measures the distances between consecutive accesses (dynamic finger), the average entropy (static optimality), or the delays between multiple accesses of an element (working set). The difficulty of proving dynamic optimality is witnessed by other highly restricted special cases that remain unresolved; one prominent example is the *traversal conjecture* [Sleator and Tarjan, 1985], which states that *preorder sequences* (whose optimum is linear) are linear-time accessed by splay trees; no online BST is known to satisfy this conjecture.

In this paper, we prove two different relaxations of the traversal conjecture for GREEDY: (i) GREEDY is almost linear for preorder traversal, (ii) if a linear-time preprocessing¹ is allowed, GREEDY is in fact linear. These statements are corollaries of our more general results that express the complexity of access sequences in terms of a *pattern avoidance* parameter k . Pattern avoidance is a well-established concept in combinatorics, and the classes of input sequences thus defined are rich, e.g. the $k = 3$ case includes preorder sequences. For any sequence X with parameter k , our most general result shows that GREEDY achieves the cost $n2^{\alpha(n)^{O(k)}}$ where α is the inverse Ackermann function. Furthermore, a broad subclass of parameter- k sequences has a natural combinatorial interpretation as *k-decomposable sequences*. For this class of inputs, we obtain an $n2^{O(k^2)}$ bound for GREEDY when preprocessing is allowed. For $k = 3$, these results imply (i) and (ii). To our knowledge, these are the first upper bounds for GREEDY that are not known to hold for any other online BST. To obtain these results we identify an *input-revealing* property of GREEDY. Informally, this means that the execution log partially reveals the structure of the access sequence. This property facilitates the use of rich technical tools from *forbidden submatrix theory*.

Further studying the intrinsic complexity of k -decomposable sequences, we make several observations. First, in order to obtain an offline optimal BST, it is enough to bound GREEDY on non-decomposable access sequences. Furthermore, we show that the optimal cost for k -decomposable sequences is $\Theta(n \log k)$, which is well below the proven performance of all known BST algorithms. Hence, sequences in this class can be seen as a “candidate counterexample” to dynamic optimality.

*Work partly done while at Saarland University.

¹The purpose of preprocessing is to bring the data structure into a state of our choice. This state is independent of the actual input. This kind of preprocessing is implicit in the Demaine et al. definition of GREEDY.

1 Introduction

The binary search tree (BST) model is one of the most fundamental and thoroughly studied data access models in computer science. In this model, given a sequence $X \in [n]^m$ of keys, we are interested in $\text{OPT}(X)$, the optimum cost of accessing X by a binary search tree. When the tree does not change between accesses, OPT is well understood: both efficient exact algorithms (Knuth [23]) and a linear time approximation (Mehlhorn [28]) have long been known. By contrast, in the dynamic BST model (where restructuring of the tree is allowed between accesses) our understanding of OPT is very limited. No polynomial time exact or constant-approximation algorithms are known for computing OPT in the dynamic BST model.

Theoretical interest in the dynamic BST model is partly due to the fact that it holds the possibility of a meaningful “instance optimality”, i.e. of a BST algorithm that is constant factor competitive (in the amortized sense) with any other BST algorithm, even with those that are allowed to see into the future. The existence of such a “dynamically optimal” algorithm has been postulated in 1985 by Sleator and Tarjan [34]; they conjectured their splay tree to have this property. Later, Lucas [25] and independently Munro [29] proposed a greedy offline algorithm as another candidate for dynamic optimality. In 2009 Demaine, Harmon, Iacono, Kane, and Pătraşcu (DHIKP) [11] presented a geometric view of the BST model, in which the Lucas-Munro offline algorithm naturally turns into an online one (simply called GREEDY). Both splay and GREEDY are considered plausible candidates for achieving optimality. However, despite decades of research, neither of the two algorithms are known to be $o(\log n)$ -competitive (the best known approximation ratio for any BST is $O(\log \log n)$ [12], but the algorithms achieving this bound are less natural). Thus, the dynamic optimality conjecture remains wide open.

Why is dynamic optimality difficult to prove (or disprove)? For a vast majority of access sequences, the optimal access cost by any dynamic BST (online or offline) is² $\Theta(m \log n)$ and any static balanced tree can achieve this bound.³ Hence, dynamic optimality is almost trivial on most sequences. However, when an input sequence has complexity $o(m \log n)$, a candidate BST must “learn” and “exploit” the specific structure of the sequence, in order to match the optimum. This observation motivates the popular research direction that aims at understanding “easy” sequences, i.e. those with $\text{OPT} = O(md)$ when d is a parameter that depends on the structure of the sequence; typically $d = o(\log n)$. Parameter d can be seen as a measure of “easiness”, and the objective is to prove that a candidate algorithm achieves the absolute bound of $O(md)$; note that this does not immediately imply that the algorithm matches the optimum.

The splay tree is known to perform very well on many “easy” sequences: It achieves the access time of $O(md)$ when d is a parameter that roughly captures the average distances between consecutive accesses [9,10], the average entropy, or the recentness of the previous accesses [34]; these properties are called *dynamic finger*, *static optimality* and *working set* bounds respectively (the latter two bounds are also satisfied by GREEDY [15]). The notorious difficulty of dynamic optimality is perhaps witnessed by the fact that many other highly restricted special cases have resisted attempts and remain unresolved. These special cases stand as a roadblock to proving (or disproving) the optimality of any algorithm: proving that \mathcal{A} is optimal requires proving that it performs well on the easy inputs; refuting the optimality of \mathcal{A} requires a special “easy” subclass on which \mathcal{A} is provably inefficient. In any case, a better understanding of easy sequences is needed, but currently lacking, as observed by several authors [11,30].

One famous example of such roadblocks is the *traversal conjecture*, proposed by Sleator and Tarjan in 1985. This conjecture states that starting with an arbitrary tree T (called *initial tree*) and accessing its keys using the splay algorithm in the order given by the preorder sequence of another tree T' takes linear time. Since the optimum of such preorder sequences is linear, any dynamically optimal algorithm must match this bound. After three decades, only very special cases of this conjecture are known to hold, namely when T' is a monotone path (the so-called *sequential access* [36]) or when $T' = T$ [8].⁴ Prior to this work, no online BST algorithm was known to be linear on such access sequences.

Motivated by well-established concepts in combinatorics [21,24,37], in this paper we initiate the study of the complexity of access sequences in terms of pattern avoidance parameters. Our main contributions are

²This observation appears to be folklore. We include a proof in Appendix G.

³As customary, we only consider successful accesses, we ignore inserts and deletes, and we assume that $m \geq n$.

⁴This result implies an offline BST with $O(n)$ cost for accessing a preorder sequence, by first rotating T to T' .

| | Structure | splay bound | GREEDY bound | Remark |
|------------------------|--|---|----------------------------------|---|
| Static optimality | low entropy | $O(\sum_{i=1}^n f_i(1 + \log \frac{m}{f_i}))$ | same as splay [6, 15] | corollaries of Access Lemma [34] |
| Working set | temporal locality | $O(m + n \log n + \sum_{t=1}^m \log(\tau_t + 1))$ | same as splay [6, 15] | |
| Dynamic finger | key-space locality | $O(m + \sum_{t=2}^m \log x_t - x_{t-1})$ | ? | |
| Deque | n updates at min/max elements | $O(n\alpha^*(n))$ [30] | $O(n2^{\alpha(n)})$ [7] | $O(n)$ for multi-splay [38] from empty tree |
| Sequential | avoid (2, 1) | $O(n)$ [31, 36] | $O(n)$ [7, 15] | |
| Traversal | avoid (2, 3, 1) | ? | $n2^{\alpha(n)O(1)}$ [* Cor 1.2] | OPT = $O(n)$ |
| Pattern-avoiding (new) | avoid (2, 3, 1) with preprocessing | ? | $O(n)$ [* Cor 1.5] | OPT = $O(n)$ |
| | avoid $(k, \dots, 1)$ (k -increasing) | ? | $n2^{O(k^2)}$ [* Thm 1.6] | |
| | avoid all simple perm. of size $> k$ (k -decomposable) with preprocessing | ? | $n2^{O(k^2)}$ [* Thm 1.4] | OPT = $O(n \log k)$ [* Cor 1.10] |
| | avoid an arbitrary perm. of size k | ? | $n2^{\alpha(n)O(k)}$ [* Thm 1.1] | Best known lower bound is at most $n2^{O(k)}$ [* Thm 1.7] |

Table 1: Easy sequences for BST and best known upper bounds. Results referenced as [*] are new. The symbol “?” means that no nontrivial bound is known. In the first rows, f_i is the number of times that i is accessed, x_t is the element accessed at time t , and τ_t is the number of distinct accesses since the last access of x_t . In the “Sequential”, “Traversal”, and “Pattern-avoiding” rows it is assumed that the access sequence is a permutation, i.e. $m = n$.

two-fold (the precise statement of the results is in § 1.1, and the results are summarized in Table ??).

- (i) We study access sequences parametrized by pattern avoidance and analyze the access time of GREEDY in terms of this parameter. As a by-product, we almost settle the traversal conjecture for GREEDY in two orthogonal directions: (a) GREEDY is “almost linear” for the traversal conjecture, and (b) if a linear-cost preprocessing is allowed, GREEDY is in fact linear. This is perhaps the first evidence in support of GREEDY as a candidate for dynamic optimality (all previously known properties of GREEDY were subsumed by splay). These results are derived via an *input-revealing property* of GREEDY, in the sense that the execution log of GREEDY partially reveals the structure of the input sequence.
- (ii) We study a decomposability property of sequences and show that a wide subclass of the pattern-avoiding class satisfies this property. This allows us to (a) identify the core difficulty of designing an offline algorithm for dynamic optimality, (b) obtain a tight bound for the optimum of this input class, and (c) derive a simple proof that Cole’s showcase sequence [10] is linear.

1.1 Our results

We first define the pattern avoidance parameters. Let $[n] = \{1, \dots, n\}$ be the set of keys. We represent a length- m sequence by an m -tuple $X = (x_1, \dots, x_m) \in [n]^m$. Throughout the paper, we assume that X is a permutation, i.e. $m = n$ and $x_i \neq x_j$ for all $i \neq j$. This is justified by Theorem 2.1, but we remark that many of our results can be extended to non-permutation access sequences.

Let S_k denote the set of all permutations in $[k]^k$. Given $X \in S_n$, we say that X *contains* a permutation pattern $\pi \in S_k$, if there are indices $1 \leq i_1 < i_2 < \dots < i_k \leq n$, such that $(x_{i_1}, \dots, x_{i_k})$ is order-isomorphic to $\pi = (\pi_1, \dots, \pi_k)$. Otherwise we say that X *avoids* π , or that X is π -free. Two sequences of the same length are *order-isomorphic* if they have the same pairwise comparisons, e.g. $(5, 8, 2)$ is order-isomorphic to $(2, 3, 1)$ (Figure 1). As examples, we mention that $(2, 1)$ -free and $(2, 3, 1)$ -free sequences are exactly the sequential, respectively, preorder access sequences.

For each permutation X , a *pattern avoidance characteristic* of X , denoted by Σ_X , is the set of all patterns not contained in X , i.e. $\Sigma_X = \{\pi \in S_\ell : \ell \geq 1 \text{ and } X \text{ is } \pi\text{-free}\}$. The *pattern avoidance parameter* of X , denoted by $k(X)$, is the minimum $k \in \mathbb{N}$ for which $S_k \cap \Sigma_X \neq \emptyset$.

For each $k \in \mathbb{N}$, the pattern avoidance class \mathcal{C}_k contains all sequences X with pattern avoidance parameter k . This characterization is powerful and universal: $\mathcal{C}_i \subseteq \mathcal{C}_{i+1}$ for all i , and \mathcal{C}_{n+1} contains all sequences.

Bounds for GREEDY in terms of k . We study the access cost of GREEDY in terms of the pattern avoidance parameter. Our most general upper bound of $n2^{\alpha(n)^{O(k)}}$ holds for any access sequence in \mathcal{C}_k (Theorem 1.1). Later, we show a stronger bound of $n2^{O(k^2)}$ for two broad subclasses of \mathcal{C}_k that have intuitive geometric interpretation: k -decomposable sequences (Theorem 1.4), which generalize preorder sequences and k -increasing sequences (Theorem 1.6), which generalize sequential access.

Theorem 1.1. *Let $X \in \mathcal{C}_k$ be an access sequence of length n . The cost of accessing X using GREEDY is at most $n2^{\alpha(n)^{O(k)}}$.*

The following corollary simply follows from the fact that all preorder sequences X belong to \mathcal{C}_3 (plugging $k = 3$ into the theorem).

Corollary 1.2 (“Almost” traversal conjecture for GREEDY). *The cost of accessing a preorder sequence of length n using GREEDY, with arbitrary initial tree is at most $n2^{\alpha(n)^{O(1)}}$.*

Next, we show sharper bounds for GREEDY when some structural property of Σ_X is assumed. First, we prove a bound when Σ_X contains all “non-decomposable” permutations of length at least $k + 1$. We show that this property leads to a decomposition that has intuitive geometric meaning, which we describe below.

Given a permutation $\sigma = (\sigma_1, \dots, \sigma_\ell) \in S_\ell$ we call a set $[a, b]$ a *block* of σ , if $\{\sigma_a, \sigma_{a+1}, \dots, \sigma_b\} = \{c, c + 1, \dots, d\}$ for some integers $c, d \in [\ell]$. In words, a block corresponds to a contiguous interval that is mapped to a contiguous interval. We say that a permutation $\sigma \in S_\ell$ is decomposable if there is a block of σ of size strictly between 1 and ℓ ; otherwise, we say that it is *non-decomposable* (a.k.a. *simple*⁵).

We say that an input sequence X is decomposable into k blocks if there exist disjoint $[a_1, b_1], \dots, [a_k, b_k]$ such that each $[a_i, b_i]$ is a block for X and $(\bigcup_i [a_i, b_i]) \cap \mathbb{N} = [n]$. The *recursive decomposition complexity* of a sequence X is at most d if one can recursively decompose X into subblocks until singletons are obtained, and each step creates at most d blocks (equivalently we say that X is d -recursively decomposable, or simply *d -decomposable*). See Figure 1. It is easy to see that preorder sequences are 2-decomposable.⁶

The following lemma (proof in Lemma D.1) connects pattern avoidance and recursive decomposition complexity of a sequence.

Lemma 1.3. *Let $X \in S_n$. Then Σ_X contains all simple permutations of length at least $k + 1$ if and only if X is k -decomposable.*

This lemma implies in particular that if X is k -decomposable, then $X \in \mathcal{C}_{k+1}$ and Theorem 1.1 can be applied. The following theorem gives a sharper bound, when a preprocessing (independent of X) is performed. This yields another relaxed version of the traversal conjecture.

Theorem 1.4. *Let X be a k -decomposable access sequence. The cost of accessing X using GREEDY, with a linear-cost preprocessing, is at most $n2^{O(k^2)}$.*

Corollary 1.5 (Traversal conjecture for GREEDY with preprocessing). *The cost of accessing a preorder sequence using GREEDY, with preprocessing, is linear.*

Our preprocessing step is done in the geometric view (explained in §2) in order to “hide” the initial tree T from GREEDY. In the corresponding tree-view, our algorithm preprocesses the initial tree T , turning it into another tree T' , independent of X , and then starts accessing the keys in X using the tree-view variant of GREEDY (as defined in [11]). We mention that DHIKP [11] define GREEDY without initial tree, and thus their algorithm (implicitly) performs this preprocessing.

⁵See [3] for a survey of this well-studied concept. The decomposition described here appears in the literature as *substitution-decomposition*.

⁶The entire set of 2-decomposable permutations is known as the set of *separable permutations* [2].

As another natural special case, we look at $\Sigma_X \supseteq \{(k, \dots, 1)\}$, for some value k . In words, X has no decreasing subsequence of length k (alternatively, X can be partitioned into at most $k - 1$ increasing subsequences⁷). Observe that the $k = 2$ case corresponds to sequential access. We call this the k -increasing property of X . Similarly, if $\Sigma_X \supseteq \{(1, \dots, k)\}$, we say that X is k -decreasing. We obtain the following generalization of the sequential access theorem (this result holds for arbitrary initial tree).

Theorem 1.6. *Let X be a k -increasing or k -decreasing sequence. The cost of accessing X using GREEDY is at most $n2^{O(k^2)}$.*

All the aforementioned results are obtained via the *input-revealing* property of GREEDY (the ideas are sketched in §1.2).

Bound for signed GREEDY in terms of k . We further show the applicability of the input-revealing technique. Signed GREEDY [11] (denoted SGREEDY) is an algorithm similar to GREEDY. It does not always produce a feasible BST solution, but its cost serves as a lower bound for the optimum. Let us denote by $\mathcal{A}(X)$ and $\text{OPT}(X)$ the cost of an algorithm \mathcal{A} on X , respectively the optimum cost of X . Then we have $\mathcal{A}(X) \geq \text{OPT}(X)$, and $\text{OPT}(X) = \Omega(\text{SGREEDY}(X))$.

Conjecture 1 (Wilber [39], using [11]). $\text{OPT}(X) = \Theta(\text{SGREEDY}(X))$ for all input sequences X .

We show that our techniques can be used to upper bound the cost of SGREEDY.

