

Prefix Discrepancy, Smoothed Analysis, and Combinatorial Vector Balancing

Nikhil Bansal ✉

University of Michigan, Ann Arbor, MI, USA

Haotian Jiang ✉

University of Washington, Seattle, WA, USA

Raghu Meka ✉

University of California, Los Angeles, CA, USA

Sahil Singla ✉

Georgia Institute of Technology, Atlanta, GA, USA

Makrand Sinha ✉

Simons Institute, Berkeley, CA, USA

University of California, Berkeley, CA, USA

Abstract

A well-known result of Banaszczyk in discrepancy theory concerns the *prefix discrepancy* problem (also known as the *signed series* problem): given a sequence of T unit vectors in \mathbb{R}^d , find \pm signs for each of them such that the signed sum vector along any prefix has a small ℓ_∞ -norm? This problem is central to proving upper bounds for the Steinitz problem, and the popular Komlós problem is a special case where one is only concerned with the final signed sum vector instead of all prefixes.

Banaszczyk gave an $O(\sqrt{\log d + \log T})$ bound for the prefix discrepancy problem. We investigate the tightness of Banaszczyk’s bound and consider natural generalizations of prefix discrepancy:

- We first consider a *smoothed analysis* setting, where a small amount of additive noise perturbs the input vectors. We show an exponential improvement in T compared to Banaszczyk’s bound. Using a primal-dual approach and a careful chaining argument, we show that one can achieve a bound of $O(\sqrt{\log d + \log \log T})$ with high probability in the smoothed setting. Moreover, this smoothed analysis bound is the best possible without further improvement on Banaszczyk’s bound in the worst case.
- We also introduce a generalization of the prefix discrepancy problem to arbitrary DAGs. Here, vertices correspond to unit vectors, and the discrepancy constraints correspond to paths on a DAG on T vertices – prefix discrepancy is precisely captured when the DAG is a simple path. We show that an analog of Banaszczyk’s $O(\sqrt{\log d + \log T})$ bound continues to hold in this setting for adversarially given unit vectors and that the $\sqrt{\log T}$ factor is unavoidable for DAGs. We also show that unlike for prefix discrepancy, the dependence on T cannot be improved significantly in the smoothed case for DAGs.
- We conclude by exploring a more general notion of vector balancing, which we call *combinatorial vector balancing*. In this problem, the discrepancy constraints are generalized from paths of a DAG to an arbitrary set system. We obtain near-optimal bounds in this setting, up to poly-logarithmic factors.

2012 ACM Subject Classification Mathematics of computing → Combinatoric problems

Keywords and phrases Prefix discrepancy, smoothed analysis, combinatorial vector balancing

Digital Object Identifier 10.4230/LIPIcs.ITCS.2022.13

Related Version *Full Version*: <https://arxiv.org/abs/2111.07049>

Funding *Nikhil Bansal*: Supported in part by the NWO VICI grant 639.023.812. The work was done when the author was at CWI, Amsterdam.

Haotian Jiang: Supported in part by the National Science Foundation, Grant Number CCF-1749609, CCF-1740551, DMS-1839116.



© Nikhil Bansal, Haotian Jiang, Raghu Meka, Sahil Singla, and Makrand Sinha; licensed under Creative Commons License CC-BY 4.0

13th Innovations in Theoretical Computer Science Conference (ITCS 2022).

Editor: Mark Braverman; Article No. 13; pp. 13:1–13:22

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Raghu Meka: Supported by NSF Grant CCF-1553605.

Sahil Singla: The work was partly done when the author was at Princeton University.

Makrand Sinha: Work done while at CWI Amsterdam and supported by the NWO VICI grant 639.023.812.

Acknowledgements The authors would like to thank the anonymous reviewers of ITCS 2022 for helpful comments.

1 Introduction

Given a sequence of T vectors v_1, \dots, v_T in the Euclidean unit ball \mathbb{B}_2^d (i.e., of ℓ_2 -length at most 1), the *prefix discrepancy* problem, also widely known as the *signed series* problem, asks to find a coloring¹ $x \in \{\pm 1\}^T$ to minimize $\max_{\tau \in [T]} \|\sum_{t \leq \tau} x_t v_t\|_\infty$. The classical Komlós problem is a special case of this problem where we want to minimize only the final discrepancy $\|\sum_{t \leq T} x_t v_t\|_\infty$. The prefix discrepancy problem was introduced by Spencer [25] who showed that there always exists a coloring with a discrepancy bound that only depends on d (i.e., independent of T). This was later improved to $O(d)$ by Barany-Grinberg [8]. Since then, the prefix-discrepancy problem has been greatly studied and has found several important applications. For instance, it implies a bound on the classical Steinitz problem [11, 28] on the rearrangement of vector sequences, it appears in online discrepancy problems where one is interested in bounding the discrepancy at all times [5, 4, 1], and it also implies the best known bound for Tusnády’s problem [21]. Additionally, the Steinitz problem itself has many applications, including faster algorithms for solving integer programs [9, 12, 16] and scheduling [7, 24].

In a remarkable result, Banaszczyk [3] showed that the prefix-discrepancy problem admits a coloring of $O(\sqrt{\log d + \log T})$ discrepancy, thereby exponentially improving the dependency on d in the Barany-Grinberg bound, but incurring a $\sqrt{\log T}$ dependence on the number of vectors T .

