

Query Complexity of Inversion Minimization on Trees

Ivan Hu*

Dieter van Melkebeek†

Andrew Morgan‡

Abstract

We consider the following computational problem: Given a rooted tree and a ranking of its leaves, what is the minimum number of inversions of the leaves that can be attained by ordering the tree? This variation of the well-known problem of counting inversions in arrays originated in mathematical psychology. It has the evaluation of the Mann–Whitney statistic for detecting differences between distributions as a special case.

We study the complexity of the problem in the comparison-query model, the standard model for problems like sorting, selection, and heap construction. The complexity depends heavily on the shape of the tree: for trees of unit depth, the problem is trivial; for many other shapes, we establish lower bounds close to the strongest known in the model, namely the lower bound of $\log_2(n!)$ for sorting n items. For trees with n leaves we show, in increasing order of closeness to the sorting lower bound:

- (a) $\log_2((\alpha(1-\alpha)n!)) - O(\log n)$ queries are needed whenever the tree has a subtree that contains a fraction α of the leaves. This implies a lower bound of $\log_2(\binom{k}{k+1} n!) - O(\log n)$ for trees of degree k .
- (b) $\log_2(n!) - O(\log n)$ queries are needed in case the tree is binary.
- (c) $\log_2(n!) - O(k \log k)$ queries are needed for certain classes of trees of degree k , including perfect trees with even k .

The lower bounds are obtained by developing two novel techniques for a generic problem Π in the comparison-query model and applying them to inversion minimization on trees. Both techniques can be described in terms of the Cayley graph of the symmetric group with adjacent-rank transpositions as the generating set, or equivalently, in terms of the edge graph of the permutahedron, the polytope spanned by all permutations of the vector $(1, 2, \dots, n)$. Consider the subgraph consisting of the edges between vertices with the same value under Π . We show that the size of any decision tree for Π must be at least:

- (i) the number of connected components of this subgraph, and
- (ii) the factorial of the average degree of the complementary subgraph, divided by n .

Lower bounds on query complexity then follow by taking the base-2 logarithm. Technique (i) represents a discrete analog of a classical technique in algebraic complexity and allows us to establish (c) and a tight lower bound for counting cross inversions, as well as unify several of the known lower bounds in the comparison-query model. Technique (ii) represents an analog of sensitivity arguments in Boolean complexity and allows us to establish (a) and (b).

Along the way to proving (b), we derive a tight upper bound on the maximum probability of the distribution of cross inversions, which is the distribution of the Mann–Whitney statistic in the case of the null hypothesis. Up to normalization the probabilities alternately appear in the literature as the coefficients of polynomials formed by the Gaussian binomial coefficients, also known as Gaussian polynomials.

*University of Wisconsin-Madison.

†University of Wisconsin-Madison.

‡University of Wisconsin-Madison.

1 Overview

The result of a hierarchical cluster analysis on a set X of items can be thought of as an unordered rooted tree T with leaf set X . To visualize the tree, or to spell out the classification in text, one needs to decide for every internal node of T in which order to visit its children. Figure 1a represents an example of a classification of eight body parts from the psychology literature [Deg82]. It is obtained by repeatedly clustering nearest neighbors where the distance between two items is given by the number of people in a survey who put the items into different classes [Mil69]. The ordering of the resulting binary tree in Figure 1a is the output produced by a particular implementation of the clustering algorithm.




Figure 1: Classification of body parts

Another ordering is given in Figure 1b; black marks the nodes whose children have been swapped from the ordering in Figure 1a. Figure 1b has the advantage over Figure 1a that the leaves now appear in an interesting global order, namely head-to-toe: ear, cheek, mouth, chest, waist, thigh, knee, toe. Indeed, Figure 1b makes apparent that the anatomical order correlates perfectly with the clustering. In general, given a tree T and a ranking ρ of its leaves, one might ask “how correlated” is T with ρ ? Degerman [Deg82] suggests evaluating the orderings of T in terms of the number of inversions of the left-to-right ranking σ of the leaves with respect to the given ranking ρ , and use the minimum number over all orderings as a measure of (non)correlation.

DEFINITION 1.1. (RANKING, INVERSION, $\text{Inv}(\cdot)$) A ranking ρ of a set X of n items is a bijection from X to $[n]$. Given two rankings σ and ρ , an inversion of σ with respect to ρ is a pair of items $x_1, x_2 \in X$ such that $\rho(x_1) < \rho(x_2)$ but $\sigma(x_1) > \sigma(x_2)$. The number of inversions is denoted by $\text{Inv}_\rho(\sigma)$. An inversion in an array A of values is an inversion of σ with respect to ρ where σ denotes the ranking by array index and ρ the ranking by value; in this setting we write $\text{Inv}(A)$ for $\text{Inv}_\rho(\sigma)$.

The minimum number of inversions can be used to compare the quality of different trees T for a given ranking ρ , or of different rankings ρ for a given tree T . This mimics the use of the number of inversions in applications like collaborative filtering in recommender systems, rank aggregation for meta searching the web, and Kendall’s test for dependencies between two random variables. In particular, the Mann–Whitney test for differences between random variables can be viewed as a special case of our optimization problem. The test is widely used because of its nonparametric nature, meaning that no assumptions need to be made about the distribution of the two variables; the distribution of the statistic in the case of the null hypothesis (both variables have the same distribution) is always the same. The test achieves this property by only considering the relative order of the samples. It takes a sequence A of a samples from a random variable Y , a sequence B of b samples from another random variable Z , and computes the statistic $U \doteq \min(\text{XInv}(A, B), \text{XInv}(B, A))$ that is the minimum of the number $\text{XInv}(A, B)$ of cross inversions from A to B , and vice versa.

DEFINITION 1.2. (CROSS INVERSIONS, $\text{XInv}(\cdot, \cdot)$) Let ρ be a ranking of X , and $A, B \subseteq X$. A cross inversion from A to B with respect to ρ is a pair $(x_1, x_2) \in A \times B$ that is out of order with respect to ρ , i.e., such that $\rho(x_1) > \rho(x_2)$. The number of cross inversions is denoted by $\text{XInv}_\rho(A, B)$. For two arrays A and B of values, a cross inversion from A to B is a cross inversion from the set of entries in A to the set of entries in B where ρ denotes the ranking by value; in this setting we write $\text{XInv}(A, B)$ for $\text{XInv}_\rho(A, B)$.




Figure 2: Mann–Whitney instance

The statistic U coincides with the optimum value of our optimization problem on input the tree T in Figure 2, where the leftmost a leaves correspond to the samples A , the rightmost b leaves to the samples B , and the ranking ρ to the value order of the combined $a + b$ samples.

We mainly study the value version of our optimization problem, which we denote by MinInv .

DEFINITION 1.3. (INVERSION MINIMIZATION ON TREES, $\text{MinInv}(\cdot, \cdot)$, Π .) *Inversion minimization on trees is the computational problem with the following specification:*

Input: A rooted tree T with leaf set X of size n , and a ranking ρ of X .

Output: $\text{MinInv}(T, \rho)$, the minimum of $\text{Inv}_\rho(\sigma)$ over all possible orderings of T , where σ denotes the left-to-right ranking of X induced by the ordering of T .

For any fixed tree T with leaf set X , we use the short-hand Π_T to denote the computational problem that takes as input a ranking ρ of X and outputs $\text{MinInv}(T, \rho)$.

Degerman [Deg82] observes that the ordering at each internal node can be optimized independently in a greedy fashion. In the setting of binary trees, for each node v , we can count the cross inversions from the leaves in the left subtree of v to the leaves in the right subtree of v . Between the two possible orderings of the children of a node v , we choose the one that yields the smaller number of cross inversions. Based on his observation, Degerman presents a polynomial-time algorithm for the case of binary trees T . A more refined implementation and analysis yields a running time of $O(d_{\text{avg}}(T) \cdot n)$, where $d_{\text{avg}}(T)$ denotes the average depth of a leaf in T . For balanced binary trees the running time becomes $O(n \log n)$. All of this can be viewed as variants of the well-known $O(n \log n)$ divide-and-conquer algorithm for counting inversions in arrays of length n .

For trees of degree $\deg(T) > 2$, the local greedy optimization at each internal node becomes more complicated, as there are many ways to order the children of each internal node. Exhaustive search results in a running time of $O((\deg(T)! + \deg(T) \cdot d_{\text{avg}}(T)) \cdot n)$, which can be improved to $O((\deg(T)^2 2^{\deg(T)} + \deg(T) \cdot d_{\text{avg}}(T)) \cdot n)$ using dynamic programming. The problem is closely related to the classical problem of minimum arc feedback set, and becomes NP-hard without any constraints on the degree. We refer to the full version of the paper for more details.

Query complexity. Rather than running time in the Turing machine model, our focus lies on query complexity in the comparison-query model. There we can only access the ranking $\rho : X \rightarrow [n]$ via queries of the form: Is $\rho(x_1) < \rho(x_2)$? For any fixed tree T , we want to determine the minimum number of queries needed to solve the problem.

The comparison-query model represents the standard model for analyzing problems like sorting, selection, and heap construction. Sorting represents the hardest problem in the comparison-query model as it is tantamount to knowing the entire ranking ρ . Its query complexity has a well-known information-theoretic lower bound of $\log_2(n!) = n \log_2(n/e) + \frac{1}{2} \log_2(n) + O(1)$. Standard algorithms such as mergesort and heapsort yield an upper bound of $\log_2(n!) + O(n)$, which has been improved to $\log_2(n!) + o(n)$ recently [Ser21]. We refer to Section 2 for an overview of results and techniques for lower bounds in the model.

Information theory only yields a very weak lower bound on the query complexity of inversion minimization on trees: $\log_2 \binom{n}{2} = 2 \log_2(n) - O(1)$. The complexity of the problem critically depends on the shape of the tree T and can be significantly lower than the one for sorting. For starters, the problem becomes trivial for trees of depth one as their leaves can be arranged freely in any order. More precisely, the trees T for which the answer

is identically zero, irrespective of the ranking ρ , are exactly those such that all root-to-leaf paths have only the root in common.

Arguably, the simplest nontrivial instances of inversion minimization are for trees T of the Mann–Whitney type in Figure 2 with $a = 1$ and $b = n - 1$. Depending on the rank r of the isolated leaf, an optimal ordering of T is either the left or the right part in Figure 3, where the label of each leaf is its rank under ρ .



Figure 3: Rank instance

As the ordering on the left has $r - 1$ inversions and the one on the right $n - r$, the answer is $\min(r - 1, n - r)$. Thus, this instance of inversion minimization on trees is essentially equivalent to rank finding, which has query complexity exactly $n - 1$.

Results. We prove that for many trees T , inversion minimization on T is nearly as hard as sorting. First, we exhibit a common structure that guarantees high complexity, namely a subtree that contains a fairly balanced fraction of the leaves. We make use of the following notation.

DEFINITION 1.4. (LEAF SET, $L(\cdot)$, AND SUBTREE) For a tree T , the leaf set of T , denoted $L(T)$, is the set of leaves of T . For a node v to T , T_v denotes the subtree of T rooted at v .

The quantitative statement references the gamma function Γ , which is a proxy for any convex real function that interpolates the factorial function on the positive integers. More precisely, we have that $\Gamma(n + 1) = n!$ for every integer $n \geq 1$.

THEOREM 1.1. (LOWER BOUND FOR GENERAL TREES) Let T be a tree with n leaves, and v a node with $|L(T_v)| = \ell$. The query complexity of inversion minimization on T is at least $\log_2(\Gamma(\frac{\ell(n-\ell)}{n} + 1))$. In particular, the complexity is at least $\log_2(\Gamma(\frac{k}{(k+1)^2} \cdot n + 1))$ where k denotes the degree of T .

For trees of constant degree, Theorem 1.1 yields a lower bound that is as strong as the one for sorting up to a constant multiplicative factor. For the important case of binary trees (like the classification trees from the motivating example), we obtain a lower bound that is only a logarithmic additive term shy of the lower bound for sorting.

THEOREM 1.2. (LOWER BOUND FOR BINARY TREES) For binary trees T with n leaves, the query complexity of inversion minimization on T is at least $\log_2(n!) - O(\log n)$.

The logarithmic loss can be reduced to a constant for certain restricted classes of trees. The full statement is somewhat technical. First, it assumes that the tree has no nodes of degree 1. This is without loss of generality, as we can short-cut all degree-1 nodes in the tree without affecting the minimum number of inversions. For example, trivial trees for inversion minimization have depth 1 without loss of generality. Second, the strength of the lower bound depends on the maximum size of a leaf child set, defined as follows.

DEFINITION 1.5. (LEAF CHILD SET, $LC(\cdot)$) The leaf child set $LC(v)$ of a vertex v in a tree T is the set $LC(v)$ of all the children of v that are leaves in T .

Most importantly, the result requires certain fragile parity conditions to hold. That said, there are interesting classes satisfying all requirements, and the bounds are very tight.

THEOREM 1.3. (LOWER BOUND FOR RESTRICTED CLASSES) Let T be a tree without nodes of degree 1 such that the leaf child sets have size at most k , at most one of them is odd, and if there exists an odd one, say $LC(v^*)$, then all ancestors of v^* have empty leaf child sets. The query complexity of inversion minimization on T is at least $\log_2(n!) - O(k \log k)$. In particular, the lower bound applies to:

- perfect trees of even degree k , and
- full binary ($k = 2$) trees with at most one leaf without a sibling leaf.

Recall that a tree of degree k is full if every node has degree 0 or k . It is perfect if it is full and all leaves have the same depth.

For the Mann–Whitney statistic, Theorem 1.1 provides an $\Omega(n \log n)$ lower bound for balanced instances, i.e., when a and b are $\Theta(n)$. For unbalanced instances there is a more efficient way to count cross inversions and thus evaluate the statistic: Sort the smaller of the two sides, and then do a binary search for each item of the larger side to find its position within the sorted smaller side so as to determine the number of cross inversions that it contributes. For $a \leq b$ the approach makes $b \log_2(a) + O(a \log a)$ comparisons. We establish a lower bound that shows the approach is optimal up to a constant multiplicative factor.

THEOREM 1.4. (LOWER BOUND FOR COUNTING CROSS INVERSIONS) *Counting cross inversions from a set A of size a to a set B of size $b \geq a$ with respect to a ranking ρ of $X \doteq A \sqcup B$ requires $\Omega((a + b) \log(a))$ queries in the comparison-query model, as does inversion minimization on the tree of Figure 2.*

Techniques. We obtain our results by developing two new query lower bound techniques for generic problems Π in the comparison-query model, and then instantiating them to the problem Π_T of inversion minimization on a fixed tree T . Both techniques follow the common pattern of lower bounding the number of distinct execution traces that any algorithm for Π needs to have.

DEFINITION 1.6. (EXECUTION TRACE, COMPLEXITY MEASURES $D(\cdot)$ AND $Q(\cdot)$) *Consider an algorithm A for a problem Π in the comparison-query model. An execution trace of A is the sequence of comparisons that A makes on some input ρ , as well as the outcomes of the comparisons. The complexity $D(\Pi)$ is the minimum over all possible algorithms for Π of the number of distinct traces the algorithm has over the set of all inputs ρ . The complexity $Q(\Pi)$ is the minimum, over all possible algorithms for Π of the maximum number of comparisons that the algorithm makes over the set of all inputs ρ .*

The complexity measure Q is what we refer to as query complexity. Since the maximum number of queries that an algorithm A makes is at least the base-2 logarithm of the number of execution traces, we have that $Q(\Pi) \geq \log_2(D(\Pi))$. Note that, in order to avoid confusion with the tree T specifying an instance of inversion minimization, we refrain from the common terminology of decision trees in the context of the complexity measure D . In those terms, we lower bound the number of leaves of any decision tree for Π , and use the fact that the depth of this binary decision tree is at least the base-2 logarithm of the number of leaves.

Both techniques proceed by considering the effect on the output of perturbations to the input ranking ρ that are hard for queries to observe. More specifically, we consider the following perturbations:

DEFINITION 1.7. (ADJACENT-RANK TRANSPOSITION, AFFECTED ITEMS) *An adjacent-rank transposition is a permutation τ of $[n]$ of the form $\tau = (r, r + 1)$, where $r \in [n - 1]$ and n denotes the number of items. Given τ and a ranking $\rho : X \rightarrow [n]$, the affected items are the two elements $x \in X$ for which $\tau(\rho(x)) \neq \rho(x)$, i.e., the items with ranks r and $r + 1$ under ρ .*

As with any permutation of the set of ranks, the effect of τ on a ranking ρ is the ranking $\tau\rho$. Adjacent-rank transpositions are the least noticeable perturbations one can apply to a ranking in the following sense: If two rankings differ by an adjacent-rank transposition, then the only query that distinguishes them is the query that compares the affected items.

