SIIM-TR-A-**05-17**

# SIIM Technical Report

# Smoothed Analysis of the Height of Binary Search Trees

**by**

Bodo Manthey and Rüdiger Reischuk

**Schriftenreihe der Institute für**

**Informatik/Mathematik**

**Serie A**

June 2, 2005

Universität zu Lübeck

Technisch-Naturwissenschaftliche Fakultät

| | | | |
|---|---|---|---|
| Email: | manthey@tcs.uni-luebeck.de | Phone: | +49-451-500-5313 |
| | | Fax: | +49-451-500-5301 |

# Smoothed Analysis of the Height of Binary Search Trees


Bodo Manthey* Rüdiger Reischuk

Universität zu Lübeck
Institut für Theoretische Informatik
Ratzeburger Allee 160, 23538 Lübeck, Germany
manthey/reischuk@tcs.uni-luebeck.de



**Abstract**

Binary search trees are one of the most fundamental data structures. While the worst case height of a binary search tree is linear, the average height of binary search trees under the uniform distribution is logarithmic and one of the best studied problems in average case complexity.

We investigate what happens between worst and average case, namely the smoothed height of binary search trees: we randomly perturb a given (adversarial) sequence and then consider the height of the binary search tree generated by the resulting sequence. As perturbation models, we consider partial permutations, partial alterations, and partial deletions.

On the one hand, we prove tight lower and upper bounds of roughly $\Theta(\sqrt{n})$ (in contrast to height $n$ in the worst case and $\Theta(\log n)$ on average) for the expected height of binary search trees under partial permutations and partial alterations. That means worst case instances break down under slight perturbations. On the other hand, we examine how much a perturbation can increase the height of a binary search tree, i.e. how much worse best case instances can become.



**Keywords:** Smoothed Analysis, Binary Search Trees, Discrete Perturbations, Permutations.

**ACM Computing Classification:** E.1 [Data Structures]: Trees; F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems—sorting and searching; G.2.2 [Discrete Mathematics] Combinatorics—permutations and combinations.



*Supported by DFG research grant Re 672/3.

# 1 Introduction

To explain the discrepancy between average case and worst case behaviour of the simplex algorithm, Spielman and Teng introduced the notion of *smoothed analysis* [25, 28]. Smoothed analysis interpolates between average case and worst case analysis: Instead of taking the worst case example or, as in average case analysis, choosing an instance completely at random, we analyse the complexity of (worst case) objects subject to slight random perturbations.

Smoothed analysis takes into account that a *typical* instance is not necessarily a *random* instance and that worst case instances are usually artificial and rarely occur in practice.

Let $C$ be some complexity measure. The worst case complexity is $\max_x C(x)$, and the average case complexity is $\mathbb{E}_{x \sim \Delta} C(x)$, where $\mathbb{E}$ denotes expectation with respect to a probability distribution $\Delta$. The smoothed complexity is defined as $\max_x \mathbb{E}_{y \sim \Delta(x,p)} C(y)$. Here, $x$ is chosen by an adversary and $y$ is randomly chosen according to some probability distribution $\Delta(x, p)$ that depends on $x$ and a parameter $p$. Usually, $\Delta(x, p)$ favours $y$ in the vicinity of $x$. The smoothing parameter $p$, which is for instance the variance of $\Delta(x, p)$, denotes how strong $x$ is perturbed. Intuitively, for $p = 0$, smoothed complexity becomes worst case complexity, while for large $p$, smoothed complexity becomes average case complexity.

Smoothed complexity can be interpreted as follows: If the smoothed complexity of an algorithm is low, then we must be unlucky to accidentally hit an instance on which our algorithm behaves poorly, even if the worst case complexity of our algorithm is bad. In this situation, worst case instances are isolated events.

Binary search trees are one of the most fundamental data structures and thus a building block for many advanced data structures. The main criteria of the "quality" of a binary search tree is its height, i.e. the length of the longest path from the root to a leaf. Unfortunately, the height equals the number of elements in the worst case, i.e. for totally unbalanced trees. On the other hand, if a binary search tree is chosen at random, then the expected height is only logarithmic in the number of elements (more details will be discussed in Section 1.1.2). Thus, there is a huge discrepancy between the worst case and the average case behaviour of binary search trees.

We analyse what happens in between, namely the *smoothed height of binary search trees*, which is the height subject to slight perturbations: given a sequence by an adversary, it will randomly be perturbed and then the height of the binary search tree from the sequence thus obtained is measured. Thus, our instances are neither adversarial nor completely random. As perturbation models, we consider *partial permutations*, *partial alterations*, and *partial deletions*. For all three, we show tight lower and upper bounds. As a byproduct, we also obtain tight bounds for the smoothed number of left-to-right maxima, which is the number of new maxima seen when scanning a sequence from the left to the right. Thus, the number of left-to-right maxima of a sequence is simply the length of the right-most path in the binary search tree grown from that sequence.

In smoothed analysis one analyses how fragile worst case instances are. We suggest to examine also another property: given a good (or best case) instance, how much can the complexity increase by slightly perturbing the instance? In other words, how stable are best case instances under perturbations? For binary search trees, we prove that there are best case instances that indeed are not stable, i.e. there are sequences yielding trees of logarithmic depth, but slightly perturbing the sequences yields trees of polynomial depth.

## 1.1 Previous Results

Since we are concerned with smoothed analysis and binary search trees, we briefly review both areas.

### 1.1.1 Smoothed Analysis

Spielman and Teng introduced smoothed analysis as a hybrid of average case and worst case complexity [25, 28]. They showed that the simplex algorithm for linear programming with the shadow vertex pivot rule has polynomial smoothed complexity. That means, the running time of the algorithm is expected to be polynomial in terms of the input size and the variance of the Gaussian perturbation. Since then, smoothed analysis has been applied to a variety of fields, e.g., several variants of linear programming [7, 27], properties of moving objects [9], online and other algorithms [5, 24], property testing [26], discrete optimisation [6, 23], graph theory [15], and computational geometry [10].

Banderier, Beier, and Mehlhorn [4] applied the concept of smoothed analysis to combinatorial problems by analysing Quicksort, left-to-right maxima, and single-source shortest path. In particular, they analysed the number of left-to-right maxima of a sequence, which is the number of maxima seen when scanning a sequence from left to right. Here the worst case is the sequence $1, 2, \dots, n$, which yields $n$ left-to-right maxima. On average we expect $H_n = \sum_{i=1}^{n} i^{-1} \approx \ln n$ left-to-right maxima. The perturbation model used by Banderier et al. are *partial permutations*, where each element of the sequence is independently selected with a given probability $p \in [0, 1]$ and then a random permutation on the selected elements is performed (see Section 3.1 for a precise definition).

Banderier et al. proved that the number of left-to-right maxima under partial permutations is expected $O(\sqrt{(n/p)\log n})$. On the other hand, they have shown a lower bound of $\Omega(\sqrt{n/p})$ that holds for $p \le \frac{1}{2}$.

### 1.1.2 Binary Search Trees

Given a sequence $\sigma = (\sigma_1, \sigma_2, \ldots, \sigma_n)$ of $n$ distinct elements from any ordered set, we obtain a binary search tree $T(\sigma)$ by iteratively inserting the elements $\sigma_1, \sigma_2, \ldots, \sigma_n$ into the initially empty tree (this is formally described in Section 2.3).

The study of binary search trees is one of the most fundamental problems in computer science since they are the building block for a large variety of data structures (see e.g. Aho et al. [1, 2] and Knuth [16]). For instance, the height of $T(\sigma)$ is just the number of levels of recursion required by Quicksort if the first element of the sequence to be sorted is chosen as the pivot (see e.g. Cormen et al. [8]).

The worst case height of a binary search tree obtained in this way is obviously $n$: just take $\sigma = (1, 2, \ldots, n)$. (In this paper, the length of a path is the number of vertices and not the number of edges it contains; however, all known and new results mentioned in this paper hold for both variants.) The expected height of the binary search tree obtained from a random permutation (with all permutations being equally likely) has been the subject of a considerable amount of research in the past. We briefly review some results. Let the random variable $H(n)$ denote the height of a binary search tree obtained from a random permutation. Robson [19] proved that $\mathbb{E}H(n) \approx c\ln(n) + o(\ln(n))$ for some $c \in [3.63, 4.3112]$ and observed that $H(n)$ does not vary much from experiment to experiment [20]. Pittel [17] proved the existence of a $\gamma > 0$ with $\gamma = \lim_{n\to\infty} \frac{\mathbb{E}H(n)}{\ln(n)}$. Devroye [11] then proved that $\lim_{n\to\infty} \frac{\mathbb{E}H(n)}{\ln(n)} = \alpha$ with $\alpha \approx 4.31107$ being the larger root of $\alpha\ln(2e/\alpha) = 1$. The variance of $H(n)$ was shown to be $O((\operatorname{llog} n)^2)$ by Devroye and Reed [12] and by Drmota [13]. Robson [21] proved that the expectation of the absolute value of the difference between the height of two random trees is constant. Thus, the height of the random trees is concentrated around the mean. A climax was the result discovered independently by Reed [18] and Drmota [14] that the variance of $H(n)$ actually is $O(1)$. Furthermore, Reed [18] proved that the expectation of $H(n)$ is $\alpha\ln n + \beta\ln(\ln n) + O(1)$ with $\beta = \frac{3}{2\ln(\alpha/2)} \approx 1.953$. Finally, Robson [22] proved strong upper bounds on the probability of large deviations from the median. His results suggest that all moments of $H(n)$ are bounded from above by a constant.

Although worst case and average case height of binary search trees are very well understood, nothing is known in between, i.e. when the sequences are not completely random, but the randomness is limited.

## 1.2 New Results

We consider the height of binary search trees subject to slight perturbations (*smoothed height*), i.e. the expected height under limited randomness. In addition to *partial permutations*, which were introduced by Banderier et al. [4], we consider *partial alterations* and *partial deletions* (both defined in Section 3) as perturbation models. Thus, we generalise the results of Banderier et al. [4]. For all three models, we prove matching lower and upper bounds for the expected height of binary search trees obtained from sequences that have been perturbed by one of the perturbation models. More precisely: for all $p \in (0,1)$ and all sequences of length $n$, the height of a binary search tree obtained via $p$-partial permutation is expected to be at most $7 \cdot (1-\alpha^{-1}p) \cdot \sqrt{n/p}$ for arbitrarily small $\alpha > 1$ and sufficiently large $n$, which is roughly $O((1-p) \cdot \sqrt{n/p})$. Previously, there was only a bound of $O(\sqrt{(n/p) \cdot \log n})$ known for the number of left-to-right maxima under partial permutations, but nothing was known about the smoothed height of binary search trees.

On the other hand, the height of a binary search tree obtained from the sorted sequence via $p$-partial permutation is at least $0.8 \cdot (1-p) \cdot \sqrt{n/p}$ in expectation. This matches the upper bound up to a constant factor and $\alpha$.

For the number of left-to-right maxima under partial permutations and partial alterations, we are able to prove an even better upper bound of $3.6 \cdot (1-p) \cdot \sqrt{n/p}$ for all sufficiently large $n$. As a lower bound, we prove that there exist sequences of $n$ elements such that after $p$-partial permutation at least $0.4 \cdot (1-p) \cdot \sqrt{n/p}$ left-to-right maxima are expected.

All these bounds hold for $p$-partial alterations as well.

Thus, under limited randomness, the behaviour of binary search trees differs completely from both the worst case and the average case.

For partial deletions, we obtain $(1-p) \cdot n$ both as lower and upper bound. This result is straight-forward. The main reason for considering partial deletions is that we can bound the expected height subject to partial alterations and permutation by the expected height subject to partial deletions. The converse holds as well, we only have to blow up the sequences quadratically. We exploit this when considering the stability of the permutation models: we prove that partial deletions and thus partial permutations and partial alterations as well are quite unstable, i.e. can cause best case instances to become much worse. More precisely: there are sequences of length $n$ that yield trees of depth $O(\log n)$, but the expected height of the tree obtained after smoothing is $\Omega(n^\epsilon)$ for some $\epsilon > 0$ that depends only on $p$.

## 1.3 Outline

In the next section, we introduce some basic notation. We define the perturbation models partial permutations, partial alterations, and partial deletions in Section 3. Then we show some basic properties of binary search trees (Section 4.1), partial permutations (Section 4.2), and partial alterations (Section 4.3). In Section 5 we show matching lower and upper bounds for the expected number of left-to-right maxima under perturbation. After that, we consider the smoothed height of binary search trees under partial permutations and partial alterations (Section 6). We prove matching lower and upper bounds for the expected height of binary search trees that hold for both perturbation models. Then we compare partial deletions with the two other models (Section 7). These results are exploited in Section 8, where we consider the stability of the perturbation models. Finally, we give some concluding remarks (Section 9).