► **Theorem 1** ([3]). *Given $v_1, \dots, v_T \in \mathbb{B}_2^d$, there always exists a signing $x \in \{\pm 1\}^T$ such that*

$$\max_{\tau \in [T]} \left\| \sum_{t \leq \tau} x_t v_t \right\|_\infty = O(\sqrt{\log d + \log T}) .$$

Given the many applications of prefix discrepancy, a natural question is whether we can improve this dependency on T in Theorem 1 while still being poly-logarithmic in d . Optimistically, one could even hope for removing the dependence on T altogether. Variants and special cases of this question have previously been asked [26, 3] and remain open (see full version). Besides being a natural (and beautiful) question at the border of combinatorics and geometry, further motivation to study prefix discrepancy comes from the Komlós problem: there is currently no known separation between the prefix discrepancy and the Komlós problems². Thus, obtaining better upper or lower bounds for prefix discrepancy will also help us understand the avenues for making progress on the Komlós conjecture.

¹ A coloring is also referred to as a signing in the literature. We will use these two words interchangeably.

² Note that for the Komlós problem one can always assume that $T \leq d$ using a standard linear algebraic trick, while this might not be possible if one considers all prefixes of the signed sum vector.

In this work, we strengthen Banaszczyk’s result for the prefix discrepancy problem in two natural ways. First, we exponentially improve the dependence on T in Banaszczyk’s prefix discrepancy bound in a *smoothed setting*. Next, we extend the prefix discrepancy problem beyond *prefixes of a path*, to *prefixes of a given DAG*, as well as to more general combinatorial settings.

1.1 Smoothed Analysis of Prefix Discrepancy

Our first result concerns the prefix discrepancy problem in a *smoothed analysis* setting. The motivation is that for several important problems, although the worst-case algorithmic bounds are bad, these worst-case instances could be brittle. The study of algorithmic problems in a smoothed model was initiated by Spielman and Teng [27]. They showed that although the popular Simplex method can take an exponential time to solve an adversarially chosen linear programming instance, a slight perturbation of this instance allows the Simplex method to run in polynomial time. Since for many applications the input already contains some measurement noise, the analysis of algorithms for slight perturbations of the input is a natural model. Since the work of Spielman and Teng on linear programs, smoothed analysis has been successfully used for analyzing several other problems such as Local Search, Mechanism Design, and computing Pareto curves (see, e.g., the book [23]).

In the smoothed setting of prefix discrepancy³, an adversary first chooses T vectors $\bar{v}_1, \dots, \bar{v}_T \in \mathbb{B}_2^d$ and then each vector \bar{v}_t is perturbed by a random independent noise $\hat{v}_t \in \mathbb{B}_2^d$ with covariance $\text{Cov}(\hat{v}_t) \succcurlyeq \epsilon^2 I_d$. The goal of the algorithm is to solve the prefix discrepancy problem on input vectors v_1, \dots, v_T where $v_t := \bar{v}_t + \hat{v}_t$. In other words, find a signing $x \in \{\pm 1\}^T$ that minimizes

$$\text{PrefixDisc}(v_1, \dots, v_T) := \min_{x \in \{\pm 1\}^T} \max_{\tau \in [T]} \left\| \sum_{t \leq \tau} x_t v_t \right\|_{\infty}.$$

This smoothed setting also generalizes stochastic discrepancy models where the input vectors are drawn i.i.d. from a given distribution, i.e., when we take all vectors $\bar{v}_t = 0$ and the noise vectors \hat{v}_t are drawn i.i.d.

In the above smoothed setting, we exponentially improve the dependency on T in Banaszczyk’s bound.

► **Theorem 2** (Smoothed analysis upper bound). *Suppose for each $t \in [T]$ the noise $\hat{v}_t \in \mathbb{B}_2^d$ is sampled independently and satisfies covariance $\text{Cov}(\hat{v}_t) \succcurlyeq \epsilon^2 I_d$ for some $\epsilon \geq 1/\text{poly}(d, \log T)$. Then, the prefix discrepancy in the smoothed setting is $O(\sqrt{\log d + \log \log T})$ with probability at least $1 - 1/\text{poly}(T)$.*

Note that the assumption on noise in Theorem 2 is quite mild and captures many natural perturbation models. For instance, the theorem holds even if each vector is *unchanged* with probability $1 - \epsilon$ for $\epsilon = 1/\text{poly}(d, \log T)$ and we perturb a *single random* coordinate by an additive $\pm \epsilon$ with probability ϵ . This is in contrast to the recent line of work that studies discrepancy minimization for stochastic inputs [13, 6, 15, 22], where the input stochasticity is used very heavily.

Theorem 2 is also interesting since it suggests if there is an adversarial $\Omega(\sqrt{\log d + \log T})$ lower bound instance for the prefix discrepancy problem then such an instance must be quite brittle: a slight perturbation of the instance will admit much better upper bounds.

³ A recent independent work has also introduced a smoothed model for discrepancy minimization [14]. However, they consider an *online* setting, and hence their model and results are incomparable.

13:4 Prefix Discrepancy, Smoothed Analysis, and Combinatorial Vector Balancing

Our proof of the theorem relies on a primal-dual approach and a careful chaining argument (see Section 1.4). We note that it does not give us an efficient algorithm since it relies on Theorem 1, which we do not know how to do algorithmically for prefixes. However, one can get a slightly weaker bound of $O(\log d + \log \log T)$ in polynomial time using a recent algorithm of [1] instead of Theorem 1.