Sensitivity. Our first technique turns this observation around to obtain a lower bound on query complexity. We adopt the terminology of sensitivity from Boolean query complexity.

DEFINITION 1.8. (SENSITIVITY, AVERAGE SENSITIVITY, $s(\cdot)$) *Let Π be a computational problem in the comparison-query model on a set X of items. For a fixed ranking ρ and adjacent-rank transposition τ , we say that Π is sensitive to τ at ρ if $\Pi(\rho) \neq \Pi(\tau\rho)$. The sensitivity of Π at ρ is the number of adjacent-rank transpositions τ such that Π is sensitive to τ at ρ . The average sensitivity of Π , denoted $s(\Pi)$, is the average sensitivity of Π at ρ when ρ is drawn uniformly at random from all rankings of X .*

On input a ranking ρ , any algorithm for Π needs to make a number of queries that is at least the sensitivity of Π at ρ . Indeed, consider an adjacent-rank transposition τ to which Π is sensitive at ρ . If the algorithm does not make the query that compares the affected items, then it must output the same answer on input $\tau\rho$ as on input ρ . Since the value of Π differs on both inputs, this means the algorithm makes a mistake on at least one of the two. It follows that the average number of queries that any algorithm for Π makes is at least the average sensitivity $s(\Pi)$. A fortiori, $Q(\Pi) \geq s(\Pi)$.

As sensitivity cannot exceed $n-1$, the best lower bound on query complexity that we can establish based on the above basic observation alone, is $n-1$. The following improvement yields a the lower bound $D(\Pi) \geq n!/n = (n-1)!$, and therefore $Q(\Pi) \geq \log_2(n!) - \log_2(n)$ for problems Π of maximum average sensitivity $s(\Pi) = n-1$. The argument hinges on an efficient encoding of rankings that share the same execution trace. See Section 3 for more details.

LEMMA 1.1. (SENSITIVITY LEMMA) *For any problem Π in the comparison-query model with n items, $D(\Pi) \geq \Gamma(s(\Pi) + 2)/n$.*

The lower bound for general trees T in Theorem 1.1 and the strengthening for binary trees in Theorem 1.2 follow from corresponding lower bounds on the average sensitivity $s(\Pi_T)$. Theorem 1.1 only requires a short analysis to establish the sensitivity lower bound needed for the application of the Sensitivity Lemma; this illustrates the power of the lemma and of the lower bound technique. Theorem 1.2 requires a more involved sensitivity analysis, but then yields a very tight lower bound. Owing to the average-case nature of the underlying measure, the technique also exhibits some degree of robustness. For the particular problem of inversion minimization on trees, we show that small changes to the tree T do not affect the average sensitivity $s(\Pi_T)$ by much. See Section 4 and Section 5.

For sorting, counting inversions, and inversion parity, the average sensitivity reaches its maximum value of $n-1$, and Lemma 1.1 recovers the standard lower bounds up to a small loss. In contrast, for selection, the average sensitivity equals 1 for ranks 1 and n , and 2 for other ranks, so the bound from Lemma 1.1 is no good. This reflects that, just like in the Boolean setting, (average) sensitivity is sometimes too rough of a measure and not always capable of proving strong lower bounds. Our second technique looks at a more delicate structural aspect, which enables it to sometimes yield stronger lower bounds.

Permutahedron graph. Before introducing our second technique, we cast our first technique in graph theoretic terms. In fact, both our techniques can be expressed naturally in subgraphs of the graph with the rankings as vertices and adjacent-rank transpositions as edges. The latter graph can be viewed as the Cayley graph of the symmetric group with adjacent-rank transpositions as the generating set. It is also the edge graph of the permutahedron, the convex polytope spanned by all permutations of the vertex $(1, 2, \dots, n)$ in \mathbb{R}^n . The permutahedron resides inside the hyperplane where the sum of the coordinates equals $\binom{n}{2}$, has positive volume inside that hyperplane, and can thus be represented naturally in dimension $n-1$; see Figure 4 for a rendering of the instance with $n=4$ [Epp07].

We think of coloring the vertices of the permutahedron with their values under Π and make use of the subgraph with the same vertex set but only containing the monochromatic edges, i.e., the edges whose end points have the same value under Π . We also consider the complementary subgraph containing all bichromatic edges.

DEFINITION 1.9. (PERMUTAHEDRON GRAPH, $G(\cdot)$, $\bar{G}(\cdot)$) *Let Π be a computational problem in the comparison-query model on a set X of items. The permutahedron graph of Π , denoted $G(\Pi)$, has the rankings of X as vertices, and an edge between two rankings ρ_1 and ρ_2 if $\Pi(\rho_1) = \Pi(\rho_2)$ and there exists an adjacent-rank transposition such that $\rho_2 = \tau\rho_1$. The complementary permutahedron graph of Π , denoted $\bar{G}(\Pi)$, is defined similarly by replacing the condition $\Pi(\rho_1) = \Pi(\rho_2)$ by its complement, $\Pi(\rho_1) \neq \Pi(\rho_2)$.*




Figure 4: Permutahedron for $n=4$ items

Our first technique looks at degrees in the complementary permutahedron graph $\overline{G}(\Pi)$, and more specifically at the *average degree* $\text{deg}_{\text{avg}}(\overline{G}(\Pi)) \doteq \mathbb{E}(\text{deg}_{\overline{G}(\Pi)}(\rho))$, where the expectation is with respect to a uniform choice of the ranking ρ . Our second technique looks at the connected components of the permutahedron graph $G(\Pi)$.

Connectivity. Our second technique is reminiscent of a result in algebraic complexity theory, where the number of execution traces of an algorithm for a problem Π in the algebraic comparison-query model is lower bounded in terms of the number of connected components that Π induces in its input space \mathbb{R}^n [BO83]. In the comparison-query setting, we obtain the following lower bound.

LEMMA 1.2. (CONNECTIVITY LEMMA) *For any problem Π in the comparison-query model, $D(\Pi)$ is at least the number of connected components of $G(\Pi)$.*

The Connectivity Lemma allows for a simple and unified exposition of many of the known lower bounds. For counting inversions and inversion parity the argument goes as follows. Every adjacent-rank transposition changes the number of inversions by exactly one (up or down), and therefore changes the output of Π , so all $n!$ vertices in $G(\Pi)$ are isolated. This means that any algorithm for Π actually needs to sort and has to make at least $\log_2(n!)$ queries. See Section 7 for a proof of the Connectivity Lemma and more applications to classical problems, including the $\Omega(n)$ lower bound for median finding.

The Connectivity Lemma also enables us to establish strong lower bounds for inversion minimization on special types of trees T , namely those of Theorem 1.2 and the Mann–Whitney instances in Theorem 1.4, closely related to counting inversions. Both theorems involve an analysis of the size of the connected component of a random ranking ρ in $G(\Pi_T)$, and Theorem 1.2 uses the delicate parity conditions of its statement to keep $G(\Pi_T)$ as sparse as possible. See Section 8 for more details.

The Mann–Whitney setting illustrates well the relative power of our techniques. In the Mann–Whitney instances of inversion minimization, the leaves are naturally split between a subtree containing a of them and a subtree containing b of them. The argument behind Theorem 1.1 yields a lower bound of $\frac{ab}{a+b}$ on the sensitivity $s(\Pi_T)$. The true sensitivity is just $O(1)$ below the one for counting cross inversions, which is $\frac{2ab}{a+b}$. The resulting lower bounds on the query complexity in case $a \leq b$ are $\Theta(a \log a)$, which roughly account for sorting the smaller side but not for the $b \log_2(a)$ comparisons used in the subsequent binary searches for counting cross inversions. Our approach based on the Connectivity Lemma yields a lower bound that includes both terms. On the other hand, it is easier to estimate and obtain the lower bound via the Sensitivity Lemma than to argue the query lower bound via the Connectivity Lemma or from scratch.

Other modes of computation. We stated our lower bounds for the standard, deterministic mode of computation. Both of our techniques provide lower bounds for the number of distinct execution traces that are needed to cover all input rankings, irrespective of whether these execution traces derive from a single algorithm. Such execution traces can be viewed as certificates or witnesses for the value of Π on a given input ρ , or as valid execution traces of a *nondeterministic* algorithm for Π . We define the minimum number of traces needed to cover all input rankings for a problem Π as the nondeterministic complexity of Π and denote it by $N(\Pi)$, along the lines of the Boolean setting [JRSW99]. All of our lower bounds on $D(\Pi)$ actually hold for $N(\Pi)$. See Remark 3.1 and Remark 7.1 for further discussion.

Since *randomized algorithms with zero error* are also nondeterministic algorithms, all of our lower bounds apply verbatim to the former mode of computation, as well. As for *randomized algorithms with bounded error*, we argue in Section 6 that our lower bounds on the query complexity of inversion minimization on trees that follow from the Sensitivity Lemma carry over modulo a small loss in strength. We do so by showing generically that high average sensitivity implies high query complexity against such algorithms.

The fact that our techniques yield lower bounds on $N(\Pi)$ and not just $D(\Pi)$ also explains why our approaches sometimes fail. For example, for the problem Π of finding the minimum of n items, a total of n certificates suffice and are needed, namely one for each possible item being the minimum. This means that our techniques cannot give a lower bound on the query complexity of Π that is better than $\log_2(n)$. In contrast, as reviewed in Section 2, $D(\Pi) = 2^{n-1}$ and the number of queries needed is $n - 1$.

Cross-inversion distribution. As a technical result in the sensitivity analysis for inversion minimization on binary trees (Theorem 1.2), we need a strong upper bound on the probability that the number of cross inversions $X\text{Inv}_\rho(A, B)$ takes on any particular value when the ranking ρ of the set $X = A \sqcup B$ is chosen uniformly at random. This is the distribution of the Mann–Whitney statistic under the null hypothesis. Mann and Whitney [MW47] argued that it converges to a normal distribution with mean $\mu = ab/2$ and variance $\sigma^2 = ab(a+b+1)/12$

as $a \doteq |A|$ and $b \doteq |B|$ grow large. Since the normal distribution has a maximum density of $1/(\sqrt{2\pi}\sigma)$, their result suggests that the maximum of the underlying probability distribution is $O(1/\sigma) = O(1/\sqrt{ab(a+b+1)})$. Takács [Tak86] managed to formally establish such a bound for all pairs (a, b) with $|a - b| = O(\sqrt{a+b})$, Stanley and Zanello [SZ16] for all pairs (a, b) with $\min(a, b)$ bounded, and Melczer, Panova, and Pemantle [MPP20] for all pairs (a, b) with $|a - b| \leq \alpha \cdot (a + b)$ for some constant $\alpha < 1$. However, these results do not cover all regimes and leave open a single bound of the same form that applies to all pairs (a, b) , which is what we need for Theorem 1.2. We establish such a bound in the full version of the paper. The counts of the rankings ρ with a particular value for $\text{XInv}_\rho(A, B)$ appear as the coefficients of the *Gaussian polynomials*. Our bound can be stated equivalently as a bound on those coefficients.

Organization. We have organized the material so as to provide a shortest route to a full proof of Theorem 1.1. Here are the sections needed for the different main results:

- Theorem 1.1 (lower bound for general trees): 3, 4.
- Theorem 1.2 (lower bound for binary trees): 3, 5.
- Theorem 1.3 (lower bound for restricted classes): 7, 8 up to 8.2 inclusive.
- Theorem 1.4 (lower bound for counting cross inversions): 7, 8 but not 8.1 nor 8.2.

In Section 2, we provide some background on known lower bounds in the comparison-query model, several of which are unified by the Sensitivity Lemma and Connectivity Lemma.

2 The Comparison-Query Model

In this section we provide an overview of known results and techniques for lower bounds in the comparison-query model. This section can be skipped without a significant loss in continuity.

Tight bounds have been established for problems like sorting, selection, and heap construction.

- We already discussed the central problem of sorting in Section 1.
- In *selection* we are told a rank r , and must identify the item with rank r . The query complexity is known to be $\Theta(n)$ [BFP⁺73, DZ99, DZ01]. There is also *multiple selection*, in which one is given multiple ranks r_1, \dots, r_k , and must identify each of the corresponding items. The query complexity of multiple selection is likewise known up to a $\Theta(n)$ gap between the upper and lower bounds [KMMS05].
- In *heap construction* we must arrange the items as nodes in a complete binary tree such that every node has a rank no larger than its children. The query complexity is known to be $\Theta(n)$.

All the problems above can be cast as instantiations of a general framework known as *partial order production* [Sch76]. Here, in addition to query access to the ranking ρ of the items, we are given n slots and regular access to a partial order $<_{\text{slot}}$ on the slots. The objective is to put each item into a slot, one item per slot, so that whenever two slots, s_1 and s_2 , are related by $s_1 <_{\text{slot}} s_2$, we also have $\rho(s_1) < \rho(s_2)$. Sorting coincides with the case where $<_{\text{slot}}$ is a total order. In selection of rank r , there is a designated slot s^* , and there are exactly $r - 1$ slots s with $s <_{\text{slot}} s^*$ and exactly $n - r$ slots s with $s^* <_{\text{slot}} s$; there are no other relations in $<_{\text{slot}}$. Multiple selection is similar. For heap construction, $<_{\text{slot}}$ matches the complete binary tree arrangement.

There is a generic lower bound for partial order production, the information-theoretic limit. For each way of putting items into slots, the number of input rankings ρ for which that way is a correct answer is bounded by $e(<_{\text{slot}})$, the number of ways to extend $<_{\text{slot}}$ to a total order. Therefore, there must be at least $n!/e(<_{\text{slot}})$ distinct execution traces. Since each execution trace is determined by the outcomes of its queries, and each query has only two outcomes, we conclude that $\lambda(<_{\text{slot}}) \doteq \log_2(n!/e(<_{\text{slot}}))$ queries are necessary to solve partial order production. Complementing this lower bound there exists an upper bound of $(1 + o(1)) \cdot \lambda(<_{\text{slot}}) + O(n)$ queries [CFJ⁺10]. One may assume without loss of generality the relationship $\lambda(<_{\text{slot}}) \geq n - 1$, in which case $O(\lambda(<_{\text{slot}}))$ queries always suffices. Thus, the complexity of partial order production is $\Theta(\lambda)$.

Not every problem of interest in the comparison model is an instance of partial order production. Here are a few examples.

- In *rank finding* there is a designated item x^* , and we have to compute its rank. The rank can be computed by comparing x^* with each of the $n - 1$ other items. Any combination of less than $n - 1$ queries leaves at least one item of which the relative ranking with x^* remains undetermined. Thus, the query complexity is exactly $n - 1$.
- In *counting inversions* the items are arranged in some known order σ and the objective is to count the number of inversions of σ with respect to ρ . As we reviewed in Section 1, counting inversions has exactly the same query complexity as sorting.
- The problem of *inversion parity* is the same as counting inversions except that one need only count the number of inversions modulo 2. This problem, as well as counting inversions modulo m for any integer $m > 1$, also has exactly the same complexity as sorting.

For each of the three problems above, information theory does not provide a satisfactory lower bound. For example, in the inversion parity problem there are only two possible outputs, which yields a lower bound of $\log_2(2) = 1$. It so happens that for each of the preceding three examples, the query complexity is known quite precisely; however, the known arguments are rather problem-specific.

Inversion minimization on trees is another example that does not fit the framework of partial order generation, and for which information theory only yields a weak lower bound: $\log_2 \binom{n}{2} = 2 \log_2(n) - \Theta(1)$. In contrast to the above examples, a strong lower bound does not seem to follow from a simple ad-hoc argument nor from a literal equivalence to sorting.

3 Sensitivity Lemma

In this section we develop Lemma 1.1. We actually prove a somewhat stronger version.

LEMMA 3.1. (STRONG SENSITIVITY LEMMA) *Consider an algorithm A in the comparison-based model with n items, color each vertex of the permutahedron with its execution trace under A , and let \overline{H} denote the subgraph with the same vertex set but only containing the bichromatic edges. The number of distinct execution traces of A is at least $g(\deg_{\text{avg}}(\overline{H}) + 1)/n$, where $g : [1, \infty) \rightarrow \mathbb{R}$ is any convex function with $g(x) = x!$ for $x \in [n]$.*

The Sensitivity Lemma follows from Lemma 3.1 because the coloring with execution traces of an algorithm A for Π is a refinement of the coloring with Π , so every edge of the permutahedron that is bichromatic under Π is also bichromatic under A , and

$$s(\Pi) \doteq \mathbb{E}(\deg_{\overline{G}(\Pi)}(\rho)) \leq \mathbb{E}(\deg_{\overline{H}}(\rho)) \doteq \deg_{\text{avg}}(\overline{H}).$$

Provided g is nondecreasing, it follows that $D(\Pi) \geq g(\deg_{\text{avg}}(\overline{H}) + 1)/n \geq g(s(\Pi) + 1)/n$.