Theorem 1.7. *The cost of SGREEDY on an access sequence $X \in \mathcal{C}_k$ is at most $n2^{O(k)}$.*

Corollary 1.8. *Let \mathcal{A} be an algorithm. If Conjecture 1 is true, then one of the following must hold:*

- (i) *The cost of accessing X using \mathcal{A} is at most $n2^{O(k)}$, for all $X \in \mathcal{C}_k$, for all k . In particular, the traversal conjecture holds for \mathcal{A} .*
- (ii) *\mathcal{A} is not dynamically optimal.*

Decomposability. Next, we prove a general theorem on the complexity of sequences with respect to the recursive decomposition. Let $X \in S_n$ be an arbitrary access sequence. We represent the process of decomposing X into blocks until singletons are obtained by a tree \mathcal{T} , where each node v in \mathcal{T} is associated with a sequence X_v (see Figure 1).

Theorem 1.9 (Decomposition theorem, informal). $\text{OPT}(X) \leq \sum_{v \in \mathcal{T}} \text{GREEDY}(X_v) + O(n)$.

Corollary 1.10. *Let X be a k -decomposable sequence of length n (with k possibly depending on n). Then X has optimum cost $\text{OPT}(X) \leq O(n \log k)$.*

This result is tight for all values of k , since for all k , there is a k -decomposable sequence whose complexity is $\Omega(n \log k)$ (Proof in Appendix H.2). This result gives a broad range of input sequences whose optimum is well below the performance of all known BST algorithms. The class can thus be useful as “candidate counterexample” for dynamic optimality. (Since we know OPT, we only need to show that no online algorithm can match it.) We remark that GREEDY asymptotically matches this bound when k is constant (Theorem 1.4).

The following observation follows from the fact that the bound in Theorem 1.9 is achieved by our proposed offline algorithm, and the fact that $\text{OPT}(X) \geq \sum_{v \in \mathcal{T}} \text{OPT}(X_v)$ (see Lemma E.4).

Corollary 1.11. *If GREEDY is constant competitive on non-decomposable sequences, then there is an offline $O(1)$ -approximate algorithm on any sequence X .*

This corollary shows that, in some sense, non-decomposable sequences serve as the “core difficulty” of proving dynamic optimality, and understanding the behaviour of GREEDY on them might be sufficient.

⁷This equivalence is proven in Lemma ??.

Finally, Cole et al. [10] developed sophisticated technical tools for proving the dynamic finger theorem for splay trees. They introduced a “showcase” sequence to illustrate the power of their tools, and showed that splay has linear cost on such a sequence. Theorem 1.9 immediately gives an upper bound on the optimum access cost of this sequence.

Corollary 1.12 (informal). *Let X be Cole’s showcase sequence. Then $\text{OPT}(X) = O(n)$.*




Figure 1: From left to right: (i) plot of permutation $\sigma = (6, 1, 3, 2, 8, 7, 4, 5)$ containing $(2, 1, 3)$ (highlighted) and avoiding $(4, 3, 2, 1)$, (ii) permutation σ , and (iii) its block decomposition tree; permutation $(3, 1, 4, 2)$ at the root is obtained by contracting the four blocks into points, (iv) simple permutation $(6, 1, 8, 4, 2, 7, 3, 5)$.

1.2 Overview of techniques

We sketch the main technical ideas of our proofs. First, we define roughly the execution log matrix of GREEDY. The log of executing GREEDY on $X \in [n]^m$ is an n -by- m matrix G_X such that $G_X(a, t) = 1$ if and only if element a is touched by GREEDY at time t . The cost of this execution is exactly the number of ones in G_X , denoted by $w(G_X)$ (we make this claim precise in §2).

A submatrix of G_X is obtained by removing some rows and columns of G_X . We say that G_X contains a matrix B if there is a submatrix B' of G_X such that $B(i, j) \leq B'(i, j)$ for all i, j (in words, any ‘one’ entry in B must also be a one in B').

We say that an algorithm \mathcal{A} is *input-revealing* if there is a constant-size matrix B such that any submatrix of an execution matrix of \mathcal{A} containing B must contain an access point. When such a matrix B exists for an algorithm, we say that B is a *capture gadget*. The existence for GREEDY of a capture gadget B allows us to observe that G_X contains the pattern⁸ $Q \otimes B$ only if the input sequence contains the pattern defined by Q . So, if we execute GREEDY on an input sequence $X \in \mathcal{C}_k$ that is Q -free for $Q \in \mathcal{S}_k$, then G_X cannot contain $Q \otimes B$. Now, we can use results from forbidden submatrix theory for bounding $w(G_X)$. In short, forbidden submatrix theory is a collection of theorems of the following type:

Theorem 1.13. *Let M and P be n -by- n and k -by- k matrices respectively, such that $k \leq n$ (mostly k should be thought of as constant). If M does not contain P , then $w(M) \leq n f_P(n, k)$.*

The function f_P depends on the matrix P that is forbidden in M . For instance, if P contains at most one 1-entry per row and column, then $f_P(n, k)$ is only a function of k . If P contains at most one 1-entry per column but possibly multiple 1s per row, then $f_P(n, k) = 2^{\alpha(n)^{O(k)}}$. The tightness of the bounds we obtain depends on the structure of the capture gadget B . This is the reason we have a weaker bound for any $X \in \mathcal{C}_k$ and stronger bounds when X is more restricted.

We also develop a different set of tools to study the intrinsic complexity of decomposable sequences. Here, in order to compute bounds on OPT , we decompose the execution matrix M_X of an algorithm in a way that mirrors the recursive decomposition of the input sequence X . Unfortunately, the behaviour of GREEDY is too capricious to afford such a clean decomposition. We therefore modify GREEDY to reduce the interference between blocks in the input sequence. The modified algorithm might perform more work locally than GREEDY, but globally it is more “robust”. It is however, an offline algorithm; whether there exists an online algorithm competitive with our robust GREEDY, is an interesting open question.

⁸ \otimes denotes the tensor (Kronecker) product, to be defined later.

Related work. For an overview of results related to dynamic optimality, we refer to the survey of Iacono [19]. Besides Tango trees, other $O(\log \log n)$ -competitive BST algorithms have been proposed, similarly using Wilber’s first bound (multi-splay [38] and chain-splay [17]). Using [13], all the known bounds can be simultaneously achieved by combining several BST algorithms. Our study of the execution log of BST algorithms makes use of the geometric view of BST introduced by DHIKP [11].

Pattern avoidance in sequences has a large literature in combinatorics [21, 37], as well as in computer science [2, 4, 24, 32, 35]. The theory of forbidden submatrices can be traced back to a 1951 problem of Zarankiewicz, and its systematic study has been initiated by Füredi [16], and Bienstock and Gyóri [1]. Our use of this theory is inspired by work of Pettie [30, 31], who applied forbidden pattern techniques to prove bounds about data structures (in particular he proves that splay trees achieve $O(n\alpha^*(n))$ cost on deque-sequences, and he reproves the sequential access theorem for splay trees). There are marked differences in our use of the techniques compared with the work of Pettie. In particular, we look for patterns directly in the execution log of a BST algorithm, without any additional encoding step. Furthermore, instead of using fixed forbidden patterns, we make use of patterns that depend on the input.

Organization of the paper. In §2 we give a short description of the geometric view of the BST model (following DHIKP [11]), and we introduce concepts used in the subsequent sections. In §3 we study the input-revealing property of GREEDY, and prove our bounds for pattern-avoiding sequences. In §4 we introduce a robust GREEDY, state the decomposition theorem and its applications. Finally, §5 is reserved for open questions. We defer the proofs of several results, as well as some auxiliary material to the Appendix.

2 Preliminaries

The main object of study in this paper are $\{0, 1\}$ -matrices (corresponding to the execution matrix of an algorithm). For geometric reasons, we write matrix entries bottom-to-top, i.e. the first row is the bottom-most. Thus, $M(i, j)$ is the value in the i -th column (left-to-right) and the j -th row (bottom-to-top) of M . We simultaneously view M as a set of integral points in the plane. For a point $p = (x, y)$, we say that $p \in M$, if $M(x, y) = 1$. We denote by $p.x$ and $p.y$ the x - and y -coordinates of point p respectively.

Given a point set (matrix) M , let p and q be points not necessarily in M . Let \square_{pq} be the closed rectangle whose two corners are p and q . Rectangle \square_{pq} is *empty* if $M \cap \square_{pq} \setminus \{p, q\} = \emptyset$. Rectangle \square_{pq} is *satisfied* if p, q have the same x or y coordinate, or there is another point $r \in M \cap \square_{pq} \setminus \{p, q\}$. A point set M is (*arborally*) *satisfied* if \square_{pq} is satisfied for any two points $p, q \in M$.

Geometric setting for online BST. We describe the geometric model for analyzing online BST algorithms (or simply online BST). This model is essentially the same as [11], except for the fact that our model allows cost analysis of BST algorithms that start accessing an input sequence from arbitrary initial tree.

For any access sequence $X \in [n]^m$, we also refer to X as an n -by- m matrix, called *access matrix*, where there is a point $(x, t) \in X$ iff an element x is accessed at time t . We call (x, t) the access point (or input point) at time t . We also refer to the x -th column of X as *element x* , and to the t -th row of X as *time t* .

For any tree T of depth d , let $d(x)$ be the depth of element x in T . We also refer to T as an n -by- d matrix, called *initial tree matrix*, where in each column x , there is a stack of points $(x, 1), \dots, (x, d - d(x) + 1) \in T$. See Figure 2 for an illustration. Observe that matrix T is satisfied.

An *online geometric* BST \mathcal{A} accessing sequence X starting with initial tree T works as follows. Given a matrix $\begin{bmatrix} X \\ T \end{bmatrix}$ as input, \mathcal{A} outputs a satisfied matrix $\begin{bmatrix} \mathcal{A}_T(X) \\ T \end{bmatrix}$, where $\mathcal{A}_T(X) \supseteq X$ is an n -by- m matrix called the *touch matrix* of \mathcal{A} on X with initial tree T . More precisely, \mathcal{A} proceeds in an online fashion: At time $t = 1, \dots, n$, the algorithm \mathcal{A} has access to T as well as to the first t rows of matrix X , and it must decide on the t -th row of the output $\mathcal{A}_T(X)$ which cannot be changed afterwards. Algorithm \mathcal{A} is *feasible* if $\begin{bmatrix} \mathcal{A}_T(X) \\ T \end{bmatrix}$ is satisfied. The matrix $\mathcal{A}_T(X)$ can be thought of as the execution log of \mathcal{A} , and points in $\mathcal{A}_T(X)$ are called *touch points*; the *cost* of \mathcal{A} is simply $w(\mathcal{A}_T(X))$, i.e. the number of points (ones) in $\mathcal{A}_T(X)$.

An online algorithm with *preprocessing* is an online algorithm that is allowed to perform differently at time $t = 0$. The linear-cost preprocessing used in this paper, as well as in DHIKP [11], is to put all elements

into a *split tree* (see [11] for the details of this data structure). The consequences of this preprocessing are that (i) in the geometric view the initial tree matrix can be removed, i.e. an algorithm \mathcal{A} only gets X as input and outputs $\mathcal{A}(X)$, and (ii) in the tree-view, an input initial tree T is turned into another tree T' independent of the input X , and then keys in X are accessed using the tree-view variant of GREEDY.

DHIKP [11] showed⁹ that online *geometric* BSTs are essentially equivalent to online BSTs (in tree-view). In particular, given any online geometric BST \mathcal{A} , there is an online BST \mathcal{A}' where the cost of running \mathcal{A}' on any sequence X with any initial tree T is a constant factor away from $w(\mathcal{A}_T(X))$. Therefore, in this paper, we focus on online geometric BSTs and refer to them from now on simply as online BSTs.

Unless explicitly stated otherwise, we assume that the access sequence X is a permutation, i.e. each element is accessed only once, and thus X and $\mathcal{A}_T(X)$ are n -by- n matrices. This is justified by the following theorem (proof in Appendix F).

Theorem 2.1. *Assume that there is a c -competitive geometric BST \mathcal{A}^p for all permutations. Then there is a $4c$ -competitive geometric BST \mathcal{A} for all access sequences.¹⁰*

GREEDY and signed GREEDY. Next, we describe the algorithms GREEDY and signed GREEDY [11] (denoted as SGREEDY).

Consider the execution of an online BST \mathcal{A} on sequence X with initial tree T just before time t . Let O_{t-1} denote the output at time $t-1$ (i.e. the initial tree T and the first $t-1$ rows of $\mathcal{A}_T(X)$). Let $\tau(b, t-1) = \max\{i \mid (b, i) \in O_{t-1}\}$ for any $b \in [n]$. In words, $\tau(b, t-1)$ is the last time *strictly before* t when b is touched, otherwise it is a non-positive number defined by the initial tree. Let $(a, t) \in X$ be the access point at time t . We define $\text{stair}_t(a)$ as the set of elements $b \in [n]$ for which $\square_{(a,t),(b,\tau(b,t-1))}$ is not satisfied. GREEDY touches element $x \in [n]$ if and only if $x \in \text{stair}_t(a)$. See Figure 2 for an illustration of stairs.




Figure 2: Tree-view and geometric view of initial tree. Search path and stair (at time 1) of element 5.

We define $\text{stairLeft}_t(a) = \{b \in \text{stair}_t(a) \mid b < a\}$ and $\text{stairRight}_t(a) = \{b \in \text{stair}_t(a) \mid b > a\}$. The algorithms GREEDYLEFT and GREEDYRIGHT are analogous to GREEDY, but they only touch $\text{stairLeft}_t(a)$, respectively $\text{stairRight}_t(a)$ for every time t ; SGREEDY simply refers to the union of the GREEDYLEFT and GREEDYRIGHT outputs, which is not necessarily satisfied. It is known [11, 18] that SGREEDY constant-approximates the *independent rectangle bound* which subsumes the two previously known lower bounds on OPT, called Wilber’s first and second bounds [39].

Theorem 2.2 ([11, 18, 39]). *For any BST \mathcal{A} (online or offline) that runs on sequence X with arbitrary initial tree T , $w(\mathcal{A}_T(X)) = \Omega(\text{SGREEDY}(X))$.*

A nice property of GREEDY and SGREEDY is that an element x can become “hidden” in some range in the sense that if we do not access in x ’s hidden range, then x will not be touched. Next, we formally define this concept, and list some cases when elements become hidden during the execution of GREEDY and SGREEDY. The concept is illustrated in Figure 4.

Definition 2.3 (Hidden). *For any algorithm \mathcal{A} , an element x is hidden in range $[w, y]$ after t if, given that there is no access point $p \in [w, y] \times (t, t']$ for any $t' > t$, then x will not be touched by \mathcal{A} in time $(t, t']$.*

⁹with some minor additional argument about initial trees from Appendix C.1.

¹⁰The statement holds unconditionally for offline algorithms, and under a mild assumption if we require \mathcal{A}^p and \mathcal{A} to be online.




Figure 3: After time t , element x is hidden in $(w, n]$ for GREEDY, and in $(w, x]$ for GREEDYRIGHT. After time t' , element x is hidden in (w, y) for GREEDY; it is easy to verify that any access outside of $[w, y]$ will never touch x .

Lemma 2.4 (Proof in Appendix B.1). *Let x be some element.*

- (i) *If there is an element $w < x$ where $\tau(w, t) \geq \tau(x, t)$, then x is hidden in $(w, n]$ and $(w, x]$ after t for GREEDY and GREEDYRIGHT respectively.*
- (ii) *If there is an element $y > x$ where $\tau(y, t) \geq \tau(x, t)$, then x is hidden in $[1, y)$ and $[x, y)$ after t for GREEDY and GREEDYLEFT respectively.*
- (iii) *If there are elements w, y where $w < x < y$, and $\tau(w, t), \tau(y, t) \geq \tau(x, t)$, then x is hidden in (w, y) after t for GREEDY.*

Forbidden submatrix theory. Let M be an n -by- m matrix. Given $I \subseteq [n]$ and $J \subseteq [m]$, a submatrix $M|_{I,J}$ is obtained by removing columns in $[n] \setminus I$ and rows in $[m] \setminus J$ from M .

Definition 2.5 (Forbidden submatrix). *Given a matrix M and another matrix P of size $c \times d$, we say that M contains P if there is a $c \times d$ submatrix M' of M such that $M'(i, j) = 1$ whenever $P(i, j) = 1$. We say that M avoids (forbids) P if M does not contain P .*

We often call an avoided (forbidden) matrix a “pattern”. For any permutation $\pi = (\pi_1, \dots, \pi_k) \in S_k$, we also refer to π as a k -by- k matrix where $(\pi_i, i) \in \pi$ for all $1 \leq i \leq k$. Permutation matrices have a single one-entry in every row and in every column. Relaxing this condition slightly, we call a matrix that has a single one in every column a *light* matrix. We make use of the following results.