We also show that Theorem 2 is tight, unless there is a breakthrough for adversarial prefix discrepancy.

► **Theorem 3** (Informal Theorem 16). *The bound in Theorem 2 is the best possible for the smoothed setting assuming the tightness of Banaszczyk’s $O(\sqrt{\log d + \log T})$ bound in the adversarial setting.*

1.2 Prefix Discrepancy for DAGs

We next consider a substantial generalization of the prefix discrepancy problem to arbitrary directed acyclic graphs (DAGs). To motivate it, let us consider a slightly different view of the prefix discrepancy problem with input vectors $v_1, v_2, \dots, v_T \in \mathbb{B}_2^d$. Consider the directed path, Path_T , on T vertices and suppose that we have a vector v_t at every vertex t on the path. The prefix discrepancy problem then corresponds to assigning signs to the vertices of Path_T and looking at the signed-sums of vectors associated with the vertices on any path starting from the root in Path_T ⁴.

We investigate the case where the *base graph* instead of being a simple path is an arbitrary DAG. In particular, given a DAG $G = ([T], E)$ with vertices $[T] := \{1, \dots, T\}$, we denote by $\text{prefix}(G)$ the following set family on the ground set $[T]$: First order the vertices of G topologically; now a set $S \subseteq [T]$ is included in $\text{prefix}(G)$ iff S is a path in G that starts at the topologically ordered root vertex. For the generalization, given T vectors in \mathbb{B}_2^d corresponding to the T vertices, our goal is to find a signing of the vertices (equivalently, the vectors) to minimize the maximum discrepancy with respect to any *prefix constraint* in $\text{prefix}(G)$. In other words, given input vectors $v_1, \dots, v_T \in \mathbb{B}_2^d$, find a signing $x \in \{\pm 1\}^T$ that minimizes

$$\text{DAGDisc}_G(v_1, \dots, v_T) := \min_{x \in \{\pm 1\}^T} \max_{S \in \text{prefix}(G)} \left\| \sum_{t \in S} x_t v_t \right\|_\infty.$$

Clearly, if G is the simple path on T vertices, we recover prefix discrepancy.

We prove the following generalization of Theorem 1 in terms of the hereditary discrepancy, HerDisc , of the set system $\text{prefix}(G)$, which is the largest possible discrepancy of any “sub-DAG” (see Definition 9). This is a natural parameter since even for $d = 1$ (i.e., scalars), HerDisc is a lower bound on DAGDisc_G if we take the vectors $\{v_t\}_{t \in [T]}$ to be 0/1-valued scalars.

► **Theorem 4** (Prefix discrepancy for DAGs). *Given a T -node DAG G and input vectors $v_1, \dots, v_T \in \mathbb{B}_2^d$, we have $\text{DAGDisc}_G(v_1, \dots, v_T) \leq \text{HerDisc}(\text{prefix}(G)) \cdot O(\sqrt{\log d + \log T})$. Additionally, if $d = 1$, then $\text{DAGDisc}_G(v_1, \dots, v_T) = O(\text{HerDisc}(\text{prefix}(G)))$.*

Theorem 4 is a generalization of Theorem 1 because $\text{HerDisc}(\text{prefix}(G)) = 1$ when G is a path on $[T]$, as in Banaszczyk’s setting. We prove the above theorem by showing a novel structural result about DAGs, which allows us to approximate the given DAG with

⁴ The restriction of involving the root is only notational and can be relaxed at the cost of a factor of two.

rooted-trees with “distance” given by $\text{HerDisc}(\text{prefix}(G))$ (see Lemma 8). Next, to handle any rooted-tree \mathcal{T} , we use techniques inspired from [3] and show that for any given vectors $v_1, \dots, v_T \in \mathbb{B}_2^d$, we have $\text{DAGDisc}_{\mathcal{T}}(v_1, \dots, v_T) = O(\sqrt{\log d + \log T})$ ⁵.

It is an intriguing open problem whether the dependence of $\sqrt{\log T}$ in Banaszczyk’s bound can be improved. We show, however, that the $\sqrt{\log T}$ in Theorem 4 cannot be improved even for rooted-trees and even in the two-dimensional case. This gives an interesting contrast between trees and paths, even though they have the same $O(\sqrt{\log d + \log T})$ upper bound.

► **Theorem 5** (Prefix Discrepancy Lower Bound). *For any integer $T > 0$, there exists a T -node rooted tree \mathcal{T} and input vectors $v_1, \dots, v_T \in \mathbb{B}_2^2$ such that $\text{DAGDisc}_{\mathcal{T}}(v_1, \dots, v_T) = \Omega(\sqrt{\log T})$.*

Note that for $d = 2$, [8] gives a bound of $O(1)$ on the prefix discrepancy problem (irrespective of T). Further, for the one-dimensional case, the discrepancy of any rooted tree is at most 1.

One might wonder if our improvement to $\sqrt{\log d + \log \log T}$ in the smoothed setting (Theorem 2) could also be generalized to DAGs. We show that such an improvement is impossible, in particular there is an $\Omega(\sqrt{\log T / \log \log T})$ lower bound even for trees in a completely stochastic setting.