In the Sensitivity Lemma we set $g(x) = \Gamma(x + 1)$. An optimal (but less elegant) choice for g is the piece-wise linear function that interpolates the prescribed values at the integral points in $[n]$, namely

$$g(x) \doteq (x - [x]) \cdot ([x]!) + (1 - (x - [x])) \cdot ([x]!).$$

For the proof of Lemma 3.1 we take intuition from a similar result in the Boolean setting [O'D14, Exercise 8.43], where the hypercube plays the role of the permutahedron in our setting.

FACT 3.1. *Let A be a query algorithm on binary strings of length n . Color each vertex of the n -dimensional hypercube by its execution trace under A , and let \overline{H} denote the subgraph with the same vertex set but only containing the bichromatic edges. Then the number of distinct execution traces is at least $2^{\deg_{\text{avg}}(\overline{H})}$.*

One way to argue Fact 3.1 is to think of assigning a weight $w(x)$ to each $x \in \{0, 1\}^n$ so as to maximize the total weight on all inputs, subject to the constraint that the total weight on each individual execution trace is at most 1. Then the number of distinct execution traces must be at least the sum of all the weights. If the weight only depends on the degree, i.e., if we can write $w(x) = f(\deg_{\overline{H}}(x))$ for some function $f : [0, \infty) \rightarrow \mathbb{R}$, then we can lower bound the number k of distinct execution traces as follows:

$$(3.1) \quad k \geq \sum_x w(x) = \sum_x f(\deg_{\overline{H}}(x)) \geq 2^n \cdot f(\mathbb{E}(\deg(x))) = 2^n \cdot f(\deg_{\text{avg}}(\overline{H})),$$

where the last inequality holds provided f is convex.

In the Boolean setting, the set R of inputs $x \in \{0, 1\}^n$ with a particular execution trace forms a subcube of dimension $n - \ell$, where ℓ denotes the length of the execution trace, i.e., the number of queries. Each $x \in R$ has degree ℓ in H ; this is because a change in a single queried position results in a different execution trace, and a change in another position does not. Therefore, a natural choice for the weight of $x \in R$ is $w(x) = f(\ell)$ where $f(x) = 1/2^{n-\ell}$. It satisfies the constraint that the total weight on R is (at most) one, and f is convex. We conclude by (3.1) that the number of distinct execution traces is at least $2^n \cdot f(\deg_{\text{avg}}(\overline{H})) = 2^{\deg_{\text{avg}}(\overline{H})}$, as desired.

Proof. [Proof of Lemma 3.1] Let k denote the number of distinct execution traces of A , and let R_1, \dots, R_k denote the corresponding sets of rankings. Following a similar strategy, we want to find a convex function $f : [0, \infty) \rightarrow \mathbb{R}$ such that the weight function $w(\rho) = f(\deg_{\overline{H}}(\rho))$ does not assign weight more than 1 to any one of the sets R_i . The following claim, to be proven later, is the crux of this.

CLAIM 3.1. *Let R denote the set of all rankings ρ that follow a particular execution trace on A , and let $d \in \{0, \dots, n - 1\}$. The number of rankings $\rho \in R$ with $\deg_{\overline{H}}(\rho) = d$ is at most $\frac{n!}{(d+1)!}$.*

Based on Claim 3.1, a natural choice for f is any convex function that satisfies $f(x) = \frac{1}{n} \frac{(x+1)!}{n!}$ for $x \in \{0, \dots, n - 1\}$. The factor of $\frac{1}{n}$ comes from the fact that there are n terms to sum together after the weights have been normalized. For every $i \in [k]$ we then have

$$\sum_{\rho \in R_i} w(\rho) = \sum_{d=0}^{n-1} |\{\rho \in R_i : \deg_{\overline{H}}(\rho) = d\}| \cdot f(d) \leq \sum_{d=0}^{n-1} \frac{1}{n} \frac{(d+1)!}{n!} \cdot \frac{n!}{(d+1)!} = \sum_{d=0}^{n-1} \frac{1}{n} = 1.$$

Similar to (3.1) we conclude

$$(3.2) \quad k \geq \sum_{i=1}^k \sum_{\rho \in R_i} w(\rho) = \sum_{\rho} w(\rho) = \sum_{\rho} f(\deg_{\overline{H}}(\rho)) \geq n! \cdot f(\mathbb{E}(\deg_{\overline{H}}(\rho))) = n! \cdot f(\deg_{\text{avg}}(\overline{H})).$$

Setting $f(x) = \frac{1}{n} \frac{g(x+1)}{n!}$ turns the requirements for f into those for g in the statement of the lemma, and yields that $k \geq n! \cdot f(\deg_{\text{avg}}(\overline{H})) = g(\deg_{\text{avg}}(\overline{H}) + 1)/n$. \square

We now turn to proving Claim 3.1. The comparisons and outcomes that constitute a particular execution trace of A can be thought of as directed edges between the items in X . We refer to the resulting digraph on the vertex set X as the comparison graph C . Since the outcomes of the comparisons are consistent with some underlying ranking, the digraph C is acyclic. The rankings in R are in one-to-one and onto correspondence with the linear orderings of the DAG C . For a given ranking $\rho \in R$, the degree $\deg_{\overline{H}}(\rho)$ equals the number of $r \in \{2, \dots, n\}$ such that swapping ranks $r - 1$ and r in ρ results in a ranking $\rho' = \tau\rho$ that is not in R , where τ denotes the adjacent-rank transposition $(r - 1, r)$. The ranking ρ' not being in R means that it is inconsistent with the combined comparisons and outcomes of the underlying execution trace, which happens exactly when there is a path in C from the item $\rho^{-1}(r - 1)$ of rank $r - 1$ in ρ to the item $\rho^{-1}(r)$ with rank r in ρ . Thus, the degree $\deg_{\overline{H}}(\rho)$ equals the number of $r \in \{2, \dots, n\}$ such that there is a path from $\rho^{-1}(r - 1)$ to $\rho^{-1}(r)$ in C . See Figure 5 for an illustration, where a squiggly edge $u \rightsquigarrow v$ denotes that there exists a path from u to v in C . We only draw squiggly edges from one position to the next, so $\deg_{\overline{H}}(\rho)$ equals the number of squiggly edges in Figure 5.




Figure 5: Ranking encoding

Our strategy is to give a *compressed encoding* of the rankings in R such that there is more compression as the number of squiggly edges increases. Our encoding is based on the well-known algorithm to compute a linear order of a DAG. Algorithm 3.1 provides pseudocode for the algorithm, which we refer to as BuildRanking.

ALGORITHM 3.1. (BUILDRANKING(C))

Input: DAG C on vertex set X

Output: ranking of X that is a linear order of C

```

1:  $T \leftarrow \emptyset$ 
2: for  $r = 1$  to  $n$  do
3:    $x \leftarrow$  arbitrary element of  $S \doteq \{v \in X \mid \text{there is no } u \in X \setminus T \text{ with } u \rightsquigarrow v \text{ in } G\}$ 
4:    $\rho^{-1}(r) \leftarrow x$ 
5:    $T \leftarrow T \cup \{x\}$ 
6: end for

```

In our formulation, BuildRanking is nondeterministic: There is a choice to make in step 3 for each $r = 1, \dots, n$. The possible executions of BuildRanking are in one-to-one and onto correspondence with the linear orders of C , and thus with the rankings in R .

Our encoding is a compressed description of how to make the decisions in BuildRanking such that the output is ρ . Note that if $\rho^{-1}(r-1) \rightsquigarrow \rho^{-1}(r)$, then the item x with rank r cannot enter the set S before iteration r . This is because before $\rho^{-1}(r-1)$ is removed from T at the end of iteration $r-1$, the edge $\rho^{-1}(r-1) \rightsquigarrow \rho^{-1}(r)$ prevents x from being in S . Thus, whenever $\rho^{-1}(r-1) \rightsquigarrow \rho^{-1}(r)$, the item $x = \rho^{-1}(r)$ is *lucky* in the sense that it gets picked in step 3 as soon as it enters the set S . In fact, the lucky items with respect to a ranking $\rho \in R$ are exactly those for which $\rho^{-1}(r-1) \rightsquigarrow \rho^{-1}(r)$ for some $r \in \{2, \dots, n\}$, as well as the item $\rho^{-1}(1)$ with rank 1. In Figure 5 the lucky items are marked black. Their number equals $\deg_{\overline{H}}(\rho) + 1$.

In order to generate a ranking ρ using BuildRanking, it suffices to know:

- (a) the lucky items (as a the set, not their relative ordering), and
- (b) the ordering of the non-lucky items (given which items they are).

This information suffices to make the correct choices in step 3 of Algorithm 3.1:

- If the set S contains a lucky item, there will be a unique lucky item in S ; pick it as the element x .
- Otherwise, pick for x the first item in the ordering of the non-lucky items that is not yet in T . Such an element will exist, and all the items that come after it in the ordering are not yet in T either.

If ρ has degree $d = \deg_{\overline{H}}(\rho)$, then there are $d + 1$ lucky items, so there are at most $\binom{n}{d+1}$ choices for (a), and at most $(n - d - 1)!$ choices for (b), resulting in a total of at most $\binom{n}{d+1} \cdot (n - d - 1)! = \frac{n!}{(d+1)!}$ choices. This proves Claim 3.1.

REMARK 3.1. *Suppose we allow an algorithm A to have multiple valid execution traces on a given input ρ , and let R_i denote the set of rankings on which the i -th execution trace is valid. The proof of Claim 3.1 carries over as it considers individually sets R_i , and only depends on the DAG that the comparisons in R_i induce. The rest of the proof of Lemma 3.1 carries through modulo the first equality in (3.2), which no longer holds as the sets R_i may overlap. However, the equality can be replaced by the inequality \geq , which does hold and is sufficient for the argument. This means that we can replace D(II) in the statement of the Sensitivity Lemma by its nondeterministic variant N(II).*

4 Sensitivity Approach for General Trees

In this section we analyze the average sensitivity of the problem Π_T of inversion minimization on a tree T with a general shape. In Section 4.1 we show that the existence of a subtree containing a fair fraction of the leaves implies high sensitivity. The lower bound on query complexity for Π_T in Theorem 1.1 then follows from the Sensitivity Lemma. In Section 4.3 we prove that the average sensitivity measure is Lipschitz continuous. For the analysis, we make use of the decomposition of the objective of inversion minimization on trees mentioned earlier. We describe the decomposition in more detail in Section 4.2; it will be helpful in later parts of this paper, as well.

4.1 Subtree-induced sensitivity We first introduce a sensitivity bound for inversion minimization based on the size of a subtree.

LEMMA 4.1. (SUBTREE-INDUCED SENSITIVITY) *Consider a tree T with $n \doteq |\mathbf{L}(T)|$ leaves, and some node v in T with $\ell \doteq |\mathbf{L}(T_v)|$ leaves. We have*

$$s(\Pi_T) \geq \frac{\ell(n - \ell)}{n} - 1.$$

Note that v is not necessarily a direct child of the root, as shown in Fig. 6.

We now prove Lemma 4.1. Let ρ be a ranking of the leaves of T , and let σ_{\min} be a tree ordering that minimizes the number of inversions with respect to ρ .

CLAIM 4.1. ρ is sensitive to the transposition $\tau = (r, r + 1)$ if $\sigma_{\min}(\rho^{-1}(r)) > \sigma_{\min}(\rho^{-1}(r + 1))$.

Proof. If $\sigma_{\min}(\rho^{-1}(r)) > \sigma_{\min}(\rho^{-1}(r + 1))$, then $\text{Inv}_{\tau\rho}(\sigma_{\min}) = \text{Inv}_{\rho}(\sigma_{\min}) - 1$. Since $\text{Inv}_{\rho}(\sigma_{\min}) = \text{MinInv}(T, \rho)$, this means that $\text{MinInv}(T, \tau\rho) < \text{MinInv}(T, \rho)$, or that ρ is sensitive to τ . \square


In the case of general trees, a tree ordering σ that minimizes the number of inversions with respect to ρ is NP-hard to find. Our strategy is to find a lower bound on the number of r for which $\sigma(\rho^{-1}(r)) > \sigma(\rho^{-1}(r + 1))$ that applies regardless of σ .

CLAIM 4.2. For any ordering σ , the number of r such that $\sigma(\rho^{-1}(r)) > \sigma(\rho^{-1}(r + 1))$ is at least one less than the number of s such that $\rho^{-1}(s) \in \mathbf{L}(T_v)$ and $\rho^{-1}(s + 1) \notin \mathbf{L}(T_v)$.


Proof. For all except at most one value of s (the maximum s for which $\rho^{-1}(s) \in \mathbf{L}(T_v)$), there exists a minimal $s' > s$ such that $\rho^{-1}(s') \in \mathbf{L}(T_v)$. We claim that at least one value of $r = s, \dots, s' - 1$ satisfies $\sigma(\rho^{-1}(r)) > \sigma(\rho^{-1}(r + 1))$. If not, then σ would rank $\rho^{-1}(s), \rho^{-1}(s + 1), \dots, \rho^{-1}(s')$ in increasing order. Because σ is a tree ordering, the leaves of $\mathbf{L}(T_v)$ must be mapped into a contiguous range by σ , as shown in Fig. 7. However, we have $\rho^{-1}(s), \rho^{-1}(s') \in \mathbf{L}(T_v)$ but $\rho^{-1}(s + 1) \notin \mathbf{L}(T_v)$, which violates this property since σ ranks a leaf outside $\mathbf{L}(T_v)$ between two leaves inside $\mathbf{L}(T_v)$.

Because each value of r is found between consecutive pairs of values in $\mathbf{L}(T_v)$, the values of r are distinct.

\square



(a) Leaves in σ -order, labeled with ρ -ranks.



(b) Corresponding plot of ρ, σ for each leaf.

Figure 7: σ maps leaves of $\mathbf{L}(T_v)$ in a contiguous range.

CLAIM 4.3. Over a uniformly random ρ , the expected number of s such that $\rho^{-1}(s) \in \mathbf{L}(T_v)$ and $\rho^{-1}(s + 1) \notin \mathbf{L}(T_v)$ is $\frac{\ell(n - \ell)}{n}$.

Proof. For $s = 1, \dots, n-1$, the probability that $\rho^{-1}(s) \in L(T_v)$ is $\frac{\ell}{n}$, and the probability that $\rho^{-1}(s+1) \notin L(T_v)$ given that $\rho^{-1}(s) \in L(T_v)$ is $\frac{n-\ell}{n-1}$. Using linearity of expectation on the indicator random variables for $\rho^{-1}(s) \in L(T_v)$ and $\rho^{-1}(s+1) \notin L(T_v)$, the expected number of s satisfying this property is

$$(n-1) \left(\frac{\ell(n-\ell)}{n(n-1)} \right) = \frac{\ell(n-\ell)}{n}.$$

□

Combining Claim 4.1, Claim 4.2, and Claim 4.3, we can conclude with Lemma 4.1.

Bounded degree. We apply our analysis to the case of trees of degree k . Observe that for fixed n , Lemma 4.1 is strongest when $\ell = n/2$. Not every tree T has a subtree with exactly $n/2$ leaves, but Lemma 4.1 still gives a useful bound for subtrees that do not contain too few or too many leaves. In the case of trees of bounded degree, there always exists a subtree T_v that contains a fairly balanced fraction of the leaves. The following quantification is folklore. We include a proof for completeness.

FACT 4.1. *If T is a tree of degree k with n leaves, there exists a node v in T such that $\ell \doteq |L(T_v)| = \alpha \cdot n$, where $\frac{1}{k+1} \leq \alpha \leq \frac{k}{k+1}$.*

Proof. Let r be the root of T and construct a sequence $v_1 = r, v_2, v_3, \dots$ such that v_i is a child of v_{i-1} that maximizes $\ell_i \doteq |L(T_{v_i})|$, with ties broken arbitrarily. Notice that $\{\ell_i\}$ is a decreasing sequence, and since T has degree k , $\ell_i \leq k\ell_{i+1}$ for all i . We claim that some v_i in this sequence satisfies the conditions of the claim. If not, then for some i , $\ell_i > \frac{k}{k+1} \cdot n$ and $\ell_{i+1} < \frac{1}{k+1} \cdot n$, which contradicts the fact that $\ell_i \leq k\ell_{i+1}$. □

By choosing a subtree satisfying Fact 4.1, we can apply Lemma 4.1 and conclude that $s(\Pi_T) \geq \frac{k}{(k+1)^2} \cdot n - 1$. The Sensitivity Lemma then gives the “in particular” part of Theorem 1.1.