Lemma 2.6. *Let M and P be n -by- n and k -by- k matrices such that M avoids P .*

- (i) (Marcus, Tardos [27], Fox [14]) *If P is a permutation, the number of ones in M is at most $n2^{O(k)}$.*
- (ii) (Nivasch [?], Pettie [?]) *If P is light, the number of ones in M is at most $n2^{\alpha(n)^{O(k)}}$.*

We define the *tensor product* between matrices as follows: $M \otimes G$ is the matrix obtained by replacing each one-entry of M by a copy of G , and replacing each zero-entry of M by an all-zero matrix equal in size with G . Suppose that M contains P . A *bounding box* B of P is a minimal submatrix $M|_{I,J}$ such that P is contained in B , and I and J are sets of consecutive columns, resp. rows. We denote $\min(I)$, $\max(I)$, $\min(J)$, $\max(J)$, as $B.xmin$, $B.xmax$, $B.ymin$, $B.ymax$, respectively.

3 GREEDY is input-revealing

In this section, we prove Theorems 1.1, 1.4, 1.6 and 1.7. We refer to § 1.1 for definitions of pattern-avoiding, k -decomposable and k -increasing permutations.

To bound the cost of algorithm \mathcal{A} , we show that $\mathcal{A}_T(X)$ (or $\mathcal{A}(X)$ if preprocessing is allowed) avoids some pattern. In the following, Inc_k , Dec_k , and Alt_k are permutation matrices of size k , and Cap is a light matrix. These matrices will be defined later.

Lemma 3.1. *Let X be an access sequence.*

- (i) If X avoids a permutation P of size k , then for any initial tree T , $\text{GREEDYRIGHT}_T(X)$, $\text{GREEDYLEFT}_T(X)$ and $\text{GREEDY}_T(X)$ avoid $P \otimes (1, 2)$, $P \otimes (2, 1)$ and $P \otimes \mathbf{Cap}$ respectively.
- (ii) If X is k -increasing (k -decreasing), then for any initial tree T , $\text{GREEDY}_T(X)$ avoids $(k, \dots, 1) \otimes \mathbf{Dec}_{k+1}$ ($(1, \dots, k) \otimes \mathbf{Inc}_{k+1}$).
- (iii) If X is k -decomposable, then $\text{GREEDY}(X)$ avoids $P \otimes \mathbf{Alt}_{k+4}$ for all simple permutations P of size at least $k + 1$.

The application of Lemma 2.6 is sufficient to complete all the proofs, together with the observation that, for any permutation matrix P of size $k \times k$, if G is a permutation (resp. light) matrix of size $\ell \times \ell$, then $P \otimes G$ is a permutation (resp. light) matrix of size $k\ell \times k\ell$.

The rest of this section is devoted to the proof of Lemma 3.1. We need the definition of the *input-revealing gadget* which is a pattern whose existence in the touch matrix (where access points are indistinguishable from touch points) gives some information about the locations of access points.

Definition 3.2 (Input-revealing gadgets). *Let X be an access matrix, \mathcal{A} be an algorithm and $\mathcal{A}(X)$ be the touch matrix of \mathcal{A} running on X .*

- **Capture Gadget.** *A pattern P is a capture gadget for algorithm \mathcal{A} if, for any bounding box B of P in $\mathcal{A}(X)$, there is an access point $p \in B$.*
- **Monotone Gadget.** *A pattern P is a k -increasing (k -decreasing) gadget for algorithm \mathcal{A} if, for any bounding box B of P in $\mathcal{A}(X)$, either (i) there is an access point $p \in B$ or, (ii) there are k access points p_1, \dots, p_k such that $B.ymin \leq p_1.y < \dots < p_k.y \leq B.ymax$ and $p_1.x < \dots < p_k.x < B.xmin$ ($p_1.x > \dots > p_k.x > B.xmax$). In words, p_1, \dots, p_k form the pattern $(1, \dots, k)$ (resp. $(k, \dots, 1)$).*
- **Alternating Gadget.** *A pattern P is a k -alternating gadget for algorithm \mathcal{A} if, for any bounding box B of P in $\mathcal{A}(X)$, either (i) there is an access point $p \in B$ or, (ii) there are k access points p_1, \dots, p_k such that $B.ymin \leq p_1.y < \dots < p_k.y \leq B.ymax$ and $B.xmax < p_i.x$ for all odd i and $p_i.x < B.xmin$ for all even i .*

Notice that if we had a permutation capture gadget G for algorithm \mathcal{A} , we would be done: If X avoids pattern Q , then $\mathcal{A}_T(X)$ avoids $Q \otimes G$ and forbidden submatrix theory applies. With arbitrary initial tree T , it is impossible to construct a permutation capture gadget for GREEDY even for 2-decomposable sequences (see Appendix H.1). Instead, we define a *light* capture gadget for GREEDY for any sequence X and any initial tree T . Additionally, we construct permutation monotone and alternating gadgets. After a preprocessing step, we can use the alternating gadget to construct a permutation capture gadget for k -decomposable sequences.

- Lemma 3.3.** (i) (12) is a capture gadget for GREEDYRIGHT and (21) is a capture gadget for GREEDYLEFT .
(ii) \mathbf{Cap} is a capture gadget for GREEDY where $\mathbf{Cap} = (\cdot \cdot \cdot)$.
(iii) \mathbf{Inc}_{k+1} (\mathbf{Dec}_{k+1}) is a k -increasing (k -decreasing) gadget for GREEDY where $\mathbf{Inc}_k = (k, \dots, 1)$ and $\mathbf{Dec}_k = (1, 2, \dots, k)$.
(iv) \mathbf{Dec}_{k+1} (\mathbf{Inc}_{k+1}) is a capture gadget for GREEDY that runs on k -increasing (k -decreasing) sequence.
(v) \mathbf{Alt}_{k+1} is a k -alternating gadget for GREEDY where $\mathbf{Alt}_k = (\lfloor (k+1)/2 \rfloor, k, 1, k-1, 2, \dots)$.
(vi) \mathbf{Alt}_{k+4} is a capture gadget for GREEDY that runs on k -decomposable sequence without initial tree.

For example, $\mathbf{Alt}_5 = \begin{pmatrix} \cdot & \cdot & \cdot \\ \cdot & & \cdot \\ & \cdot & \cdot \end{pmatrix}$. We emphasize that only result (vi) requires a preprocessing.

Given Lemma 3.3, and the fact that k -increasing, k -decreasing, and k -decomposable sequences avoid $(k, \dots, 1)$, $(1, \dots, k)$, resp. all simple permutations of size at least $k + 1$, Lemma 3.1 follows immediately.

Proof. We only show the proofs for (i) and (ii) in order to convey the flavor of the arguments. The other proofs are more involved. See Appendix C.2 for proofs of (iii),(iv), and Appendix D.2 for proofs of (v),(vi).

(i) We only prove for GREEDYRIGHT . Let a, b be touch points in $\text{GREEDYRIGHT}_T(X)$ such that $a.y < b.y$ and form $(1, 2)$ with a bounding box B . As $a.y = \tau(a.x, a.y) \geq \tau(b.x, a.y)$, by Lemma 2.4 (i), $b.x$ is hidden in $(a.x, b.x]$ after $a.y$. Suppose that there is no access point in $B = \square_{ab}$, then b cannot be touched, which is a contradiction.

(ii) Let a, b, c be touch points in $\text{GREEDY}_T(X)$ such that $a.x < b.x < c.x$ and a, b, c form **Cap** with bounding box B . As $\tau(a.x, a.y), \tau(c.x, a.y) \geq \tau(b.x, a.y)$, by Lemma 2.4 (iii), $b.x$ is hidden in $(a.x, c.x)$ after $a.y$. Suppose that there is no access point in B , then b cannot be touched, which is a contradiction. ■

4 Decomposition theorem and applications

In this section, we take a deeper look at decomposable permutations and investigate whether **GREEDY** is “decomposable”, in the sense that one can bound the cost of **GREEDY** on input sequences by summing over local costs on blocks. Proving such a property for **GREEDY** seems to be prohibitively difficult (we illustrate this in Appendix H.1). Therefore, we propose a more “robust” but offline variant of **GREEDY** which turns out to be decomposable. We briefly describe the new algorithm, and state the main theorems and three applications, leaving the details to Appendix E.

Recursive decomposition. We refer the reader to §1.1 for the definitions of blocks and recursive decomposition. Let S be a k -decomposable permutation access sequence, and let \mathcal{T} be its (not necessarily unique) decomposition tree (of arity at most k). Let P be a block in the decomposition tree \mathcal{T} of S . Consider the subblocks P_1, P_2, \dots, P_k of P , and let \tilde{P} denote the unique permutation of size $[k]$ that is order-isomorphic to a sequence of arbitrary representative elements from the blocks P_1, P_2, \dots, P_k . We denote $P = \tilde{P}[P_1, P_2, \dots, P_k]$ a *deflation* of P , and we call \tilde{P} the *skeleton* of P . Visually, \tilde{P} is obtained from P by contracting each block P_i to a single point. We associate with each block P (i.e. each node in \mathcal{T}) both the corresponding rectangular region (in fact, a square), and the corresponding skeleton \tilde{P} (Figure 1). We denote the set of leaves and non-leaves of the tree \mathcal{T} by $L(\mathcal{T})$ and $N(\mathcal{T})$ respectively.

Consider the k subblocks P_1, \dots, P_k of a block P in this decomposition, numbered by increasing y -coordinate. We partition the rectangular region of P into k^2 rectangular regions, whose boundaries are aligned with those of P_1, \dots, P_k . We index these regions as $R(i, j)$ with i going left to right and j going bottom to top. In particular, the region $R(i, j)$ is aligned (same y -coordinates) with P_j .

A robust version of **GREEDY.** The main difficulty in analyzing the cost of **GREEDY** on decomposable access sequences is the “interference” between different blocks of the decomposition. Our **RGREEDY** algorithm is an extension of **GREEDY**. Locally it touches at least those points that **GREEDY** would touch; additionally it may touch other points, depending on the recursive decomposition of the access sequence (this a priori knowledge of the decomposition tree is what makes **RGREEDY** an offline algorithm). The purpose of touching extra points is to “seal off” the blocks, limiting their effect to the outside.

We now define the concept of *topwing*. The *topwing* of a rectangular region R at a certain time in the execution of **RGREEDY** contains a subset of the points in R (at most one in each column). A point is in the topwing if it forms an empty rectangle with the top left or the top right corner of R (this definition includes the topmost touch point of the leftmost and rightmost columns of R , if such points exist in R). When accessing an element at time t , **RGREEDY** touches the projections to the timeline t of the topwings of all regions $R(i, j)$ aligned with a block P_j ending at time t (this includes the region of P_j itself). Observe that multiple nested blocks might end at the same time, in this case we process them in increasing order of their size. See Figure 3 for an illustration of **RGREEDY**.

Theorem 4.1 (Decomposition theorem, Theorem E.2). *For any decomposition tree \mathcal{T} of a permutation P , the total cost of **RGREEDY** is bounded as*

$$\text{RGREEDY}(P) \leq 4 \cdot \sum_{\tilde{P} \in N(\mathcal{T})} \text{GREEDY}(\tilde{P}) + \sum_{P \in L(\mathcal{T})} \text{GREEDY}(P) + 3n.$$

Application 1: Cole et al.’s “showcase” sequences. This sequence can be defined as a permutation $P = \tilde{P}[P_1, \dots, P_{n/\log n}]$ where each P_j is the permutation $(1, \dots, \log n)$, i.e. a sequential access. It is known that **GREEDY** is linear on sequential access, so $\text{GREEDY}(P_j) = O(\log n)$. Now, by applying Theorem 4.1, we




Figure 4: Illustration of RGREEDY. Left: region with points and *topwing* highlighted. Right: a sample run of RGREEDY on a 2-decomposable input. Access points are dark circles, points touched by GREEDY are gray circles, points touched by the RGREEDY augmentation are gray squares. The access point in the topmost line completes the black square and also the enclosing gray square. The topwing of the black square consists of the black circle at position (1,6) and the gray circle at (3,5). Therefore the augmentation of the black square adds the gray square at (3,6). The topwing of $R(2,2)$ consists of the gray circle at (4,4) and hence the gray square at (4,6) is added. The topwing of the gray square consists of points (1,6) and (6,3); the point (4,4) does not belong to the topwing because (4,6) has already been added. Therefore, the gray square at (6,6) is added.

can say that the cost of RGREEDY on P is $\text{RGREEDY}(P) \leq 4 \cdot \text{GREEDY}(\tilde{P}) + \sum_j \text{GREEDY}(P_j) + O(n) \leq O(\frac{n}{\log n} \cdot \log(\frac{n}{\log n})) + \frac{n}{\log n} O(\log n) + O(n) = O(n)$. This implies the linear optima of this class of sequences.

Application 2: Simple permutations as the “core” difficulty of BST inputs. We prove that, if GREEDY is c -competitive on simple permutations, then RGREEDY is $O(c)$ -approximate for all permutations. To prove this, we consider a block decomposition tree \mathcal{T} in which the permutation corresponding to every node is simple. Using Theorem 4.1 and the hypothesis that GREEDY is c -competitive for simple permutations:

$$\text{RGREEDY}(P) \leq 4c \cdot \sum_{\tilde{P} \in N(\mathcal{T})} \text{OPT}(\tilde{P}) + c \cdot \sum_{P \in L(\mathcal{T})} \text{OPT}(P) + 3n.$$

Now we decompose OPT into subproblems the same way as done for RGREEDY. We prove in Lemma E.4 that $\text{OPT}(P) \geq \sum_{P \in \mathcal{T}} \text{OPT}(P) - 2n$. Combined with Theorem 4.1 this yields $\text{RGREEDY}(P) \leq 4c \cdot \text{OPT}(P) + (8c + 3)n$. Using Theorem 2.1 we can extend this statement to arbitrary access sequences.

Application 3: Recursively decomposable permutations. Let X be a k -decomposable permutation. Then there is a decomposition tree \mathcal{T} of X in which each node is a permutation of size at most k . Each block $P \in L(\mathcal{T}) \cup N(\mathcal{T})$ has $\text{GREEDY}(P) = O(|P| \log |P|) = O(|P| \log k)$. By standard charging arguments, the sum $\sum_{P \in N(\mathcal{T}) \cup L(\mathcal{T})} |P| = O(n)$, so the total cost of RGREEDY is $O(n \log k)$; see the full proof in Appendix E. We remark that the logarithmic dependence on k is optimal (Proof in Appendix H.2), and that GREEDY asymptotically matches this bound when k is constant (Theorem 1.4).

5 Discussion and open questions

Besides the long-standing open question of resolving dynamic optimality, our work raises several new ones in the context of pattern avoidance. We list some that are most interesting to us.

Can the bound of Theorem 1.1 be improved? While our input-revealing techniques are unlikely to yield a linear bound (Appendix H.1), a slight improvement could come from strengthening Lemma 2.6(ii) for the special kind of light matrices that arise in our proof.

An intriguing question is whether there is a natural characterization of all sequences with linear optimum. How big is the overlap between “easy inputs” and inputs with pattern avoidance properties? We have shown that avoidance of a small pattern makes sequences easy. The converse does not hold: There is a permutation $X \in S_n$ with $k(X) = \sqrt{n}$ but for which $\text{GREEDY}(X) = O(n)$; see Appendix H.3. Note that our pattern-avoiding properties are incomparable with the earlier parametrizations (e.g. dynamic finger); see Appendix H.3. Is there a parameter that subsumes both pattern-avoidance and dynamic finger?

A question directly related to our work is to close the gap between $\text{OPT} = O(n \log k)$ and $n2^{O(k^2)}$ by GREEDY on k -decomposable sequences (when $k = \omega(1)$). Matching the optimum (if at all possible) likely requires novel techniques: splay is not even known to be linear on preorder sequences with preprocessing, and with forbidden-submatrix-arguments it seems impossible to obtain bounds beyond¹¹ $O(nk)$.

We proved that if GREEDY is optimal on simple permutations, then RGREEDY is optimal on all access sequences. Can some property of simple permutations¹² be algorithmically exploited? Can RGREEDY be made online? Can our application for Cole’s showcase sequence be extended in order to prove the dynamic finger property for GREEDY ?

Finally, making further progress on the traversal conjecture will likely require novel ideas. We propose a simpler question that captures the barrier for three famous conjectures: *traversal*, *deque*, and *split* (Appendix H.4). Given any initial tree T , access a preorder sequence defined by a path P (a BST where each non-leaf node has a single child). Prove a linear bound for your favorite online BST!

¹¹An n -by- n matrix avoiding *all* permutations of size at least k can contain $\Omega(nk)$ ones.

¹²Note that simple permutations can have linear cost. In Appendix H.3 we show such an example.