# 2 Preliminaries

## 2.1 Notations

We denote by log and ln the logarithm to base 2 and $e$, respectively, while exp denotes the exponential function to base $e$. We abbreviate the twice iterated logarithm $\log \circ \log$ by llog. For any $x \in \mathbb{R}$, let $[x] = \{x - i \mid i \in \mathbb{N}, x - i > 0\}$. For instance, $[n] = \{1, 2, \ldots, n\}$ and $[n - \frac{1}{2}] = \{\frac{1}{2}, \frac{3}{2}, \ldots, n - \frac{1}{2}\}$ for $n \in \mathbb{N}$.

Let $\sigma = (\sigma_1, \ldots, \sigma_n) \in S^n$ for some ordered set $S$. We call $\sigma$ a **sequence**. Usually, we assume that all elements of $\sigma$ are distinct, i.e. $\sigma_i \neq \sigma_j$ for all $i \neq j$. The length of $\sigma$ is $n$. In most cases, $\sigma$ will simply be a permutation of $[n]$. We denote the sorted sequence $(1, 2, \ldots, n)$ by $\sigma^n_{\text{sort}}$. When considering partial alterations, we have $\sigma^n_{\text{sort}} = (0.5, 1.5, \ldots, n - 0.5)$ instead (this will be clear from the context).

Let $\tau = (\tau_1, \ldots, \tau_t)$. We call $\tau$ a **subsequence of $\boldsymbol{\sigma}$** if there are number $i_1 < i_2 < \ldots < i_t$ with $\tau_j = \sigma_{i_j}$ for all $j \in [t]$. Let $\mu = \{i_1, \ldots, i_t\} \subseteq [n]$. Then $\sigma_\mu = (\sigma_{i_1}, \ldots, \sigma_{i_t})$ denotes the subsequence consisting of all elements of $\sigma$ at positions in $\mu$. For instance, $\sigma_{[k]}$ denotes the prefix of length $k$ of $\sigma$. By abusing notation, we sometimes consider $\sigma_\mu$ as the set of elements at positions in $\mu$, i.e. in this case $\sigma_\mu = \{\sigma_i \mid i \in \mu\}$. However, whether we consider $\sigma_\mu$ as a sequence or as a set will always be clear from the context. For $\mu \subseteq [n]$, we define $\overline{\mu} = [n] \setminus \mu$.

## 2.2 Probability Theory

We denote probabilities by $\mathbb{P}$ and expectations by $\mathbb{E}$. To bound large deviations, we will frequently use Chernoff bounds [3, Corollary A.7]. Let $p \in (0, 1)$ and let $X_1, X_2, \ldots, X_n$ be mutually independent random variables with $\mathbb{P}(X_i = 1) = 1 - \mathbb{P}(X_i = 0) = p$ and $X = \sum_{i=1}^n X_i$. Clearly, $\mathbb{E}(X) = pn$. The probability that $X$ deviates by more than $a$ from its expectation is bounded from above by

$$\mathbb{P}(|X - p \cdot n| > a) < 2 \cdot \exp\left(-\frac{2a^2}{n}\right) . \tag{2.1}$$

We will frequently use the following lemma.

**Lemma 2.1.** *Let $k \in \mathbb{N}$, $\alpha > 1$ and $p \in [0, 1]$. Assume that we have mutually independent random variables $X_1, \ldots, X_k$ as above. Then*

$$\mathbb{P}\left((X > \alpha pk) \vee (X < \alpha^{-1} pk)\right) \leq 2 \cdot \exp\left(-2(1 - \alpha^{-1})^2 p^2 k\right) .$$

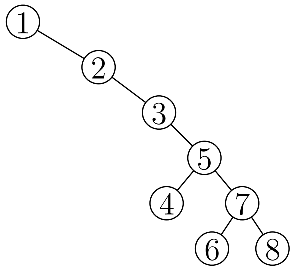


Figure 1: The binary search tree $T(\sigma)$ obtained from $\sigma = (1, 2, 3, 5, 7, 4, 6, 8)$. We have $\mathrm{height}(\sigma) = 6$.

*Proof.* Since $\alpha - 1 \geq 1 - \alpha^{-1}$ for all $\alpha > 1$, let $a = (1 - \alpha^{-1})pk$. Then we apply Formula 2.1 and get

$$\begin{aligned}\mathbb{P}\left((X > \alpha pk) \vee (X < \alpha^{-1}pk)\right) &\leq \mathbb{P}\left(|X - pk| > (1 - \alpha^{-1})pk\right) \\ &< 2 \cdot \exp\left(-\frac{2(1 - \alpha^{-1})p^2k^2}{k}\right) \\ &= 2 \cdot \exp\left(-2(1 - \alpha^{-1})^2 p^2 k\right) .\end{aligned}$$

□

## 2.3 Binary Search Trees and Left-to-right Maxima

Let $\sigma = (\sigma_1, \ldots, \sigma_n)$ be a sequence. We obtain a **binary search tree $T(\sigma)$** from a sequence $\sigma$ by iteratively inserting the elements $\sigma_1, \sigma_2, \ldots, \sigma_n$ into the initially empty tree as follows:

- The root of $T(\sigma)$ is the first element $\sigma_1$ of $\sigma$.
- Let $\sigma_< = \sigma_{\{i \mid \sigma_i < \sigma_1\}}$ be $\sigma$ restricted to elements smaller than $\sigma_1$. Then the left subtree of the root $\sigma_1$ of $T(\sigma)$ is obtained inductively from $\sigma_<$.

  Analogously, let $\sigma_> = \sigma_{\{i \mid \sigma_i > \sigma_1\}}$ be $\sigma$ restricted to elements greater than $\sigma_1$. Then the right subtree of the root $\sigma_1$ of $T(\sigma)$ is the tree obtained inductively from $\sigma_>$.

Figure 1 shows an example. We denote the height of $T(\sigma)$ by **height($\sigma$)**, i.e., $\mathrm{height}(\sigma)$ is the number of nodes on the longest path from the root to a leaf. (We consider a single node as a tree of height one.)

The element $\sigma_i$ is called a **left-to-right maximum** of $\sigma$ if $\sigma_i > \sigma_j$ for all $j \in [i-1]$. Let **ltrm($\sigma$)** denote the number of left-to-right maxima of $\sigma$. We have $\mathrm{ltrm}(\sigma) \leq \mathrm{height}(\sigma)$ since the number of left-to-right maxima of a sequence is just the length of the right-most path in the tree $T(\sigma)$.

# 3 Perturbation Models for Permutations

All three perturbation models considered start with randomly marking elements: given $p \in [0, 1]$, every element of a sequence $\sigma = (\sigma_1, \ldots, \sigma_n)$ is marked independently of the others with probability $p$. To be more formally: the random variable $M_p^n$ is a random subset of $[n]$ with $\mathbb{P}(i \in M_p^n) = p$ for all $i \in [n]$. For any $\mu \subseteq [n]$ we have $\mathbb{P}(M_p^n = \mu) = p^{|\mu|} \cdot (1-p)^{|\overline{\mu}|}$.

Let $\mu \subseteq [n]$ be the set of positions marked. If $i \in \mu$, then we say that position $i$ and element $\sigma_i$ are marked. Thus, $\sigma_\mu$ is the set (or sequence) of all marked elements.

We denote by $\textbf{height-perm}_{\boldsymbol{p}}\boldsymbol{(\sigma)}$, $\textbf{height-alter}_{\boldsymbol{p}}\boldsymbol{(\sigma)}$, and $\textbf{height-del}_{\boldsymbol{p}}\boldsymbol{(\sigma)}$ the expected height of the binary search tree $T(\sigma')$ originated from the sequence $\sigma'$ obtained by performing a $p$-partial permutation, alteration, and deletion, respectively, on $\sigma$ (all three perturbation models will be defined in the subsequent sections). Analogously, we denote by $\textbf{ltrm-perm}_{\boldsymbol{p}}\boldsymbol{(\sigma)}$, $\textbf{ltrm-alter}_{\boldsymbol{p}}\boldsymbol{(\sigma)}$, and $\textbf{ltrm-del}_{\boldsymbol{p}}\boldsymbol{(\sigma)}$ the expected number of left-to-right maxima of the sequence $\sigma'$ obtained from $\sigma$ via $p$-partial permutation, alteration, and deletion, respectively.

## 3.1 Partial Permutations

The notion of ***p*-partial permutations** has been introduced by Banderier et al. [4]. A $p$-partial permutation of $\sigma$ is obtained as follows. Every element of $\sigma$ is marked independently with probability $p$. The elements thus marked are permuted according to a permutation drawn uniformly at random.

Let $\sigma = (\sigma_1, \ldots, \sigma_n)$ and $\mu \subseteq [n]$. Then the sequence $\sigma' = \Pi(\sigma, \mu)$ is a random variable with the following properties:

- $\Pi$ chooses a permutation $\pi$ of $\mu$ uniformly at random and
- sets $\sigma'_{\pi(i)} = \sigma_i$ for all $i \in \mu$ and $\sigma'_i = \sigma_i$ for all $i \notin \mu$.

Thus, a $p$-partial permutation $\Pi(\sigma, M_p^n)$ of $\sigma$ consists of two steps: randomly mark elements of $\sigma$ as described above, i.e. randomly create a set $\mu = M_p^n \subseteq [n]$ of marked elements, and then randomly permute all the marked elements, i.e. the elements at positions in $\mu$. Note that $i \in \mu$ does not necessarily mean that $\sigma_i$ is at a position different from $i$ in $\Pi(\sigma, \mu)$; the random permutation can of course map $\pi(i) = i$.

**Example 3.1.** *Figure 2 shows an example.*

By choosing $p$, we can interpolate between average and worst case: for $p = 0$, no element is marked and $\sigma' = \sigma$, while for $p = 1$, all elements are marked and thus $\sigma'$ is a random permutation of the elements of $\sigma$ with all permutations being equally likely.

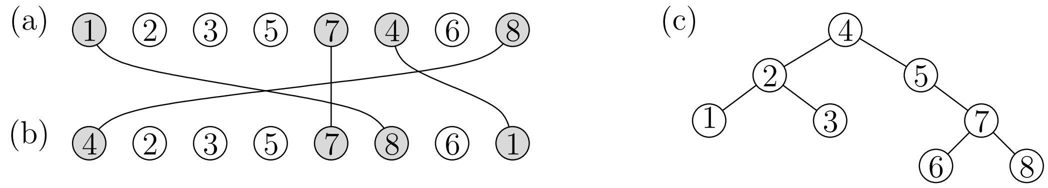


Figure 2: A partial permutation. (a) The sequence $\sigma = (1, 2, 3, 5, 7, 4, 6, 8)$ (Figure 1 shows $T(\sigma)$). The first, fifth, sixth, and eighth element is (randomly) marked, thus $\mu = M_p^n = \{1, 5, 6, 8\}$. (b) The marked elements are randomly permuted. The result is the sequence $\sigma' = \Pi(\sigma, \mu)$, in this case $\sigma' = (4, 2, 3, 5, 7, 8, 6, 1)$. (c) $T(\sigma')$ with $\mathrm{height}(\sigma') = 4$.

## 3.2 Partial Alterations

Let us now introduce ***p*-partial alterations**. For this perturbation model, we restrict the sequences of length $n$ to be a permutation of $[n - \frac{1}{2}]$.

We again start with a random marking $M_p^n$. Then we replace every marked element by a real number drawn uniformly and independently at random from $[0, n)$ to obtain a sequence $\sigma'$. With probability one, all elements in $\sigma'$ are distinct.

Instead of considering permutations of $[n - \frac{1}{2}]$, we can also consider permutations of $[n]$ and draw the random values from $[\frac{1}{2}, n + \frac{1}{2})$. This does not change the results. Another possibility is to consider permutations of $[n]$ and draw the random values from $[0, n + 1)$. This does not change the results by much. However, for technical reasons we consider partial alterations as introduced above.

**Example 3.2.** *Let* $\sigma = (0.5, 1.5, 2.5, 4.5, 6.5, 3.5, 5.5, 7.5)$ *(which is the sequence of Example 3.1 with* 0.5 *subtracted from each element) and* $\mu = \{1, 5, 6, 8\}$*. Then we obtain the sequence* $(X_1, 1.5, 2.5, 4.5, X_2, X_3, 5.5, X_4)$*, where* $X_1, X_2, X_3, X_4$ *are drawn independently and uniformly from* $[0, 8)$*.*

Like partial permutations, partial alterations interpolate between worst case ($p = 0$) and average case ($p = 1$). Partial alterations are somewhat easier to analyse: the majority of the results on the average case height of binary search trees (see for instance Pittel [17] and Devroye [11]) is obtained not via random permutations but the binary search trees are grown from a sequence of $n$ random variables that are uniformly and independently drawn from $[0, 1)$. There is no difference between partial permutations and partial alterations for $p = 1$. It turns out that this seems to hold for all $p$; the bounds obtained for partial permutations and partial alterations are equal for all $p$.