4.2 Decomposition of the objective function For use in this section as well as later parts of the paper, we now explain how the objective of inversion minimization on trees decomposes. We introduce the notion of root inversion along the way, and observe the effect of adjacent-rank transpositions on the decomposition.

The objective $\text{MinInv}(T, \rho)$ can be written as the sum of contributions from each of the individual nodes. A node v contributes those inversions that reside in the subtree T_v and go through the root v of T_v . We refer to them as the root inversions in T_v .

DEFINITION 4.1. (ROOT INVERSIONS, $\text{RInv}(\cdot, \cdot, \cdot)$, $\text{MinRInv}(\cdot, \cdot)$) *Given a tree T , a ranking ρ of the leaves of T , and an ordering σ of T , a root inversion of σ with respect to ρ is an inversion (ℓ_1, ℓ_2) of σ with respect to ρ for which the lowest common ancestor $\text{LCA}(\ell_1, \ell_2)$ is the root of T . The number of root inversions of σ with respect to ρ in T is denoted by $\text{RInv}(T, \rho, \sigma)$. The minimum number of root inversion in T with respect to ρ is denoted*

$$(4.3) \quad \text{MinRInv}(T, \rho) \doteq \min_{\sigma} \text{RInv}(T, \rho, \sigma),$$

where σ ranges over all possible orderings of T .

The only aspect of the ordering σ of T_v that affects $\text{RInv}(T_v, \rho, \sigma)$ is the relative order of the children of v . For a node v with k children u_1, \dots, u_k , by abusing notation and using σ to also denote the ranking of the children induced by the ordering of the tree, we have

$$(4.4) \quad \text{RInv}(T, \rho, \sigma) \doteq \sum_{1 \leq i < j \leq k} \text{XInv}_{\rho}(L_{\sigma(i)}, L_{\sigma(j)}),$$

where L_i is a short-hand for the leaf set $L(T_{u_i})$. The contributions of the nodes can be optimized independently:

$$(4.5) \quad \text{MinInv}(T, \rho) = \sum_v \text{MinRInv}(T_v, \rho),$$

where v ranges over all nodes of T with degree $\deg_T(v) > 1$.

When we apply an adjacent-rank transposition τ to a ranking ρ , at most one of terms in the decomposition (4.5) can change, and the change is at most one unit. We capture this observation for future reference as it will be helpful in several sensitivity analyses.

PROPOSITION 4.1. *Let ρ be a ranking of the leaf set X of a tree T , τ an adjacent-rank transposition, and ℓ_1 and ℓ_2 be the affected leaves. Then*

$$\text{MinRInv}(T_v, \rho) = \text{MinRInv}(T_v, \tau\rho)$$

for all nodes v in T except possibly $v = \text{LCA}(\ell_1, \ell_2)$. Moreover, the difference is at most 1 in absolute value.

Proof. Since the ranks of ℓ_1 and ℓ_2 under ρ are adjacent, for any leaf ℓ other than ℓ_1 and ℓ_2 , the relative order of ℓ under ρ is the same with respect to ℓ_1 as it is with respect to ℓ_2 . This means that the adjacent-rank transposition τ does not affect whether a pair of leaves constitutes an inversion unless that pair equals $\{\ell_1, \ell_2\}$. As a result, the only term on the right-hand side of (4.5) that can be affected by the transposition τ is the one corresponding to the node v , and it can change by at most one unit. \square

4.3 Lipschitz continuity Average-case notions typically do not change much under small changes to the input. This is indeed the case for the average sensitivity when “small” is interpreted as affecting few of the subtrees. The following lemma quantifies the property and can be viewed as a form of Lipschitz continuity.

LEMMA 4.2. *Given a tree T , if a subtree T_{v^*} with ℓ leaves is replaced with a tree T'_{v^*} with the same number of leaves, resulting in the tree T' , then*

$$|s(\Pi_T) - s(\Pi_{T'})| \leq \frac{\ell(\ell - 1)}{n}.$$

Proof. We think of the leaf sets of T and T' as being the same set $X = L(T) = L(T')$, and fix a ranking ρ of X . Consider an ordering of T and the ranking σ of X that it induces. Outside of T'_{v^*} we can order T' in the same way as T . Irrespective of how we order T' inside T'_{v^*} , the induced ranking σ' of X agrees with σ on all leaves in X except possibly those in $Y \doteq L(T_{v^*}) = L(T'_{v^*})$. Moreover, under both σ and σ' , the set Y gets mapped to the same contiguous interval. It follows that for all pairs (ℓ_1, ℓ_2) of distinct leaves of which at least one lies outside of Y , (ℓ_1, ℓ_2) constitutes an inversion of σ with respect to ρ if and only if (ℓ_1, ℓ_2) constitutes an inversion of σ' with respect to ρ . For any node v outside of T_{v^*} , root inversions in T_v cannot involve leaves that are both in $Y \doteq L(T_{v^*})$. See Figure 8 for an illustration. Thus, for such nodes v , $\text{RInv}(T_v, \rho, \sigma) = \text{RInv}(T'_v, \rho, \sigma')$. By taking the minimum over all orderings, we conclude:




Figure 8: Effects of changing T_{v^*}

CLAIM 4.4. *$\text{MinRInv}(T_v, \rho) = \text{MinRInv}(T'_v, \rho)$ holds for every node v outside of T_{v^*} (or equivalently, outside of T'_{v^*}).*

Consider a ranking ρ and an adjacent-rank transposition $\tau = (r, r + 1)$. We claim that, unless $(\ell_1, \ell_2) \doteq (\rho^{-1}(r), \rho^{-1}(r + 1)) \in Y \times Y$, Π_T is sensitive to τ at ρ if and only if $\Pi_{T'}$ is sensitive to τ at ρ . This is because by Proposition 4.1 the only term in the decomposition (4.5) of $\text{MinInv}(T, \rho)$ that can be affected by τ is the contribution $\text{MinRInv}(T_v, \rho)$ for $v = \text{LCA}(\ell_1, \ell_2)$. If at least one of ℓ_1 or ℓ_2 is not inside T_{v^*} , then v is not inside T_{v^*} either, so by Claim 4.4, $\text{MinRInv}(T_v, \rho) = \text{MinRInv}(T'_v, \rho)$. By the same token, $\text{MinRInv}(T_v, \tau\rho) = \text{MinRInv}(T'_v, \tau\rho)$. It follows that $\text{MinInv}(T, \rho) \neq \text{MinInv}(T, \tau\rho)$ if and only if $\text{MinInv}(T', \rho) \neq \text{MinInv}(T', \tau\rho)$.

We bound the expected number of values of r for which $(\rho^{-1}(r), \rho^{-1}(r + 1)) \in Y \times Y$ with $Y \doteq L(T_v)$ when ρ is chosen uniformly at random. For $r \in [n - 1]$, the probability that $\rho^{-1}(r) \in Y$ is $\frac{\ell}{n}$, and the probability that $\rho^{-1}(r + 1) \in Y$ given that $\rho^{-1}(r) \in Y$ is $\frac{\ell - 1}{n - 1}$. Using linearity of expectation on the indicators, the expected number of said r is

$$(n - 1) \left(\frac{\ell(\ell - 1)}{n(n - 1)} \right) = \frac{\ell(\ell - 1)}{n}.$$

\square

Lemma 4.2 helps to extend query lower bounds based on average sensitivity to larger classes. Suppose we have established a good lower bound on the sensitivity $s(\Pi_T)$ for a class C of trees. Consider a class C' obtained

by taking a tree T in class C and replacing some of the subtrees T_v by other subtrees T'_v on the same number of leaves. For this new class C' the same lower bound on the sensitivity of inversion minimization applies modulo the Lipschitz loss. For example, Theorem 1.2 holds by virtue of a lower bound of the form $s(\Pi_T) \geq (n-1) - c \log(n)$ for every binary tree T with n leaves, where c is a universal constant. If we allow some of the subtrees of T to be replaced by, say freely arrangeable ones on the same leaves, applying Lemma 4.2 for each of the modified subtrees in sequence shows that the resulting new tree T' has

$$s(\Pi_{T'}) \geq s(\Pi_T) - \frac{\alpha n(\alpha n - 1)}{n} \geq (n-1) - c \log(n) - \alpha^2(n-1) = (1-\alpha^2)(n-1) - c \log(n),$$

where α denotes the fraction of leaves that belong to one of the replaced subtrees.

5 Refined Sensitivity Approach for Binary Trees

In this section we show how to refine the sensitivity approach for lower bounds on the query complexity of the problem Π_T of inversion minimization on trees in the important special case of binary trees T . In Section 5.1 we first develop a criterion for when a particular ranking ρ is sensitive to a particular adjacent-rank transposition τ . We then analyze the root sensitivity of binary trees in Section 5.2 and finally establish a strong lower bound on the average sensitivity in Section 5.3. An application of the Sensitivity Lemma then yields Theorem 1.2.

5.1 Sensitivity criterion Recall the decomposition of the objective function $\text{MinInv}(T, \rho)$ into contributions attributed to each node v of degree $\deg_T(v) > 1$, as given by (4.5) in Section 4.2. In the case of binary trees, the contribution of node v can be calculated simply as

$$(5.6) \quad \text{MinRInv}(T_v, \rho) = \min(\text{XInv}_\rho(L_1, L_2), \text{XInv}_\rho(L_2, L_1)),$$

where u_1 and u_2 denote the two children of v , and $L_1 \doteq L(T_{u_1})$ and $L_2 \doteq L(T_{u_2})$ their leaf sets. This simplicity makes a precise analysis of sensitivity feasible, as we will see next.

For a given ranking ρ of T and a given adjacent-rank transposition τ , we would like to figure out the effect of τ on the objective $\text{MinInv}(T, \cdot)$, in particular when $\text{MinInv}(T, \tau\rho) = \text{MinInv}(T, \rho)$. Let ℓ_{lo} and ℓ_{hi} denote the two leaves that are affected by the transposition τ on the ranking ρ , where the subscript “lo” indicates the lower of the two leaves with respect to ρ , and “hi” the higher of the two. Let v be the lowest common ancestor $\text{LCA}(\ell_{\text{lo}}, \ell_{\text{hi}})$. We use the same subscripts “lo” and “hi” for the two children of v : u_{lo} denotes the child whose subtree contains ℓ_{lo} , and u_{hi} its sibling. Similarly, we denote by L_{lo} the leaf set of $T_{u_{\text{lo}}}$, and by L_{hi} the leaf set of $T_{u_{\text{hi}}}$. See Figure 9 for the subsequent analysis.




Figure 9: Sensitivity analysis for binary trees

By Proposition 4.1, the situation before and after the application of τ is as follows, where $x \doteq \text{XInv}_\rho(L_{\text{lo}}, L_{\text{hi}})$ and $y \doteq \text{XInv}_\rho(L_{\text{hi}}, L_{\text{lo}})$.

ranking	$\text{RInv}(T_v, \cdot, \sigma(\ell_{\text{lo}}) < \sigma(\ell_{\text{hi}}))$	$\text{RInv}(T_v, \cdot, \sigma(\ell_{\text{hi}}) < \sigma(\ell_{\text{lo}}))$	$\text{MinRInv}(T_v, \cdot)$
ρ	x	y	$\min(x, y)$
$\tau\rho$	$y - 1$	$x + 1$	$\min(y - 1, x + 1)$

The objective function remains the same iff $\min(x, y) = \min(y - 1, x + 1)$, which happens iff $x - y = -1$, or equivalently iff

$$(5.7) \quad \text{DInv}_\rho(L_{\text{lo}}, L_{\text{hi}}) = \text{DInv}_\rho(\{\ell_{\text{lo}}\}, \{\ell_{\text{hi}}\}),$$

where we introduce the following short-hand:

DEFINITION 5.1. (CROSS INVERSION DIFFERENCE, $\text{DInv}(\cdot, \cdot)$) For a ranking ρ of a set X , and two subsets $A, B \subseteq X$,

$$\text{DInv}_\rho(A, B) \doteq \text{XInv}_\rho(A, B) - \text{XInv}_\rho(B, A).$$

We can split L_{lo} as $L_{\text{lo}} = \{\ell_{\text{lo}}\} \sqcup A = A_{<} \sqcup \{\ell_{\text{lo}}\} \sqcup A_{>}$, where $A_{<}$ contains all leaves in L_{lo} that ρ ranks before ℓ_{lo} , and $A_{>}$ contains all the leaves in L_{lo} that ρ ranks after ℓ_{lo} . We similarly split L_{hi} , as indicated in Figure 9. We have that

$$\text{DInv}_\rho(L_{\text{lo}}, L_{\text{hi}}) = \text{DInv}_\rho(\{\ell_{\text{lo}}\}, \{\ell_{\text{hi}}\}) + \text{DInv}_\rho(\{\ell_{\text{lo}}\}, B) + \text{DInv}_\rho(A, \{\ell_{\text{hi}}\}) + \text{DInv}_\rho(A, B).$$

Since the ranks of ℓ_{lo} and ℓ_{hi} under ρ are adjacent, we have that $\text{DInv}_\rho(\{\ell_{\text{lo}}\}, B) = |B_{<}| - |B_{>}|$ and $\text{DInv}_\rho(A, \{\ell_{\text{hi}}\}) = |A_{>}| - |A_{<}|$. Plugging everything into (5.7) we conclude:

PROPOSITION 5.1. Let T be a binary tree, ρ a ranking of the leaves of T , τ an adjacent-rank transposition, ℓ_{lo} and ℓ_{hi} the two leaves affected by τ under ρ such that ρ ranks ℓ_{lo} before ℓ_{hi} . Referring to the notation in Figure 9, we have that

$$(5.8) \quad \text{MinInv}(T, \rho) = \text{MinInv}(T, \tau\rho) \Leftrightarrow \text{DInv}_\rho(L_{\text{lo}}, L_{\text{hi}}) = \text{DInv}_\rho(\{\ell_{\text{lo}}\}, \{\ell_{\text{hi}}\})$$

$$(5.9) \quad \Leftrightarrow \text{DInv}_\rho(A, B) = |A_{<}| - |A_{>}| + |B_{>}| - |B_{<}|.$$

5.2 Root sensitivity Given a ranking ρ and an adjacent-rank transposition τ , we know by Proposition 4.1 that at most one of the terms in the decomposition (4.5) of $\text{MinInv}(T, \rho)$ is affected by the transposition, namely $\text{MinRInv}(T_v, \rho)$ where v is the lowest common ancestor of the affected leaves ℓ_{lo} and ℓ_{hi} . It follows that we can write the average sensitivity of $\Pi_T \doteq \text{MinInv}(T, \cdot)$ as the following convex combination:

$$(5.10) \quad \begin{aligned} s(\Pi_T) &= (n - 1) \cdot \mathbb{P}[\text{MinInv}(T, \rho) \neq \text{MinInv}(T, \tau\rho)] \\ &= (n - 1) \sum_v \mathbb{P}[v = \text{LCA}(\ell_{\text{lo}}, \ell_{\text{hi}})] \cdot \mathbb{P}[\text{MinRInv}(T_v, \rho) \neq \text{MinRInv}(T_v, \tau\rho) \mid v = \text{LCA}(\ell_{\text{lo}}, \ell_{\text{hi}})], \end{aligned}$$

where the probability is over a uniformly random choice of the ranking ρ and the adjacent-rank transposition τ , and ℓ_{lo} and ℓ_{hi} denote the affected leaves. The conditional probability on the right-hand side of (5.10) only depends on the subtree T_v . The ranking ρ of all leaves induces a ranking ρ' of the leaves of T_v that is uniform under the conditioning. Similarly, the adjacent-rank transposition τ for ρ induces an adjacent-rank transposition τ' for ρ' ; the distribution of τ' under the conditioning is independent of ρ' and uniform among all adjacent-rank transpositions such that the affected leaves live in subtrees of different children of v . Thus, the probability on the right-hand side of (5.10) coincides with the following notion for the subtree T_v .