References

- [1] Daniel Bienstock and Ervin Györi. An extremal problem on sparse 0-1 matrices. *SIAM J. Discrete Math.*, 4(1):17–27, 1991. 5
- [2] Prosenjit Bose, Jonathan F Buss, and Anna Lubiw. Pattern matching for permutations. *Information Processing Letters*, 65(5):277–283, 1998. 3, 5
- [3] Robert Brignall. A survey of simple permutations. *CoRR*, abs/0801.0963, 2008. 3, 20, 25
- [4] Marie-Louise Bruner and Martin Lackner. The computational landscape of permutation patterns. *CoRR*, abs/1301.0340, 2013. 5
- [5] Mihai Bădoiu, Richard Cole, Erik D. Demaine, and John Iacono. A unified access bound on comparison-based dynamic dictionaries. *Theoretical Computer Science*, 382(2):86–96, August 2007. Special issue of selected papers from the 6th Latin American Symposium on Theoretical Informatics, 2004. 32
- [6] P. Chalermsook, M. Goswami, L. Kozma, K. Mehlhorn, and T. Saranurak. A global geometric view of splaying. *To Appear*, CoRR abs/1503.03105, 2015. 2
- [7] P. Chalermsook, M. Goswami, L. Kozma, K. Mehlhorn, and T. Saranurak. Greedy is an almost optimal deque. *To Appear*, 2015. 2
- [8] R. Chaudhuri and H. Höft. Splaying a search tree in preorder takes linear time. *SIGACT News*, 24(2):88–93, April 1993. 1
- [9] R. Cole. On the dynamic finger conjecture for splay trees. part ii: The proof. *SIAM Journal on Computing*, 30(1):44–85, 2000. 1, 2
- [10] Richard Cole, Bud Mishra, Jeanette Schmidt, and Alan Siegel. On the dynamic finger conjecture for splay trees. part i: Splay sorting log n-block sequences. *SIAM J. Comput.*, 30(1):1–43, April 2000. 1, 2, 4, 25
- [11] Erik D. Demaine, Dion Harmon, John Iacono, Daniel M. Kane, and Mihai Patrascu. The geometry of binary search trees. In *SODA 2009*, pages 496–505, 2009. 1, 3, 4, 5, 6, 7, 18
- [12] Erik D. Demaine, Dion Harmon, John Iacono, and Mihai Patrascu. Dynamic optimality - almost. *SIAM J. Comput.*, 37(1):240–251, 2007. 1
- [13] Erik D. Demaine, John Iacono, Stefan Langerman, and Özgür Özkan. Combining binary search trees. In *Automata, Languages, and Programming - 40th International Colloquium, ICALP 2013, Riga, Latvia, July 8-12, 2013, Proceedings, Part I*, pages 388–399, 2013. 5
- [14] Jacob Fox. Stanley-Wilf limits are typically exponential. *CoRR*, pages –1–1, 2013. 7
- [15] Kyle Fox. Upper bounds for maximally greedy binary search trees. In *WADS 2011*, pages 411–422, 2011. 1, 2
- [16] Zoltán Füredi. The maximum number of unit distances in a convex n -gon. *J. Comb. Theory, Ser. A*, 55(2):316–320, 1990. 5
- [17] George F. Georgakopoulos. Chain-splay trees, or, how to achieve and prove loglogn-competitiveness by splaying. *Inf. Process. Lett.*, 106(1):37–43, 2008. 5
- [18] Dion Harmon. *New Bounds on Optimal Binary Search Trees*. PhD thesis, Massachusetts Institute of Technology, 2006. 7

- [19] John Iacono. In pursuit of the dynamic optimality conjecture. In Andrej Brodnik, Alejandro Lopez-Ortiz, Venkatesh Raman, and Alfredo Viola, editors, *Space-Efficient Data Structures, Streams, and Algorithms*, volume 8066 of *Lecture Notes in Computer Science*, pages 236–250. Springer Berlin Heidelberg, 2013. 5
- [20] Balázs Keszegh. On linear forbidden submatrices. *Journal of Combinatorial Theory, Series A*, 116(1):232 – 241, 2009. 7
- [21] S. Kitaev. *Patterns in Permutations and Words*. Monographs in Theoretical Computer Science. An EATCS Series. Springer, 2011. 1, 5
- [22] Martin Klazar. The Füredi-Hajnal conjecture implies the Stanley-Wilf conjecture. In Daniel Krob, Alexander A. Mikhalev, and Alexander V. Mikhalev, editors, *Formal Power Series and Algebraic Combinatorics*, pages 250–255. Springer Berlin Heidelberg, 2000. 7
- [23] D.E. Knuth. Optimum binary search trees. *Acta Informatica*, 1(1):14–25, 1971. 1
- [24] Donald E. Knuth. *The Art of Computer Programming, Volume I: Fundamental Algorithms*. Addison-Wesley, 1968. 1, 5
- [25] Joan M. Lucas. Canonical forms for competitive binary search tree algorithms. *Tech. Rep. DCS-TR-250, Rutgers University*, 1988. 1
- [26] Joan M. Lucas. On the competitiveness of splay trees: Relations to the union-find problem. *Online Algorithms, DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, 7:95–124, 1991. In Lyle A. McGeoch and Daniel D. Sleator, editors. 34
- [27] Adam Marcus and Gábor Tardos. Excluded permutation matrices and the Stanley-Wilf conjecture. *Journal of Combinatorial Theory, Series A*, 107(1):153 – 160, 2004. 7
- [28] Kurt Mehlhorn. Nearly optimal binary search trees. *Acta Informatica*, 5(4):287–295, 1975. 1
- [29] J.Ian Munro. On the competitiveness of linear search. In Mike S. Paterson, editor, *Algorithms - ESA 2000*, volume 1879 of *Lecture Notes in Computer Science*, pages 338–345. Springer Berlin Heidelberg, 2000. 1
- [30] Seth Pettie. Splay trees, Davenport-Schinzel sequences, and the deque conjecture. In *Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '08*, pages 1115–1124, Philadelphia, PA, USA, 2008. Society for Industrial and Applied Mathematics. 1, 2, 5
- [31] Seth Pettie. Applications of forbidden 0-1 matrices to search tree and path compression-based data structures. In Moses Charikar, editor, *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2010, Austin, Texas, USA, January 17-19, 2010*, pages 1457–1467. SIAM, 2010. 2, 5
- [32] Vaughan R. Pratt. Computing permutations with double-ended queues, parallel stacks and parallel queues. In *Proceedings of the Fifth Annual ACM Symposium on Theory of Computing, STOC '73*, pages 268–277, New York, NY, USA, 1973. ACM. 5
- [33] James H. Schmerl and William T. Trotter. Critically indecomposable partially ordered sets, graphs, tournaments and other binary relational structures. *Discrete Mathematics*, 113(13):191 – 205, 1993. 20
- [34] Daniel Dominic Sleator and Robert Endre Tarjan. Self-adjusting binary search trees. *J. ACM*, 32(3):652–686, July 1985. 1, 2, 34
- [35] Robert Endre Tarjan. Sorting using networks of queues and stacks. *J. ACM*, 19(2):341–346, April 1972. 5

- [36] Robert Endre Tarjan. Sequential access in splay trees takes linear time. *Combinatorica*, 5(4):367–378, 1985. 1, 2, 34
- [37] Vincent Vatter. Permutation classes. *CoRR*, abs/1409.5159, 2014. 1, 5
- [38] Chengwen C. Wang, Jonathan C. Derryberry, and Daniel D. Sleator. $O(\log \log n)$ -competitive dynamic binary search trees. In *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithm*, SODA '06, pages 374–383, Philadelphia, PA, USA, 2006. Society for Industrial and Applied Mathematics. 2, 5
- [39] R. Wilber. Lower bounds for accessing binary search trees with rotations. *SIAM Journal on Computing*, 18(1):56–67, 1989. 4, 7

A Warmup: A direct proof for preorder sequences

In this section we show that the cost of GREEDY on an arbitrary preorder access sequence is linear. In other words, we give a direct proof of Corollary 1.5. We assume the preprocessing discussed in §2, or equivalently, in geometric view we assume that GREEDY runs with no initial tree¹³, i.e. initially all stairs are empty. The proof is intended as a warmup, to develop some intuition about the geometric model and to familiarize with the concepts of *wings* and *hidden elements* used in the subsequent proofs. We remark that an alternative proof of the same result (up to the same constants) can be obtained using forbidden submatrix arguments.

Preorder sequences. The class PreOrd_n of preorder sequences of length n can be defined recursively as follows. A permutation sequence $X = (x_1, \dots, x_n)$ belongs to PreOrd_n if (i) $|X| = n = 1$ or (ii) there is a number $n' < n$, such that

- $L(X) = (x_2, \dots, x_{n'}) \in \text{PreOrd}_{n'-1}$,
- $R(X) = (x_{n'+1}, \dots, x_n) \in \text{PreOrd}_{n-n'}$,
- $\max\{L(X)\} < x_1 < \min\{R(X)\}$.

We remark again the well-known fact that preorder sequences are exactly the $(2, 3, 1)$ -free permutations, and the easily verifiable fact that preorder sequences are 2-decomposable (see Figure 5 for an illustration).




Figure 5: Preorder sequence with sets L and R , with $wing_L(a)$ highlighted. Observe that point a together with L can form a single block.

We prove the following theorem.

Theorem A.1. *Let $X \in \text{PreOrd}_n$. Then $\text{GREEDY}(X) \leq 4n$.*

A.1 Wings and hidden elements

Let $X \in \text{PreOrd}_n$. For an arbitrary element $a \in X$, let t_a denote its access time. Let L_a denote the maximal contiguous subsequence of X immediately after a , that consists of elements smaller than a . Let $r_a \in X$ be the smallest element that is larger than a and precedes a in X . By convention, r_a is ∞ if there are no elements larger than a that precede a in X . Let R_a denote the maximal contiguous subsequence after L_a of elements larger than a and smaller than r_a .

Let \mathcal{T}_a be the concatenation of L_a and R_a . Let $\text{lca}(b, c)$ denote the last element a , such that $b, c \in \mathcal{T}_a$.

We define $wing_L(a)$ to contain all elements $b \in L_a$ such that the rectangle with corners (b, t_b) , (a, t_a) contains no other access point. Symmetrically, we define $wing_R(a)$ to contain all elements $b \in R_a$ such that the rectangle with corners (b, t_b) , (a, t_a) contains no other access point (Figure 5).

Lemma A.2. *The sets $\{wing_L(a)\}_{a \in [n]}$ are disjoint. Similarly, the sets $\{wing_R(a)\}_{a \in [n]}$ are disjoint.*

¹³curiously in the case of GREEDY and preorder sequences this has the same effect as if the initial tree would be exactly the tree that generated the input sequence.

Proof. Assume for contradiction that there are two elements $a, a' \in [n]$ such that $wing_L(a) \cap wing_L(a') \neq \emptyset$. Assume w.l.o.g. that $a < a'$. By definition, $wing_L(a)$ only contains elements in L_a , so there must be an element $b \in L_a \cap L_{a'}$, such that $b < a < a'$ and $t_b > t_a$ and $t_b > t'_a$. There are two possibilities: either $t_a > t'_a$, contradicting the fact that the rectangle with corners (a', t'_a) , (b, t_b) is empty, or $t_a < t'_a$, in which case the subsequence (a, a', b) is order-isomorphic to $(2, 3, 1)$, contradicting the fact that X is a preorder sequence. ■

The reader might find it instructive to see the corresponding situation in tree-view. Figure 6 shows a proof by picture in the tree corresponding to the preorder sequence X .




Figure 6: BST corresponding to a preorder sequence. Left wings of elements indicated as shaded rectangles.

Corollary A.3. $\sum_a (|wing_L(a)| + |wing_R(a)|) \leq 2n$.

Next we restate and prove Lemma 2.4 about hidden elements. Recall that $\tau(x, t - 1)$ denotes the last time strictly before t when x is touched.

Definition A.4 (Definition 2.3). *For any algorithm \mathcal{A} , an element x is hidden in range $[w, y]$ after t if, given that there is no access point $a \in [w, y] \times (t, t']$ for any $t' > t$, then x will not be touched by \mathcal{A} in time $(t, t']$.*

Lemma A.5 (Lemma 2.4). *Let x be some element.*

- (i) *If there is an element $w < x$ where $\tau(w, t) \geq \tau(x, t)$, then x is hidden in $(w, n]$ and $(w, x]$ after t for GREEDY and GREEDYRIGHT respectively.*
- (ii) *If there is an element $y > x$ where $\tau(y, t) \geq \tau(x, t)$, then x is hidden in $[1, y)$ and $[x, y)$ after t for GREEDY and GREEDYLEFT respectively.*
- (iii) *If there are elements w, y where $w < x < y$, and $\tau(w, t), \tau(y, t) \geq \tau(x, t)$, then x is hidden in (w, y) after t for GREEDY.*

Proof. (i) Consider any $t' > t$. For the case of GREEDY, assume that there is no access point in $(w, n] \times (t, t']$. Suppose x is touched in the time interval $(t, t']$, and let (p, t_p) be the first access point in this time interval that causes the touching of x . Then $p \in [1, w]$, and $t_p \in (t, t']$. As x is not touched in the time interval (t, t_p) by the choice of p , we have that $\tau(x, t_p - 1) = \tau(x, t)$. If $\tau(w, t) \geq \tau(x, t)$, then the rectangle with corners (p, t_p) , and $(x, \tau(x, t_p - 1))$ contains the point $(w, \tau(w, t))$, and thus it is satisfied before accessing p . Therefore, the accessing of p via GREEDY does not touch x , a contradiction. It follows that x is hidden in $(w, n]$ for GREEDY after t .

For the case of GREEDYRIGHT, assume that there is no access point in $(w, x] \times (t, t']$. Suppose x is touched in the time interval $(t, t']$, and let (p, t_p) be the first access point in this time interval that causes the touching of x . Then $p \in [1, w]$, and $t_p \in (t, t']$. (The case $p \in (x, n]$ is not possible for GREEDYRIGHT, since $p < x$ must hold.) The remainder of the argument is the same as for GREEDY.

- (ii) The argument is symmetric to (i).

- (iii) Consider any $t' > t$. Assume that there is no access point in $(w, y) \times (t, t']$. Suppose x is touched in the time interval $(t, t']$, and let (p, t_p) be the first access point in this time interval that causes the touching of x . There are two cases. If $p \in [1, w]$, we use the argument of (i). If $p \in [y, n]$, we use the argument of (ii). ■

A.2 Bounding the cost of GREEDY

Lemma A.6. *For any element $c \in L_a \setminus \text{wing}_L(a)$, a is not touched when accessing c using GREEDY. Similarly for $c \in R_a \setminus \text{wing}_R(a)$.*

Proof. There must be an access point b , such that $c < b < a$ and $t_a < t_b < t_c$, for otherwise c would be in $\text{wing}_L(a)$. Since $\tau(b, t_b) \geq \tau(a, t_b)$, it follows that a is hidden in $(b, n]$ after t_b . All accesses in the interval $(t_b, t_c]$ are in $[c, b)$. Hence, a is not touched. ■

Corollary A.7. *Let a be any element. GREEDY touches a when accessing elements in \mathcal{T}_a at most $|\text{wing}_L(a)| + |\text{wing}_R(a)|$ times.*

Lemma A.8. *Let a be any element. GREEDY touches a when accessing elements outside \mathcal{T}_a at most once.*

Proof. No element is touched before it is accessed. Since the elements in \mathcal{T}_a are accessed consecutively, we only need to consider accesses after all elements in \mathcal{T}_a have been accessed. All such accesses are to elements to the right of \mathcal{T}_a . Let c be the first element accessed after the last access to an element in \mathcal{T}_a and let $b = \text{lca}(a, c)$. Then $a \in L_b$, and c is the first element in R_b .

Since \mathcal{T}_c contains all elements in $(b, c]$, GREEDY does not touch any element in $(b, c]$ before time t_c . Let d be the element in $(a, b]$ with largest time $\tau(d, t_c - 1)$. If there are several such d , let d be the rightmost such element. Since no element in $(b, c]$ is accessed before time t_c , d is also the rightmost element in $(a, c]$ with largest time $\tau(d, t_c - 1)$. The rectangle with corners $(d, \tau(d, t_c - 1))$ and (c, t_c) contains no other point and hence d is touched at time t_c .

Clearly, $\tau(a, t_c) \leq t_c = \tau(d, t_c)$. Thus a is hidden in $[1, d)$ after time t_c and hence a is not touched after time t_c . ■

Corollary B.7 and Lemma B.8 together imply that each element a is only touched at most $|\text{wing}_L(a)| + |\text{wing}_R(a)| + 1$ times, so the total cost is at most $\sum_{a \in [n]} (|\text{wing}_L(a)| + |\text{wing}_R(a)| + 2) \leq 4n$. This concludes the proof of Theorem B.1.

B Access sequences with initial trees

B.1 Initial trees.

In the BST model, the *initial tree* is the state of the BST before the first element is accessed. In the tree-view of BST execution, we know the structure of the tree exactly in every time step, and hence the structure of the initial tree as well. This is less obvious, however, in the geometric view. In fact, according to the reduction of DHIKP, the initial tree is always a *split tree* [11], hidden from us in geometric view. We show how to “force” an arbitrary initial tree into the geometric view of BST execution.

First, we observe a simple fact. Given a binary search tree T , we denote the search path of an element a in T as $\text{path}(a, T)$. It is easy to show that two BSTs T_1 and T_2 , containing $[n]$, have the same structure iff $\text{path}(a, T_1) = \text{path}(a, T_2)$ as a set, for all $a \in [n]$.

The stair of an element is, in some sense, the geometric view of the search path. Indeed, $\text{path}(a, T) = \{b \mid b \text{ is touched if } a \text{ is accessed in } T\}$. Similarly, $\text{stair}_t(a) = \{b \mid b \text{ is touched if } a \text{ is accessed at time } t\}$. Thus, to model an initial tree T into the geometric view, it is enough to enforce that $\text{stair}_1(a) = \text{path}(a, T)$.

We can achieve this in a straightforward way, by putting stacks of points in each column below the first row, corresponding to the structure of the initial tree T . More precisely, in column x we put a stack of height

$d - d(x) + 1$, where $d(x)$ is the depth of element x in T , and d is the maximum depth of an element (Figure 2). One can easily show that $\text{stair}_1(\mathbf{a}) = \text{path}(\mathbf{a}, T)$ for all $\mathbf{a} \in [n]$.