► **Theorem 6** (Prefix Discrepancy Lower Bound in the Random Setting). *Let the dimension $d = 2$. There exists a T -node rooted tree \mathcal{T} such that when unit vectors $v_t \in \mathbb{R}^d$ are drawn i.i.d. uniformly from \mathbb{S}^{d-1} , then w.h.p. every ± 1 -coloring satisfies $\text{DAGDisc}_{\mathcal{T}}(v_1, \dots, v_T) = \Omega(\sqrt{\log T / \log \log T})$.*

1.3 Combinatorial Vector Balancing

Finally, we generalize the vector balancing problem beyond prefixes and introduce the *Combinatorial Vector Balancing* problem – here the discrepancy constraints come from arbitrary set systems. In particular, given a set family \mathcal{S} on the ground set $[T]$ and input vectors $v_1, \dots, v_T \in \mathbb{B}_2^d$, find a signing $x \in \{\pm 1\}^T$ that minimizes

$$\text{CombDisc}_{\mathcal{S}}(v_1, \dots, v_T) := \min_{x \in \{-1, +1\}^n} \max_{S \in \mathcal{S}} \left\| \sum_{i \in S} x_i v_i \right\|_{\infty}.$$

Notice that $\text{CombDisc}_{\mathcal{S}}(v_1, \dots, v_T)$ captures $\text{DAGDisc}_G(v_1, \dots, v_T)$ when the set system \mathcal{S} is given by $\text{prefix}(G)$, and therefore also generalizes the prefix discrepancy problem. Moreover, $\text{CombDisc}_{\mathcal{S}}(v_1, \dots, v_T)$ can be viewed as a natural generalization of the combinatorial discrepancy of the set system \mathcal{S} , where v_1, \dots, v_T are all scalars and equal to 1. Combinatorial discrepancy of set systems has been well-studied in many different contexts and we refer the reader to the books [10, 19] for an overview.

We also define the quantity

$$\text{CombDisc}(\mathcal{S}, d) := \sup_{v_1, \dots, v_T \in \mathbb{B}_2^d} \text{CombDisc}_{\mathcal{S}}(v_1, \dots, v_T),$$

which can be viewed as a generalization of hereditary discrepancy of the set system \mathcal{S} . In particular, for one dimensional input vectors, we have that $\text{HerDisc}(\mathcal{S}) \leq \text{CombDisc}(\mathcal{S}, 1)$ by taking the corresponding scalars $\{v_t\}_{t \in [T]}$ to be 0/1-valued.

⁵ It is easy to see that for any rooted tree \mathcal{T} , we have that $\text{HerDisc}(\text{prefix}(\mathcal{T})) = 1$, as in the case of prefixes of a path.

13:6 Prefix Discrepancy, Smoothed Analysis, and Combinatorial Vector Balancing

Our results from the previous section imply that if $\mathcal{S} = \text{prefix}(G)$ for some DAG G , then $\text{CombDisc}(\mathcal{S}, 1) = \Theta(\text{HerDisc}(\mathcal{S}))$ and $\text{CombDisc}(\mathcal{S}, d) \leq \text{HerDisc}(\mathcal{S}) \cdot O(\sqrt{\log d + \log T})$. A natural question that comes up is whether an analogous relation also holds for a general set system \mathcal{S} . We answer this question, up to poly-logarithmic factors.

► **Theorem 7 (Combinatorial Vector Balancing Upper Bound).** *Let \mathcal{S} be a set system on the ground set $[T]$. Given input vectors $v_1, \dots, v_T \in \mathbb{B}_2^d$, we have that*

$$\text{CombDisc}_{\mathcal{S}}(v_1, \dots, v_T) \leq \text{HerDisc}(\mathcal{S}) \cdot \log(\min\{|\mathcal{S}|, T\}) \cdot \sqrt{\log d + \log |\mathcal{S}|}.$$

Moreover, the upper bound is constructive in the sense that there is a polynomial time algorithm to find such a coloring.

Note that the above also implies that $\text{CombDisc}(\mathcal{S}, d)$ is bounded by the right hand side in the above theorem.

Our proof of Theorem 7 is via the γ_2 -norm, which is used in [18] as a tool to approximate $\text{HerDisc}(\mathcal{S})$. We show that since the vectors v_1, \dots, v_T are in \mathbb{B}_2^d , the γ_2 norm of the left hand side (viewed as a $d|\mathcal{S}| \times T$ matrix) is at most $\gamma_2(A_{\mathcal{S}})$, where $A_{\mathcal{S}}$ is the incidence matrix of the set system \mathcal{S} . This allows us to apply the connection between γ_2 -norm and hereditary discrepancy in [18] to prove Theorem 7. More details on the proof can be found in Section 5.

1.4 Technical Overview

1.4.1 Smoothed Analysis of Prefix Discrepancy

To achieve the $\sqrt{\log \log T}$ dependence in Theorem 2 in the smoothed setting, we group the input vectors into consecutive blocks of n vectors, with a total of T/n blocks, where $n = \text{poly}(d, \log T)$. Now Banaszczyk's result in Theorem 1 implies that the prefix discrepancy within each block (viewed as an instance with n input vectors) is $\Delta = O(\sqrt{\log d + \log n}) = O(\sqrt{\log d + \log \log T})$. However, the discrepancy for different blocks might add up arbitrarily. Our idea to get around this issue is to color the vectors block-by-block while using the current block to cancel any previously accumulated discrepancy.