DEFINITION 5.2. (ROOT SENSITIVITY) Let T be a tree. The root sensitivity of T is the probability that $\text{MinRInv}(T_v, \rho) \neq \text{MinRInv}(T_v, \tau\rho)$ when ρ is a uniform random ranking of $\text{L}(T)$, and τ a uniform random adjacent transposition with the condition that the affected leaves are in subtrees of different children of the root of T .

Note that the only nodes v that need to be considered in the sum on the right-hand side of (5.10) are those that can appear as lowest common ancestor of two leaves, and such that T_v is not freely arrangeable. In the case of binary trees, this means that we only need to consider nodes v of degree 2 such that T_v contains more than 2 leaves. In this section we prove a strong lower bound on the root sensitivity of such trees T_v .

Consider the binary tree T with root v in Figure 9. The distribution underlying Definition 5.2 can be generated as follows: Pick a leaf on each side of the root v uniformly at random, and let ρ be a ranking of the leaves of T that is uniform random with the condition that the selected leaves receive adjacent ranks; τ then is the adjacent-rank

transposition that swaps the two selected leaves. The root sensitivity of T is the complement of the probability that the right-hand side of (5.9) holds under this distribution. Let us analyze the left-hand side of (5.9) further. As $A = A_{<} \sqcup A_{>}$ and $B = B_{<} \sqcup B_{>}$, we have that

$$\text{DInv}_\rho(A, B) = \text{DInv}_\rho(A_{<}, B_{<}) + \text{DInv}_\rho(A_{<}, B_{>}) + \text{DInv}_\rho(A_{>}, B_{<}) + \text{DInv}_\rho(A_{>}, B_{>}).$$

By the defining properties of the sets involved (see Figure 9), we know that $\text{DInv}_\rho(A_{<}, B_{>}) = -a_{<} b_{>}$ and $\text{DInv}_\rho(A_{>}, B_{<}) = a_{>} b_{<}$, where $a_{<} \doteq |A_{<}|$, $a_{>} \doteq |A_{>}|$, $b_{<} \doteq |B_{<}|$, and $b_{>} \doteq |B_{>}|$. Thus, we can rewrite criterion (5.9) as:

$$(5.11) \quad \text{DInv}_\rho(A_{<}, B_{<}) + \text{DInv}_\rho(A_{>}, B_{>}) = a_{<} b_{>} - a_{>} b_{<} + a_{<} - a_{>} + b_{>} - b_{<}.$$

A critical observation that helps us to bound the probability of (5.11) is that, conditioned on all four values $a_{<}$ and $b_{>}$, the right-hand side of (5.11) is fixed, but the left-hand side still contains a lot of randomness. In fact, under the conditioning stated, the ranking that ρ induces on $A_{<} \sqcup B_{<}$ is still distributed uniformly at random, the same holds for the ranking that ρ induces on $A_{>} \sqcup B_{>}$, and both distributions are independent. This means that, under the same conditioning, the left-hand side of (5.11) has the same distribution as the sum $X_{a_{<}, b_{<}} + X_{a_{>}, b_{>}}$ of two independent random variables of the following type:

DEFINITION 5.3. (CROSS INVERSION DISTRIBUTION, $X_{\cdot, \cdot}$) *For nonnegative integers a and b , $X_{a,b}$ denotes the random variable $\text{XInv}(A, B)$ that counts the number of cross inversions from A to B , where A is an array of length a , B an array of length b , and the concatenation AB is a random permutation of $[a + b]$.*

In the full version of the paper, we establish the following upper bound on the probability that the number of cross inversions takes on any specific value.

LEMMA 5.1. *There exists a constant C such that for all integers $a, b \geq 1$ and $0 \leq k \leq ab$,*

$$(5.12) \quad \mathbb{P}[X_{a,b} = k] \leq \frac{C}{\sqrt{ab(a+b)}}.$$

Using Lemma 5.1 we can establish an upper bound of the same form as the right-hand side of (5.12) for the probability that (5.11) holds: For some constant C'

$$(5.13) \quad \mathbb{P}[X_{a_{<}, b_{<}} + X_{a_{>}, b_{>}} = a_{<} b_{>} - a_{>} b_{<} + a_{<} - a_{>} + b_{>} - b_{<}] \leq \frac{C'}{\sqrt{ab(a+b)}},$$

where $a \doteq a_{<} + a_{>} \geq 1$ and $b \doteq b_{<} + b_{>} \geq 1$. We consider several cases based on the relative sizes of $a_{<}$ vs $a_{>}$, and $b_{<}$ vs $b_{>}$.

- (i) In case both $a_{<} \geq \eta a$ and $b_{<} \geq \eta b$, the bound (5.13) follows from (5.12) as long as $C' \geq C/\eta^{3/2}$.
- (ii) By reverting the order in (i), the same holds true in case both $a_{>} \geq \eta a$ and $b_{>} \geq \eta b$.
- (iii) In case $a_{>} \leq \eta a$ and $b_{<} \leq \eta b$, the left-hand side of (5.11) is at most $2\eta ab$, whereas the right-hand side is at least

$$(1 - \eta)^2 ab - \eta^2 ab + (1 - \eta)a - \eta a + (1 - \eta)b - \eta b = (1 - 2\eta)(ab + a + b).$$

As long as $2\eta \leq 1 - 2\eta$, or equivalently, $\eta \leq 1/4$, this case cannot occur.

- (iv) By reverting the role of A and B in (iii), the same holds true in case $b_{>} \leq \eta b$ and $a_{<} \leq \eta a$.

As long as $\eta \leq 1/2$, it holds that either $a_{<} \geq \eta a$ or $a_{>} \geq \eta a$, and either $b_{<} \geq \eta b$ or $b_{>} \geq \eta b$. Distributing the “and” over the “or”, we obtain the four cases we considered, which are therefore exhaustive. We conclude that (5.13) holds for $C' = 4^{3/2}C$ whenever $a \doteq |A| \geq 1$ and $b \doteq |B| \geq 1$.

In the case where $a = 0$ and $b \geq 1$, the right-hand side of (5.9) vanishes, as do $|A_{<}|$ and $|A_{>}|$, so (5.9) holds if and only if $|B_{<}| = |B_{>}|$, or equivalently, the leaf ℓ_{hi} is ranked exactly in the middle of the leaf set L_{hi} . As the ranking ρ is chosen uniformly at random, this happens with probability $1/(b + 1)$ where $b \doteq |B| = |L_{\text{hi}}| - 1$. The case where $a \geq 1$ and $b = 0$ is symmetric. The remaining case, $a = b = 0$, is one we do not need to consider as the tree T then only has two leaves. In all other cases we obtain a strong upper bound on the probability that (5.9) holds, and by complementation a strong lower bound on the root sensitivity. We capture the lower bound in the following single expression that holds for all cases under consideration.

LEMMA 5.2. *There exists a constant c such that for every binary tree T with at least 3 leaves and a root of degree 2, the root sensitivity of T is at least*

$$(5.14) \quad 1 - \frac{c}{\sqrt{n_1 n_2 (n_1 + n_2)}},$$

where n_1 and n_2 denote the number of leaves in the subtrees rooted by the two children of the root.

5.3 Average sensitivity We are now ready to establish that, except for trivial cases, the average sensitivity of a binary tree is close to maximal. The trivial cases are those where the tree has at most two leaves, in which case the sensitivity is zero.

LEMMA 5.3. *The average sensitivity of Π_T for binary trees T with $n \geq 2$ leaves is $n - O(1)$.*

Proof. We use the expression (5.10) for the average sensitivity of Π_T , where v ranges over all nodes of degree 2 such that T_v contains at least two leaves. Consider a node v of degree 2 such that T_v contains $n_{v,1}$ leaves on one side and $n_{v,2}$ leaves on the other side, where $n_{v,1} + n_{v,2} \geq 3$. If we choose the ranking ρ and the adjacent-rank transposition τ uniformly at random, each of the $\binom{n}{2}$ pairs of leaves are equally likely to be the affected pair. As there are $n_{v,1} \cdot n_{v,2}$ choices that result in v as their lowest common ancestor, we have that $\mathbb{P}[v = \text{LCA}(\ell_{10}, \ell_{hi})] = \frac{2n_{v,1}n_{v,2}}{n(n-1)}$. Combining this with the root sensitivity lower bound given by (5.14), we have that

$$\begin{aligned} s(\Pi_T) &\geq (n-1) \sum_v \frac{2n_{v,1}n_{v,2}}{n(n-1)} \cdot \left(1 - \frac{c}{\sqrt{n_{v,1}n_{v,2}(n_{v,1} + n_{v,2})}}\right) \\ &= (n-1) - \frac{2c}{n} \sum_v \sqrt{\frac{n_{v,1}n_{v,2}}{n_{v,1} + n_{v,2}}}. \end{aligned}$$

The following claim then completes the proof. \square

CLAIM 5.1. *There is a constant c' such that for all binary trees T with n leaves*

$$(5.15) \quad \sum_v \sqrt{\frac{n_{v,1}n_{v,2}}{n_{v,1} + n_{v,2}}} \leq c'n,$$

where the sum ranges over all nodes v of degree 2 such that T_v contains at least 3 leaves.

Proof. [Proof of Claim 5.1] We use structural induction to prove a somewhat stronger claim, namely that

$$(5.16) \quad \sum_v \sqrt{\frac{n_{v,1}n_{v,2}}{n_{v,1} + n_{v,2}}} \leq c'n - d'\sqrt{n}$$

for some constants c' and d' to be determined. As the base case we consider binary trees T with at most two leaves. In this case, the left-hand side of (5.16) is zero and the right-hand side is non-negative provided $c' \geq d'$, so (5.16) holds.

For the inductive step, the case where the root of T has degree 1 immediately follows from the inductive hypothesis for the subtree T_u rooted by the child u of the root of T . The remaining case is where the root of T has degree 2. Let u_1 and u_2 be the two children for the root, $n_1 = |L(T_{u_1})|$, and $n_2 = |L(T_{u_2})|$. The sum on the left-hand side of (5.16) has three contributions: $\sqrt{\frac{n_1 n_2}{n_1 + n_2}}$ from the root, and the contributions from T_{u_1} and T_{u_2} , to which we can individually apply the inductive hypothesis. This gives us an upper bound of

$$\sqrt{\frac{n_1 n_2}{n_1 + n_2}} + (c'n_1 - d'\sqrt{n_1}) + (c'n_2 - d'\sqrt{n_2}) = \sqrt{\frac{n_1 n_2}{n_1 + n_2}} + c'(n_1 + n_2) - d'(\sqrt{n_1} + \sqrt{n_2}),$$

which we want to upper bound by

$$c'n - d'\sqrt{n} = c'(n_1 + n_2) - d'\sqrt{n_1 + n_2}.$$

Writing $n_1 = \alpha n$ for some $\alpha \in [0, 1]$ and rearranging terms, the upper bound holds if and only if

$$\sqrt{\alpha(1-\alpha)} \leq d'(\sqrt{\alpha} + \sqrt{1-\alpha} - 1).$$

We claim that the upper bound holds for $d' = (\sqrt{2} + 1)/2$. Let

$$F(\alpha) \doteq d'(\sqrt{\alpha} + \sqrt{1-\alpha} - 1) - \sqrt{\alpha(1-\alpha)}.$$

It suffices to show that $F(\alpha) \geq 0$. Since F is continuous on $[0, 1]$, it attains a minimum on $[0, 1]$. On $(0, 1)$, F is differentiable. It can be verified that F' has a unique zero in $(0, 1/2)$, which needs to be a maximum as F is increasing at $\alpha = 0$. By the symmetry $F(\alpha) = F(1-\alpha)$, it follows that the minimum of F on $[0, 1]$ is attained at the midpoint $\alpha = 1/2$ or at one of the endpoint $\alpha = 0$ or $\alpha = 1$. At all three points $F(\alpha) = 0$. We conclude that (5.16) holds for any constants $d' \geq (\sqrt{2} + 1)/2$ and $c' \geq d'$. \square

6 Sensitivity Approach for Bounded Error

In this section, we apply the sensitivity approach to obtain lower bounds on the query complexity of problems in the comparison-query model against randomized algorithms with bounded error. We derive a generic result that query lower bounds against deterministic algorithms that are based on the Sensitivity Lemma, also hold against bounded-error randomized algorithms with a small loss in strength. The approach works particularly well when we have linear lower bounds on the average sensitivity, in which case there is only a constant-factor loss in the strength of the query lower bound. Among others, this applies to the $\Omega(n \log n)$ query lower bound for inversion minimization on trees of bounded degree.

Generic lower bound. Our approach is based on Yao's minimax principle [Yao77], which lower bounds worst-case complexity against randomized algorithms with bounded error by average-case complexity against deterministic algorithms with bounded distributional error. We view a deterministic algorithm with small distributional error for a problem Π as an exact deterministic algorithm for a slightly modified problem Π' . The idea is to then apply the sensitivity approach to Π' , and capitalize on the closeness of the average sensitivities of Π and Π' to obtain a lower bound in terms of the sensitivity of Π . By using the Sensitivity Lemma as a black-box, the approach yields a lower bound on the query complexity of bounded-error algorithms that is worst-case with respect to the input and with respect to the randomness, i.e., the lower bound holds for some input and some computation path on that input. By delving into the proof of the Sensitivity Lemma, we are able to obtain a lower bound that is worst-case with respect to the input but average-case with respect to the randomness, i.e., the lower bound holds for the expected number of queries on some input.¹

We first define the notions of randomized complexity and distributional complexity.

DEFINITION 6.1. (RANDOMIZED AND DISTRIBUTIONAL QUERY COMPLEXITY, $\text{RQ}(\cdot)$ AND $\text{DistQ}(\cdot, \cdot)$) *Let Π be a problem in the comparison-query model and $\varepsilon \in [0, 1]$.*

A randomized algorithm R for Π is said to have error ε if on every input ρ , the algorithm outputs $\Pi(\rho)$ with probability at least $1 - \varepsilon$. The query complexity of R is the maximum, over all inputs ρ , of the expected number of queries that R makes on input ρ . The ε -error randomized query complexity of Π , denoted $\text{RQ}_\varepsilon(\Pi)$, is the minimum query complexity of R over all ε -error randomized algorithms R for Π .

Let \mathcal{D} be a probability distribution on the inputs ρ . A deterministic algorithm A for Π has error ε with respect to \mathcal{D} if the probability that $A(\rho) = \Pi(\rho)$ is at least $1 - \varepsilon$ where the input ρ is chosen according to \mathcal{D} . The query complexity of A with respect to \mathcal{D} is the expected number of queries that A makes on input ρ when ρ is chosen according to \mathcal{D} . The ε -error distributional query complexity of Π with respect to \mathcal{D} , denoted $\text{DistQ}_\varepsilon(\Pi, \mathcal{D})$, is the minimum query complexity of A with respect to \mathcal{D} over all deterministic algorithms A for Π that have error ε with respect to \mathcal{D} .

The relationship between randomized complexity and distributional complexity is described by Yao's principle.

¹In fact, the approach yields a lower bound that is average-case with respect to the input (chosen uniformly at random) as well as the randomness. This follows because the proof of Yao's minimax principle allows us to replace the left-hand side of (6.17) by the average of the expected number of queries with respect to the distribution \mathcal{D} , which we pick to be uniform in our application of the principle.

LEMMA 6.1. (YAO'S MINIMAX PRINCIPLE [YAO77]) *Let Π be a problem in the comparison-query model, $\varepsilon \in [0, 1/2]$, and \mathcal{D} a distribution on the inputs ρ .*

$$(6.17) \quad \text{RQ}_\varepsilon(\Pi) \geq \frac{1}{2} \text{DistQ}_{2\varepsilon}(\Pi, \mathcal{D}).$$

We now prove lower bounds on the distributional query complexity, and thus on randomized query complexity, of comparison-query problems Π based on average sensitivity bounds. For these bounds, we always set \mathcal{D} to be the uniform distribution, the distribution underlying the notion of average sensitivity.

We start by studying average-case query complexity, i.e., zero-error distributional query complexity, and its relationship to the average sensitivity. We follow a strategy similar to the one in the proof of the Sensitivity Lemma. Whereas a bound on deterministic complexity Q follows purely from the number of execution traces D , here, the execution traces are weighted by their depth and their probability of occurring.

Recall that g in the statement of the Strong Sensitivity Lemma denotes any convex function $g : [1, \infty) \rightarrow \mathbb{R}$ with $g(x) = x!$ for $x \in [n]$; for deriving the Sensitivity Lemma from the Strong Sensitivity Lemma we also need g to be nondecreasing. One such function is $g(x) = \Gamma(x+1)$. To prove a lower bound on the zero-error distributional complexity, we need the function g to be not only convex, but *log-convex*, i.e., $\log_2 g(x)$ needs to be convex. The function $g(x) = \Gamma(x+1)$ satisfies this constraint, as well.