We remark that the points representing the initial tree are not counted in the cost.

B.2 k -increasing sequences

We prove Lemma 3.3 (iii) and (iv).

Lemma B.1 (Lemma 3.3(iii)). *$\text{Inc}_{k+1}(\text{Dec}_{k+1})$ is a k -increasing (k -decreasing) gadget for GREEDY where $\text{Inc}_k = (k, \dots, 1)$ and $\text{Dec}_k = (1, 2, \dots, k)$.*

Proof. We only prove that Inc_{k+1} is a k -increasing gadget. Let q_0, \dots, q_k be touch points in $\text{GREEDY}_T(X)$ such that $q_i.y < q_{i+1}.y$ for all i and form Inc_{k+1} with bounding box B . Note that, for all $1 \leq i \leq k$, we can assume that there is no touch point q'_i where $q'_i.x = q_i.x$ and $q_{i-1}.y < q'_i.y < q_i.y$, otherwise we can choose q'_i instead of q_i . So, for any $t \in [q_{i-1}.y, q_i.y]$ we have that $\tau(q_{i-1}.x, t) \geq \tau(q_i.x, t)$ for all $1 \leq i \leq k$. Suppose that there is no access point in B .

By Lemma 2.4 (ii), $q_1.x$ is hidden in $[1, q_0.x]$ after $q_0.y$. So there must be an access point $p_1 \in [1, B.xmin) \times (q_0.y, q_1.y]$, otherwise q_1 cannot be touched. We choose such p_1 where $p_1.x$ is maximized. If $p_1.y < q_1.y$, then $\tau(p_1.x, p_1.y), \tau(q_0.x, p_1.y) \geq \tau(q_1.x, p_1.y)$. So by Lemma 2.4 (iii), $q_1.x$ is hidden in $(p_1.x, q_0.x]$ after $p_1.y$ and hence q_1 cannot be touched, which is a contradiction. So we have $p_1.y = q_1.y$ and $p_1.x < B.xmin$.

Next, we prove by induction from $i = 2$ to k that, given the access point p_{i-1} where $p_{i-1}.y = q_{i-1}.y$ and $p_{i-1}.x < B.xmin$, there must be an access point p_i where $p_i.y = q_i.y$ and $p_{i-1}.x < p_i.x < B.xmin$. Again, by Lemma 2.4 (iii), there must be an access point $p_i \in (p_{i-1}.x, q_{i-1}.x) \times (q_{i-1}.y, q_i.y]$ otherwise q_i cannot be touched. We choose the point p_i where $p_i.x$ is maximized. If $p_i.y < q_i.y$, then $\tau(p_i.x, p_i.y), \tau(q_{i-1}.x, p_i.y) \geq \tau(q_i.x, p_i.y)$. So by Lemma 2.4 (iii), $q_i.x$ is hidden in $(p_i.x, q_{i-1}.x]$ after $p_i.y$ and hence q_i cannot be touched. So we have $p_i.y = q_i.y$ and $p_{i-1}.x < p_i.x < B.xmin$. Therefore, we get $p_1.x < \dots < p_k.x < B.xmin$ which concludes the proof. ■

Lemma B.2 (Lemma 3.3(iv)). *$\text{Dec}_{k+1}(\text{Inc}_{k+1})$ is a capture gadget for GREEDY that runs on k -increasing (k -decreasing) sequence.*

Proof. We only prove that Inc_{k+1} is a capture gadget for GREEDY running on k -decreasing sequence X . Since Inc_{k+1} is an increasing gadget, if Inc_{k+1} appears in $\text{GREEDY}_T(X)$ with bounding box B , then either there is an access point in B or there are access points forming $(1, 2, \dots, k)$. The latter case cannot happen because X avoids $(1, 2, \dots, k)$. So Inc_{k+1} is a capture gadget. ■

C Access sequences without initial trees: k -decomposable permutations

C.1 Pattern avoidance and recursive decomposability

We restate and prove Lemma 1.3 that connects the k -decomposability of a permutation with pattern avoidance properties.

Lemma C.1 (Extended form of Lemma 1.3). *Let P be a permutation, and let k be an integer. The following statements are equivalent:*

- (i) P is k -decomposable,
- (ii) P avoids all simple permutations of size at least $k + 1$,
- (iii) P avoids all simple permutations of size $k + 1$ and $k + 2$.

Proof. (ii) \implies (iii) is obvious.

(iii) \implies (ii) follows from the result of Schmerl and Trotter [33] that every simple permutation of length n contains a simple permutation of length $n - 1$ or $n - 2$, and the simple observation that if P avoids Q , then P also avoids all permutations containing Q .

(ii) \implies (i) follows from the observation that a permutation P contains all permutations in its block decomposition tree \mathcal{T}_P . Further, it is known [3] that every permutation P has a block decomposition tree \mathcal{T}_P in which all nodes are simple permutations. If P contains no simple permutation of size $k + 1$ or more, it must have a block decomposition tree in which all nodes are simple permutations of size at most k , it is therefore, k -decomposable.

(i) \implies (ii): we show the contrapositive $\neg(ii) \implies \neg(i)$. Indeed, if P contains a simple permutation Q of size at least $k + 1$, then any k -decomposition of P would induce a k -decomposition of Q , contradicting the fact that Q is simple. \blacksquare

C.2 k -decomposable permutations

We prove Lemma 3.3 (v) and (vi). We refer to §3 for the necessary definitions. We show an example of the gadget Alt_k :

$$\text{Alt}_7 = \left(\begin{array}{c} \cdot \\ \cdot \quad \cdot \\ \cdot \quad \cdot \quad \cdot \\ \cdot \quad \cdot \\ \cdot \end{array} \right).$$

Lemma C.2 (Lemma 3.3(v)). *Alt_{k+1} is a k -alternating gadget for GREEDY where $\text{Alt}_k = (\lfloor (k + 1)/2 \rfloor, k, 1, k - 1, 2, \dots)$.*

Proof. Let q_0, \dots, q_k be touch points in $\text{GREEDY}(X)$ such that $q_i.y < q_{i+1}.y$ for all i and form Alt_{k+1} with bounding box B . Suppose that there is no access point in B .

We prove for $i = 1, \dots, k$, that there exists an access point $p_i \in (B.xmax, n] \times (q_{i-1}.y, q_i.y]$ for odd i , and $p_i \in [1, B.xmin) \times (q_{i-1}.y, q_i.y]$ for even i .

For odd i , by Lemma 2.4 (i), $q_i.x$ is hidden in $(q_{i-1}.x, n]$ after $q_{i-1}.y$. So there must be an access point $p_i \in (B.xmax, n] \times (q_{i-1}.y, q_i.y]$, otherwise q_i cannot be touched.

For even i , by Lemma 2.4 (ii), $q_i.x$ is hidden in $[1, q_{i-1}.x)$ after $q_{i-1}.y$. So there must be an access point $p_i \in [1, B.xmin) \times (q_{i-1}.y, q_i.y]$, otherwise q_i cannot be touched.

Observe that p_1, \dots, p_k satisfy $B.ymin \leq p_1.y < \dots < p_k.y \leq B.ymax$ and $B.xmax < p_i.x$ for all odd i and $p_i.x < B.xmin$ for all even i , fulfilling the condition of Definition 3.2. Hence, Alt_{k+1} is a k -alternating gadget. \blacksquare

Lemma C.3 (Lemma 3.3(vi)). *Alt_{k+4} is a capture gadget for GREEDY that runs on k -decomposable sequence without initial tree.*

Proof. Let X be a k -decomposable permutation. Since Alt_{k+4} is a $(k + 3)$ -alternating gadget, if Alt_{k+4} appears in $\text{GREEDY}(X)$ with bounding box B , then either there is an access point in B or there are access points p_1, p_2, \dots, p_{k+3} , satisfying the conditions $B.ymin \leq p_1.y < \dots < p_{k+3}.y \leq B.ymax$ and $B.xmax < p_i.x$ for odd i and $p_i.x < B.xmin$ for even i . Suppose there is no access point in B .

Let q_0, q_1, \dots, q_{k+3} denote the points that form Alt_{k+4} (where $q_0.y < q_1.y < \dots < q_{k+3}.y$), and let us denote $\mathcal{P} = \{p_1, \dots, p_{k+3}\}$. Let \mathcal{T} be a block decomposition tree of X of arity at most k . We look at each block $P \in \mathcal{T}$ as a minimal rectangular region in the plane that contains all access points in the block.

Let P^* be the *smallest* block in \mathcal{T} that contains two distinct access points $p_i, p_j \in \mathcal{P}$ such that i is odd, and j is even. (Observe that p_i is to the right of B , and p_j is to the left of B .)

Observe first that the bounding box of P^* must contain or intersect both vertical sides of B , otherwise it could not contain points on both sides of B . Furthermore, the bottom horizontal side of B must be contained entirely in the bounding box of P^* : If that were not the case, then there would be no accesses below B ,

and q_0 (the lowest point of Alt_{k+4}) would not be touched by GREEDY (since there is no initial tree). The following is the key observation of the proof.

Lemma C.4. *Let k' be the largest integer such that P^* contains $q_0, \dots, q_{k'}$. Then, $k' < k + 1$. In words, P^* contains at most the $k + 1$ lowest points of Alt_{k+4} .*

Proof. Let $P^* = \tilde{P}(P_1, \dots, P_k)$ be the decomposition of P^* into at most k blocks. We claim that each block P_j can contain at most one point from \mathcal{P} . First, by the minimality of P^* , each block P_j never contains two access points from \mathcal{P} from different sides of B ; and P_j cannot contain two points from \mathcal{P} from the same side of B either, because assuming otherwise that P_j contains two points from \mathcal{P} from the left of B , it would follow that there is another access point $(a, t) \in \mathcal{P}$ on the right of B outside of P_j such that $(P_j).y_{\min} < t < (P_j).y_{\max}$, contradicting that P_j is a block.

Observe further that, except for the block containing p_1 , each block P_j cannot overlap with B . Since we have at most k such blocks in the decomposition, and because $q_i.y \geq p_i.y$ for all i , it follows that the top boundary of P^* must be below q_{k+1} . ■

Since the bounding box of P^* contains (at best) the first $k+1$ points of Alt_{k+4} , the three remaining points of Alt_{k+4} need to be touched after the last access in P^* . In the following we show that this is impossible.

Let (L, t_L) be the topmost access point inside P^* to the left of B , such that there is a touch point inside B at time t_L . Let (R, t_R) be the topmost access point inside P^* to the right of B , such that there is a touch point inside B at time t_R . Let L' be the rightmost element inside B touched at time t_L , and let R' be the leftmost element inside B touched at time of t_R .

Lemma C.5. *After time t_L , within the interval $[B.x_{\min}, L']$, only L' can be touched. Similarly, after time t_R , within the interval $[R', B.x_{\max}]$, only R' can be touched.*

Proof. Let L'' be the rightmost touched element in $[L, B.x_{\min})$ at time t_L , and let R'' be the leftmost touched element in $(B.x_{\max}, R]$ at time t_R .

We have that $\tau(L'', t_L), \tau(L', t_L) \geq \tau(x, t_L)$, for any $x \in [B.x_{\min}, L')$. Thus, any $x \in [B.x_{\min}, L')$ is hidden in (L'', L') after t_L .

Above P^* there can be no access to an element in $[L'', L']$, since that would contradict the fact that P^* is a block. Within P^* after t_L there can be no access to an element in $(L'', B.x_{\min})$, as the first such access would necessarily cause the touching of a point inside B , a contradiction to the choice of (L, t_L) . An access within B is ruled out by our initial assumption.

Thus, there is no access in $(L'', L') \times (t_L, n]$, and from Lemma 2.4(iii) it follows that any $x \in [B.x_{\min}, L')$ can not be touched after time t_L .

We argue similarly on the right side of B . ■

As a corollary of Lemma D.5, note that above P^* only elements within $[L', R']$ can be touched within B . If $L' = R'$, then only one element can be touched, and we are done. Therefore, we assume that L' is strictly to the left of R' .

We assume w.l.o.g. that $t_R > t_L$ (the other case is symmetric). Let (Q, t_Q) be the first access point left of B with $t_Q > t_R$ such that there is a touch point inside B at time t_Q . Observe that (Q, t_Q) is outside of P^* by our choice of (L, t_L) . If no such Q exists, we are done, since we contradict the fact that Alt_{k+4} is a $(k+3)$ -alternating gadget.

Let (Z, t_Z) denote the last touch point with the property that $Z \in [(P^*).x_{\min}, L')$, and $t_Z \in [t_L, t_Q)$. If there are more such points in the same row, pick Z to be the leftmost. Note that (Z, t_Z) might coincide with (L, t_L) . We have $t_Z \leq t_R$ because otherwise (Q, t_Q) cannot exist. Note that $t_Z > t_R$ would imply that elements in $[L', R']$ are hidden in (Z, R) after t_R .

Next, we make some easy observations.

Lemma C.6. *The following ranges are empty:*

- (i) $[B.xmin, B.xmax] \times [t_R + 1, t_Q - 1]$,
- (ii) $[1, Z - 1] \times [t_Z + 1, t_R]$,
- (iii) $[Q, Z - 1] \times [t_Z + 1, t_Q - 1]$,
- (iv) $[Q, R' - 1] \times [t_R + 1, t_Q - 1]$.

Proof. (i) It is clear that there can be no access point inside B by assumption. Suppose there is a touch point in $[B.xmin, B.xmax] \times [t_R + 1, t_Q - 1]$, and let $(Q', t_{Q'})$ be the first (lowest) such point. Denote the access point at time $t_{Q'}$ as $(Q'', t_{Q''})$. Clearly $t_{Q'} > (P^*).ymax$ must hold, otherwise the choice of (L, t_L) or (R, t_R) would be contradicted. Also $Q'' > (P^*).xmax$ must hold, otherwise the choice of (Q, t_Q) would be contradicted.

But this is impossible, since $Q' \in [L', R']$, and $\tau(Q', t_{Q'} - 1) \leq t_R$, and thus the rectangle with corners $(Q'', t_{Q''})$ and $(Q', \tau(Q', t_{Q'} - 1))$ contains the touch point (R, t_R) , contradicting the claim that GREEDY touches Q' at time $t_{Q'}$.

- (ii) All elements in $[1, Z - 1]$ are hidden in $[1, Z - 1]$ after t_Z . There can be no access points on the left of P^* , due to the structure of the block decomposition, and there is no access point in $[(P^*).xmin, Z - 1] \times [t_Z + 1, t_Q - 1]$ due to the choice of Z . Hence there cannot be a touch points in $[1, Z - 1] \times [t_Z + 1, t_R]$.
- (iii) First there is no access point in $[Q, Z - 1] \times [t_Z + 1, (P^*).ymax]$ due to the choice of (Z, t_Z) and the structure of block decomposition. Also, there is no access point in $[Q, Z - 1] \times ((P^*).ymax, t_Q - 1]$: Assume there were such an access point $(Q', t_{Q'})$. Then it must be the case that $Q < Q' < (P^*).xmin$ and $(P^*).ymax < t_{Q'} < t_Q$. Any rectangle formed by (Q, t_Q) and a point in $P^* \cap B$ would have contained $(Q', t_{Q'})$, a contradiction to the fact that GREEDY touches a point inside B at time t_Q .

Next, we argue that there is no non-access touch point $(a, t) \in [Q, Z - 1] \times [t_Z + 1, t_Q - 1]$. There are three cases.

- $(a, t) \in [(P^*).xmin, Z - 1] \times [t_Z + 1, t_Q - 1]$ contradicts the choice of (Z, t_Z) .
 - $(a, t) \in [Q, (P^*).xmin] \times [t_Z + 1, (P^*).ymax]$ contradicts the fact that all elements in $[1, Z]$ are hidden in $[1, Z]$ after t_Z , and there is no access in $[1, Z]$ in the time interval $[t_Z + 1, (P^*).ymax]$, since P^* is a block.
 - $(a, t) \in [Q, (P^*).xmin] \times ((P^*).ymax, t_Q - 1]$, contradicts the claim that at time t_Q GREEDY touches a point inside B , since any rectangle formed by (Q, t_Q) and a point in $P^* \cap B$ would have contained (a, t) .
- (iv) Given the previous claims, it remains only to prove that there is no touch point $(a, t) \in [Z, B.xmin] \times t \in [t_R + 1, t_Q - 1]$. There cannot be such a touch point for $t \leq (P^*).ymax$ due to the choice of (Z, t_Z) . For $t > (P^*).ymax$, there cannot be such an access point due to the structure of block decomposition. Remains the case when (a, t) is a non-access touch point in $[Z, B.xmin] \times ((P^*).ymax, t_Q - 1]$. This is also impossible, as any rectangle formed by (Q, t_Q) and a point in $P^* \cap B$ would have contained (a, t) , contradicting the choice of (Q, t_Q) . ■

Lemma C.7. *At time t_Q we touch R' . Let L^* be the leftmost element touched in $[L', R']$ at time t_Q (it is possible that $L^* = R'$). After time t_Q , only L^* and R' can be touched within $[L', R']$.*

Proof. The first part follows from the emptiness of rectangles in Lemma D.6. That is, the lemma implies that the rectangle formed by (Q, t_Q) and (R', t_R) is empty, so GREEDY touches R' at time t_Q .