## 3.3 Partial Deletions

As third perturbation model, let us introduce ***p*-partial deletions**: again, we have a random marking $M_p^n$ as in Section 3.1. Then we delete all marked elements and obtain the sequence $\sigma_{\overline{M_p^n}}$.

**Example 3.3.** *The sequence $\sigma$ and the marking $\mu$ as in Example 3.1 yield the sequence* $(2, 3, 5, 6)$.

Partial deletions are not really perturbing a sequence: any ordered sequence remains ordered even if elements are deleted. The main reason for considering partial deletions is that they are easy to analyse when considering the stability of perturbation models (Section 8). The results obtained for partial permutations then carry over to partial permutations and partial alterations since the expected height with respect to these three models is closely related (Section 7).

# 4 Basic Properties

In this section, we state some basic properties of binary search trees (Section 4.1), partial permutations (Section 4.2), and partial alterations (Section 4.3) that we will exploit in subsequent sections.

## 4.1 Properties of Binary Search Trees

We start by introducing a new measure for the height of binary search trees. Let $\mu \subseteq [n]$ and $\sigma$ be a sequence of length $n$. The **$\boldsymbol{\mu}$-restricted height of $\boldsymbol{T(\sigma)}$**, denoted by **$\mathbf{height}(\boldsymbol{\sigma}, \boldsymbol{\mu})$**, is the maximum number of elements of $\sigma_\mu$ on a root-to-leaf path in $T(\sigma)$.

**Lemma 4.1.** *For all sequences $\sigma$ of length $n$ and $\mu \subseteq [n]$, we have*

$$\begin{array}{rcl} \mathrm{height}(\sigma) & \leq & \mathrm{height}(\sigma, \mu) + \mathrm{height}(\sigma, \overline{\mu}) \quad \textit{and} \\ \mathrm{height}(\sigma, \mu) & \leq & \mathrm{height}(\sigma_\mu)\,. \end{array}$$

*Proof.* Consider any path of maximum length from the root to a leaf in $T(\sigma)$. This path consists of at most $\mathrm{height}(\sigma, \mu)$ elements of $\sigma_\mu$ and at most $\mathrm{height}(\sigma, \overline{\mu})$ elements of $\sigma_{\overline{\mu}}$, which proves the first part.

For the second part, let $a$ and $b$ be elements of $\sigma_\mu$ that do not lie on the same path from the root to a leaf in $T(\sigma_\mu)$. Assume that $a < b$. Then there exists a $c$ prior to $a$ and $b$ in the sequence $\sigma_\mu$ with $a < c < b$. Thus, $a$ and $b$ do not lie on the same root-to-leaf path in the tree $T(\sigma)$ as well. Consider now any root-to-leaf path of $T(\sigma)$ with $\mathrm{height}(\sigma, \mu)$ elements of $\sigma_\mu$. Then all these elements from $\sigma_\mu$ lie on the same root-to-leaf path in $T(\sigma_\mu)$, which proves the second part of the lemma. $\square$

Of course $\mathrm{height}(\sigma, \mu) \leq \mathrm{height}(\sigma)$ for all $\sigma$ and $\mu$. But $\mathrm{height}(\sigma_\mu) \leq \mathrm{height}(\sigma)$, which would imply $\text{height-del}_p(\sigma) \leq \mathrm{height}(\sigma)$, does not hold in general: Consider $\sigma = (c, a, b, d, e)$ (we use letters and their alphabetical ordering instead of numbers for readability) and $\mu = \{2, 3, 4, 5\}$, then $\sigma_\mu = (a, b, d, e)$. Thus, $\mathrm{height}(\sigma) = 3$ and height$(\sigma_\mu) = 4$. This will further be investigated in Section 8, when we consider the stability of the perturbation models.

For bounding the smoothed height from above, we will use the following lemma, which is an immediate consequence of Lemma 4.1.

**Lemma 4.2.** *For all sequences $\sigma$ of length $n$ and $\mu \subseteq [n]$, we have*

$$\operatorname{height}(\sigma) \leq \operatorname{height}(\sigma_\mu) + \operatorname{height}(\sigma, \overline{\mu}) .$$

*Proof.* According to Lemma 4.1, we have $\operatorname{height}(\sigma) \leq \operatorname{height}(\sigma, \mu) + \operatorname{height}(\sigma, \overline{\mu}) \leq \operatorname{height}(\sigma_\mu) + \operatorname{height}(\sigma, \overline{\mu})$. □

For left-to-right maxima, we can state equivalent lemmas. Let $\sigma$ be a sequence of length $n$ and $\mu \subseteq [n]$. Then **ltrm$(\boldsymbol{\sigma}, \boldsymbol{\mu})$** denotes the **$\boldsymbol{\mu}$-restricted number of left-to-right maxima** of $\sigma$, i.e. the number of elements $\sigma_i$ for $i \in \mu$ such that $\sigma_i > \sigma_j$ for all $j \in [i-1]$. We omit the proof of the following lemma since it is almost equal to the proofs of the lemmas above.

**Lemma 4.3.** *Let $\sigma$ be a sequence of length $n$ and $\mu \subseteq [n]$. Then*

$$\begin{array}{lcl} \operatorname{ltrm}(\sigma) & \leq & \operatorname{ltrm}(\sigma, \mu) + \operatorname{ltrm}(\sigma, \overline{\mu}) , \\ \operatorname{ltrm}(\sigma, \mu) & \leq & \operatorname{ltrm}(\sigma_\mu) \text{ , and} \\ \operatorname{ltrm}(\sigma) & \leq & \operatorname{ltrm}(\sigma_\mu) + \operatorname{ltrm}(\sigma, \overline{\mu}) . \end{array}$$

## 4.2 Properties of Partial Permutations

Let us now prove some properties of partial permutations. The two lemmas proved in this section are crucial for estimating the smoothed height under partial permutations. In the next section, we prove counterparts of these lemmas for partial alterations that will play a similar role in estimating the height under partial alterations.

We start by proving that the expected height under partial permutations merely depends on the elements that are left unmarked. The marked elements contribute at most height $O(\log n)$. Thus, when estimating the expected height in the subsequent sections, we can restrict ourselves to considering the elements that are left unmarked.

**Lemma 4.4.** *Let $\sigma$ be a sequence of length $n$ and $p \in (0,1)$. Let $\mu \subseteq [n]$ be the random set of marked positions and $\sigma' = \Pi(\sigma, \mu)$ be the random sequence obtained from $\sigma$ via $p$-partial permutation. Then*

$$\text{height-perm}_p(\sigma) = \mathbb{E}(\operatorname{height}(\sigma')) \leq \mathbb{E}\left(\operatorname{height}(\sigma', \overline{\mu})\right) + O(\log n) .$$

*Proof.* We have $\operatorname{height}(\sigma_\mu) \in O(\log n)$ since the elements at positions in $\mu$ are randomly permuted. Then the lemma follows from Lemma 4.2. □

And again we obtain an equivalent lemma for left-to-right maxima.

**Lemma 4.5.** *Under the assumptions of Lemma 4.4, we have*

$$\text{ltrm-perm}_p(\sigma) \leq \mathbb{E}\left(\text{ltrm}(\sigma', \overline{\mu})\right) + O(\log n) .$$

The following lemma bounds the probability from above that no element of a fixed set of elements is permuted to a position of a fixed set of positions.

**Lemma 4.6.** *Let $p \in (0,1)$, $\alpha > 1$, $n \in \mathbb{N}$ be sufficiently large, and $\sigma$ be a sequence of length $n$ with elements from $[n]$. Let $\sigma' = \Pi(\sigma, M_p^n)$.*

*Let $\ell = a\sqrt{n/p}$ and $k = b\sqrt{n/p}$ with $a, b \in \Omega((\text{polylog}\, n)^{-1}) \cap O(\text{polylog}\, n)$. Let $A = \sigma'_{[\ell]}$ be the set of the first $\ell$ elements of $\sigma'$. Let $B \subseteq [n]$ be any subset with $|B| = k$.*

*Then $\mathbb{P}(A \cap B = \emptyset) \leq \exp(-ab/\alpha)$.*

*Proof.* We choose $\beta$ with $1 < \beta^3 < \alpha$ arbitrarily. According to Lemma 2.1, the probability $P$ that

- $|M_p^n \cap [\ell]| < \beta^{-1}p\ell$, i.e. that too few of the first $\ell$ positions are marked,
- $|\sigma_{M_p^n} \cap B| < \beta^{-1}pk$, i.e. that too few of the elements of $B$ are marked, or
- $|M_p^n| > \beta pn$, i.e. that too many positions are marked overall

is $O(\exp(-n^\epsilon))$ for fixed $p \in (0,1)$, $\beta > 1$, and appropriately chosen $\epsilon > 0$. This holds since $a, b \in \Omega((\text{polylog}\, n)^{-1})$.

From now on, assume that at least $\beta^{-1}p\ell$ of the first $\ell$ positions of $\sigma$ are marked, at least $\beta^{-1}pk$ elements in $B$ are marked, and at most $\beta pn$ positions are marked overall. The probability that then no element from $B$ is in $A$ is at most

$$\begin{aligned}
\left(\frac{\beta pn - \beta^{-1}p\ell}{\beta pn}\right)^{\beta^{-1}pk} &= \left(1 - \frac{\ell}{\beta^2 n}\right)^{\beta^{-1}pk} \\
= \left(\left(1 - \frac{\ell}{\beta^2 n}\right)^{\frac{\beta^2 n}{\ell}}\right)^{\frac{\ell}{\beta^2 n}\cdot\beta^{-1}pk} &\leq \exp\left(-\frac{\ell}{\beta^2 n} \cdot \beta^{-1}pk\right) = \exp\left(-\frac{ab}{\beta^3}\right) .
\end{aligned}$$

Overall, $\mathbb{P}(A \cap B = \emptyset) \leq \exp(-ab/\beta^3) + P \leq \exp(-ab/\alpha)$ for sufficiently large $n$ since $a, b \in O(\text{polylog}\, n)$. $\square$

## 4.3 Properties of Partial Alterations

Partial alterations fulfil roughly the same properties as partial permutations. We state the lemmas and restrict ourselves to pointing out the differences in the proofs.

**Lemma 4.7.** *Let $\sigma$ be a sequence of length $n$ with elements from $[n-\frac{1}{2}]$ and $p \in (0,1)$. Let $\sigma'$ be the random sequence obtained from $\sigma$ via $p$-partial alteration and $\mu$ be the random set of marked positions. Then*

$$\begin{aligned}
\text{height-alter}_p(\sigma) &\leq \mathbb{E}(\text{height}(\sigma', \overline{\mu})) + O(\log n) \quad \textit{and} \\
\text{ltrm-alter}_p(\sigma) &\leq \mathbb{E}(\text{ltrm}(\sigma', \overline{\mu})) + O(\log n) .
\end{aligned}$$

The following lemma is the counterpart of Lemma 4.6 above.

**Lemma 4.8.** *Let $p \in (0,1)$, $\alpha > 1$, and $n \in \mathbb{N}$ be sufficiently large, and $\sigma$ be a sequence with elements from $[n - \frac{1}{2}]$. Let $\sigma'$ be the random sequence obtained from $\sigma$ by performing a p-partial alteration.*

*Let $\ell = a\sqrt{n/p}$ and $k = b\sqrt{n/p}$ with $a, b \in \Omega((\operatorname{polylog} n)^{-1}) \cap O(\operatorname{polylog} n)$. Let $A = \sigma'_{[\ell]}$ and $B = [x, x+k) \subseteq [0, n)$.*

*Then $\mathbb{P}(A \cap B = \emptyset) \leq \exp(-ab/\alpha)$.*

*Proof.* The proof is similar to the proof of Lemma 4.6. Choose $\beta$ arbitrarily with $1 < \beta < \alpha$. Assume that at least $\beta^{-1}p\ell$ of the first $\ell$ positions of $\sigma$ are marked. Then the probability that no element in $A$ assumes a value of $B$ is at most

$$\left(\frac{n-k}{n}\right)^{\beta^{-1}p\ell} = \left(\left(1 - \frac{k}{n}\right)^{\frac{n}{k}}\right)^{ab/\beta} \leq \exp(-ab/\beta) .$$

The remainder of the proof is as in the proof of Lemma 4.6. □

# 5 Tight Bounds for the Smoothed Number of Left-To-Right Maxima

## 5.1 Partial Permutations

**Theorem 5.1.** *Let $p \in (0,1)$. Then for all sufficiently large $n$ and for all sequences $\sigma$ of length $n$, we have*

$$\text{ltrm-perm}_p(\sigma) \leq 3.6 \cdot (1-p) \cdot \sqrt{n/p} .$$

*Proof.* According to Lemma 4.5, it suffices to show

$$\mathbb{E}(\text{ltrm}(\sigma', \overline{\mu})) \leq C \cdot (1-p) \cdot \sqrt{n/p}$$

for some $C < 3.6$, where $\mu \subseteq [n]$ is the random set of marked positions and $\sigma'$ is the sequence obtained via randomly permuting the elements of $\sigma_\mu$. Then $\text{ltrm-perm}_p(\sigma) \leq C(1-p)\sqrt{n/p} + O(\log n) \leq 3.6(1-p)\sqrt{n/p}$. We assume without loss of generality that $\sigma$ is a permutation of $[n]$.