PROPOSITION 6.1. *Let Π be a problem in the comparison-query model with n items, \mathcal{D} the uniform distribution on the inputs ρ , and $g : [1, \infty) \rightarrow \mathbb{R}$ a nondecreasing log-convex function with $g(x) = x!$ for $x \in [n]$.*

$$\text{DistQ}_0(\Pi, \mathcal{D}) \geq \log_2(g(s(\Pi) + 1)/n)$$

Proof. Let k be the number of distinct execution traces of a deterministic algorithm A for Π , and let R_1, \dots, R_k denote the corresponding sets of rankings. Interpreting A as a binary decision tree, let $d(R_i)$ be the depth of the execution trace corresponding to R_i . By Kraft's inequality,

$$\sum_{i=1}^k 2^{-d(R_i)} \leq 1.$$

Let $f(x) = \frac{1}{n} \frac{g(x+1)}{n!}$ and define the weight function $w(\rho) = f(\text{deg}_{\overline{H}}(\rho))$, where \overline{H} refers to the notation of the Strong Sensitivity Lemma: \overline{H} denotes the subgraph of the full permutahedron that only consists of the bichromatic edges when the vertices are colored with their execution trace under A . By Claim 3.1, the sum of the weights of all rankings ρ in R_i is at most 1. Therefore,

$$\sum_{\rho} 2^{-d(\rho)} w(\rho) \leq 1.$$

Dividing both sides by $n!$ and taking the logarithm of both sides, we get that

$$(6.18) \quad \log_2 \mathbb{E} \left[2^{-d(\rho)} w(\rho) \right] \leq \log_2(1/n!),$$

where the expectation is with respect to a uniform distribution over the inputs ρ . By Jensen's inequality, since \log is concave, we get

$$\log_2 \mathbb{E} \left[2^{-d(\rho)} w(\rho) \right] \geq \mathbb{E} \left[\log_2 \left(2^{-d(\rho)} w(\rho) \right) \right] = \mathbb{E}[-d(\rho)] + \mathbb{E}[\log_2 w(\rho)],$$

which, in combination with (6.18), implies

$$\mathbb{E}[\log_2 w(\rho)] \leq \mathbb{E}[d(\rho)] + \log_2(1/n!).$$

Note that since g is log-convex, so is f . By applying Jensen's inequality again,

$$\mathbb{E}[\log_2 w(\rho)] = \mathbb{E}[\log_2 f(\text{deg}_{\overline{H}}(\rho))] \geq \log_2 f(\mathbb{E}[\text{deg}_{\overline{H}}(\rho)]),$$

implying

$$\log_2 \left(\frac{1}{n} \cdot \frac{g(\mathbb{E}[\deg_{\overline{H}}(\rho)] + 1)}{n!} \right) \leq \mathbb{E}[d(\rho)] + \log_2(1/n!),$$

or equivalently,

$$\mathbb{E}[d(\rho)] \geq \log_2(g(\mathbb{E}[\deg_{\overline{H}}(\rho)] + 1)/n).$$

The result follows since A is an arbitrary deterministic algorithm for Π , $\mathbb{E}[d(\rho)]$ equals the query complexity of A with respect to the uniform distribution \mathcal{D} , $\mathbb{E}[\deg_{\overline{H}}(\rho)] \geq \mathbb{E}[\deg_{\overline{G}(\Pi)}(\rho)] = s(\Pi)$, and g is nondecreasing. \square

Proposition 6.1 allows us to prove a lower bound on the ε -error distributional query complexity of Π with respect to the uniform distribution. In order to do so, we view a deterministic algorithm with distributional error ε for Π as an exact deterministic algorithm for a modified problem Π' , apply Proposition 6.1, and lower bound the sensitivity of Π' in terms of the sensitivity of Π .

PROPOSITION 6.2. *Let Π be a problem in the comparison-query model with n items, \mathcal{D} the uniform distribution on the inputs ρ , $\varepsilon \in [0, 1]$, and $g : [1, \infty) \rightarrow \mathbb{R}$ a nondecreasing log-convex function with $g(x) = x!$ for $x \in [n]$.*

$$(6.19) \quad \text{DistQ}_\varepsilon(\Pi, \mathcal{D}) \geq \log_2 \left(\frac{g(s(\Pi) + 1 - 2(n-1)\varepsilon)}{n} \right)$$

Proof. Consider any algorithm A with error ε for Π , or in other words, $\mathbb{P}[A(\rho) \neq \Pi(\rho)] \leq \varepsilon$. Let Π_A be the problem of determining the output of A . We prove that

$$s(\Pi_A) \geq s(\Pi) - 2(n-1)\varepsilon,$$

which implies the desired result by Proposition 6.1, since A is a deterministic algorithm for Π_A and g is nondecreasing.

Let G denote the full permutahedron graph for n items. We use the fact that $s(\Pi) = (n-1) \cdot \mathbb{P}_{e \in G}[e \in G(\Pi)]$, and similarly, $s(\Pi_A) = (n-1) \cdot \mathbb{P}_{e \in G}[e \in G(\Pi_A)]$, where all the underlying distributions are uniform. Suppose the endpoints of e are ρ_1 and ρ_2 . Note that if $e \in G$ is picked uniformly at random, then the marginal distributions of both ρ_1 and ρ_2 are also uniform. If $A(\rho_1) = \Pi(\rho_1)$, $A(\rho_2) = \Pi(\rho_2)$, and $e \in G(\Pi_A)$, then $e \in G(\Pi)$, as well. By a union bound, the probability that $A(\rho_1) \neq \Pi(\rho_1)$ or $A(\rho_2) \neq \Pi(\rho_2)$ is at most 2ε .

$$\begin{aligned} \mathbb{P}_{e \in G}[e \in G(\Pi_A)] &\geq \mathbb{P}_{e \in G}[e \in G(\Pi)] - \mathbb{P}_{e \in G}[A(\rho_1) \neq \Pi(\rho_1) \text{ or } A(\rho_2) \neq \Pi(\rho_2)] \\ &\geq \mathbb{P}_{e \in G}[e \in G(\Pi)] - 2\varepsilon. \end{aligned}$$

Multiplying both sides by $n-1$ gives $s(\Pi_A) \geq s(\Pi) - 2(n-1)\varepsilon$. \square

Since $s(\Pi) \leq n-1$, Proposition 6.2 only yields nontrivial lower bounds for small ε . In order to establish lower bounds for the standard $\varepsilon = 1/3$, we first reduce the error using standard techniques. Doing so such that the argument of g on the right-hand side of (6.19) remains $\Omega(s(\Pi))$, and picking $g(x) = \Gamma(x+1)$, we conclude:

LEMMA 6.2. (BOUNDED-ERROR SENSITIVITY LEMMA) *For any problem Π in the comparison-query model with n items,*

$$\text{RQ}_{1/3}(\Pi) = \Omega \left(\frac{s \log s}{\log(2n/s)} \right),$$

where $s \doteq s(\Pi)$.

Proof. By taking the majority vote of multiple independent runs and a standard analysis, e.g, based on Chernoff bounds, we have that $\text{RQ}_\varepsilon(\Pi) = O(\log(1/\varepsilon)) \text{RQ}_{1/3}(\Pi)$ for any $\varepsilon \leq 1/3$. Combining this with Lemma 6.1 and Proposition 6.2, we have:

$$\text{RQ}_{1/3}(\Pi) = \Omega \left(\frac{\text{RQ}_\varepsilon(\Pi)}{\log(1/\varepsilon)} \right) = \Omega \left(\frac{\text{DistQ}_{2\varepsilon}(\Pi)}{\log(1/\varepsilon)} \right) = \Omega \left(\frac{\log(g(s+1-4(n-1)\varepsilon)/n)}{\log(1/\varepsilon)} \right).$$

Setting ε such that $4n\varepsilon = s/2$ yields

$$\text{RQ}_{1/3}(\Pi) = \Omega\left(\frac{\log(g(s/2 + 1)/n)}{\log(8n/s)}\right).$$

Picking $g(x) = \Gamma(x + 1)$ and using the fact that $\Gamma(x) \geq \sqrt{2\pi x} \left(\frac{x}{e}\right)^x$, we obtain

$$\text{RQ}_{1/3}(\Pi) = \Omega\left(\frac{(s/2) \log(s/(2e)) - \log n}{\log(8n/s)}\right) = \Omega\left(\frac{s \log s}{\log(2n/s)}\right),$$

where the simplification can be verified by considering the cases of large s (say $s \geq \sqrt{n}$) and small s separately. \square

We can apply Lemma 6.2 to the sensitivity lower bounds of Lemma 4.1 and produce randomized lower bounds for inversion minimization on bounded-degree trees. Using Fact 4.1 we obtain:

THEOREM 6.1. (LOWER BOUND AGAINST BOUNDED-ERROR FOR INVERSION MINIMIZATION ON TREES) *Let T be a tree with $\deg(T) \leq k$. The query complexity of Π_T for bounded-error randomized algorithms is $\Omega\left(\frac{n \log(n/k)}{k \log(k)}\right)$.*

7 Connectivity Lemma

In this section we establish Lemma 1.2 and use it to present some of the known lower bounds in a unified framework. We actually prove the following somewhat stronger result.

LEMMA 7.1. (STRONG CONNECTIVITY LEMMA) *Consider an algorithm A in the comparison-based model, color each vertex of the permutahedron with its execution trace under A , and let H denote the subgraph with the same vertex set but only containing the monochromatic edges. The number of distinct execution traces of A equals the number of connected components of H .*

The Connectivity Lemma follows from Lemma 7.1 because the coloring with execution traces of an algorithm A for Π is a refinement of the coloring with Π . Note that the counterpart of Lemma 7.1 in the Boolean setting is trivial. This is because an execution trace in the Boolean setting is specified by values for a subset of the input bits, so the set of inputs that follow a particular execution trace form a subcube of the hypercube, the Boolean counterpart of the permutahedron. Subcubes are trivially connected inside the hypercube. In the comparison-query model, the sets of inputs that follow a particular execution trace can be more complicated, and their connectedness is no longer trivial but still holds.

Proof. [Proof of Lemma 7.1] Two rankings ρ_1 and ρ_2 that have distinct execution traces under A cannot be connected because any path between them needs to contain at least one bichromatic edge. For the remainder of the proof, we consider two rankings ρ_1 and ρ_2 that have the same execution trace under A , and construct a path from ρ_1 to ρ_2 in H .

If $\rho_1 = \rho_2$, we do not need to make any move and use an empty path.

Otherwise, there exists a rank $r < n$ such that ρ_1 and ρ_2 agree on ranks less than r and disagree on rank r . We have the following situation, where the item y_r with rank r under ρ_2 , has rank $s > r$ under ρ_1 .

rank	1	...	$r - 1$	r	...	$s - 1$	s	...	n
ρ_1^{-1}	x_1	...	x_{r-1}	x_r	...	x_{s-1}	$x_s = y_r$...	
	=	=	=	≠					
ρ_2^{-1}	y_1	...	y_{r-1}	y_r	

Considering ranking ρ_1 , we have that $\rho_1(x_{s-1}) = s - 1 < s = \rho_1(x_s)$. Considering ranking ρ_2 , since x_{s-1} differs from $y_i = x_i$ for every $i \in [r - 1]$ and also differs from y_r , we have that $\rho_2(x_{s-1}) > r = \rho_2(y_r) = \rho_2(x_s)$. Thus, the relative ranks of x_{s-1} and x_s under ρ_1 and ρ_2 differ. As ρ_1 and ρ_2 have the same execution trace, this means that the algorithm does not compare x_{s-1} and x_s on either input, and on ρ_1 in particular. Let ρ'_1 be the ranking obtained from ranking ρ_1 by applying the adjacent-rank transposition $\tau = (s - 1, s)$. Since the algorithm does not compare the affected items, the execution trace for ρ'_1 and ρ_1 are the same, so the edge from ρ_1 to ρ_1 is

monochromatic and in H . We use this edge as the first on the path from ρ_1 to ρ_2 in H . What remains is to find a path from ρ'_1 to ρ_2 in H . The situation is the same as the one depicted above but with r increased by one in case $s = r + 1$, and with the same r and s decreased by one, otherwise. The proof then follows by induction on the ordered pair $(r, n - s)$. \square

REMARK 7.1. *Suppose we allow an algorithm A to have multiple valid execution traces on a given input ρ , and let R denote the set of rankings on which a particular execution trace is valid. The construction in the proof of Lemma 7.1 yields a path in the permutahedron between any two rankings in R such that the path entirely stays within R . This means that we can replace $D(\Pi)$ in the statement of the Connectivity Lemma by its nondeterministic variant $N(\Pi)$.*

We already explained in Section 1 how the Connectivity Lemma shows that counting inversions and inversion parity amount to sorting, and require at least $\log(n!)$ queries. We now illustrate its use for a classical problem that is easier than sorting, namely median finding.

Let Π denote the selection problem with rank $r = \lceil n/2 \rceil$. For any ranking, the adjacent-rank transpositions τ that change the item with rank r are the two that involve rank r : $\tau = (r - 1, r)$ and $\tau = (r, r + 1)$. Those transpositions are the ones that correspond to missing edges in the permutahedron graph $G(\Pi)$. As a result, for any two rankings, there exists a path between them in $G(\Pi)$ if and only if they have the same median as well as the same set of items with rank less than r (and also the same set of items with rank greater than r). As there are n possibilities for the median and, for each median, $\binom{n-1}{r-1}$ possibilities for the set of items that have rank less than r , $G(\Pi)$ has $n \cdot \binom{n-1}{r-1}$ connected components. It follows that any algorithm for Π has at least $n \cdot \binom{n-1}{r-1} = \Omega(\sqrt{n} \cdot 2^n)$ distinct execution paths, and therefore needs to make at least $n + \frac{1}{2} \log(n) - O(1)$ queries.

As a side note, this example clarifies a subtlety in the equivalence between ordinary selection and the instantiation of partial order production that is considered equivalent to selection. Whereas selection of rank r ordinarily requires outputting only the item of rank r , the instantiation of partial order production additionally requires partitioning the remaining items according to whether their ranks are less than or greater than r . The above analysis implies that it is impossible for the algorithm to know the item of rank r without also knowing how to partition the remaining items into those of rank less than and greater than r . It follows that, in the comparison-based model, ordinary selection and the instantiation of partial order production are equivalent.

8 Connectivity Approach

In this section we develop the connectivity approach to obtain query lower bounds for problems Π in the comparison-query model. Our main focus is the problem Π_T of inversion minimization on certain types of trees T , for which we derive the very strong query lower bounds of Theorem 1.3. While developing the application, we observe that some parts carry through for a broader class of problems. This allows us to apply the same ideas to the problem of counting cross inversions, for which we obtain the query lower bound of Theorem 1.4, as well as the closely related problem of inversion minimization on the Mann–Whitney trees of Figure 2.

In order to obtain good lower bounds on $D(\Pi_T)$ using the Connectivity Lemma, it is sufficient to find good upper bounds on the size of the connected components of a typical vertex in $G(\Pi_T)$. Assume without loss of generality that T has no internal nodes of degree 1, i.e., no nodes with exactly one child. Π_T is always insensitive to an adjacent-rank transposition τ at a ranking ρ when the affected leaves are siblings in T . Thus, the corresponding edges from the permutahedron are always present in $G(\Pi_T)$. From the perspective of ensuring small connected components in $G(\Pi_T)$, the ideal situation would be if there were no other edges in $G(\Pi_T)$. That is to say, Π_T is *sensitive* at ρ to *every* adjacent-rank transposition τ *except* when the affected leaves are siblings. We will investigate conditions on T that guarantee this situation in the next subsections. For now, let us investigate the size of the connected components of $G(\Pi_T)$ when T is of the desired type.

To facilitate the analysis, recall our notation for the maximal sets of sibling leaves.

DEFINITION 8.1. (LEAF CHILD SET, LC) *The leaf child set $LC(v)$ of a vertex v in a tree T is the set $LC(v)$ of all the children of v that are leaves in T .*

Let LC_1, LC_2, \dots denote the leaf child sets in T . See Figure 10a for an example. As the only steps we can take in $G(\Pi_T)$ correspond to adjacent-rank transpositions τ that swap leaves in the same leaf child set LC_i , the sets LC_i remain invariant, irrespective of the ranking ρ we start from.