The elements $a \in (L^*, R')$ cannot be touched because they are hidden in (L^*, R') after t_Q , and there is no access point in this range after t_Q , due to the structure of the block decomposition. Suppose that some element $a \in [L', L^*)$ is touched (for the first time) at some time $t > t_Q$ by accessing element x . So the rectangle formed by (x, t) and $(a, \tau(a, t - 1))$ is empty, and we know that $\tau(a, t - 1) < t_R$. Notice that




Figure 7: Illustration of the proof of Lemma 3.3(vi).

$x < (P^*).x_{min}$ for otherwise (R, t_R) would have been in the rectangle formed by (x, t) and $(a, \tau(a, t - 1))$, a contradiction. Furthermore, $\tau(a, t - 1) > t_Z$ for otherwise, (Z, t_Z) would have been in the rectangle formed by (x, t) and $(a, \tau(a, t - 1))$, a contradiction. Now, since $\tau(a, t - 1) > t_Z$, Lemma D.6 implies that the rectangle formed by (Q, t_Q) and $(a, \tau(a, t - 1))$ must be empty, therefore a is touched at time t_Q ; this contradicts the choice of L^* . ■

Lemma D.4, D.5, and D.7 together imply that in total at best the $k + 3$ lowest points of Alt_{k+4} can be touched (out of the $k + 4$ needed). This means that our assumption that B is free of access points was false. Therefore Alt_{k+4} is a capture gadget, finishing the proof of Lemma 3.3(vi). See Figure 7 for an illustration. ■

D Decomposition theorem and applications

D.1 Preliminaries

Given a permutation $\sigma : [n] \rightarrow [n]$, we call a set $[a, b] \times [c, d]$, with the property that $\{\sigma(a), \sigma(a+1), \dots, \sigma(b)\} = [c, d]$, a *block* of σ induced by the interval $[a, b]$. Let P be a block induced by interval I . Consider a partitioning of I into k intervals, denoted (left-to-right) as I_1, I_2, \dots, I_k , such that each interval induces a block, denoted respectively as P_1, P_2, \dots, P_k . Consider the unique permutation \tilde{P} of size $[k]$ that is order-isomorphic to a sequence of arbitrary representative elements from the blocks P_1, P_2, \dots, P_k . We denote $P = \tilde{P}[P_1, P_2, \dots, P_k]$ a *deflation* of P , and we call \tilde{P} the *skeleton* of P . Visually, \tilde{P} is obtained from P by contracting each block P_i to a single point (Figure 1). Deflation is analogously defined on permutations. Every permutation σ has the trivial deflations $\sigma = \sigma[1, \dots, 1]$ and $\sigma = 1[\sigma]$. Permutations that have only the trivial deflations are called *simple permutations*.

A *block decomposition tree* \mathcal{T} of σ is a tree in which each node is a block of σ . We define \mathcal{T} recursively as follows: The root of \mathcal{T} is the block $[n] \times [n]$. A node P of \mathcal{T} is either a leaf, or it has children P_1, P_2, \dots corresponding to a nontrivial deflation $P = \tilde{P}[P_1, P_2, \dots]$. To every non-leaf block P we associate its skeleton

\tilde{P} . For convenience, we assume that the blocks P_1, \dots, P_k are ordered by their y -coordinates from bottom to top. We partition P into k^2 rectangular regions whose boundaries are aligned with P_1, \dots, P_k , and index them $R(i, j)$ with i going left to right and j going bottom to top. In particular, the region $R(i, j)$ is aligned with P_j (same y -coordinates). To every leaf block $P = [a, b] \times [c, d]$ we associate the unique permutation that is order-isomorphic to $(\sigma(a), \sigma(a+1), \dots, \sigma(b))$. When there is no chance of confusion, we refer interchangeably to a block and its associated permutation. We denote the leaves and non-leaves of the tree \mathcal{T} by $L(\mathcal{T})$ and $N(\mathcal{T})$ respectively.

A permutation is k -decomposable if it has a block decomposition tree \mathcal{T} in which every node has an associated permutation of size at most k . As an example, we observe that preorder sequences of binary search trees are 2-decomposable.

D.2 A robust GREEDY

We augment GREEDY with some additional steps and we obtain an offline algorithm (called RGREEDY) with several desirable properties. We remark that in this section we always assume the preprocessing discussed in §2 or equivalently, in geometric view we assume that RGREEDY is executed without initial tree (since RGREEDY is offline, this is a non-issue).

Before describing the algorithm, we give some simple definitions.

We denote the top left, respectively top right corner point of a rectangular region R as $opleft(R)$, respectively $topright(R)$. We define $topwing(R)$ to be a subset of the points touched in R , that are the last touch points of their respective columns. A point q is in $topwing(R)$ if at least one of the following properties is satisfied:

- The rectangle with corners q and $opleft(R)$ is empty,
- The rectangle with corners q and $topright(R)$ is empty,
- q is in the leftmost or rightmost column of R .

Note that the set $topwing(R)$ depends on the time point where we are in the execution of the algorithm, but for simplicity of notation, we omit the reference to a specific time (it will always be clear from the context).

We can now describe how RGREEDY processes an input point (a, t) . First, it acts as GREEDY would, i.e. for every empty rectangle with corners (a, t) and (b, t') , it touches the point (b, t) on the current timeline. Let Y_t be the points touched in this way (including the point (a, t)). Second, it touches additional points on the timeline. Let $v_1, v_2 = parent(v_1), \dots, v_\ell = parent(v_{\ell-1})$ be the nodes of the decomposition tree such that the blocks associated with these nodes end at time t . Note that v_1 is a leaf. Let $v_{\ell+1} = parent(v_\ell)$; $v_{\ell+1}$ exists except when the last input is processed. For $h \in [2, \ell+1]$ in increasing order, we consider the deflation $\tilde{P}[P_1, \dots, P_k]$ associated with v_h . Let (a, t) lie in P_j ; note that $j = k$ if $h \leq \ell$ and $j < k$ for $h = \ell + 1$. For every sibling region $R(i, j)$ of P_j , RGREEDY touches the projection of $topwing(R(i, j))$ on the timeline t .

Let $Y_{h,t}$ be the points touched. Figure 3 illustrates the definition of RGREEDY.

We need to argue that the augmentation step does not violate the feasibility of the solution.

Lemma D.1. *The output of RGREEDY is feasible.*

Proof. Suppose for contradiction that at time t we have created an unsatisfied rectangle r with corners (a, t) and (x, y) , where $y < t$. The point (a, t) is the result of projecting (a, t') onto the horizontal line t , where $(a, t') \in topwing(R)$, for some region R processed at time t . Clearly $t' < t$. Assume that before adding (a, t) , the rectangle formed by the points (a, t') and $opleft(R)$ is empty. A symmetric argument works if (a, t') formed an empty rectangle with $topright(R)$. We have $y > t'$, for otherwise the point (a, t') would satisfy r . Let $R = R(i, j)$. Since (x, y) lies in the time window of R , we have $(x, y) \in R(i', j)$ for some i' .

Suppose first that $i = i'$. If $x \leq a$, then the rectangle with corners (a, t') and $opleft(R)$ contains (x, y) , contradicting the fact that (a, t') was in $topwing(R)$ before the augmentation. Therefore $x > a$ holds. But then, if r is not satisfied, then the rectangle with corners (x, y) and $opleft(R)$ had to be empty before the augmentation, and therefore (x, t) must have been touched in the augmentation step, making r satisfied.

The case $i \neq i'$ remains. If $i' < i$, then r contains the rectangle formed by (x, y) and the top right corner of $R(i', j)$. Either this rectangle is nonempty, or (x, y) was in $topwing(R(i', j))$ before the augmentation, in

which case (x, t) was touched in the augmentation step. In both cases r must be satisfied. The case when $i' > i$ is symmetric, finishing the proof. ■

Observe that RGREEDY processes the input and emits the output line-by-line, in increasing order of time. It is, however, not an online algorithm, as it has a priori access to a block decomposition tree of the input.

D.3 Decomposition theorem

Our main technical contribution in this section is the following result:

Theorem D.2 (Decomposition theorem for permutations). *For any block decomposition tree \mathcal{T} of a permutation P , the total cost of RGREEDY is bounded as*

$$\text{RGREEDY}(P) \leq 4 \cdot \sum_{\tilde{P} \in N(\mathcal{T})} \text{GREEDY}(\tilde{P}) + \sum_{P \in L(\mathcal{T})} \text{GREEDY}(P) + 3n.$$

We discuss applications of Theorem E.2.

Application 1: Cole et al.’s “showcase” sequences. In Cole et al. [10], a showcase sequence is defined in order to introduce necessary technical tools for proving the dynamic finger theorem. This sequence can be defined (in our context) as a permutation $P = \tilde{P}[P_1, \dots, P_{n/\log n}]$ where each P_j is a sequence for which GREEDY has linear linear access cost, i.e. $|P_j| = \log n$, and $\text{GREEDY}(P_j) = O(\log n)$. Now, by applying Theorem E.2, we can say that the cost of RGREEDY on P is $\text{RGREEDY}(P) \leq 4 \cdot \text{GREEDY}(\tilde{P}) + \sum_j \text{GREEDY}(P_j) + O(n) \leq O(\frac{n}{\log n} \cdot \log(\frac{n}{\log n})) + \frac{n}{\log n} O(\log n) + O(n) = O(n)$. This implies that our algorithm has linear access cost on this sequence.

Application 2: Simple permutations as the “core” difficulty of BST inputs. We argue that whether RGREEDY performs well on all sequences depends only on whether GREEDY does well on simple permutation sequences. More formally, we prove the following theorem.

Theorem D.3. *If GREEDY is c -competitive on simple permutations, then RGREEDY is an $O(1)$ approximation algorithm for $\text{OPT}(X)$ for all permutations X .*

Using Theorem 2.1 we can extend this statement from all permutations to all access sequences.

To prove Theorem E.3 we use of the known easy fact [3] that every permutation P has a block decomposition tree \mathcal{T} in which all permutations in $L(\mathcal{T})$ and all permutations in $N(\mathcal{T})$ are simple. Applying Theorem E.2 and the hypothesis that GREEDY is c -competitive for simple permutations, we get the bound:

$$\text{RGREEDY}(P) \leq 4c \cdot \sum_{\tilde{P} \in N(\mathcal{T})} \text{OPT}(\tilde{P}) + c \cdot \sum_{P \in L(\mathcal{T})} \text{OPT}(P) + 3n.$$

Now we need to decompose OPT into subproblems the same way as done for RGREEDY. The following lemma finishes the proof (modulo the proof of Theorem E.2), since it implies that $\text{RGREEDY}(P) \leq 4c \cdot \text{OPT}(P) + (8c + 3)n$.

Lemma D.4. $\text{OPT}(P) \geq \sum_{P \in \mathcal{T}} \text{OPT}(P) - 2n$.

Proof. We show that for an arbitrary permutation $P = \tilde{P}[P_1, \dots, P_k]$, it holds that $\text{OPT}(P) \geq \text{OPT}(\tilde{P}) + \sum_{i=1}^k \text{OPT}(P_i) - k$. The result then follows simply by iterating this operation for every internal node of the block decomposition tree \mathcal{T} of P and observing that the numbers of children of each internal node add up to at most $2n$.

Let Y^* be the set of points of the optimal solution for P . We partition Y^* into sets Y_i^* contained inside each block P_i , and we let $Y_{out}^* = Y^* \setminus \bigcup_{i=1}^k Y_i^*$. Each Y_i^* must clearly be feasible for P_i , so we have $|Y_i^*| \geq \text{OPT}(P_i)$. It remains to show that $|Y_{out}^*| \geq \text{OPT}(\tilde{P}) - k$.

Again, consider the partitioning of P into k^2 rectangular regions $R(i, j)$, aligned with the blocks P_1, \dots, P_k . Consider the contraction operation that replaces each region $R(i, j)$ with a point (i, j) if and only if $Y^* \cap R(i, j)$ is nonempty. We obtain a point set $Y' \in [k] \times [k]$, and clearly $|Y'| \leq |Y_{out}^*| + |\tilde{P}|$. We claim that the obtained point set Y' is a satisfied superset of \tilde{P} , and therefore $|Y_{out}^*| \geq \text{OPT}(\tilde{P}) - k$.

Indeed, suppose for contradiction that there is an unsatisfied rectangle with corner points $p, q \in Y'$. Suppose p is to the upper-left of q (the other cases are symmetric). Let R_p and R_q be the regions before the contraction that were mapped to p , respectively q . Since $Y^* \cap R_p$ is nonempty, we can consider p' to be the lowest point in $Y^* \cap R_p$ and if there are several such points, we take p' to be the rightmost. Similarly, since $Y^* \cap R_q$ is nonempty, we can take q' to be the topmost point in $Y^* \cap R_q$ and if there are several such points, we take q' to be the leftmost. Now, if the rectangle with corners p and q is unsatisfied in Y' , then the rectangle with corners p' and q' is unsatisfied in Y^* , contradicting the claim that Y^* is a feasible solution. \blacksquare

Application 3: Recursively decomposable permutations. We generalize preorder access sequences to k -decomposable permutations, and show that **RGREEDY** has linear access cost for all such permutations. Let k be a constant such that there is a tree decomposition \mathcal{T} in which each node is a permutation of size at most k . Notice that a preorder sequence is 2-decomposable.

Theorem D.5. *The cost of **RGREEDY** on any k -decomposable access sequence is at most $O(n \log k)$.*

Proof. Applying Theorem E.2, we obtain that $\text{RGREEDY}(P) \leq 4 \sum_{\tilde{P} \in N(\mathcal{T})} \text{GREEDY}(\tilde{P}) + \sum_{P \in L(\mathcal{T})} \text{GREEDY}(P) + O(n)$. First, notice that the elements that belong to the leaf blocks are induced by disjoint intervals, so we have $\sum_{P \in L(\mathcal{T})} |P| = n$. For each such leaf block, since **GREEDY** is known to have logarithmic amortized cost, we have $\text{GREEDY}(P) \leq O(|P| \log |P|) = O(|P| \log k)$. Summing over all leaves, we have $\sum_{P \in L(\mathcal{T})} \text{GREEDY}(P) \leq O(n \log k)$.

Next, we bound the cost of non-leaf blocks. Since the numbers of children of non-leaf blocks add up to at most $2n$, we have $\sum_{\tilde{P} \in N(\mathcal{T})} |\tilde{P}| \leq 2n$. Again, using the fact that **GREEDY** has amortized logarithmic cost, it follows that $\sum_{\tilde{P} \in N(\mathcal{T})} \text{GREEDY}(\tilde{P}) \leq \sum_{\tilde{P} \in N(\mathcal{T})} O(|\tilde{P}| \log k) \leq O(n \log k)$. This completes the proof. \blacksquare

We mention that the logarithmic dependence on k is essentially optimal, since there are simple permutations of size n with optimum cost $\Omega(n \log n)$. In fact, for arbitrary k , the bound is tight, as shown in Appendix H.2. In Theorem 1.4 we proved that for arbitrary constant k , the linear upper bound for k -decomposable permutations is asymptotically matched by **GREEDY**.

D.4 Proof of Theorem E.2

We are now ready to prove the decomposition theorem. Let Y be the output of **RGREEDY** on a permutation of P of size n , and let \mathcal{T} be the block decomposition tree of P . We wish to bound $|Y|$.

For each leaf block $P \in L(\mathcal{T})$, let Y_P be the points in $P \cap Y$. For each non-leaf block $\tilde{P} \in N(\mathcal{T})$, with children P_1, \dots, P_k , let $Y_{\tilde{P}}$ be the set of points in $Y \cap (\tilde{P} \setminus \bigcup_j P_j)$. In words, these are the points in block \tilde{P} that are not in any of its children (we do not want to overcount the cost). It is clear that we have already accounted for all points in Y , i.e. $Y = \bigcup_{P \in \mathcal{T}} Y_P$, so it suffices to bound $\sum_{P \in \mathcal{T}} |Y_P|$.

For each non-leaf block $\tilde{P} \in N(\mathcal{T})$, we show later that $|Y_{\tilde{P}}| \leq 4 \cdot \text{GREEDY}(\tilde{P}) + 5 \text{degree}(\tilde{P}) + \text{junk}(\tilde{P})$, where $\text{degree}(\tilde{P})$ is the number of children of \tilde{P} , and the term $\text{junk}(\tilde{P})$ will be defined such that $\sum_{\tilde{P} \in N(\mathcal{T})} \text{junk}(\tilde{P}) \leq 2n$. For each leaf block P , notice that apart from the **RGREEDY** augmentation in the last row, the cost is exactly $\text{GREEDY}(P)$. Including the extra points touched by **RGREEDY** we have that $|Y_P| = \text{GREEDY}(P) + |\text{topwing}(P)|$. Since leaf blocks are induced by a partitioning of $[n]$ into disjoint intervals, we have that $\sum_{P \in L(\mathcal{T})} |\text{topwing}(P)| \leq n$. So in total, the cost can be written as $4 \cdot \sum_{\tilde{P} \in N(\mathcal{T})} \text{GREEDY}(\tilde{P}) + \sum_{P \in L(\mathcal{T})} \text{GREEDY}(P) + 5n$ as desired. It remains to bound $\sum_{\tilde{P} \in N(\mathcal{T})} |Y_{\tilde{P}}|$.