Let $K_c = c\sqrt{n/p}$ for $c \in [\log n]$. In this and the following proofs, we assume that $K_c$ is a natural number for the sake of readability. If $K_c$ is not a natural number, then we can replace $K_c$ by $\lceil K_c \rceil$. The proofs remain valid.

Choose $\beta$ and $\gamma$ with $1 < \gamma < \beta < \alpha$ and $\gamma < 1.001$. Let $P$ denote the probability that less than $\gamma^{-1}pK_c$ of the first $K_c$ positions are marked or that less than $\gamma^{-1}pK_c$ of the $K_c$ largest elements are marked for some $c \in [\log n]$ or that overall more than $\gamma pn$ elements are marked. $P$ tends exponentially fast to zero as $n$ increases by Lemma 2.1.

From now on, we assume that for all $c \in [\log n]$, at least $\gamma^{-1} p K_c$ of the $K_c$ first positions and of the $K_c$ largest elements are marked. In this case, we say that the partial permutation is **partially successful**. If a partial permutation is not partially successful, we bound the expected number of left-to-right maxima by $n$.

We call $\sigma'$ ***c*-successful** for $c \in [\log n]$ if one of the largest $K_c$ elements $n, n-1, \ldots, n-K_c+1$ of $\sigma$ is among the first $K_c$ elements in $\sigma'$.

Assume that $\sigma'$ is $c$-successful and that $x \in \{n-K_c+1, \ldots, n\}$ is among the first $K_c$ elements of $\sigma'$. Then only the unmarked among the first $K_c$ positions and the unmarked among elements larger than $x$ can contribute to $\mathrm{ltrm}(\sigma', \overline{\mu})$. All other unmarked elements are smaller than $x$ and located after $x$ in $\sigma'$. Thus, they are no left-to-right maxima. The expected number of unmarked elements larger than $n - K_c$ plus the expected number of unmarked positions among the first $K_c$ positions is at most $2 \cdot (1-p) \cdot K_c = Q_c$. Thus, we have $\mathbb{E}(\mathrm{ltrm}(\sigma', \overline{\mu})) \leq Q_c$ if $\sigma'$ is $c$-successful.

The probability that a partially successful partial permutation is not $c$-successful for $c \in O(\log n)$ is bounded from above by $\exp(-c^2/\gamma)$ according to Lemma 4.6. Particularly, the probability that $\sigma'$ is not $\log n$-successful is at most $P' = \exp(-(\log n)^2/\gamma)$. If $\sigma'$ is not $\log n$-successful, we bound the number of left-to-right maxima by $n$.

Thus, restricted to partially successful partial permutations, we have

$$\mathbb{P}(\mathrm{ltrm\text{-}perm}_p(\sigma) > Q_c) \leq \exp(-c^2/\gamma) .$$

Hence, we can bound $\mathrm{ltrm}(\sigma', \overline{\mu})$ from above by

$$\begin{aligned}
& \sum_{c=0}^{\log n} Q_{c+1} \cdot \mathbb{P}(\sigma' \text{ is not } c\text{-successful but } (c+1)\text{-successful}) + n \cdot (P + P') \\
\leq \;& 2 \cdot (1-p) \cdot \sqrt{n/p} \cdot \underbrace{\sum_{c \in \mathbb{N}} (c+1) \cdot e^{-\frac{c^2}{\gamma}}}_{< 1.8 \text{ for } \gamma < 1.001} + n \cdot (P + P') \\
\leq \;& C \cdot (1-p) \cdot \sqrt{n/p}
\end{aligned}$$

for some $C < 3.6$, which proves the theorem. □

The following lemma is an improvement of the lower bound proof for the number of left-to-right maxima under partial permutations presented by Banderier et al. [4]. This way we get a lower bound with a much larger constant that holds for all $p \in (0, 1)$; the lower bound provided by Banderier et al. holds only for $p \leq \frac{1}{2}$.

**Lemma 5.2.** *Let $p \in (0, 1)$, $\alpha > 1$, and $c > 0$. Then for all sufficiently large $n$, there exist sequences $\sigma$ of length $n$ with*

$$\mathrm{ltrm\text{-}perm}_p(\sigma) \geq \exp(-c^2 \alpha) \cdot c \cdot (1-p) \cdot \sqrt{n/p} \,.$$

*Proof.* Let $K_c = c\sqrt{n/p}$. Let $\sigma = (n - K_c + 1, n - K_c + 2, \ldots, n, 1, 2, \ldots, n - K_c)$. Choose $\beta$ arbitrarily with $1 < \beta^3 < \alpha$. Let $P$ denote the probability that more than $\beta p K_c$ of the first $K_c$ elements or less than $\beta^{-1} pn$ of the remaining elements are selected. $P$ tends exponentially fast to zero as $n$ increases (Lemma 2.1).

Let $\mu$ be the set of marked positions and $\mu_c = \mu \cap [K_c]$ be the set of marked positions among the first $K_c$ positions. Let $y = |\mu \setminus \mu_c|$ and $\mu_c = \{i_1, \ldots, i_x\}$ with $i_1 < \ldots < i_x$, i.e. $|\mu_c| = x$. Let $f$ be a random permutation of $\mu$. We say that $f$ is **successful** if $f(i) > i$ for all $i \in \mu_c$. Thus, under a successful permutation, all marked elements of $\{n - K_c + 1, \ldots, n\}$ are moved further to the back.

If $f$ is successful, then all $K_c - x$ unmarked elements of $\{n - K_c + 1, \ldots, n\}$ are left-to-right maxima. Provided that at most $\beta p K_c$ of the first $K_c$ elements are marked, i.e. $x \leq \beta p K_c$, the expectation of $K_c - x$ is at least $(1 - p)K_c$.

Let us bound the probability from below that the random permutation of $\mu$ is successful for a given $\mu$: For $i_x$, there are $y$ positions allowed and $x$ positions not allowed, for $i_{x-1}$, there are $y$ positions allowed (all in $\mu \setminus \mu_c$ plus $i_x$ minus $f(i_x)$) and $x - 1$ positions not allowed, ..., for $i_1$, there are $y$ positions allowed and one position not allowed. Thus, the probability that the random permutation is successful is at least

$$\left(\frac{y}{y+x}\right)^x = \big(\underbrace{(1 - \frac{x}{y+x})^{\frac{y+x}{x}}}_{\geq e^{-1}\cdot(1-\frac{x}{y+x})}\big)^{\frac{x^2}{y+x}} \geq \exp\left(\left(\ln\left(1 - \frac{x}{y+x}\right) - 1\right) \cdot \frac{x^2}{y+x}\right).$$

Provided that $x \leq \beta p K_c$ and $x + y \geq y \geq \beta^{-1} pn$, we obtain that the probability that the random function is successful is at least

$$\begin{aligned} & \exp\left(\left(\ln\left(1 - \frac{\beta p K_c}{\beta^{-1} pn}\right) - 1\right) \cdot \frac{\beta^2 p^2 K_c^2}{\beta^{-1} pn}\right) \\ = \; & \exp\left(\left(\ln\left(1 - \frac{\beta^2 c}{\sqrt{pn}}\right) - 1\right) \cdot \beta^3 c^2\right) = Q \cdot \exp(-\beta^3 c^2) \end{aligned}$$

for $Q = \left(1 - \frac{\beta^2 c}{\sqrt{pn}}\right)^{\beta^3 c^2}$, which tends to one as $n$ increases. Thus, with probability at least $(1-P) \cdot Q \cdot \exp(-\beta^3 c^2)$, all unmarked elements of $\{K_c + 1, \ldots, n\}$ are left-to-right maxima. Furthermore, we have $(1 - P) \cdot Q \cdot \exp(-\beta^3 c^2) \geq \exp(-c^2 \alpha)$ for sufficiently large $n$. Since the expectation of the number of unmarked elements among the first $K_c$ elements is at least $(1 - p)K_c$, the lemma is proved. □

By choosing $\alpha$ sufficiently close to 1 and $c = \sqrt{1/2}$, we immediately get the following theorem from Lemma 5.2.

**Theorem 5.3.** *For all $p \in (0, 1)$ and all sufficiently large $n$, there exists a sequence $\sigma$ of length $n$ with*

$$\text{ltrm-perm}_p(\sigma) \geq 0.4 \cdot (1 - p) \cdot \sqrt{n/p}\,.$$

Theorem 5.3 would also yield the same lower bound for height-perm$_p(\sigma)$ since the number of left-to-right maxima of a sequence bounds the height of the binary search tree obtained from that sequence from below. However, for the smoothed height of binary search trees, we can proof a stronger lower bound (Theorem 6.3).

Another consequence of Lemma 5.2 is that there does not exist a constant $c$ such that the number of left-to-right maxima is at most $c \cdot (1-p) \cdot \sqrt{n/p}$ with high probability, i.e. with probability at least $1 - n^{-O(1)}$. Thus, the bounds proved for the expectation of the tree height or the number of left-to-right maxima cannot be generalised to bounds that hold with high probability. A bound that holds with high probability can directly be obtained from Lemma 4.6: Let $\sigma'$ be the sequence obtained from $\sigma$ via $p$-partial permutation. Then $\text{height}(\sigma') \in O(\sqrt{(n/p) \cdot \log n})$ with probability at least $1 - n^{-\Omega(1)}$. The same holds for $\text{ltrm}(\sigma')$.

## 5.2 Partial Alterations

As for the height of binary search trees, we obtain the same upper bound for the expected number of left-to-right maxima under partial alterations.

**Theorem 5.4.** *Let $p \in (0,1)$. Then for all sufficiently large $n$ and for all sequences $\sigma$ of length $n$ (where $\sigma$ is a permutation of $[n - \frac{1}{2}]$), we have*

$$\text{ltrm-alter}_p(\sigma) \leq 3.6 \cdot (1-p) \cdot \sqrt{n/p}\,.$$

*Proof.* The mere difference between the proof of this theorem and the proof of Theorem 5.1 is that we have to use Lemma 4.8 instead of Lemma 4.6. For the remaining differences, we refer to the proof of Theorem 6.4, which differs from the proof of Theorem 6.1 in the same way as the proof of this theorem differs from the proof of Theorem 5.1. □

Let us now prove the counterpart for partial alterations of Lemma 5.2.

**Lemma 5.5.** *Let $p \in (0,1)$, $\alpha > 1$, and $c > 0$. Then for all sufficiently large $n$, there exist sequences $\sigma$ of length $n$ with*

$$\text{ltrm-alter}_p(\sigma) \geq \exp(-c^2\alpha) \cdot c \cdot (1-p) \cdot \sqrt{n/p}\,.$$

*Proof.* Let $K_c = c\sqrt{n/p}$. Let $\sigma = (K_c + 1, K_c + 2, \ldots, n, 1, 2, \ldots, K_c)$. Choose $\beta$ arbitrarily with $1 < \beta < \alpha$. Let $P$ denote the probability that more than $\beta p K_c$ are marked. $P$ tends exponentially fast to zero as $n$ increases (Lemma 2.1).

Let $\mu$ be the set of marked positions and $\mu_c = \mu \cap [K_c]$ be the set of marked positions among the first $K_c$. Let $\mu_c = \{i_1, \ldots, i_x\}$ with $i_1 < \ldots < i_x$, i.e. $|\mu_c| = x$. We have $\sigma_{i_j} = n - K_c + i_j - \frac{1}{2}$ for all $j \in [x]$. Let $\sigma'$ be the sequence obtained from $\sigma$ by replacing all marked elements by random numbers from $[0, n)$. We say that $\sigma'$

is successful if $\sigma'_{i_j} \leq n - K_c$. If $\sigma'$ is successful, then all $K_c - x$ unmarked elements among the first $K_c$ elements of $\sigma$ are left-to-right maxima.