Coloring Strategy for a Block. We color the first block of n vectors by using Theorem 1 and for the subsequent blocks, we adopt the following strategy. Suppose that we have assigned signs to the previous blocks and let $M \in \mathbb{R}^{d \times n}$ be the matrix whose columns are the current block of vectors that need to be colored. To find a coloring $x \in \{\pm 1\}^n$ for a given block, we proceed in two phases: first we find a *fractional* signing $x \in [-1, 1]^n$ of the columns of M so that

1. We cancel out the discrepancy vector w of the previous blocks, i.e., $Mx + w = 0$, and
2. The intermediate discrepancy $\|M_I x\|_{\infty} \leq O(\Delta)$ for every prefix interval $I \subseteq [n]$ and $M_I \in \mathbb{R}^{d \times n}$ is the matrix obtained from M by zeroing out columns not in I .

If we can find such a fractional solution, then a variant of the *bit-by-bit rounding procedure* of [17] allows us to round x to a full coloring $x^* \in \{-1, 1\}^n$ so that the prefix discrepancy for block M due to the rounding is at most $\Delta := \Theta(\sqrt{\log d + \log \log T})$ (Lemma 18). In particular, we show that if the previous discrepancy vector w satisfies $\|w\|_{\infty} \leq \Delta$, then the discrepancy vector after block M satisfies $\|w + Mx^*\|_{\infty} = \|M(x^* - x)\|_{\infty} \leq \Delta$, which allows us to proceed with coloring the next block inductively. The central part of the argument is to find a fractional signing x satisfying the above two properties for a given block, and we describe the ideas behind it next.

Primal-Dual Approach for Feasibility. We need to show that for any vector $\|w\|_\infty \leq \Delta$ (w is the discrepancy vector and the assumption on norm is the inductive hypothesis), with probability $1 - 1/\text{poly}(T)$, we can find a fractional signing $x \in [-1, 1]^n$ satisfying (a) $Mx + w = 0$, and (b) $\|M_I x\|_\infty \leq O(\Delta)$ for every prefix interval I . These linear constraints are naturally captured by the feasibility of the following *stochastic* linear program $\text{LP}(\mathcal{I})$, where \mathcal{I} is the set of all prefixes of $[n]$ (see Section 3.1.2).

$$\begin{aligned} (Mx)_i &= -w_i && \forall i \in [d] \\ -2\Delta &\leq (M_I x)_i \leq 2\Delta && \forall I \in \mathcal{I}, \forall i \in [d] \\ -1 &\leq x_j \leq 1 && \forall j \in [n] \end{aligned} \quad \text{LP}(\mathcal{I})$$

Our goal is to show that the above linear program is feasible with high probability over M . At a high-level this is reminiscent of the approach of [6] who also use a stochastic LP to get cancellations, but the complexities are quite different: their LP is simpler and they work in the *completely random* setting as opposed to the more challenging smoothed-analysis setting as we do. Recall that for us, the columns of M are generated in the smoothed analysis model: starting with an arbitrary matrix, we perturb the columns independently with random vectors whose covariance is lower bounded (in the PSD ordering) by $\epsilon^2 I_d$. Using a primal-dual approach, we derive the following sufficient condition for the feasibility of $\text{LP}(\mathcal{I})$: For all dual variables $(y, \{\alpha_I\}_{I \in \mathcal{I}})$ ($y, \alpha_I \in \mathbb{R}^d$) satisfying $\|y\|_1 = d$ and $\sum_{I \in \mathcal{I}} \|\alpha_I\|_1 \leq d/2$, we must have

$$\|y^\top M - \sum_{I \in \mathcal{I}} \alpha_I^\top M_I\|_1 \geq \Delta \cdot d. \quad (1)$$

For intuition on why the above might be true, consider the case of a fixed choice of dual variables $(y, \{\alpha_I\}_{I \in \mathcal{I}})$. Denoting v_j as the j th column of M , the left hand side of (1) can be written as $\sum_{j=1}^n |(y - \sum_{I: j \in I} \alpha_I)^\top v_j|$. The latter is a random variable (because of the perturbations in v_j 's) that is a sum of n independent terms, each with variance at least $\epsilon^2 \cdot \|(y - \sum_{I: j \in I} \alpha_I)\|_2^2$, while the right hand side of Equation 1 is only of order $\sqrt{\log n}$. Therefore, picking $n = \text{poly}(d, \log T)$ to be large enough, Chernoff bound shows that the left hand side is larger than the right hand side with probability $1 - \exp(-n)$, for any fixed setting of the dual variables $(y, \{\alpha_I\}_{I \in \mathcal{I}})$ that satisfies $\|y\|_1 = d$ and $\sum_{I \in \mathcal{I}} \|\alpha_I\|_1 \leq d/2$. However, there are n prefix intervals in \mathcal{I} and each α_I needs to be discretized to multiples of $1/\text{poly}(n, d)$, so the total number of choices of $(y, \{\alpha_I\}_{I \in \mathcal{I}})$ is $\text{poly}(n, d)^n \gg \exp(n)$. Thus, even though, inequality in Equation 1 holds with high probability for a fixed dual solution, we cannot use a union bound over all the choices of $(y, \{\alpha_I\}_{I \in \mathcal{I}})$ to show that $\text{LP}(\mathcal{I})$ is feasible with high probability.