Consider any non-leaf block \tilde{P} . Define the rectangular regions $R(i, j)$ as before. We split $|Y_{\tilde{P}}|$ into three separate parts. Let $Y_{\tilde{P}}^1$ be the subset of $Y_{\tilde{P}}$ consisting of points that are *in the first row* of some region $R(i, j)$. Let $Y_{\tilde{P}}^3$ be the subset of $Y_{\tilde{P}}$ consisting of points that are *in the last row* of some region $R(i, j)$. Observe that any points possibly touched in an RGREEDY augmentation step fall into $Y_{\tilde{P}}^3$. Let $Y_{\tilde{P}}^2 = Y_{\tilde{P}} - (Y_{\tilde{P}}^1 \cup Y_{\tilde{P}}^3)$ be the set of touch points that are not in the first or the last row of any region $R(i, j)$.

Denote by $left(R)$ and $right(R)$ the leftmost, respectively rightmost column of a region R . Similarly, denote by $top(R)$ and $bottom(R)$ the topmost, respectively bottommost y -coordinate of a region R .

Lemma D.6. *If region $R(i, j)$ has no block P_q below it, RGREEDY touches no point in $R(i, j)$. If $R(i, j)$ has a block P_q below it, RGREEDY touches only points in columns $left(R(i, j))$ and $right(R(i, j))$.*

Proof. If $R(i, j)$ has no block P_q below it, then it is never touched. Assume therefore that $R(i, j)$ is above (and aligned with) block P_q for some $q < j$. When RGREEDY accessed the last element in P_q , it touched the left and right upper corners of P_q in the augmentation step (since the topmost element in the leftmost and rightmost columns belong to $topwing(P_q)$). As there are no more accesses after $top(P_q)$ inside $[left(P_q) + 1, right(P_q) - 1]$, all elements in this interval are hidden in $[left(P_q) + 1, right(P_q) - 1]$, and therefore not touched again. ■

Let $m_{i,j}$ be the indicator variable of whether the region $R(i, j)$ is touched by RGREEDY. Consider the execution of GREEDY on permutation \tilde{P} and define $n_{i,j}$ to be the indicator variable of whether the element (i, j) is touched by GREEDY on input \tilde{P} . Notice that $\sum_{i,j} n_{i,j} = \text{GREEDY}(\tilde{P})$. We refer to the execution of GREEDY on permutation \tilde{P} as the “execution on the contracted instance”, to avoid confusion with the execution of RGREEDY on block \tilde{P} .

Lemma D.7. *Before the augmentation of the last row of \tilde{P} by $topwing(\tilde{P})$ the following holds: A rectangular area $R(i, j)$ in block \tilde{P} is touched by RGREEDY iff GREEDY touches (i, j) in the contracted instance \tilde{P} ; in other words, $m_{i,j} = n_{i,j}$ for all i, j .*

Lemma E.6 and Lemma E.7 together immediately imply the bounds $|Y_{\tilde{P}}^1| \leq 2 \cdot \text{GREEDY}(\tilde{P})$, and $|Y_{\tilde{P}}^3| \leq 2 \cdot \text{GREEDY}(\tilde{P}) + |topwing(\tilde{P})|$, since in the first or last row of a region $R(i, j)$ that is not one of the P_j 's at most two points are touched and no point is touched if GREEDY does not touch the region. In the case of $Y_{\tilde{P}}^3$, RGREEDY also touches the elements in $topwing(\tilde{P})$ after the last access to a point in \tilde{P} .

Proof. We prove by induction on j that after P_j is processed, the rectangular area $R(i, j)$ is touched by RGREEDY if and only if (i, j) is touched by GREEDY in the contracted instance.

We consider the processing of P_j and argue that this invariant is maintained. When $j = 1$, the induction hypothesis clearly holds. There are no other touch points below the block \tilde{P} (due to the property of the decomposition), so when points in P_1 are accessed, no points in $\tilde{P} \setminus P_1$ are touched. Similarly, in the contracted instance when the first point in \tilde{P} is accessed by GREEDY, no other points are touched.

First we prove that GREEDY touching (i, j) in the contracted instance implies that RGREEDY touches $R(i, j)$. Suppose GREEDY touches point (i, j) . Let p_j be the access point in the contracted instance at time j . There has to be an access or touch point (i, j') where $j' < j$ such that the rectangle with corners (i, j') , and p' had no other points before adding (i, j) (the point (i, j') was on the stair of p_j at time j). The induction hypothesis guarantees that $R(i, j')$ is touched by RGREEDY, so some point q is touched or accessed in this region. Choose q to be the topmost such point, and if there are more points in the same row, choose q to be as close to block P_j as possible.

We claim that q is in the stair of the first access point p^* of P_j , therefore touching $R(i, j)$: If q was not in the stair, there would be another point q' in the rectangle formed by p^* and q . Observe that q' cannot be in $R(i, j')$, because it would contradict the choice of q being closest to p^* . So q' is in some region $R(i'', j'')$ that lies inside the minimal rectangle embracing $R(i, j')$, and the region immediately below P_j . By induction hypothesis, the fact that $R(i'', j'')$ contains point q' implies that (i'', j'') is touched by GREEDY in the contracted instance, contradicting the claim that the rectangle with corners (i, j') , p_j is empty.

Now we prove the other direction. Suppose RGREEDY touches region $R(i, j)$. This also implies that GREEDY (without augmentation step) already touches $R(i, j)$. But then GREEDY must touch a point in

$R(i, j)$ when accessing the first element p^* in P_j . Assume w.l.o.g. that $R(i, j)$ is to the right of P_j . It means that there is a point q in the stair of p^* , such that q is in some region $R(i, j')$ for $j' < j$. If in the contracted instance (i, j') is in the stair of the access point p_j at time j , we would be done (because GREEDY would add point (i, j)), so assume otherwise that this is not the case. Let (i'', j'') be the point that blocks (i, j') , so the corresponding region $R(i'', j'')$ lies inside the minimal rectangle embracing both P_j and $R(i, j')$. Also note that the region $R(i'', j'')$ cannot be vertically aligned with P_j because \tilde{P} is a permutation. Thus $i < i''$. The fact that (i'', j'') is touched by GREEDY implies, by induction hypothesis, the existence of a point $q' \in R(i'', j'')$, and therefore the augmentation step of RGREEDY touches some point q^* on the top row of region $R(i'', j'')$. This point would prevent q from being on the stair of p^* , a contradiction. ■

Lemma D.8. $|topwing(\tilde{P})| \leq 2 \cdot degree(\tilde{P})$.

Proof. $topwing(\tilde{P})$ contains at most 2 points in block $R(i, k)$ for every k . ■

The only task left is to bound $|Y_{\tilde{P}}^2|$, the number of touch points in non-first and non-last rows of a region. This will be the *junk* term mentioned earlier, that sums to $2n$ over all blocks $\tilde{P} \in N(\mathcal{T})$. In particular we show that every access point can contribute at most 2 to the entire sum $\sum_{\tilde{P} \in N(\mathcal{T})} |Y_{\tilde{P}}^2|$.

Consider any access point $X = (x, t_x)$, and let P^* denote the smallest block that contains X , and in which X is *not* the first access point (such a block exists for all but the very first access point of the input). Further observe, that every point touched at time t_x *inside* P^* is accounted for already in the sum $\sum_{\tilde{P} \in N(\mathcal{T})} |Y_{\tilde{P}}^1|$, since these points are in the first row of some region inside P^* . It remains to show that at time t_x at most two points are touched outside of P^* , excluding a possible RGREEDY augmentation step (which we already accounted for). Let p^* denote the first access point in P^* , i.e. at time $bottom(P^*)$. At the time of accessing p^* several points might be touched. Notice that all these points must be outside of P^* , since no element is touched before being accessed. Let p_{left} and p_{right} be the nearest of these touch points to p^* on its left, respectively on its right.

We claim that for all the access points in P^* other than the first access point p^* , only p_{left} and p_{right} can be touched outside of P^* . Since, by definition X is a non-first element of P^* , it follows that at time t_x at most two points are touched outside of P^* , exactly what we wish to prove.

Indeed, elements in the interval $[1, p_{left} - 1]$, and in the interval $[p_{right} + 1, n]$ are hidden in $[1, p_{left} - 1]$, respectively in $[p_{right} + 1, n]$, in the time interval $[bottom(P^*) + 1, top(P^*)]$ and therefore cannot be touched when accessing non-first elements in P^* . The case remains when p_{left} or p_{right} or both are nonexistent. In this case it is easy to see that if on the corresponding side an element outside of P^* were to be touched at the time of a non-first access in P^* , then it would have had to be touched at the time of p^* , a contradiction.

This concludes the proof.

E Proof of Theorem 2.1

In this section, we prove the following theorem.

Theorem E.1 (Theorem 2.1). *Assume that there is a c -competitive online geometric BST \mathcal{A}^P for all permutations. Then there is a $4c$ -competitive online geometric BST \mathcal{A} for all access sequences.*

Given a c -competitive online BST \mathcal{A}^P that can serve a permutation access sequence, we construct a $4c$ -competitive online BST \mathcal{A} that can serve any (not necessarily permutation) access sequences. We first sketch the outline of the proof.

Let $X \in [n]^m$ be a (possibly non-permutation) access sequence, and let T be the initial tree on $[n]$. First we show that given any online BST \mathcal{B} that can handle accesses, there is a canonical way to augment \mathcal{B} and obtain an *augmented* \mathcal{B}^{ins} , that can also handle insertions of elements.

Then, we show that given any access-only sequence $X \in [n]^m$, initial tree T and an online augmented BST $\mathcal{A}^{P, ins}$ that can serve permutation access and insert sequences, we can construct a permutation sequence (comprising of both accesses and insertions) $X^{P, ins} \in [m]^m$ to be fed to $\mathcal{A}^{P, ins}$ in an online manner. Then

there are two steps. First, we show that the execution of $\mathcal{A}^{p,ins}$ can define the execution of \mathcal{A} that runs on X with initial tree T . Second, we show that the cost of $\mathcal{A}^{p,ins}$ can be bounded by $O(\text{OPT}(X))$. Now we formally describe the proof.

Handling Insertion Given an online BST \mathcal{B} , let \mathcal{B}^{ins} be an algorithm that works as follows. To access element x , just simulate \mathcal{B} . To insert element x , try searching for x and miss, i.e. touching the search path of predecessor and successor of x in the tree. Then add x as a leaf. Then simulate \mathcal{B} as if \mathcal{B} is accessing the leaf x .

E.1 Constructing the algorithm \mathcal{A}

In the following, we work on $\mathbb{R} \times \mathbb{Z}$ instead of the integer grid $\mathbb{Z} \times \mathbb{Z}$ in order to simplify the algorithm and the proof. We extend some notations and definitions accordingly. We use the words “ y -coordinate” and “row” interchangeably. For any $X, T \subset \mathbb{R} \times \mathbb{Z}$, by $\lceil \frac{X}{T} \rceil$, we mean the point set obtained by “putting X above T ” while the x -coordinate of each element does not change. A point set $X \subset \mathbb{R} \times \mathbb{Z}$ is a *permutation* if there is at most one point in X for any fixed x -coordinate and any fixed y -coordinate. A satisfied set is defined in the usual no-empty-rectangle way. An online BST \mathcal{B} , given $\lceil \frac{X}{T} \rceil \subset \mathbb{R} \times \mathbb{Z}$, outputs a satisfied $\lceil \frac{\mathcal{B}_T(X)}{T} \rceil$ in an online manner.

To describe the algorithm \mathcal{A} for Theorem 2.1, we define two operations: *split* and *merge*. Given any access sequence X , we define *split*(X) as follows. Let ϵ be an extremely small positive constant. For each point $(x, y) \in X$, we have $(x + \epsilon y, y) \in \text{split}(X)$ except if (x, y) is the first access point of element x (i.e. there is no $(x', y) \in X$ where $x' < x$), in which case $(x, y) \in \text{split}(X)$. The intention is to “tilt” the non-first accesses of all elements in the simplest way. Note that *split*(X) is a permutation. We define the reverse operation $\text{merge}(S) = \{(\lfloor x \rfloor, y) \mid (x, y) \in S\}$ for any set S . For any (non-permutation) sequence $X \in \mathbb{Z} \times \mathbb{Z}$, let $X^{p,ins}, X^p = \text{split}(X)$ as a point set. All points in X^p are access points. However, in $X^{p,ins}$, only points with integral x -coordinate are access points, the remaining points are insertion points. That is, the first points of each element of X “become” access points in $X^{p,ins}$, the remaining points “become” insertion points.

Given $\lceil \frac{X}{T} \rceil \in \mathbb{Z} \times \mathbb{Z}$ as an input to \mathcal{A} , we define:

$$\mathcal{A}_T(X) = \text{merge}(\mathcal{A}_T^{p,ins}(X^{p,ins}))$$

Lemma E.2. *If \mathcal{A}^p is online, then \mathcal{A} is an online process.*

Proof. Observe that *split* and *merge* fix y -coordinate. If \mathcal{A}^p is online, then $\text{merge}(\mathcal{A}_T^{p,ins}(X^{p,ins}))$ can be computed row by row. ■

Lemma E.3. *For any (non-permutation) access sequence X , and initial tree T it holds that $\mathcal{A}_T(X) \supseteq X$ and $\lceil \frac{\mathcal{A}_T(X)}{T} \rceil$ is satisfied*

Proof. To show that $\mathcal{A}_T(X) \supseteq X$, observe that $X = \text{merge}(\text{split}(X))$ as ϵ is very small. Next, $X^{p,ins} \subseteq \mathcal{A}_T^{p,ins}(X^{p,ins})$ because $\mathcal{A}^{p,ins}$ is an online BST. Finally, as *merge* is monotone under inclusion, we have

$$X = \text{merge}(\text{split}(X)) = \text{merge}(X^{p,ins}) \subseteq \text{merge}(\mathcal{A}_T^{p,ins}(X^{p,ins})) = \mathcal{A}_T(X).$$

To show that $\lceil \frac{\mathcal{A}_T(X)}{T} \rceil$ is satisfied, we claim that *merge* preserves satisfiability. Therefore, as $\lceil \frac{\mathcal{A}_T^{p,ins}(X^{p,ins})}{T} \rceil$ is satisfied, $\text{merge}(\lceil \frac{\mathcal{A}_T^{p,ins}(X^{p,ins})}{T} \rceil) = \lceil \frac{\mathcal{A}_T(X)}{T} \rceil$ is satisfied.

It remains to prove the claim. Let $S \subset \mathbb{R} \times \mathbb{Z}$ be any satisfied set and $S' = \text{merge}(S)$. Consider any points $p', q' \in S'$ where $p'.x < q'.x$ and choose the rightmost (leftmost) associated point p (q) of p' (q') from S . As S is satisfied, there is a point $r \in S \setminus \{p, q\}$ where $r \in \square_{pq}$. In particular, $p.x \leq r.x \leq q.x$. Let $r' = (\lfloor r.x \rfloor, r.y) \in S'$ be the image of r under *merge*, then we have $r' \in \square_{p'q'}$ because $\lfloor p.x \rfloor \leq \lfloor r.x \rfloor \leq \lfloor q.x \rfloor$ and *merge* does not affect y -coordinate. Moreover, $r' \neq p', q'$, otherwise p is not the rightmost associated




Figure 8: Illustration of S' in the proof of Lemma F.5. The symbols \times denote accesses and circles are touched points.

point of p' , or similarly for q' . So we have identified $r' \in S' \setminus \{p', q'\}$ where $r' \in \square_{p'q'}$ and therefore S' is satisfied. \blacksquare

By Lemma F.2 and Lemma F.3, \mathcal{A} is an online process where $\mathcal{A}_T(X) \supseteq X$ and $\mathcal{A}_T(X)$ is satisfied. Therefore, \mathcal{A} is an online (geometric) BST for accessing arbitrary (non-permutation) sequences.

E.2 Competitiveness

The idea is to bound $\mathcal{A}_T^{p,ins}(X^{p,ins})$, which is the cost of the online BST which does both access and insert by the cost of the online BST running on X that only supports access.

Let T^p be a binary search tree with m elements corresponding to the distinct x -coordinates of points in X^p . The elements with integral values in T^p form a BST with exactly same structure as T . Between any two integral value elements v and $v + 1$, there will be a right path (each node has only right child) of size equal to the number of points from X^p with x -coordinate from $(v, v + 1)$. This tree is defined so that the following claim holds.

Lemma E.4. $\mathcal{A}_T^{p,ins}(X^{p,ins}) = \mathcal{A}_{T^p}^p(X^p)$.

Proof. We assume that \mathcal{A}^p does not touch elements in the initial tree that have never been in any of the search paths for access or insert. By induction on time step t , if the point $(x, t) \in X^{p,ins}$ is an access point, then $\mathcal{A}^{p,ins}$ and \mathcal{A}^p do the same thing by definition. If the point $(x, t) \in X^{p,ins}$ is an insertion point, then after $\mathcal{A}^{p,ins}$ adds x as a leaf, the structure of the tree excluding elements which are never touched is same. So $\mathcal{A}^{p,ins}$ and \mathcal{A}^p do the same thing again. \blacksquare

Lemma E.5. For any access sequence (no insertions) X , and the corresponding access sequence $split(X)$, $OPT(split(X)) \leq 4OPT(X)$.