Depending on ρ , there may be more structure inside each leaf child set LC_i ; the leaf child set may be broken up into smaller sets that are each invariant. Figure 10b illustrates how this happens for a particular leaf child set containing 7 leaves. We list the elements of the child leaf set in order of increasing rank under ρ , and include an edge between elements that have successive ranks. We introduce the term “successor graph” to capture this structure, viewed as a graph with the ranks as vertices.

DEFINITION 8.2. (SUCCESSOR GRAPH, $S(\cdot, \cdot)$) Let Π be a computational problem in the comparison-query model on a set X of n items, and ρ a ranking of X . The successor graph of Π on ρ , denoted $S(\Pi, \rho)$, has vertex set $[n]$ and contains all edges of the form $(r, r + 1)$ with $r \in [n - 1]$ such that $\Pi(\rho) = \Pi(\tau\rho)$, where τ denotes the adjacent-rank transposition $(r, r + 1)$.




Figure 10: Connected component analysis

The connected components of the successor graph $S(\Pi, \rho)$ in Figure 10b correspond to subsets of LC_i that each remain invariant. Within each of the subsets, every possible ordering can be reached, independently for each subset. This is because for any adjacent-rank transposition τ , the successor graphs $S(\Pi, \rho)$ and $S(\Pi, \tau\rho)$ are the same, and every ordering can be realized by a sequence of swaps of adjacent elements. Thus, if the sizes of the connected components of $S(\Pi, \rho)$ over all possible leaf child sets LC_i are n_1, n_2, \dots , the size of the connected component of ρ in $G(\Pi_T)$ equals $\prod_j (n_j!)$.

For future reference, we abstract the property that we need for the above analysis the carry through.

DEFINITION 8.3. (PARTITION PROPERTY) A computational problem Π in the comparison-query model on a set X of n items has the partition property if the set X can be partitioned into sets X_i such that for any ranking ρ of X and adjacent-rank transposition $\tau = (r, r + 1)$ with $r \in [n - 1]$, $\Pi(\rho) = \Pi(\tau\rho)$ if and only if $\rho^{-1}(r)$ and $\rho^{-1}(r + 1)$ belong to the same partition class X_i .

In the case of the problem Π_T , the partition classes X_i correspond to the leaf child sets LC_i . We have shown:

PROPOSITION 8.1. Let Π be a computational problem in the comparison-query model on the set X , and let ρ be a ranking of X . If Π has the partition property, then the connected component of ρ in $G(\Pi)$ has size $\prod_j (n_j!)$, where the n_j 's denote the sizes of the connected components of $S(\Pi, \rho)$.

Returning to instances of the problem Π_T with the partition property, if each leaf child set LC_i has size at most k , then each of the connected components of $S(\Pi, \rho)$ has size $n_j \leq k$, irrespective of ρ . The maximum value that $\prod_j (n_j!)$ can take under the constraints $\sum_j n_j = n$ and $n_j \leq k$ is no more than $(k!)^{n/k}$. By the Connectivity Lemma, we conclude that $D(\Pi_T) \geq n! / (k!)^{n/k}$, and that the query complexity is at least $\log_2(n!) - O(n \log(k))$.

We can do better by observing that, for a random ranking ρ , the number of adjacent-rank transpositions τ that do not jump from one leaf child set to another one, is not much larger than the average size of the leaf child sets. More precisely, for any rank $r \in [n - 1]$, the probability that $\rho^{-1}(r)$ and $\rho^{-1}(r + 1)$ belong to the same leaf child set equals $\sum_i \frac{|\text{LC}_i|}{n} \frac{|\text{LC}_i| - 1}{n - 1}$, which is at most $\frac{k-1}{n-1}$ when each leaf child set LC_i is of size at most k . It follows that the expected number of adjacent-rank transpositions τ that do not change leaf child sets, is at most $k - 1$, so for a fraction at least half of the rankings ρ the number is at most $2(k - 1)$.

The number of adjacent-rank transpositions τ that do not change leaf child sets for a given ranking ρ equals the number of edges in the successor graph $S(\Pi_T, \rho)$. This is the number of edges in Figure 10b, summed over

all leaf child sets LC_i . In terms of the sizes n_j , the number equals $\sum_j (n_j - 1)$. We are considering rankings ρ for which the sum is at most $2(k - 1)$. The maximum of $\prod_j (n_j!)$ under the constraints that $\sum_j (n_j - 1) \leq 2(k - 1)$ and that each individual $n_j \leq k$, is reached when two of the n_j 's equal k and the rest are 1. Thus, if each of the leaf child sets LC_i are of size at most k , for a fraction at least half of the rankings ρ , the size of the connected component of ρ in $G(\Pi_T)$ is at most $(k!)^2$. It follows that the number of connected components of $G(\Pi)$ is at least $n!/(2(k!)^2)$. By the Connectivity Lemma, we conclude:

LEMMA 8.1. *Let T be a tree with n leaves. If Π_T satisfies the partition property, then $D(\Pi) \geq n!/(2(k!)^2)$.*

This yields a lower bound of $\log(n!) - O(k \log(k))$ on the query complexity of Π_T whenever Π_T satisfies the partition property.

Next we find sufficient conditions on the tree T that guarantee that Π_T has the partition property. For didactic reasons we first develop the conditions for binary trees, and then generalize them to arbitrary trees.

8.1 Binary trees In the case of binary trees T , the sensitivity analysis of Section 5.1 leads to a simple sufficient condition for the partition property to hold for Π_T . Recall that we are assuming without loss of generality that T has no internal nodes of degree 1, which in the case of binary trees is equivalent to saying that the tree is full: Every internal node has the maximum degree of 2.

Consider criterion 5.8 in Proposition 5.1. The right-hand side is always -1 . As for the left-hand side, we know the following.

FACT 8.1. *For all disjoint sets $A, B \subseteq X$ and any ranking ρ of X , $D\text{Inv}_\rho(A, B) = |A| \cdot |B| \bmod 2$.*

Proof. As every pair in $A \times B$ constitutes a cross-inversion for either A to B , or B to A , we have $X\text{Inv}_\rho(A, B) + X\text{Inv}_\rho(B, A) = |A| \cdot |B|$. Thus,

$$\begin{aligned} D\text{Inv}_\rho(A, B) &\doteq X\text{Inv}_\rho(A, B) - X\text{Inv}_\rho(B, A) \\ &= (X\text{Inv}_\rho(A, B) + X\text{Inv}_\rho(B, A)) - 2 X\text{Inv}_\rho(B, A) \\ (8.20) \qquad &= |A| \cdot |B| - 2 X\text{Inv}_\rho(B, A). \end{aligned}$$

As $X\text{Inv}_\rho(B, A)$ in an integer, the claim follows. \square

Fact 8.1 implies that whenever at least one of the leaf sets L_{lo} or L_{hi} is of even cardinality, then (5.8) fails to hold, and Π_T is sensitive to the underlying τ at ρ . Thus, we can guarantee that Π_T satisfies the partition property provided that for any two siblings u_1 and u_2 in T that are not both leaves, at least one of $|L(T_{u_1})|$ or $|L(T_{u_2})|$ is even. We refer to the latter condition as the *product condition*. In trees without nodes of degree 1, the product condition can be expressed alternately in terms of the leaf child sets. We state and prove the result for arbitrary trees as it will help us in the next subsection to generalize the analysis.

PROPOSITION 8.2. *Let T be a tree without nodes of degree 1. The following two conditions are equivalent:*

- (a) *For any two siblings u_1 and u_2 that are not both leaves, at least one of $|L(T_{u_1})|$ or $|L(T_{u_2})|$ is even.*
- (b) *At most one leaf child set is odd, and if there exists a node v^* with an odd leaf child set $LC(v^*)$, then all ancestors of v^* have an empty leaf child set.*

In the case of binary trees, (b) can be simplified to: At most one leaf has a non-leaf sibling.

Proof. We establish the two directions of implication separately.

\Rightarrow : We argue the contrapositive. Suppose that at least two of the leaf child sets are odd. Start with the root of T as the node v , and iterate the following: If v has a child u such that T_u contains at least two nodes with an odd leaf child set, replace v by such a child u . When the process ends, one of the following situations applies:

- There are two distinct children u_1 and u_2 of v that are not leaves and each contain a single node with an odd leaf child set. In this case both T_{u_1} and T_{u_2} contain an odd number of leaves, violating (a).

- There exists a unique child u_1 of v that is not a leaf and contains a single node with an odd leaf child set, and v itself has an odd number of leaf children. In this case, setting u_2 to any one leaf child of v (which exists as their number is odd), leads to a violation of (a).

Next, suppose that there exists a unique node v^* that has an odd leaf child set, and that an ancestor v of v^* has a leaf child u_1 . Setting u_2 to the child of v that contains v^* in its subtree, yields a violation of (a) as T_{u_2} contains an odd number of leaves.

⇐: If neither u_1 nor u_2 are leaves, the first condition of (b) guarantees that at most one of T_{u_1} or T_{u_2} contains an odd number of leaves. If u_1 is a leaf and u_2 is not, then the second condition of (b) implies that T_{u_2} cannot contain a node with an odd leaf child set, and therefore has an even number of leaves.

In the case of binary trees, the first condition of (b) implies the second one, which can therefore be dropped from the equivalence statement. Moreover, for binary trees the first condition of (b) can be expressed as: At most one leaf has a non-leaf sibling. \square

By Proposition 8.2, in a binary tree the condition that at most one leaf has no sibling is equivalent to the product condition, which implies the partition property of Π_T , so the lower bound of Lemma 8.1 applies. This establishes Theorem 1.3 in the case of binary trees.

As a side note, Fig. 11 shows the simplest example of a full binary tree T and a ranking ρ for which there exists an adjacent-rank transposition τ to which Π_T is insensitive at ρ while the affected leaves are not siblings. The tree has two leaves without siblings, namely 1 and 2. The adjacent-rank transposition $(3, 4)$ acts on nodes that are not siblings, but leaves the minimum number of inversions at 4.




Figure 11: Tree insensitive to non-sibling transposition

8.2 General trees For general trees T the sensitivity analysis of Π_T becomes more complicated than for binary trees, and we do not know of a simple sensitivity criterion like Proposition 5.1, but we can nevertheless extend the result for binary trees to arbitrary trees with similar constraints. For a given ranking ρ of $L(T)$ and a given adjacent-rank transposition τ , we would like to figure out the effect of τ on the objective $\text{MinInv}(T, \cdot)$, in particular when $\text{MinInv}(T, \tau\rho) = \text{MinInv}(T, \rho)$. Recall the decomposition (4.5) of $\text{MinInv}(T, \cdot)$ from Section 4.2. By Proposition 4.1 the only term on the right-hand side of (4.5) that can be affected by the transposition τ is

$$\text{MinRInv}(T_v, \rho) \doteq \min_{\sigma} \text{RInv}(T, \rho, \sigma)$$

corresponding to the node v that is the least common ancestor $\text{LCA}(\ell_{\text{lo}}, \ell_{\text{hi}})$ of the two leaves ℓ_{lo} and ℓ_{hi} that are affected by τ under ρ . In Section 5 we considered the two possible relative orderings σ_1 and σ_2 of the children of v , and derived a criterion for when the lowest cost does not change under τ . More precisely, when

$$(8.21) \quad \min(\text{RInv}(T, \rho, \sigma_1), \text{RInv}(T, \rho, \sigma_2)) = \min(\text{RInv}(T, \tau\rho, \sigma_1), \text{RInv}(T, \tau\rho, \sigma_2)).$$

There are two complications in generalizing this approach from binary to general trees.

- The expression (4.4) for $\text{RInv}(T, \rho, \sigma)$ involves multiple terms instead of just one as in (5.6). This complicates probabilistic analyses like the one we did in Section 5 because the difference in cost of the two relative orderings of two children is also affected by parts of the tree outside of their combined subtrees. The issue did not matter for the analysis in Section 4. We will be able to manage it here, as well.
- There now are not just two but multiple possible orderings σ , and it is not clear what pairs (σ_1, σ_2) we need to impose (8.21) on in order to guarantee that $\text{MinRInv}(T_v, \rho) = \text{MinRInv}(T_v, \tau\rho)$ but no more.

In Section 4 we circumvented the second issue by only considering sensitivities that decrease the objective function, and establishing a lower bound on their occurrence independent of the ordering σ . Here we are also able to handle the second issue by shooting for a sufficient condition for sensitivity rather than a criterion. We do so by requiring

that for no pair of distinct orderings σ_1 and σ_2 , condition (8.21) holds (unless the two affected leaves are siblings). Similar to the case of binary trees, we guarantee that (8.21) fails based on parity considerations given the product condition.

For the analysis we again assume without loss of generality that T has no nodes of degree 1. We use the same notation as in Section 5: Let ℓ_{lo} denote the affected leaf that is smaller with respect to ρ , and ℓ_{hi} the other affected leaf. Let L_i denote the leaf set $L_i \doteq \mathbf{L}(T_{u_i})$, where u_1, \dots, u_k are the children of $v = \text{LCA}(\ell_{\text{lo}}, \ell_{\text{hi}})$. We also write u_{lo} for the child of v that contains ℓ_{lo} in its subtree, and L_{lo} for the leaf set of the subtree rooted at u_{lo} , and define ℓ_{hi} , u_{hi} , and L_{hi} similarly. See Figure 12 for a sketch of the setting.




Figure 12: Sensitivity for general trees

We slightly abuse notation and let σ denote both the ordering of the entire tree T as well as the ranking of the children of v . By the analysis of Section 5, we have that for two orderings σ_1 and σ_2 the situation (8.21) can only occur if $\sigma_1(u_{\text{lo}}) < \sigma_1(u_{\text{hi}})$, $\sigma_2(u_{\text{lo}}) > \sigma_2(u_{\text{hi}})$, and

$$(8.22) \quad \text{RInv}_\rho(T_v, \rho, \sigma_1) - \text{RInv}_\rho(T_v, \rho, \sigma_2) = -1 = \text{DInv}_\rho(\{\ell_{\text{lo}}\}, \{\ell_{\text{hi}}\}).$$

For any two disjoint sets of leaves, (8.20) lets us write

$$(8.23) \quad \text{XInv}_\rho(A, B) = \frac{1}{2} \text{DInv}_\rho(A, B) + \frac{1}{2} |A| \cdot |B|.$$

Applying (8.23) to all the terms involved in (4.4), we have

$$\begin{aligned} \text{RInv}_\rho(T_v, \rho, \sigma) &= \sum_{1 \leq i < j \leq k} \text{XInv}_\rho(L_i, L_j) \cdot \mathbb{I}[\sigma(i) < \sigma(j)] \\ &\quad + \sum_{1 \leq i < j \leq k} \text{XInv}_\rho(L_j, L_i) \cdot \mathbb{I}[\sigma(i) > \sigma(j)] \\ &= \frac{1}{2} \sum_{1 \leq i < j \leq k} \left(\text{DInv}_\rho(L_i, L_j) \cdot (-1)^{\mathbb{I}[\sigma(i) > \sigma(j)]} + |L_i| \cdot |L_j| \right) \\ \text{RInv}_\rho(T_v, \rho, \sigma_1) - \text{RInv}_\rho(T_v, \rho, \sigma_2) &= \sum_{\substack{1 \leq i < j \leq k \\ \sigma_1(i) < \sigma_2(j) \\ \sigma_2(i) > \sigma_2(j)}} \text{DInv}_\rho(L_i, L_j) + \sum_{\substack{1 \leq i < j \leq k \\ \sigma_1(i) > \sigma_2(j) \\ \sigma_2(i) < \sigma_2(j)}} \text{DInv}_\rho(L_i, L_j) \end{aligned}$$

By combining the last equation with (8.22) and separating out the term for $(i, j) = (\text{lo}, \text{hi})$, we obtain the following necessary condition for (8.21) to hold:

$$(8.24) \quad \text{DInv}_\rho(L_{\text{lo}}, L_{\text{hi}}) - \text{DInv}_\rho(\{\ell_{\text{lo}}\}, \{\ell_{\text{hi}}\}) = - \sum_{\substack{1 \leq i < j \leq k \\ \sigma_1(i) < \sigma_2(j) \\ \sigma_2(i) > \sigma_2(j) \\ (i, j) \neq (\text{lo}, \text{hi})}} \text{DInv}_\rho(L_i, L_j) + \sum_{\substack{1 \leq i < j \leq k \\ \sigma_1(i) > \sigma_2(j) \\ \sigma_2(i) < \sigma_2(j)}} \text{DInv}_\rho(L_i, L_j).$$

In order for Π_T to have the partition property, it suffices to ensure that (8.24) fails whenever u_{lo} and u_{hi} are not both leaves. By (8.20) each of the terms $\text{DInv}_\rho(L_i, L_j)$ in (8.24) has the same parity as $|L_i| \cdot |L_j|$. Since

$\text{DInv}_\rho(\{\ell_{\text{lo}}\}, \{\ell_{\text{hi}}\})$ is odd, it follows that (8.24) fails whenever at most one of the leaf sets L_i involved is odd, which is condition (a) in Proposition 8.2. Switching to the equivalent condition (b) from Proposition 8.2 allows us to conclude via Lemma 8.1:

THEOREM 8.1. *Let T be a tree without nodes of degree 1 such that the leaf child sets have size at most k , at most one of them is odd, and if there exists an odd one, say $\text{LC}(v^*)$, then all ancestors of v^* have empty leaf child sets. Then $\text{D}(\Pi_T) \geq n!/(2(k!)^2)$.*

Theorem 1.3 follows by taking the base-2 logarithm of the bound.