Block Decomposition and Chaining. To overcome the difficulty above, we work with a different family of intervals I that we call a *block decomposition* $\mathcal{I}_{\text{block}}$. These intervals satisfy:

1. Every prefix interval of $[n]$ can be written as a disjoint union of $I_1, I_2 \in \mathcal{I}_{\text{block}}$ which allows us to deduce the feasibility of $\text{LP}(\mathcal{I})$ from the feasibility of $\text{LP}(\mathcal{I}_{\text{block}})$.
2. There are only $o(n)$ “long” intervals in $\mathcal{I}_{\text{block}}$. We are able to bypass the union bound by employing a two-step union bound, best thought of as a *chaining* argument, to show that for any fixed choice of $(y, \{\alpha_I\}_{I_{\text{long}}})$, with probability at least $1 - \exp(-n)$, inequality (1) holds for **all** choices of $\{\alpha_I\}_{I_{\text{short}}}$. We once again cannot do a union bound over all choices of $\{\alpha_I\}_{I_{\text{short}}}$, but use a more subtle argument based on the values of α_I to prove this claim (see Lemma 20). Once we do so, we can take a naive union bound only over the $\exp(o(n))$ possible choices of $(y, \{\alpha_I\}_{I_{\text{long}}})$, thus showing that $\text{LP}(\mathcal{I}_{\text{block}})$ is almost always feasible. The fractional signing $x \in [-1, 1]^n$ can then be found by solving the corresponding LP feasibility problem.

1.4.2 Prefix Discrepancy for Trees and DAGs

The upper bound. The proof of Theorem 4 has two parts. First, we show that Banaszczyk’s result in Theorem 1 can be generalized to prefixes of *rooted trees* (which can be naturally viewed as DAGs by orienting the edges from the root to leaves). In particular we show that the same $O(\sqrt{\log d + \log T})$ bound holds for prefixes of any rooted tree (Lemma 21). This proof is inspired from Banaszczyk’s original proof of Theorem 1: we define intermediate auxiliary convex bodies inductively from the leaves to the root using a symmetrization operation for convex bodies defined in [3] (see Definition 12). We can show that all these intermediate bodies have large Gaussian measure and thus invoking the classic result of [2] gives us a coloring of the vectors with the above bound.


The second part is the crux of our proof, where we reduce the case of an arbitrary DAG G to that of a rooted tree. This is where the notion of hereditary discrepancy of the set system $\text{prefix}(G)$ comes in. In particular, we show:

► **Lemma 8 (Reducing DAGs to Trees).** *Given a DAG $G = ([T], E)$, there exists a (rooted) spanning tree $\mathcal{T} \subseteq G$ such that every directed path \mathcal{P} in G has at most $4 \cdot \text{HerDisc}(\text{prefix}(G))$ non-tree edges.*

Since the removed edges break each path $\mathcal{P} \in \text{prefix}(G)$ into $O(\text{HerDisc}(\text{prefix}(G)))$ sub-paths in the tree \mathcal{T} , the discrepancy along \mathcal{P} is at most $\text{HerDisc}(\text{prefix}(G)) \cdot O(\sqrt{\log d + \log T})$ using the bound for prefixes for trees (since any sub-path in the tree can be written as a difference of two prefixes). Below we describe the implications and ideas behind the above combinatorial characterization of hereditary discrepancy for prefixes of a DAG.

Hereditary Discrepancy for DAGs. Lemma 8 also gives us a combinatorial characterization of $\text{HerDisc}(\text{prefix}(G))$ as the proximity of DAG G to a tree \mathcal{T} , where the “distance” is the maximum number of non-tree edges along any path in G (see Corollary 23). This characterization might be of independent interest.

For the proof of Lemma 8, we construct such a tree \mathcal{T} by going through the vertices of the DAG in the reverse topological order and adding edges inductively. However, it’s not immediately clear what induction hypothesis we want to maintain throughout the process and how the number of non-tree edges can be related to $\text{HerDisc}(\text{prefix}(G))$. The key observation, the one that brings $\text{HerDisc}(\text{prefix}(G))$ into the picture is that $\text{HerDisc}(\text{prefix}(G))$ is large if and only if G contains a long *chain structure*. This is a subgraph C that looks like the DAG shown in Figure 1, and it has the critical property that the set family $\text{prefix}(C)$ is of exponential in ℓ size where ℓ is the length of the chain, implying that hereditary discrepancy of $\text{prefix}(C) = \Omega(\ell)$.



■ **Figure 1** An example of a chain structure of length $\ell - 1$. Note that there are two vertex disjoint paths between a_i and a_{i+1} for each $i \in [\ell - 1]$ and one of these contain at least one additional vertex apart from a_i and a_{i+1} .

Using a careful induction, we are able to show that whenever during the construction of the tree \mathcal{T} in the reverse topological order, the maximum number of non-tree edges increases on some path, then the length of the longest chain and thus the hereditary discrepancy (of the prefixes of the sub-DAG) must also increase.