The probability that $\sigma'$ is successful is at least

$$\left(\frac{n-K_c}{n}\right)^x = \Big(\underbrace{\big(1 - \frac{K_c}{n}\big)^{\frac{n}{K_c}}}_{\geq e^{-1}\cdot(1-\frac{K_c}{n})}\Big)^{\frac{xK_c}{n}} \geq \exp\left(\left(\ln\left(1 - \frac{K_c}{n}\right) - 1\right) \cdot xK_c/n\right) .$$

Provided that $x \leq \beta p K_c$, we obtain that the probability that the random function is successful is at least

$$\begin{aligned} & \exp\left(\left(\ln\left(1 - \frac{\beta p K_c}{n}\right) - 1\right) \cdot \beta p K_c^2/n\right) \\ = \ & \exp\left(\left(\ln\left(1 - \frac{\beta c}{\sqrt{pn}}\right) - 1\right) \cdot \beta c^2\right) = Q \cdot \exp(-\beta c^2) \end{aligned}$$

for $Q = \left(1 - \frac{\beta c}{\sqrt{pn}}\right)^{\beta c^2}$, which tends to one as $n$ increases. Thus, with probability at least $(1-P)\cdot Q \cdot \exp(-\beta c^2)$, all unmarked among the first $K_c$ elements are left-to-right maxima. The expectation of the number of unmarked elements among the first $K_c$ elements is at least $(1-p)K_c$. Furthermore, for sufficiently large $n$, we have $(1-P)\cdot Q\cdot \exp(-\beta c^2) \geq \exp(-\alpha c^2)$, which proves the lemma. □

From the above lemma, we obtain the same lower bounds for the number of left-to-right maxima as for partial permutations.

**Theorem 5.6.** *For all $p \in (0,1)$ and all sufficiently large $n$, there exists a sequence $\sigma$ of length $n$ with*

$$\text{ltrm-alter}_p(\sigma) \geq 0.4 \cdot (1-p) \cdot \sqrt{n/p} \,.$$

As for partial permutations, a consequence of Lemma 5.5 is that we cannot achieve a bound of $O((1-p)\cdot\sqrt{n/p})$ that holds with high probability for the number of left-to-right maxima or the height of binary search trees. But again, we obtain from Lemma 4.8 that for all sequences, the height and the number of left-to-right maxima under partial alterations is in $O(\sqrt{(n/p)\cdot \log n})$ with probability at least $1 - n^{-\Omega(1)}$.

# 6 Tight Bounds for the Smoothed Height of Binary Search Trees

In this section, we consider the smoothed height of binary search trees under the perturbation models partial permutation and partial alteration.

## 6.1 Partial Permutations

Let us now prove one of the main theorems of this work, namely an upper bound for the expected height of binary search trees obtained from sequences under partial permutations.

**Theorem 6.1.** *Let $p \in (0,1)$ and $\alpha > 1$. Then for all sufficiently large $n$ and all sequences $\sigma$ of length $n$, we have*

$$\text{height-perm}_p(\sigma) \leq 7 \cdot (1 - \alpha^{-1} p) \cdot \sqrt{n/p}\,.$$

*Proof.* According to Lemma 4.4, it suffices to show

$$\mathbb{E}(\text{height}(\sigma', \overline{\mu})) \leq 7 \cdot (1 - \beta^{-1} p) \cdot \sqrt{n/p}$$

for some $\beta$ with $1 < \beta < \alpha$, where $\mu \subseteq [n]$ is the random set of marked positions and $\sigma'$ is the sequence obtained via randomly permuting the elements of $\sigma_\mu$. Then $\text{height-perm}_p(\sigma) \leq 7\cdot(1-\beta^{-1}p)\cdot\sqrt{n/p}+O(\log n) \leq 7\cdot(1-\alpha^{-1}p)\cdot\sqrt{n/p}$ for sufficiently large $n$. Choose $\gamma$ and $\beta$ arbitrarily with $1 < \gamma < \beta < \alpha$ and $\gamma < 1.01$. Without loss of generality, we assume that $\sigma$ is a permutation of $[n]$.

Let $c \in [\log n]$ and $K_c = c\sqrt{n/p}$. We define $D(d) = \sum_{i=1}^{d-1} i^2 = \frac{1}{3}\cdot(d-1)\cdot(d-\frac{1}{2})\cdot d$. Then $D(d) \geq \frac{1}{8}d^3$ for $d \geq 2$.

We divide the sequence $\sigma$ into blocks $B_1, B_2, \ldots, B_{(\log n)^2}$. The block $B_d$ consists of $d^2 K_c$ elements: $B_1$ contains the elements of $\sigma$ at the first $K_c$ positions, $B_2$ contains the elements of $\sigma$ at the next $4K_c$ positions, and so on. Thus,

$$B_d = \sigma_{[D(d+1)\cdot K_c]} \setminus \sigma_{[D(d)\cdot K_c]}\,.$$

Let $B = \bigcup_{d=1}^{(\log n)^2} B_d$ be the set of elements that are contained in any $B_d$. We have $|B| = D((\log n)^2 + 1) \cdot K_c \geq \frac{1}{8}(\log n)^6 \cdot K_c$.

Every block $B_d$ is further divided into $d^4$ subsets $A_d^1, \ldots, A_d^{d^4}$ of elements as follows: $A_d^1$ contains the $d^{-2}K_c$ smallest elements of $B_d$, $A_d^2$ contains the $d^{-2}K_c$ second smallest elements of $B_d$, and so on. The subset $A_d^1$ contains the $d^{-2}K_c$ smallest elements of $B_d$, $A_d^2$ the $d^{-2}K_c$ second smallest elements of $B_d$, $\ldots$, and $A_d^{d^4}$ contains the $d^{-2}K_c$ largest elements of $B_d$. Figure 3(a) illustrates the division of $\sigma$ into blocks $B_1, B_2, \ldots, B_{(\log n)^2}$ and subsets $A_d^i$ ($d \in [(\log n)^2]$, $i \in [d^2]$).

Finally, we divide the numbers in $[n]$ into $\log n \cdot \sqrt{np}$ subsets $C_1, \ldots, C_{\log n\cdot\sqrt{np}}$ with

$$C_j = \left\{ \frac{\sqrt{n/p}}{\log n} \cdot (j-1) + 1, \ldots, \frac{\sqrt{n/p}}{\log n} \cdot j \right\}.$$

Thus, $C_1$ contains the $(\log n)^{-1} \cdot \sqrt{n/p}$ smallest numbers of $[n]$, $C_2$ contains the $(\log n)^{-1} \cdot \sqrt{n/p}$ second smallest numbers of $[n]$, and so on.

Let $\mu$ be the random set of marked positions and $\sigma'$ be the random sequence obtained from $\sigma$ via partial permutation. We say that $\mu$ and $\sigma'$ are **partially successful** if the following properties are fulfilled:

- for all $c \in [\log n]$, $d \in [(\log n)^2]$, and $i \in [d^4]$, at least $\gamma^{-1} c d^{-2} \sqrt{n/p}$ and at most $\gamma c d^{-2} \sqrt{n/p}$ elements of $A_d^i$ are marked, and
- for all $j \in [\log n \sqrt{np}]$, at least $\gamma^{-1}(\log n)^{-1}\sqrt{n/p}$ and at most $\gamma(\log n)^{-1}\sqrt{n/p}$ elements of $C_j$ are marked.

The probability that $\mu$ is not partially successful is at most

$$\begin{aligned} &\sum_{c=1}^{\log n} \sum_{d=1}^{(\log n)^2} d^4 \cdot 2 \cdot \exp(-2(1-\gamma^{-1})^2 p^2 d^{-2} c \sqrt{n/p}) \\ + \quad & \log n \sqrt{np} \cdot 2 \cdot \exp(-2(1-\gamma^{-1})^2 p^2 (\log n)^{-1} \sqrt{n/p}) \end{aligned}$$

by Lemma 2.1. For sufficiently large $n$, this can be bounded from above by $\exp(-n^{\epsilon})$ for some constant $\epsilon > 0$. Let $P$ denote this probability. If $\mu$ is not partially successful, we bound the height of $\sigma'$ by $n$.

From now on, we assume that $\mu$ is partially successful.

We call a subset $A_d^i$ for $d \geq 2$ and $i \in [d^4]$ ***c*-successful** if at least one element of $A_d^i$ is permuted to one of the $D(d)c\sqrt{n/p}$ positions that precede $B_d$. Thus, the probability that a fixed $A_d^i$ is not successful is at most $\exp(-c^2 D(d) d^{-2} \gamma^{-1}) \leq \exp(-c^2 d/(8\gamma))$ according to Lemma 4.6: there are $d^{-2} c \sqrt{n/p}$ elements in $A_d^i$ and $D(d)c\sqrt{n/p}$ positions that precede $B_d$.

We call a block $B_d$ for $d \geq 2$ ***c*-successful** if all subsets $A_d^1, \ldots, A_d^{d^4}$ of $B_d$ are $c$-successful. The probability that $B_d$ is not $c$-successful is at most $d^4 \cdot \exp(-c^2 d/(8\gamma))$ since there are $d^4$ subsets $A_d^1, \ldots, A_d^{d^4}$ of $B_d$. Figures 3(a) and 3(b) illustrate $c$-success.

Let $d' = (\log n)^2 + 1$ and $D' = D(d') \geq (\log n)^6/8$. A subset $C_j$ is called ***c*-successful** if at least one element of $C_j$ is among the first $D' c \sqrt{n/p}$ positions of $\sigma'$. The probability that a fixed $C_j$ is not $c$-successful is at most $\exp(-\frac{cD'}{\gamma \log n}) \leq \exp(-\frac{c(\log n)^5}{8\gamma})$. The probability that any $C_j$ is not $c$-successful is bounded from above by

$$\log n \cdot \sqrt{np} \cdot \exp\left(-\frac{c(\log n)^5}{8\gamma}\right) \leq d'^4 \cdot \exp\left(-\frac{c^2 d'}{8\gamma}\right) \tag{6.1}$$

for sufficiently large $n$.

Finally, we say that ***σ′* is *c*-successful** if

- all blocks $B_1, B_2, \ldots, B_{(\log n)^2}$ are $c$-successful and
- all subsets $C_1, \ldots, C_{\log n \sqrt{np}}$ are $c$-successful.

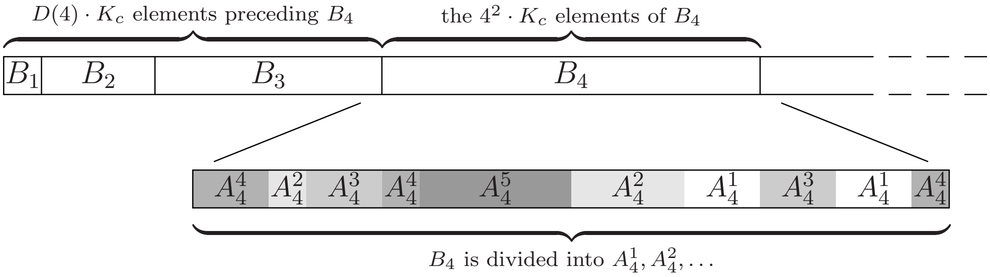


(a) The division of the first $D' \cdot K_c$ elements of $\sigma$ into blocks $B_1, \ldots, B_{(\log n)^2}$. For instance, the block $B_4$ is further divided into subsets $A_4^1, \ldots, A_4^{16}$, where $A_4^1$ contains the $K_c/4$ smallest elements of $B_4$, $\ldots$, and $A_4^{16}$ contains the $K_c/4$ largest elements of $B_4$. (For readability, $B_4$ is divided into only five subsets in the illustration.)

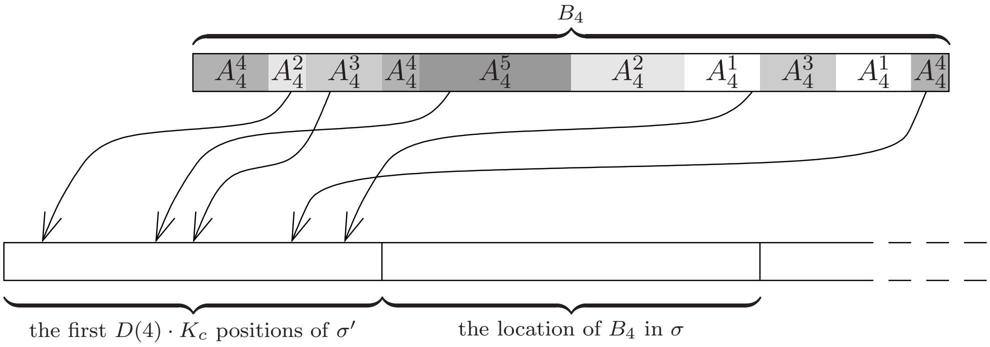


(b) A subset $A_4^i$ is $c$-successful if at least one element of $A_4^i$ is among the first $D(4) \cdot K_c$ elements of $\sigma'$. The block $B_4$ is $c$-successful if all $A_4^i$ are $c$-successful.

Figure 3: The division of $\sigma$ into blocks and subsets (shown here for $B_4$).