8.3 Counting cross inversions and evaluating the Mann–Whitney statistic We now apply the connectivity approach that we captured in Proposition 8.1 to the problem Π_{XInv} of computing the number of cross inversions between two disjoint sets A and B with respect to a ranking ρ of $X = A \sqcup B$. Note that this problem is a refinement of evaluating the Mann–Whitney statistic, or equivalently, of inversion minimization on the tree T of Figure 2: Any algorithm that solves Π_{XInv} with q queries, can be transformed into an algorithm for Π_T with q queries, namely by transforming the output y of the algorithm for Π_{XInv} to $\min(y, |A| \cdot |B| - y)$. Viewed in the contrapositive, a lower bound for Π_{XInv} is easier to obtain than one for Π_T on the tree T of Figure 2. We first establish a lower bound for Π_{XInv} and then see how it extends to Π_T .

One can think of Π_{XInv} as inversion minimization on the Mann–Whitney tree without allowing swapping the two children of the root. As a result, the problem Π_{XInv} is sensitive to *every* adjacent-rank transposition between non-siblings, and therefore automatically satisfies the partition property (with A and B being the partition classes), so Proposition 8.1 applies. In contrast, the problem Π_T of minimizing inversions on the Mann–Whitney tree may not have the partition property. This is why analyzing Π_{XInv} is a bit simpler, and why we handle it first.

Let $a \doteq |A|$ and $b \doteq |B|$. By the partition property, the average sensitivity of Π_{XInv} equals $\frac{2ab}{a+b}$, which via the Sensitivity Lemma yields a query lower bound of $\Omega(a \log(a))$ for $a \leq b$. To obtain the stronger lower bound of $\Omega((a+b) \log(a))$ we need a more detailed analysis of the connectivity of the permutahedron graph $G(\Pi_{\text{XInv}})$.

For a given ranking ρ , let x_1, \dots, x_a be the elements of A listed in increasing order, and similarly for y_1, \dots, y_b for the elements of B . We define m_1, \dots, m_{b+1} such that for each i , m_i is the number of elements of A between y_{i-1} and y_i . (Here, y_0 and y_{b+1} serve as sentinels with an infinitely low and infinitely high rank.) Similarly, we define n_1, \dots, n_{a+1} as the number of elements in B between successive elements of A . The numbers m_i and n_i are the sizes of the connected components of the successor graph $S(\Pi_{\text{XInv}}, \rho)$ (possibly with some additional zeroes). By Proposition 8.1, the connected component of ρ in $G(\Pi_{\text{XInv}})$ has size

$$(8.25) \quad (m_1)! \cdots (m_{b+1})! (n_1)! \cdots (n_{a+1})!$$

Depending on the values of $m_1, \dots, m_{b+1}, n_1, \dots, n_{a+1}$, some connected components may be much larger than others. We apply the Connectivity Lemma in a similar way as in Lemma 8.1 and only count the rankings ρ that are in small connected components, which are the rankings for which $m_1, \dots, m_{b+1}, n_1, \dots, n_{a+1}$ are bounded. Let m^* and n^* be the minimum integers for which

$$\mathbb{P}[m_1, \dots, m_{b+1} \leq m^*] \geq \frac{3}{4} \quad \text{and} \quad \mathbb{P}[n_1, \dots, n_{a+1} \leq n^*] \geq \frac{3}{4}.$$

By a union bound, the probability that both of these events hold is at least $1/2$. In other words, there are least $(a+b)!/2$ rankings ρ for which $m_1, \dots, m_{b+1} \leq m^*$ and $n_1, \dots, n_{a+1} \leq n^*$.

PROPOSITION 8.3.

$$(8.26) \quad \text{D}(\Pi_T) \geq \frac{(a+b)!}{2(m^*)^{a/m^*} (n^*)^{b/n^*}}.$$

Proof. We consider the rankings for which $m_1, \dots, m_{b+1} \leq m^*$ and $n_1, \dots, n_{a+1} \leq n^*$. We first argue the following upper bound on the size (8.25) of the connected component in $G(\Pi_{\text{XInv}})$ of any such ranking ρ :

$$(8.27) \quad (m_1)! \cdots (m_{b+1})! (n_1)! \cdots (n_{a+1})! \leq (m^*)^{a/m^*} (n^*)^{b/n^*}.$$

For nonnegative integers n , $n^{1/n}$ is increasing. This can be seen by noticing that $\log(n^{1/n})$ is the average of $\log(1), \dots, \log(n)$. As a result, for $i \in [b+1]$, $(m_i)!^{1/m_i} \leq (m^*)!^{1/m^*}$, or equivalently, $(m_i)! \leq (m^*)^{m_i/m^*}$. Using the fact that $m_1 + \dots + m_{b+1} = a$,

$$(m_1)! \cdots (m_{b+1})! \leq (m^*)^{(m_1 + \dots + m_{b+1})/m^*} = (m^*)^{a/m^*}.$$

We can apply similar reasoning to get $(n_1)! \cdots (n_{a+1})! \leq (n^*)^{b/n^*}$. From this, we conclude that the size of the connected components among the rankings under consideration is at most the right-hand side of (8.27).

Since there are at least $(a+b)!/2$ of the rankings under consideration, we derive the stated bound on $D(\Pi_T)$ by applying the Connectivity Lemma. \square

Now, we find concrete bounds on m^* and n^* .

PROPOSITION 8.4.

$$\max\left(1, \frac{a}{b+1}\right) \leq m^* \leq \frac{a+b}{b} \ln(4(b+1)).$$

Symmetrically,

$$\max\left(1, \frac{b}{a+1}\right) \leq n^* \leq \frac{a+b}{a} \ln(4(a+1)).$$

Proof. We first prove the upper bound on m^* . Let k be a positive integer. We compute the probability, over an average ranking ρ , that $m_i > k$ for a specific i . Notice that there is a one-to-one correspondence between ranking ρ and the corresponding sequence of nonnegative integers m_1, \dots, m_{b+1} such that $m_1 + \dots + m_{b+1} = a$, because the ranks of B can be uniquely recovered as $m_1 + 1, m_1 + m_2 + 2, \dots$, and the remaining ranks form A . By stars and bars, there are $\binom{a+b}{b}$ such sequences. Now, if $m_i > k$, then $m_i - (k+1)$ is an arbitrary nonnegative integer, and $m_1 + \dots + (m_i - (k+1)) + \dots + m_{b+1} = a - k - 1$. By stars and bars, there are $\binom{a+b-k-1}{b}$ such sequences. Therefore, the probability that $m_i > k$ is $\binom{a+b-k-1}{b} / \binom{a+b}{b}$. Continuing,

$$\mathbb{P}[m_i > k] = \frac{(a+b-k-1) \cdots (a-k)}{(a+b) \cdots (a+1)} \leq \left(\frac{a+b-k-1}{a+b}\right)^b \leq \exp\left(-b \cdot \frac{k+1}{a+b}\right),$$

where the last step uses the bound $1+x \leq \exp(x)$. By a union bound and taking the complement,

$$(8.28) \quad \mathbb{P}[m_1, \dots, m_{b+1} \leq k] \geq 1 - (b+1) \exp\left(-b \cdot \frac{k+1}{a+b}\right).$$

If k is such that the right-hand side of (8.28) is at least $3/4$, we know that $m^* \leq k$. Solving for k yields the stated bound on m^* .

For the lower bounds, $m^* \geq 1$ because at least one of m_1, \dots, m_{b+1} is at least 1, and $m^* \geq \frac{a}{b+1}$ because $m_1 + \dots + m_{b+1} = a$, and m^* is greater than or equal to the average term in the sum. \square

The first part of Theorem 1.4 now comes from taking the logarithm of $D(\Pi_{\text{XInv}})$ in Proposition 8.3 and using the bounds in Proposition 8.4.

Proof. [Proof of the first part of Theorem 1.4] We mainly make use of the following approximation based on Stirling's formula.

$$(8.29) \quad \ln(n!) = \left(n + \frac{1}{2}\right) \ln(n) - n + O(1).$$

In order to estimate $\ln D(\Pi_{\text{XInv}})$, we need to estimate $\ln((a+b)!) - \frac{a}{m^*} \ln(m^*) - \frac{b}{n^*} \ln(n^*)$. By (8.29),

$$(8.30) \quad \ln((a+b)!) = \left(a+b + \frac{1}{2}\right) \ln(a+b) - (a+b) + O(1)$$

$$(8.31) \quad \frac{a}{m^*} \ln((m^*)!) = \left(a + \frac{a}{2m^*}\right) \ln(m^*) - a + \frac{a}{m^*} O(1)$$

$$(8.32) \quad \frac{b}{n^*} \ln((n^*)!) = \left(b + \frac{b}{2n^*}\right) \ln(n^*) - b + \frac{b}{n^*} O(1)$$

We can use the lower bounds in Proposition 8.4, namely that $m^* \geq 1$ and $n^* \geq \frac{b}{a+1}$, to simplify the occurrences of m^* and n^* in the denominators of (8.31) and (8.32). Therefore,

$$\begin{aligned} \ln D(\Pi_{\text{XInv}}) &\geq a \ln \left(\frac{a+b}{m^*} \right) + b \ln \left(\frac{a+b}{n^*} \right) - \frac{a}{2} \ln(m^*) - \frac{a+1}{2} \ln(n^*) + \frac{1}{2} \ln(a+b) - O(a) \\ &= a \ln \left(\frac{a+b}{m^* \sqrt{m^* n^*}} \right) + \left(b + \frac{1}{2} \right) \ln \left(\frac{a+b}{n^*} \right) - O(a). \end{aligned}$$

Using the upper bounds in Proposition 8.4, absorbing low-order terms, and using the condition that $a \leq b$, we get

$$\log D(\Pi_{\text{XInv}}) \geq \Omega \left(a \log \left(\frac{b\sqrt{ab}}{a+b} \right) + b \log(a) \right) \geq \Omega(a \log(\sqrt{ab}/2) + b \log(a)) = \Omega((a+b) \log(a)).$$

□

Evaluating the Mann–Whitney statistic. We now argue how the second part of Theorem 1.4 follows, i.e., that the lower bound of $\Omega((a+b) \log(a))$ holds for inversion minimization on the Mann–Whitney tree T of Figure 2 with $a \leq b$. We do so by tweaking our lower bound argument for Π_{XInv} to the setting of Π_T .

How does the permutahedron graph for Π_T relate to the one for Π_{XInv} ? The problem Π_T is a coarsening of the problem Π_{XInv} : Output values y and $ab - y$ for Π_{XInv} are both mapped to $\min(y, ab - y)$ under Π_T . This means that all edges present in $G(\Pi_{\text{XInv}})$ are also present in $G(\Pi_T)$, but there may be more, and some of the connected components in $G(\Pi_{\text{XInv}})$ corresponding to output value y , may be merged in $G(\Pi_T)$ with some of the connected components of $G(\Pi_{\text{XInv}})$ corresponding to output value $ab - y$. However, by the reasoning behind Proposition 4.1, edges in $G(\Pi_{\text{XInv}})$ can only go between rankings whose value under Π_{XInv} differ by at most one. This means that the above merging of connected components can only happen if the difference between y and $ab - y$ is 1, i.e., for the values $\lfloor ab/2 \rfloor$ and $\lceil ab/2 \rceil$, and only if ab is odd. In fact, this is exactly the situation that we analyzed in Figure 9, where v coincides with the root of T .

If we ignore the rankings with value $\lfloor ab/2 \rfloor$ or $\lceil ab/2 \rceil$ under Π_{XInv} , our lower bound argument for Π_{XInv} carries over verbatim to Π_T , except that on the right-hand side of Proposition 8.3 the factor of $\frac{1}{2}$ is replaced by $\frac{1}{2} - W$, where W represents the fraction of rankings with value $\lfloor ab/2 \rfloor$ or $\lceil ab/2 \rceil$ under Π_{XInv} . Lemma 5.1 tells us that $W \leq 2C/\sqrt{ab(a+b)}$, where C denotes the constant from the lemma. Thus, we obtain a lower bound for $D(\Pi_T)$ that is a negligible fraction smaller than the one for $D(\Pi_{\text{XInv}})$. Taking logarithms, we obtain the same lower bound for the query complexity up to an additive term. In particular, we obtain a query lower bound of $\Omega((a+b) \log(a))$ for Π_T in case $a \leq b$. This is the second part of Theorem 1.4.

Acknowledgements

We would like to thank Greta Panova, Robin Pemantle, and Richard Stanley for pointers regarding Gaussian polynomials, Stasys Jukna for answering questions about the complexity measure N , and the anonymous reviewers for helpful suggestions.

References

- [BFP⁺73] Manuel Blum, Robert W. Floyd, Vaughan Pratt, Ronald L. Rivest, and Robert E. Tarjan. Time bounds for selection. *Journal of Computer and System Sciences*, 7(4):448–461, 1973.
- [BO83] Michael Ben-Or. Lower bounds for algebraic computation trees. *Proceedings of the 15th Annual Symposium on Theory of Computing (STOC)*, pages 80–86, 01 1983.
- [CFJ⁺10] Jean Cardinal, Samuel Fiorini, Gwenaël Joret, Raphaël M. Jungers, and J. Ian Munro. An efficient algorithm for partial order production. *SIAM Journal on Computing*, 39(7):2927–2940, 2010.
- [Deg82] Richard Degerman. Ordered binary trees constructed through an application of Kendall’s tau. *Psychometrika*, 47(4):523–527, 1982.
- [DZ99] Dorit Dor and Uri Zwick. Selecting the median. *SIAM Journal on Computing*, 28(5):1722–1758, 1999.
- [DZ01] Dorit Dor and Uri Zwick. Median selection requires $(2+\epsilon)n$ comparisons. *SIAM Journal on Discrete Mathematics*, 14(3):312–325, 2001.

- [Epp07] David Eppstein. A permutohedron. Wikipedia User Gallery, 2007.
- [JRSW99] Stasys Jukna, Alexander A. Razborov, Petr Savický, and Ingo Wegener. On P versus NP cap co-np for decision trees and read-once branching programs. *Computational Complexity*, 8:357–370, 1999.
- [KMMS05] Kanela Kaligosi, Kurt Mehlhorn, J. Ian Munro, and Peter Sanders. Towards optimal multiple selection. In *International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 103–114, 2005.
- [Mil69] George A. Miller. Psychological method to investigate verbal concepts. *Journal of Mathematical Psychology*, 6:169–191, 1969.
- [MPP20] Stephen Melczer, Greta Panova, and Robin Pemantle. Counting partitions inside a rectangle. *SIAM Journal on Discrete Mathematics*, 34(4):2388–2410, 2020.
- [MW47] Henry B. Mann and Donald R. Whitney. On a test of whether one of two random variables is stochastically larger than the other. *The Annals of Mathematical Statistics*, 18(1):50–60, 1947.
- [O’D14] Ryan O’Donnell. *Analysis of Boolean Functions*. Cambridge University Press, 2014.
- [Sch76] Arnold Schönhage. The production of partial orders. *Astérisque*, 38–39:229–246, 1976.
- [Ser21] Igor S. Sergeev. On the upper bound of the complexity of sorting. *Computational Mathematics and Mathematical Physics*, 61(2):329–346, 2021.
- [SZ16] Richard P. Stanley and Fabrizio Zanello. Some asymptotic results on q -binomial coefficients. *Annals of Combinatorics*, 20(3):623–634, 2016.
- [Tak86] Lajos Takács. Some asymptotic formulas for lattice paths. *Journal of Statistical Planning and Inference*, 14(1):123–142, 1986.
- [Yao77] Andrew Chi-Chin Yao. Probabilistic computations: Toward a unified measure of complexity. In *18th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 222–227, 1977.