Proof. We show that given $OPT(X)$, we can construct a satisfied set S' of $split(X)$ where $|S'| \leq 2OPT(X) + 2m$.

We iteratively modify $OPT(X)$ to obtain S' . Consider any column c_i of $OPT(X)$ which contains more than one access point, say m_i access points. Suppose that c_i contains $p_i \geq m_i$ touch points (thus $p_i - m_i$ non-access touch points), and lies between c_{i-1} and c_{i+1} . We “split” c_i into m_i columns $c_{i_1}, c_{i_2}, \dots, c_{i_{m_i}}$ lying between c_{i-1} and c_{i+1} . Columns c_{i_1} and $c_{i_{m_i}}$ have p_i points in the same rows as c_i ; in c_{i_1} the lowest of them is an access point, and all others above are touched, whereas in $c_{i_{m_i}}$ the highest is an access point, and all others below are touched. The remaining columns $c_{i_2}, \dots, c_{i_{m_i-1}}$ have two points each in them, one access and one touch point. The access point in column c_{i_j} is in the same row as the j th access in original column c_i , and the touch point in column c_{i_j} is in the same row as the $(j + 1)$ th access in c_i . See Figure 8.

It is easy to verify that, after placing these points, the set is still satisfied. The number of points placed in these columns is exactly $2p_i + 2(m_i - 2)$. Thus after applying this modification to every column c_i , there are at most $2OPT(X) + 2m \leq 4OPT(X)$ points. \blacksquare

Now we are ready to prove the main theorem.

Proof of Theorem 2.1. First, we have that $\mathcal{A}_T(X) \leq \mathcal{A}_T^{p,ins}(X^{p,ins})$ because *merge* never increases the number of points. We conclude that

$$\mathcal{A}_T(X) \leq \mathcal{A}_T^{p,ins}(X^{p,ins}) = \mathcal{A}_{T^p}^p(X^p) \leq c\text{OPT}(\text{split}(X)) \leq 4c\text{OPT}(X),$$

where the second inequality follows from Lemma F.4, the third inequality follows from the assumption that \mathcal{A}^p is c -competitive for only accessing on only permutations, and the last inequality follows from Lemma F.5, and hence \mathcal{A} is a $4c$ -competitive online BST for any access sequence. \square

F Random permutation is hard

We prove that the cost of most permutations is $\Omega(n \log n)$ for *all* BST algorithm. A similar result has been shown by Wilber [39] for random access sequences (that might not be permutations). More formally, we prove the following theorem.

Theorem F.1. *There is a positive constant c such that for any $L > 0$ only a fraction $1/2^L$ of the permutations X have $\text{OPT}(X) \leq (\log n! - L - n - 1)/c$.*

Proof. We use Kolmogorov complexity. Observe that the shape of a tree on k nodes can be encoded by $O(k)$ bits, e.g., by the preorder traversal. If the keys stored in the tree are clear from the context, the assignment of the keys to the nodes is unique. Let T be the initial tree, let X be any (permutation) access sequence, and let C_X be the total cost of OPT on X . We encode the shape of the initial tree with $O(n)$ bits. Consider now an access of cost k . It replaces an initial segment of the current search tree, the before-tree, by another tree on the same set of nodes, the after-tree. The before-tree must include the element accessed and has size $O(k)$. We encode the before-tree and the after-tree with $O(k)$ bits and use a special marker to designate the accessed element in the before-tree. In total, we need $c \cdot (n + C_X)$ bits for the encoding for some small constant c . Note that X can be reconstructed from the encoding. We first reconstruct the shape of the initial tree. Its labelling by keys is unique. Consider now any access. Since the position of the accessed element in the before-tree is marked and since the before- and the after-tree are on the same set of keys, the search tree after the access is determined. We conclude that we have an injective mapping from permutations X to bitstrings of length $c \cdot (n + C_X)$. Assume $C_x \leq (\log n! - L - n - 1)/c$. Then the encoding of X has at most $\log n! - L - 1$ bits and hence X is in a set of at most $n!/2^L$ permutations. \blacksquare

G Other claims

G.1 Counterexamples for straightforward attempts

Avoiding the same pattern is not enough.

It might be tempting to conjecture that if X avoids P , then the GREEDY touch matrix $\text{GREEDY}(X)$ avoids P as well (P refers both to a permutation pattern and its matrix representation). While this holds for the simple case $P = (2, 3, 1)$ in the “no initial tree”-case (proof left as an exercise for the reader), in general it is not true. Figure 9 shows a counterexample ($P = (3, 2, 1)$).

GREEDY is not decomposable.

One could prove a decomposition theorem similar to Theorem E.2 for GREEDY, if GREEDY were *decomposable*. By this we mean that for a k -decomposable input sequence, considering a certain level $P = \tilde{P}[P_1, \dots, P_k]$ of the decomposition tree, GREEDY touches a region $R(i, j)$, iff the corresponding point (i, j) is touched by GREEDY, when executed on the contracted instance \tilde{P} . Alas, this property does not hold for GREEDY, as shown on Figure 10.




Figure 9: Access points are black circles, points touched by GREEDY are gray circles. Input avoids (3,2,1), but output contains (3,2,1).




Figure 10: Access points are black circles, points touched by GREEDY are gray circles. Left: block decomposition of input, and GREEDY execution. Top left region is touched. Right: GREEDY execution on contracted instance. Top left point not touched.

No permutation input-revealing gadget with initial trees.

A permutation input-revealing gadget for GREEDY would allow us to strengthen the bound in Theorem 1.1 from quasilinear to linear. In Figure 11 we sketch a construction that shows that such a gadget can not exist, if initial tree is allowed, even if the input is the preorder sequence of a path.

G.2 Tightness of $O(n \log k)$ for k -decomposable permutations

In this section, we show that the upper bound of $\text{OPT}(X) = O(n \log k)$ for k -decomposable permutations is tight, by arguing that there is a sequence that matches this bound. We will use the following claim, which is proved in §E.4, and restated here.

Lemma G.1. *Let X be any input sequence and \mathcal{T} be its tree decomposition. Then*

$$\text{OPT}(X) \geq \sum_{P \in \mathcal{T}} \text{OPT}(P) - 2n$$

Let P' be a sequence of size k where $\text{OPT}(P') = \Omega(k \log k)$. Then we construct the tree \mathcal{T} which is a complete degree- k tree of depth d , with each node $v \in \mathcal{T}$ having the template $P_v = P'$. Then $n = k^d$, and each non-leaf node $v \in N(\mathcal{T})$ has $\text{OPT}(P_v) = \Omega(k \log k)$. There are $k^{d-1} = n/k$ non-leaf nodes, so we have $\text{OPT}(X) \geq \Omega(k \log k) \cdot \frac{n}{k} = \Omega(n \log k)$.

G.3 Decomposition of k -increasing sequences

Given a permutation $X = (X_1, \dots, X_n)$, a subsequence of X of size k is a sequence $(X_{i_1}, \dots, X_{i_k})$, where $1 \leq i_1 < \dots < i_k \leq n$.

Lemma G.2. *Let $X \in S_n$ be an arbitrary permutation. The following statements are equivalent:*

- (i) X is $(k, \dots, 1)$ -free,

(ii) There exist pairwise disjoint subsequences Y_1, Y_2, \dots, Y_{k-1} of X , such that Y_i is increasing for all $i = 1, \dots, k-1$, and $|Y_1| + \dots + |Y_{k-1}| = n$.

Proof. (ii) \implies (i):

Let Y_1, \dots, Y_{k-1} be increasing subsequences of X , as in the statement of the lemma, and suppose there exists a subsequence $X' = (X_{i_1}, \dots, X_{i_k})$ of X order-isomorphic to the pattern $(k, \dots, 1)$. Since X' is decreasing, no two elements of X' can be in the same subsequence Y_i . This is a contradiction, since we have only $k-1$ subsequences. Therefore, such an X' cannot exist, and X is $(k, \dots, 1)$ -free.

(i) \implies (ii):

Assume that X is $(k, \dots, 1)$ -free. We construct the decomposition of X into increasing subsequences Y_1, \dots, Y_{k-1} as follows. We look at the plot of permutation X , and we refer interchangeably to elements of X and points in the plot of X . Let Y_1 be the “wing”, i.e. the points in the plot of X that form an empty rectangle with the top left corner (by this definition, we have $X_1 \in Y_1$). Clearly, Y_1 forms an increasing subsequence of X . We now remove the elements of Y_1 from X and similarly find Y_2 as the wing of the remaining points. We repeat the process, thereby finding Y_1, Y_2, \dots . We claim that we run out of points before reaching Y_k , thus constructing a decomposition of X as required by the lemma.

Suppose otherwise that we have reached Y_k , and consider an arbitrary point $X_{i_k} \in Y_k$. Since X_{i_k} was not chosen in Y_{k-1} , the rectangle with corners X_{i_k} and the top-left corner has some point from Y_{k-1} . Pick such a point, and denote it as $X_{i_{k-1}}$. Observe that $X_{i_{k-1}} > X_{i_k}$, and $i_{k-1} < i_k$. Since $X_{i_{k-1}}$ was not chosen in Y_{k-2} , we can pick a suitable $X_{i_{k-2}}$ in Y_{k-2} . Continuing in this fashion, we construct the subsequence X_{i_1}, \dots, X_{i_k} of X that forms the forbidden pattern $(k, \dots, 1)$, a contradiction. \blacksquare

G.4 Comparing easy sequences

Easy sequences not captured by any known bounds.

We show a sequence X where $UB(X) = \Omega(n \log n)$, $k(X) = \sqrt{n}$, and $OPT(X) = O(n)$. By UB we denote the “unified bound” [5] that subsumes both the working set, dynamic finger and entropy bounds. At this time, however, UB is not known to be matched by any online BST. UB can be seen as the sum of logs of “ L_1 distances”. In other words, X is an easy sequence that is not captured by the known classes of easy sequences studied in the literature.

Consider a point set \mathcal{P} that is the perturbation of the \sqrt{n} -by- \sqrt{n} grid: Let $\ell = \sqrt{n}$, and

$$\mathcal{P} = \{(i\ell + (j-1), j\ell + i-1) : i, j \in [\ell]\}.$$

Denote the point $(i\ell + (j-1), j\ell + i-1)$ by $p(i, j)$. It is easy to check that there are no two points aligning on x or y coordinates: If $p(i, j)$ agrees with $p(i', j')$ on one coordinate, it implies that $i = i'$ and $j = j'$. Therefore, this point set corresponds to a permutation $X \in S_n$. Figure 12(i) illustrates the construction. The following lemma shows that $k(X) = \sqrt{n}$.

Lemma G.3. X contains all patterns σ of size at most ℓ .

Proof. Pattern σ corresponds to point set $\{(\sigma_t, t)\}_{t=1}^k$, this is order-isomorphic to $p(\sigma_1, 1), p(\sigma_2, 2), \dots, p(\sigma_k, k)$. \blacksquare

Proposition G.4. $UB(X) = \Omega(n \log n)$.

Finally, we show that $OPT(X) = O(n)$ by describing an offline algorithm that achieves this bound. Our algorithm is similar to GREEDY, where we augment some extra steps. First, partition $[n]$ into I_1, \dots, I_ℓ such that $I_j = \{(j-1)\ell + 1, j\ell\}$. Notice that the access sequence starts by touching elements $p(1, 1) \in I_1, p(2, 1) \in I_2, \dots, p(\ell, 1) \in I_\ell$ and then back to $p(2, 1) \in I_1$ and so on. This occurs for ℓ rounds.

The algorithm proceeds just like GREEDY, except that, when accessing $p(i, j)$, after touching the minimal elements, it also touches the two boundaries of the interval I_i , i.e. $(i-1)\ell + 1$ and $i\ell$.

Proposition G.5. The algorithm produces a feasible solution.

Proposition G.6. *The cost of this algorithm is $O(n)$.*

A slightly more involved inductive argument shows that GREEDY too has linear cost on X .

Proposition G.7. $\text{GREEDY}(X) = O(n)$.

Pattern avoidance and dynamic finger.

We show examples that $k(X)$ and the dynamic finger bound (denoted by $DF(X)$) are not comparable. Figure 12(ii) illustrates a sequence X where $DF(X) = O(n)$ and $k(X) = \Omega(\sqrt{n}/\log n)$. First, the sequence X starts with the perturbed grid construction of size $f(n)$ -by- $f(n)$ where $f(n) = \sqrt{n}/\log n$; then the sequence continues with a sequential access of length $n - f(n)^2$. Notice that, in the first phase of the sequence, the distances between consecutive points are $|X_{i+1} - X_i| = O(f(n))$, except for at most $f(n)$ pairs, which have distance $O(f(n)^2)$. In the second phase, the distances are one. So $DF(X) = O(f(n)^2) \log f(n) + O(n) = O(n)$. On the other hand, the $f(n)$ -by- $f(n)$ grid contains all permutations of size $f(n)$, so we have $k(X) = \Omega(\sqrt{n}/\log n)$.

Figure 12(iii) illustrates a 2-recursively decomposable sequence (thus having $k(X) = 3$) whose dynamic finger bound is $\Omega(n \log n)$; to see this, notice that $|X_{i+1} - X_i| \geq n/2$ for $i = 1, \dots, n/2$, so we have $DF(X) \geq \frac{n}{2} \log(n/2)$.

Easy non-decomposable sequence.

Figure 12(iv) shows the non-decomposable input sequence $X = (n/2, n, 1, n-1, 2, \dots)$ on which GREEDY achieves linear cost (to see this, observe that GREEDY touches at most three elements in each row).

G.5 Why is path the roadblock to deque, traversal, split

The deque-, traversal- and split conjectures are originally stated for splay trees. We first state the conjectures for any online BST \mathcal{A} .

Conjecture 2 (Deque conjecture [36]). *Starting with any initial tree T with n elements, the cost of \mathcal{A} for inserting or deleting the current minimum or maximum elements m times is $O(m + n)$*

Conjecture 3 (Traversal conjecture [34]). *Starting with any initial tree T with n elements, the cost of accessing a sequence X defined by the preorder sequence of some BST T' is $O(n)$.*

Conjecture 4 (Split conjecture [26]). *Starting with any initial tree T with n elements, the cost of \mathcal{A} for splitting T by any sequence of n splits is $O(n)$. A split at element x is to delete x and to obtain two trees whose elements are smaller resp. larger than x , each subject to further splittings.*

Next, we state an easier conjecture which is implied by all three conjectures.

Conjecture 5 (Path conjecture). *Starting with any initial tree T with n elements, the cost of \mathcal{A} for accessing a sequence X defined by the preorder sequence of some BST T' which is a path is $O(n)$.*

To express it in terms of pattern avoidance, the access sequences in the path conjecture avoid both $(2, 3, 1)$ and $(2, 1, 3)$, while the ones in traversal conjecture only avoid $(2, 3, 1)$.

Intuitively, the sequence in path conjecture is a sequence that starts from minimum or maximum and moves inwards towards the middle, for example, $(1, 2, 3, 4, 10, 9, 5, 6, 8, 7)$.

Theorem G.8. *Let \mathcal{A} be an online BST.*

- (i) *If \mathcal{A} satisfies traversal conjecture, then \mathcal{A} satisfies path conjecture.*
- (ii) *If \mathcal{A} satisfies deque conjecture, then there is an online BST \mathcal{A}' that satisfies path conjecture.*
- (iii) *If \mathcal{A} satisfies split conjecture, then there is an online BST \mathcal{A}' that satisfies path conjecture.*

Proof. (i) This is clear from the statement.

(ii) If \mathcal{A} satisfies deque conjecture, then \mathcal{A} can delete the current minimum or maximum element n times with cost $O(n)$. Observe that the sequence of these deletions can be defined by a BST which is a path. Given \mathcal{A} , we construct \mathcal{A}' that satisfies path conjecture. To access $x = 1$ or n , \mathcal{A}' simulates \mathcal{A} by deleting x , but then before finishing, \mathcal{A}' brings x back by making x the root.

To access other elements x from the “minimum” (“maximum”) side, \mathcal{A}' also simulates \mathcal{A} , but we can see that it must touch $x - 1$ ($x + 1$) as well. Again, instead of deleting x , \mathcal{A}' makes x as a root. For any time, the cost of \mathcal{A}' is at most the cost of \mathcal{A} plus one.

(iii) Splitting is equivalent to deleting when applied to the minimum or maximum element. Thus, if \mathcal{A} satisfies split conjecture, then \mathcal{A} satisfies deque conjecture without inserts. We can then use the argument of (ii). ■




Figure 11: A GREEDY execution with initial tree. Access points and initial tree are shown with black circles, points touched by GREEDY are gray circles. Black rectangle marks a region in which all permutations are contained, but there is no access point inside. Gray rectangle marks initial tree.




Figure 12: From left to right: (i) “perturbed grid” construction when $\ell = 5$ with GREEDY execution, (ii) example with low dynamic finger bound, but high pattern avoidance parameter, (iii) example with low pattern avoidance parameter and high dynamic finger bound, with GREEDY execution, (iv) non-decomposable permutation with linear cost GREEDY execution.