Let $c \geq 5$. The probability that $\sigma'$ is not $c$-successful is at most

$$\begin{aligned} & \sum_{d=2}^{(\log n)^2} d^4 \cdot \exp(-c^2 d/(8\gamma)) + \mathbb{P}(\text{some } C_j \text{ is not } c\text{-successful}) \\ \leq \; & \sum_{d=2}^{\infty} d^4 \cdot \exp(-c^2 d/(8\gamma)) \;\leq\; \sum_{d=2}^{\infty} \big(\exp(-c^2/(16\gamma))\big)^d \\ = \; & \frac{\exp(-c^2/(16\gamma))^2}{1-\exp(-c^2/(16\gamma))} = E(c,\gamma)\,. \end{aligned} \tag{6.2}$$

The first inequality holds due to Formula 6.1, the second inequality holds since $c \geq 5$. If $\sigma'$ is not $\log n$-successful, which happens with probability at most $E(\log n, \gamma) \leq \exp(-(\log n)^2/(16\gamma))$, we bound the height of $T(\sigma')$ by $n$.

Let $Q_c = (c \cdot \pi^2/3 + \frac{2}{\log n}) \cdot (1-\gamma^{-1}p) \cdot \sqrt{n/p}$.

**Claim 6.2.** *If $\sigma'$ is $c$-successful, then* $\text{height}(\sigma', \overline{\mu}) \leq Q_c$.

*Proof of Claim 6.2.* Consider any path from the root to a leaf in $T(\sigma')$. This path cannot contain unmarked elements from both $A_d^{i-1}$ and $A_d^{i+1}$ for $d \geq 2$ and $2 \leq i \leq d^4-1$ since there is at least one element of $A_d^i$ that stands before all unmarked elements of $A_d^{i-1}$ and $A_d^{i+1}$.

It is possible that unmarked elements from $A_d^i$ and $A_d^{i+1}$ are on the same root-to-leaf path in $T(\sigma')$. For every $d$ and $i$, there are at most $(1-\gamma^{-1}p)cd^{-2}\sqrt{n/p}$ unmarked elements in $A_d^i$ since $\sigma'$ is partially successful. Thus, for every $d$, at most $2(1-\gamma^{-1}p)cd^{-2}\sqrt{n/p}$ elements of $B_d$ are on the same root-to-leaf path in $T(\sigma')$.

Let $\overline{B} = [n] \setminus B$ be the set of elements of $\sigma$ that are not contained in any $A_d^i$. There cannot be unmarked elements from both $C_{j-1} \cap \overline{B}$ and $C_{j+1} \cap \overline{B}$ on the same root-to-leaf path in $\sigma'$ since there is at least one element of $C_j$ among the first $D'c\sqrt{n/p}$ elements of $\sigma'$. Thus, there are at most $2(1-\gamma^{-1}p)\frac{\sqrt{n/p}}{\log n}$ elements of $\overline{B} \cap \bigcup_{i=1}^{\log n \cdot \sqrt{np}} C_j$ on the same root-to-leaf path in $T(\sigma')$.

Overall, the maximum number of elements on any root-to-leaf path in $T(\sigma')$ can be bounded from above by

$$
\begin{aligned}
& \sum_{d=1}^{(\log n)^2} 2(1-\gamma^{-1}p)cd^{-2}\sqrt{n/p} + 2(1-\gamma^{-1}p)(\log n)^{-1}\sqrt{n/p} \\
\leq\ & \left(2c \cdot \sum_{d=1}^{\infty} \frac{1}{d^2} + \frac{2}{\log n}\right) \cdot (1-\gamma^{-1}p) \cdot \sqrt{n/p} \\
=\ & \left(c \cdot \frac{\pi^2}{3} + \frac{2}{\log n}\right) \cdot (1-\gamma^{-1}p) \cdot \sqrt{n/p} = Q_c\,,
\end{aligned}
$$

which proves the claim. □

Now we can complete the proof of the theorem. According to Claim 6.2 and Formula 6.2, we have

$$\mathbb{P}\left(\text{height}(\sigma', \overline{\mu}) > Q_c\right) \leq E(c, \gamma)$$

for $5 \leq c \leq \log n$. Hence, we can bound the expectation of $\text{height}(\sigma', \overline{\mu})$ from above by

$$
\begin{aligned}
& Q_5 + \sum_{c=5}^{\log n} Q_{c+1} \cdot \mathbb{P}(\sigma' \text{ is not } c\text{-successful but } (c+1)\text{-successful}) \\
& +n \cdot (P + E(\log n, \gamma)) \\
\leq\ & \underbrace{(1-\gamma^{-1}p) \cdot \sqrt{n/p}}_{=\Gamma} \cdot \left(5 + \sum_{c=5}^{\infty} \left(\frac{\pi^2}{3}(c+1) + \frac{2}{\log n}\right) \cdot E(c,\gamma)\right) + o(\Gamma) \\
=\ & \Gamma \cdot \left(\underbrace{5 + \frac{\pi^2}{3} \cdot \overbrace{\sum_{c \geq 5} (c+1) \cdot E(c, \gamma)}^{< 0.5 \text{ for } \gamma < 1.01}}_{< 7 \text{ for } \gamma < 1.01}\right) + o(\Gamma) \\
\leq\ & 7 \cdot (1-\beta^{-1}p)\sqrt{n/p}
\end{aligned}
$$

for sufficiently large $n$ since $\gamma < \beta$ and $\gamma < 1.01$. This completes the proof. □

As a counterpart to the above theorem, we prove the following lower bound. Interestingly, the lower bound is obtained for the sorted sequence, which is not a worst case for the expected number of left-to-right maxima; the expected number of left-to-right maxima of the sequence obtained by partially permuting the sorted sequence is only logarithmic [4].

**Theorem 6.3.** *For all $p \in (0,1)$ and all sufficiently large $n \in \mathbb{N}$, we have*

$$\text{ltrm-perm}_p(\sigma_{\text{sort}}^n) \geq 0.8 \cdot (1-p) \cdot \sqrt{n/p}\,.$$

*Proof.* Let $c > 0$ be any constant and $K_c = c\sqrt{n/p}$. Let $\sigma'$ be the sequence obtained from $\sigma_{\text{sort}}^n$ via $p$-partial permutation. We say that $\sigma'$ is **c-successful** if all marked elements among the first $K_c$ elements of $\sigma_{\text{sort}}^n$ are permuted further to the back. According to Lemma 5.2, we have

$$\mathbb{P}(\sigma' \text{ is } c\text{-successful}) \geq \exp(-c^2\alpha)$$

for arbitrarily chosen $\alpha > 1$ and sufficiently large $n$. If $\sigma'$ is $c$-successful and $x$ elements among the first $K_c$ elements are unmarked, then $\text{height}(\sigma') \geq x$. Let $Q = (1-p)\cdot\sqrt{n/p}$ for short. Analogously to Lemma 5.2, we obtain

$$\mathbb{P}\left(\text{height}(\sigma') \geq cQ\right) \geq \exp(-c^2\alpha)$$

for $\alpha > 1$ and sufficiently large $n$. The idea is now to consider $c$-success for all $c \in \{0.1, 0.2, \ldots, 9.9, 10\} = C$. Thus,

$$\begin{aligned}
\mathbb{E}(\text{height}(\sigma')) &\geq Q \cdot \sum_{c \in C} c \cdot \mathbb{P}(cQ \leq \text{height}(\sigma') < (c+1)Q) \\
&\geq Q \cdot \sum_{c \in C} 0.01 \cdot \mathbb{P}(\text{height}(\sigma') \geq cQ) \\
&\geq Q \cdot \sum_{c \in C} 0.01 \cdot \exp(-c^2\alpha) \;\geq 0.8 \cdot Q
\end{aligned}$$

for sufficiently large $n$ and $\alpha < 1.01$, which proves the theorem. □

## 6.2 Partial Alterations

The following theorem is obtained via a proof similar to the proof of Theorem 6.1.

**Theorem 6.4.** *Let $p \in (0,1)$ and $\alpha > 1$. Then for all sufficiently large $n$ and all sequences $\sigma$ of length $n$ (where $\sigma$ is a permutation of $[n-\frac{1}{2}]$), we have*

$$\text{height-alter}_p(\sigma) \leq 7 \cdot (1-\alpha^{-1}p) \cdot \sqrt{n/p}\,.$$

*Proof.* The mere difference between the proof of this theorem and the proof of Theorem 6.1 is that we have to use Lemma 4.8 instead of Lemma 4.6. The blocks $B_d$ and $C_j$ and the subsets $A_d^i$ are defined in the same way. Now we have for each subset $A_i^d$ numbers $a_d^i = \lfloor \min A_i^d \rfloor$ and $b_d^i = \lceil \max A_i^d \rceil$. We say that $A_i^d$ is ***c*-successful** if among the first $D(d)c\sqrt{n/p}$ elements there is at least one element of the interval $[a_d^i, b_d^i)$. The term $c$-successful for blocks $B_d$ is defined in the same way as in the previous proof. For subsets $C_j$, the term $c$-successful is defined just like for $A_i^d$. The remainder of the proof goes along the same lines as the proof of Theorem 6.1. □

We also get the same lower bound for the height of the binary search trees under partial alterations. Again, the lower bound is obtained for the sorted sequence.

**Theorem 6.5.** *For all $p \in (0,1)$ and all sufficiently large $n \in \mathbb{N}$, we have*

$$\text{ltrm-alter}_p(\sigma_{\text{sort}}^n - 0.5) \geq 0.8 \cdot (1-p) \cdot \sqrt{n/p}\,.$$

*Proof.* The proof is almost equal to the proof of Theorem 6.3. The only difference is that we have to use Lemma 5.5 instead of Lemma 5.2. □

# 7 Comparing Partial Deletions with Permutations and Alterations

In this section, we justify the consideration of partial deletions. Partial deletions turn out to be the worst of the three models: Trees are usually expected to be higher under partial deletions than under partial permutations or alterations, although containing less elements. Thus, the expected height under partial deletions yields upper bounds (up to an additional $O(\log n)$ term) for the expected height under partial permutations and alterations. On the other hand, we prove that lower bounds for the expected height under partial deletions yield slightly weaker lower bounds for permutations and alterations. The main advantage of partial deletions over partial permutations and partial alterations is that partial deletions are much easier to analyse.

For the sake of completeness, we start by providing matching upper and lower bounds for the height of binary search trees under partial deletions.

**Proposition 7.1.** *For all $p \in [0,1]$, $n \in \mathbb{N}$, and sequences $\sigma$ of length $n$, we have*

$$\text{height-del}_p(\sigma) \leq (1-p) \cdot n\,.$$

*Moreover,*

$$\text{height-del}_p(\sigma_{\text{sort}}^n) = (1-p) \cdot n\,.$$

*Proof.* After $p$-partial deletion, the sequence $\sigma$ is expected to have $(1-p) \cdot n$ elements. The number of elements is clearly an upper bound for the height. The second proposition holds obviously. □

The following lemma is an immediate consequence of Lemmas 4.4, 4.5, and 4.7, we therefore omit its proof.

**Lemma 7.2.** *For all sequences $\sigma$ of length $n$ and $p \in [0,1]$,*

$$\begin{array}{lcl} \text{height-perm}_p(\sigma) & \leq & \text{height-del}_p(\sigma) + O(\log n) \ \ \textit{and} \\ \text{ltrm-perm}_p(\sigma) & \leq & \text{ltrm-del}_p(\sigma) + O(\log n) \ . \end{array}$$

*If $\sigma$ is a permutation of $[n - \frac{1}{2}]$, then*

$$\begin{array}{lcl} \text{height-alter}_p(\sigma) & \leq & \text{height-del}_p(\sigma) + O(\log n) \ \ \textit{and} \\ \text{ltrm-alter}_p(\sigma) & \leq & \text{ltrm-del}_p(\sigma) + O(\log n) \ . \end{array}$$

Thus, we can bound the expected height under partial permutations or alterations from above by the expected height under partial deletions. The reverse is not true (Theorems 6.1 and 6.4 and Proposition 7.1). But we can bound the expected height under partial deletions by the expected height under partial permutations or alterations by padding the sequences considered.

**Lemma 7.3.** *Let $p \in (0,1)$ be fixed and let $\sigma$ be a sequence of length $n$ with* $\text{height}(\sigma) = d$ *and* $\text{height-del}_p(\sigma) = d'$. *Then there exists a sequence $\tilde{\sigma}$ of length $O(n^2)$ with* $\text{height}(\tilde{\sigma}) = d + O(\log n)$ *and* $\text{height-perm}_p(\tilde{\sigma}) \in \Omega(d')$.

*Proof.* Without loss of generality, we assume that $\sigma$ is a permutation of $[n]$. The idea is to attach a tail of sufficiently many elements greater than $n$ to the sequence such that all marked elements that are greater than or equal to $n$ will be permuted to this tail. Thus, the overall structure of the remaining elements from $[n]$ will be as via partial deletions.