The lower bounds. Both our lower bounds in Theorems 5 and 6 make use of a canonical construction. In particular, if one is considering prefixes of a complete binary tree with T nodes, then there is a canonical construction of two-dimensional vectors $\{v_t\}_{t \in [T]}$ such that for any coloring $x \in \{\pm 1\}^T$, there always exists a path so that the next vector on that path is always orthogonal to the current discrepancy vector. Since the square of the ℓ_2 -norm of the discrepancy vector always increases as one goes down the path, this implies that the prefix discrepancy is at least $\Omega(\sqrt{h})$ where $h \approx \log T$ is the height of the tree. For our lower bound in the stochastic setting, we show that for any ℓ -ary tree where $\ell = \text{polylog}(T)$, if one chooses vectors $\{v_t\}_{t \in [T]}$ randomly, then with high probability, such a canonical instance is embedded in this random instance.

2 Preliminaries

Notation. Throughout this paper, we use $[k]$ to denote the set $\{1, 2, \dots, k\}$ and $T \in \mathbb{Z}_{>0}$ to denote the number of vectors, which we assume to be d -dimensional. We use \mathbb{B}_p^d to denote the unit ℓ_p -ball in \mathbb{R}^d and \mathbb{S}^{d-1} to denote the unit Euclidean sphere in \mathbb{R}^d . We denote by μ_d the d -dimensional standard Gaussian measure on \mathbb{R}^d . For a random vector $v \in \mathbb{R}^d$, let $\text{Cov}(v)$ denote the $d \times d$ covariance matrix of the distribution of v . For any two positive semidefinite matrices A and B , we use $A \succcurlyeq B$ to denote that $A - B$ is positive semidefinite. Given two DAGs G_1 and G_2 on the same vertex set, we write $G_1 - G_2$ to denote the DAG which contains an edge e iff $e \in G_1$ but $e \notin G_2$. We also write $G_2 \subseteq G_1$ if G_2 is a subgraph of G_1 .

Discrepancy. Let $\mathcal{F} = \{F_1, \dots, F_m\}$ be a family of subsets over a ground set $[T]$, where $T \geq 1$ is a positive integer. The discrepancy of \mathcal{F} is defined as

$$\text{Disc}(\mathcal{F}) := \min_{x \in \{-1, 1\}^T} \text{Disc}(\mathcal{F}, x),$$

where $\text{Disc}(\mathcal{F}, x) = \max_{i \in [m]} |\sum_{j \in F_i} x_j|$. A vector $x \in \{-1, 1\}^T$ is called a coloring (or signing) of $[T]$.

► **Definition 9** (Hereditary discrepancy). *The hereditary discrepancy of \mathcal{F} is defined as*

$$\text{HerDisc}(\mathcal{F}) := \max_{J \subseteq [T]} \text{Disc}(\mathcal{F}|_J),$$

where $\mathcal{F}|_J := \{F_1 \cap J, \dots, F_m \cap J\}$ is the restriction of the set system \mathcal{F} to J .

An important measure for approximating the hereditary discrepancy of a set system is the γ_2 -norm (or the factorization norm), which is well-studied in mathematics.

► **Definition 10** (γ_2 -norm). *Let $A \in \mathbb{R}^{m \times n}$ be a matrix. The γ_2 -norm of A is defined as*

$$\gamma_2(A) := \min \{r(B) \cdot c(C) \mid A = BC\},$$

where $r(B)$ is the largest ℓ_2 -norm of the rows of B , and $c(C)$ is the largest ℓ_2 -norm of the columns of C , and the minimum is taken over all real matrices B and C such that $A = BC$.

13:10 Prefix Discrepancy, Smoothed Analysis, and Combinatorial Vector Balancing

The following relationship between γ_2 -norm and HerDisc was proved in [18].

► **Lemma 11** ([18]). *Let $A \in \mathbb{R}^{m \times T}$ with rank r . Then, we have $\frac{\gamma_2(A)}{\log r} \leq \text{HerDisc}(A) \leq \sqrt{\log m} \cdot \gamma_2(A)$.*

Banaszczyk's Symmetrization. The following definition from [2] is crucially used to obtain the bound in Theorem 1. Let $K \subseteq \mathbb{R}^d$ be a convex body and $u \in \mathbb{R}^d$ be a non-zero vector. Let Z be the cylinder (possibly empty) consisting of all lines ℓ parallel to u such that the Euclidean length of the segment $K \cap \ell$ is greater than or equal to $2\|u\|_2$. If $u = 0$, then Z is defined to be equal to K .

► **Definition 12** (Banaszczyk's symmetrization). *Let $K \subseteq \mathbb{R}^d$ be a convex body, $u \in \mathbb{R}^d$ a vector, and Z as defined above. We define $K * u$ as*

$$K * u := [(K + u) \cup (K - u)] \cap Z .$$

In particular, it follows that $K * u \subseteq (K + u) \cup (K - u)$ and that if K is symmetric, then so is $K * u$. The following lemma is central to the proof of Theorem 1.

► **Lemma 13** ([2]). *Let $K \in \mathbb{R}^d$ be a symmetric convex body with Gaussian measure $\mu_d(K) \geq \frac{1}{2}$, and let $u \in \frac{1}{5} \cdot \mathbb{B}_2^d$, then we have $\mu_d(K * u) \geq \mu_d(K)$.*

Basic Probabilistic Lemmas. Here we collect some basic probabilistic facts that will be used later.