Choose $K = n^2p$ and construct $\tilde{\sigma}$ from $\sigma$ as follows: the first $n$ items of $\tilde{\sigma}$ are just $\sigma$, which we call the **head** of $\tilde{\sigma}$. The last $K - n$ items of $\tilde{\sigma}$, which we call the **tail** of $\tilde{\sigma}$, are numbers greater than $n$ such that these numbers build a tree of height $O(\log(K - n)) = O(\log n)$. With constant probability (see the proof of Lemma 5.2), say $c$, all elements marked in the head are permuted into the tail.

Consider the tree obtained from the first $n$ elements after partial permutation under the restriction that all marked head elements are now in the tail. This tree merely equals the tree obtained via partial deletion from $\sigma$, when the same elements are marked. The only difference are some elements greater than $n$, which only elongate the right-most path. Thus, $\text{height-perm}_p(\tilde{\sigma})$ is at least $cd'$, which proves the lemma. □

The following is the analogue of the above lemma for partial alterations. Since its proof is similar to the proof of the previous lemma (the mere difference is that we have to use the proof of Lemma 5.5 instead of Lemma 5.2), we omit it.

**Lemma 7.4.** *Let $p \in (0,1)$ be fixed and $\sigma$ be a sequence of length $n$ with elements of $[n - \frac{1}{2}]$. Let* $\text{height}(\sigma) = d$ *and* $\text{height-del}_p(\sigma) = d'$. *Then there exists a sequence $\tilde{\sigma}$ of length $O(n^2)$ with* $\text{height}(\tilde{\sigma}) = d + O(\log n)$ *and* $\text{height-alter}_p(\tilde{\sigma}) \geq \Omega(d')$.

# 8 The (In-)Stability of Perturbations

Having shown that worst case instances become much better by smoothing, we now provide a family of best case instances for which smoothing results in an exponential increase in height.

We consider the following family of sequence:

- $\sigma^{(1)} = (1)$.
- $\sigma^{(k+1)} = (2^k, \sigma^{(k)}, 2^k + \sigma^{(k)})$, where $c + \sigma = (c + \sigma_1, \ldots, c + \sigma_n)$ for a sequence $\sigma$ of length $n$.

For instance, $\sigma^{(2)} = (2, 1, 3)$ and $\sigma^{(3)} = (4, 2, 1, 3, 6, 5, 7)$. Let $n = 2^k - 1$. Then $\sigma^{(k)}$ contains the numbers $1, 2, \ldots, n$, and we have $\text{height}(\sigma^{(k)}) = \text{ltrm}(\sigma^{(k)}) = k \in \Theta(\log n)$.

Let us estimate the expected number of left-to-right maxima after partial deletion, bounding the expected height of the binary search tree from below. The main reason for the following theorem and its corollary is that deleting the first element of $\sigma^{(k)}$ roughly doubles the height of the resulting tree and the number of left-to-right maxima in the resulting sequence. Figure 4 illustrates this.

**Theorem 8.1.** *Let $p \in (0, 1)$. Then for all $k \in \mathbb{N}$,*

$$\text{ltrm-del}_p(\sigma^{(k)}) = \frac{1-p}{p} \cdot \left((1+p)^k - 1\right) .$$

*Proof.* Let $\ell(k) = \text{ltrm-del}_p(\sigma^{(k)})$ for short. The root $2^{k-1}$ is deleted with probability $p$. Then the expected number of left-to-right maxima is just the expectation for the left subtree plus the expectation for the right subtree since all elements in the left subtree are smaller and occur earlier than all elements in the right subtree. Both expectations are $\ell(k-1)$. If the root is not deleted, we expect $1 + \ell(k-1)$ left-to-right maxima: One is the root and $\ell(k-1)$ are expected in the right subtree. We do not get any new maximum from the left subtree since all elements in the left subtree are smaller than the root. We have $\ell(1) = 1 - p$ since with probability $p$, the single element will be deleted. Overall, we have

$$\begin{aligned}
\ell(k) &= 2p \cdot \ell(k-1) + (1-p) \cdot (1 + \ell(k-1)) \\
&= (1+p) \cdot \ell(k-1) + (1-p) = (1-p) \cdot \sum_{i=0}^{k-1} (1+p)^i \\
&= \frac{1-p}{p} \cdot \left((1+p)^k - 1\right) .
\end{aligned}$$

$\square$

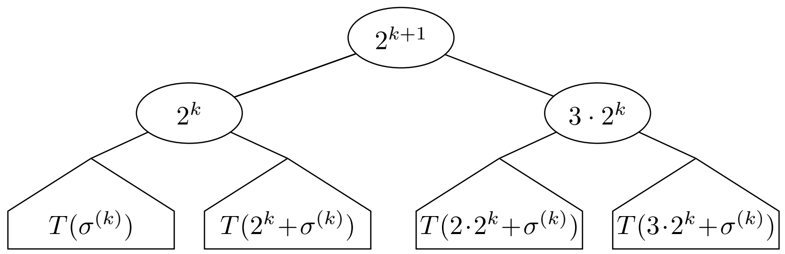


(a) $T(\sigma^{(k+2)})$.

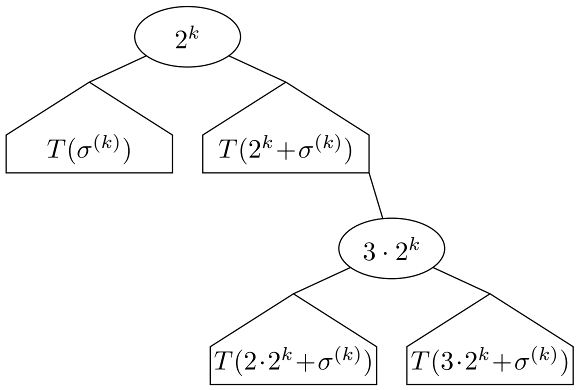


(b) Removing the root $2^{k+1}$ roughly doubles the height.

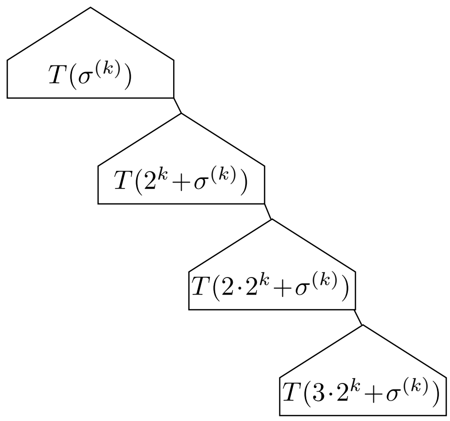


(c) Removing further the roots $2^k$ and $3 \cdot 2^k$ of $T(\sigma^{(k+1)})$ and $T(2^{k+1} + \sigma^{(k+1)})$, respectively, quadruplicates the height.

Figure 4: Removing root elements increases the height and the number of left-to-right maxima.

Since the number of left-to-right maxima of a sequence is a lower bound for the height of the binary search tree obtained from the same sequence, we get the following result.

**Corollary 8.2.** *For all $p \in (0,1)$ and all $k \in \mathbb{N}$,*

$$\text{height-del}_p(\sigma^{(k)}) \geq \frac{1-p}{p} \cdot \left((1+p)^{k+1} - 1\right).$$

We conclude that there are some best case instances that are quite fragile under partial deletions: From logarithmic height they "jump" to height $\Omega(n^\epsilon)$, for constant $\epsilon > 0$, via smoothing. Thus, the height increases exponentially.

We can transfer this result to partial permutations and partial alterations due to Lemmas 7.3 and 7.4. Therefore, we consider sequences $\tilde{\sigma}^{(k)}$ which are obtained from $\sigma^{(k)}$ as described in the proof of Lemma 7.3.

**Corollary 8.3.** *For any fixed $p \in (0,1)$, we have*

$$\begin{array}{lcl} \text{height}(\tilde{\sigma}^{(k)}) & \in & O(\log n)\ , \\ \text{height-perm}_p(\tilde{\sigma}^{(k)}) & \in & \Omega\big(n^{\epsilon}\big)\ ,\ \textit{and} \\ \text{height-alter}_p(\tilde{\sigma}^{(k)} - \frac{1}{2}) & \in & \Omega\big(n^{\epsilon}\big) \end{array}$$

*for some fixed $\epsilon > 0$. (We have to subtract $\frac{1}{2}$ from $\tilde{\sigma}^{(k)}$ in case of partial alterations to be consistent with the model.)*

For the sake of completeness, let us mention that the number of left-to-right-maxima is, at least asymptotically for any fixed $p$, as fragile as possible: There are sequences with one left-to-right maximum for which the expected number of left-to-right maxima after partial permutation is $\Omega(\sqrt{n})$. The same holds for partial alterations. For partial deletions, the number can jump from 1 to $\Omega(n)$. The proofs are straight-forward: take an adversarial sequence of length $n-1$ for proving lower bounds for the expected number of left-to-right maxima under any of these perturbation models and add an $n$ at the front of the sequence. With constant probability, this $n$ will be marked and moved behind the first $\Theta(\sqrt{n/p})$ elements in case of partial permutations. For the other two models, the proof is similar.

# 9 Conclusions

We have analysed the height of binary search trees obtained from perturbed sequences and obtained asymptotically tight lower and upper bounds of roughly $\sqrt{n}$ for the height under partial permutations and alterations. This stands in contrast to both the worst case and the average case height of $n$ and $\Theta(\log n)$, respectively. Thus, the height of binary search trees under limited randomness differs significantly from both the average and the worst case. One direction for future work is of course improving the constants of the bounds and getting rid of the $\alpha$ in the upper bound.

Interestingly, the sorted sequence $\sigma^n_{\text{sort}}$ turns out to be a worst case for the smoothed height of binary search trees in the sense that the lower bounds are obtained for $\sigma^n_{\text{sort}}$ (Theorems 6.3 and 6.5). This contrasts the fact that the expected number of left-to-right maxima of $\sigma^n_{\text{sort}}$ under $p$-partial permutations is roughly $O(\log n)$ [4]. We believe that for the height of binary search trees, $\sigma^n_{\text{sort}}$ is indeed a worst case.

**Conjecture 9.1.** *For all $p \in [0,1]$, all $n \in \mathbb{N}$, and every sequence $\sigma$ of length $n$,*

$$\begin{array}{lcl} \text{height-perm}_p(\sigma) & \le & \text{height-perm}_p(\sigma^n_{\text{sort}})\ \ \textit{and} \\ \text{height-alter}_p(\sigma) & \le & \text{height-alter}_p(\sigma^n_{\text{sort}})\ . \end{array}$$

We performed experiments to estimate the constants in the bounds for the height of binary search trees. For all $n \in \{1\,000, 2\,000, \ldots, 500\,000\}$ and $p \in \{0.1, 0.25\}$, we performed $5\,000$ partial permutations of $\sigma^n_{\text{sort}}$. Furthermore, we did the same for all $n \in \{100\,000, 200\,000, \ldots, 500\,000\}$ and $p \in \{0.01, 0.02, \ldots, 0.99\}$. (See Appendix A for more details.) The results lead to the following conjecture. Provided that Conjecture 9.1 holds, this would yield an improved upper bound for the height of binary search trees under partial permutations.

**Conjecture 9.2.** *For $p \in (0,1)$ and sufficiently large $n$,*

$$\text{height-perm}_p(\sigma^n_{\text{sort}}) = (\gamma + o(1)) \cdot (1-p) \cdot \sqrt{n/p}$$

*for some constant $\gamma \approx 1.8$.*

Throughout the work, the bounds obtained for partial permutations and partial alterations are equal. Moreover, the proofs used for obtaining these bounds are almost identical. We suspect that this always holds for binary search trees.

**Conjecture 9.3.** *For all $p \in (0,1)$ and $\sigma$,*

$$\text{height-perm}_p(\sigma) \approx \text{height-alter}_p(\sigma) .$$

Beyond partial permutations and alterations, we are interested in other perturbation models for sequences. From a more abstract point of view, a future research direction is to characterise the properties of perturbation models that lead to upper or lower bounds that are asymptotically different from the average or worst case.

Apart from lower and upper bounds, we examined the stability of perturbations, i.e. how much higher a tree can become if the underlying sequence is perturbed. It turned out that all three perturbation models are unstable. Again, we are interested in examining other perturbation models. Are there any such models that break down worst cases but do not increase the height of good cases by much? More abstractly, what are the properties of a perturbation model that does increase the height of trees by much?

Finally, we are interested in generalising these results to other problems based on permutations, like sorting algorithms (Quicksort under partial permutations has already been investigated by Banderier et al. [4]), routing algorithms, and other data structures. Hopefully, this will shed some light on the discrepancy between worst case and average case behaviour of these problems.

## Acknowledgements

We thank Jan Arpe for valuable discussions and comments.

## References

[1] Alfred V. Aho, John E. Hopcroft, and Jeffrey D. Ullman. *The Design and Analysis of Computer Algorithms.* Addison-Wesley, 1974.

[2] Alfred V. Aho, John E. Hopcroft, and Jeffrey D. Ullman. *Data Structures and Algorithms*. Addison-Wesley, 1983.

[3] Noga Alon, Joel H. Spencer, and Paul Erdős. *The Probabilistic Method*. John Wiley & Sons, 1992.

[4] Cyril Banderier, René Beier, and Kurt Mehlhorn. Smoothed analysis of three combinatorial problems. In Branislav Rovan and Peter Vojtás, editors, *Proc. of the 28th Int. Symp. on Mathematical Foundations of Computer Science (MFCS)*, volume 2747 of *Lecture Notes in Computer Science*, pages 198–207. Springer, 2003.

[5] Luca Becchetti, Stefano Leonardi, Alberto Marchetti-Spaccamela, Guido Schäfer, and Tjark Vredeveld. Average case and smoothed competitive analysis of the multi-level feedback algorithm. In *Proc. of the 44th Ann. IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 462–471. IEEE Computer Society, 2003.

[6] René Beier and Berthold Vöcking. Typical properties of winners and losers in discrete optimization. In *Proc. of the 36th Ann. ACM Symp. on Theory of Computing (STOC)*, pages 343–352. ACM Press, 2004.

[7] Avrim L. Blum and John D. Dunagan. Smoothed analysis of the perceptron algorithm for linear programming. In *Proc. of the 13th Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA)*, pages 905–914. SIAM, 2002.

[8] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to Algorithms*. MIT Press, 2nd edition, 2001.

[9] Valentina Damerow, Friedhelm Meyer auf der Heide, Harald Räcke, Christian Scheideler, and Christian Sohler. Smoothed motion complexity. In Giuseppe Di Battista and Uri Zwick, editors, *Proc. of the 11th Ann. European Symp. on Algorithms (ESA)*, volume 2832 of *Lecture Notes in Computer Science*, pages 161–171. Springer, 2003.

[10] Valentina Damerow and Christian Sohler. Extreme points under random noise. In Susanne Albers and Tomasz Radzik, editors, *Proc. of the 12th Ann. European Symp. on Algorithms (ESA)*, volume 3221 of *Lecture Notes in Computer Science*, pages 264–274. Springer, 2004.

[11] Luc Devroye. A note on the height of binary search trees. *Journal of the ACM*, 33(3):489–498, 1986.

[12] Luc Devroye and Bruce Reed. On the variance of the height of random binary search trees. *SIAM Journal on Computing*, 24(6):1157–1162, 1995.

[13] Michael Drmota. An analytic approach to the height of binary search trees. *Algorithmica*, 29(1–2):89–119, 2001.

[14] Michael Drmota. An analytic approach to the height of binary search trees II. *Journal of the ACM*, 50(3):333–374, 2003.

[15] Abraham D. Flaxman and Alan M. Frieze. The diameter of randomly perturbed digraphs and some applications. In Klaus Jansen, Sanjeev Khanna, José D. P. Rolim, and Dana Ron, editors, *Proc. of the 8th Int. Workshop on Randomization and Computation (RANDOM)*, volume 3122 of *Lecture Notes in Computer Science*, pages 345–356. Springer, 2004.

[16] Donald E. Knuth. *Sorting and Searching*, volume 3 of *The Art of Computer Programming*. Addison-Wesley, 2nd edition, 1998.

[17] Boris Pittel. On growing random binary trees. *Journal of Mathematical Analysis and Applications*, 103(2):461–480, 1984.

[18] Bruce Reed. The height of a random binary search tree. *Journal of the ACM*, 50(3):306–332, 2003.

[19] John Michael Robson. The height of binary search trees. *The Australian Computer Journal*, 11(4):151–153, 1979.

[20] John Michael Robson. The asymptotic behaviour of the height of binary search trees. Technical Report TR-CS-81-15, The Australian National University, Department of Computer Science, Canberra, 1981.

[21] John Michael Robson. On the concentration of the height of binary search trees. In Pierpaolo Degano, Roberto Gorrieri, and Alberto Marchetti-Spaccamela, editors, *Proc. of the 24th Int. Coll. on Automata, Languages and Programming (ICALP)*, volume 1256 of *Lecture Notes in Computer Science*, pages 441–448. Springer, 1997.

[22] John Michael Robson. Constant bounds on the moments of the height of binary search trees. *Theoretical Computer Science*, 276(1–2):435–444, 2002.

[23] Heiko Röglin and Berthold Vöcking. Smoothed analysis of integer programming. In Michael Jünger and Volker Kaibel, editors, *Proc. of the 11th Int. Conf. on Integer Programming and Combinatorial Optimization (IPCO)*, volume 3509 of *Lecture Notes in Computer Science*, pages 276–290. Springer, 2005.

[24] Guido Schäfer and Naveen Sivadasan. Topology matters: Smoothed competitiveness of metrical task systems. In Volker Diekert and Michel Habib, editors, *Proc. of the 22th Int. Symp. on Theoretical Aspects of Computer Science (STACS)*, volume 2996 of *Lecture Notes in Computer Science*, pages 489–500. Springer, 2004.

[25] Daniel A. Spielman and Shang-Hua Teng. Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. In *Proc. of the 33rd Ann. ACM Symp. on Theory of Computing (STOC)*, pages 296–305. ACM Press, 2001.

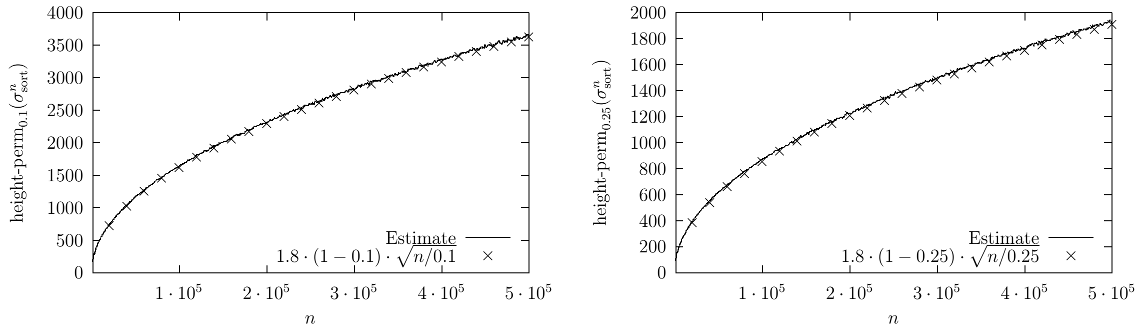


Figure 5: Experimental data for $\sigma^n_{\text{sort}}$ for $n \in \{1\,000, 2\,000, \ldots, 500\,000\}$ and $p = 0.1$ compared to $1.8 \cdot (1-p) \cdot \sqrt{n/p}$.

[26] Daniel A. Spielman and Shang-Hua Teng. Smoothed analysis: Motivation and discrete models. In Frank Dehne, Jörg-Rüdiger Sack, and Michiel Smid, editors, *Proc. of the 8th Workshop on Algorithms and Data Structures (WADS)*, volume 2748 of *Lecture Notes in Computer Science*, pages 256–270. Springer, 2003.

[27] Daniel A. Spielman and Shang-Hua Teng. Smoothed analysis of termination of linear programming algorithms. *Mathematical Programming, Series B*, 97(1–2):375–404, 2003.

[28] Daniel A. Spielman and Shang-Hua Teng. Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. *Journal of the ACM*, 51(3):385–463, 2004.

# A Experimental Results

For all $n \in \{1\,000, 2\,000, \ldots, 500\,000\}$ and $p \in \{0.1, 0.25\}$, we randomly performed $5\,000$ $p$-partial permutations on $\sigma^n_{\text{sort}}$. Then we estimated height-perm$_p(\sigma^n_{\text{sort}})$ as the average height of the trees generated by the sequences thus obtained. Figure 5 shows the results compared to $1.8 \cdot (1-p)\sqrt{n/p}$.

We performed the same experiment for all $n \in \{100\,000, 200\,000, \ldots, 500\,000\}$ and $p \in \{0.01, 0.02, \ldots, 0.99\}$. Figure 6 shows the results, again compared to $1.8 \cdot (1 - p)\sqrt{n/p}$.

These experiments lead us to Conjecture 9.2, i.e. that height-perm$_p(\sigma^n_{\text{sort}})$ is roughly $1.8 \cdot (1-p)\sqrt{n/p}$. Together with Conjecture 9.1, this would lead to considerably better upper bounds for the smoothed height of binary search trees under partial permutations.

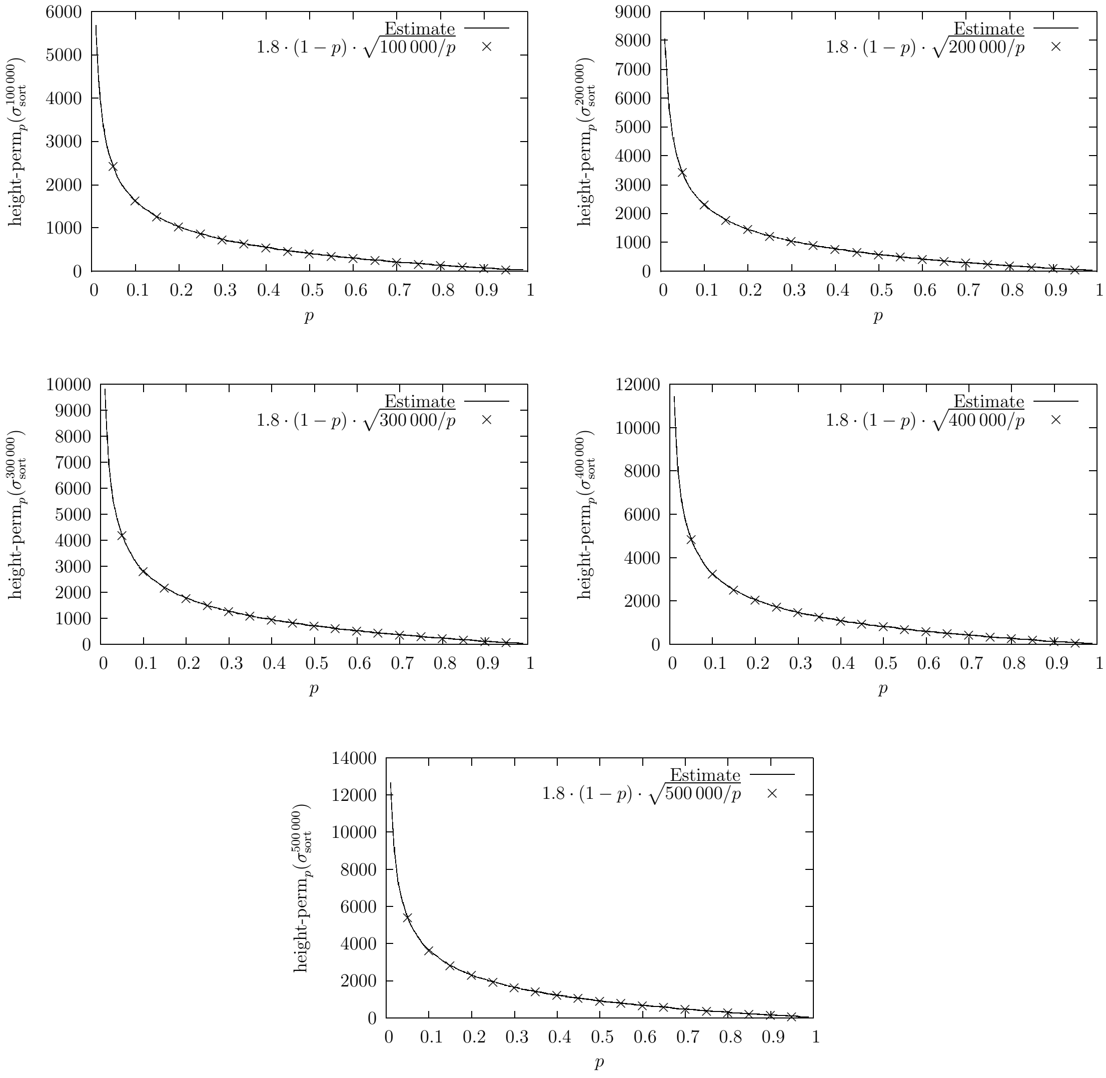


Figure 6: Experimental data for $\sigma_{\text{sort}}^n$ for $n \in \{100\,000, \ldots, 500\,000\}$ and $p \in \{0.01, 0.02, \ldots, 0.99\}$ compared to $1.8 \cdot (1-p) \cdot \sqrt{n/p}$.