► **Proposition 14.** *Let $d \geq 2$, and $\epsilon \leq 1/d$, and $X \in \mathbb{R}^d$ be sampled uniformly from the unit sphere \mathbb{S}^{d-1} . Then, for any fixed vector $u \in \mathbb{S}^{d-1}$, we have*

$$\mathbb{P}[\|X - u\|_2 \leq \epsilon] \geq \exp(-10d \log(1/\epsilon)) .$$

Proof. Since $\|u\|_2 = 1$, the event $\|X - u\|_2 \leq \epsilon$ is equivalent to saying that $X^\top u \geq 1 - \epsilon$. Thus, the probability of this event is exactly equal to the ratio $\text{Vol}(\mathcal{S})/\text{Vol}(\mathbb{S}^{d-1})$ where Vol denotes the surface area and $\mathcal{S} = \{x \in \mathbb{S}^{d-1} \mid x^\top u \geq 1 - \epsilon\}$ is the spherical cap. Using Lemma 4.1 in [20], we get that this ratio is at least $\epsilon^{10d} = \exp(-10d \log(1/\epsilon))$. ◀

► **Proposition 15.** *Let $X \in 2 \cdot \mathbb{B}_2^d$ be a random vector such that $\text{Cov}(X) \succcurlyeq \epsilon^2 I$, then for any $u \in \mathbb{R}^d$,*

$$\mathbb{P} \left[|u^\top X| \geq \frac{1}{2} \cdot \epsilon \|u\|_2 \right] \geq \frac{\epsilon^2}{10} .$$

Proof. We may assume that $\|u\|_2 = 1$. Let us denote the mean and covariance of X by μ and Σ . Note that

$$\mathbb{E}[|u^\top X|^2] = u^\top \mathbb{E}[XX^\top] u = u^\top \Sigma u + |u^\top \mu|^2 \geq \epsilon^2 ,$$

since $\Sigma = \mathbb{E}[XX^\top] - \mu\mu^\top$. Since $\|X\|_2 \leq 2$ by assumption, $|u^\top X| \leq 2$ always, and a straightforward averaging argument yields the statement of the proposition. ◀

3 Smoothed Analysis of Prefix Discrepancy

In Section 3.1 by showing an upper bound of $O(\sqrt{\log d + \log \log T})$ for prefix discrepancy in the smoothed setting. Recall that in this setting the input vectors $v_t = \bar{v}_t + \hat{v}_t$, where \bar{v}_t are worst-case vectors and \hat{v}_t are small random perturbations.

► **Theorem 2** (Smoothed analysis upper bound). *Suppose for each $t \in [T]$ the noise $\hat{v}_t \in \mathbb{B}_2^d$ is sampled independently and satisfies covariance $\text{Cov}(\hat{v}_t) \succeq \epsilon^2 I_d$ for some $\epsilon \geq 1/\text{poly}(d, \log T)$. Then, the prefix discrepancy in the smoothed setting is $O(\sqrt{\log d + \log \log T})$ with probability at least $1 - 1/\text{poly}(T)$.*

Next in Section 3.2, we show that the bound in the above theorem cannot be improved even for the stochastic setting (which is a special case of smoothed setting), without further progress on Banaszczyk's result on prefix discrepancy for adversarial vectors (Theorem 1). In particular, we show that if unit vectors $v_1, \dots, v_T \in \mathbb{S}^{d-1}$ are drawn uniformly, then the bound from Theorem 2 is attained with high probability assuming the tightness of Banaszczyk's bound.

► **Theorem 16** (Conditional Lower Bound for Smoothed/Stochastic Setting). *Let $d \geq 2$ and $n \geq \log d$ be integers. Suppose that there is an adversarial configuration of unit vectors $w_1, \dots, w_n \in \mathbb{S}^{d-1}$ satisfying $\text{PrefixDisc}(w_1, \dots, w_n) = \Omega(\sqrt{\log d + \log n})$. Let $T = n \cdot \exp(100nd \log n)$ and let unit vectors $v_1, \dots, v_T \in \mathbb{S}^{d-1}$ be drawn uniformly, then with probability at least $1 - e^{-\sqrt{T}}$, we have that $\text{PrefixDisc}(v_1, \dots, v_T) = \Omega(\sqrt{\log d + \log \log T})$.*

3.1 The Upper Bound

To prove Theorem 2, let us start by grouping the input vectors v_1, \dots, v_T into consecutive blocks of n vectors, with a total of T/n blocks (without loss of generality, assume here that T/n is an integer). We will set $n = \text{poly}(d, \log T)$ with the exact parameters to be specified later. Let $\Delta = \Delta(d, n) = \Theta(\sqrt{\log d + \log n})$ be twice the prefix discrepancy upper bound given by Banaszczyk's Theorem 1 for coloring n vectors. Our strategy will work in rounds and in round $r \in [T/n]$ it will give signs to the vectors in the r^{th} block. We maintain the following invariants at the end of round r :

Invariant 1: Discrepancy at the end of r^{th} block is $\left\| \sum_{t \leq rn} x_t v_t \right\|_\infty \leq \Delta$

Invariant 2: Discrepancy of any prefix of the first r blocks is $\max_{\tau \in [rn]} \left\| \sum_{t \leq \tau} x_t v_t \right\|_\infty \leq 6\Delta$. The first condition in the invariant allows us to inductively proceed with coloring the next block of vectors while the second condition gives us the desired bound on the maximum discrepancy of any prefix.

The first round of our strategy gives signs to the first block of n vectors using the result of [3], which maintains both invariants. Each other round $1 < r < T/n$ will have two phases: in the first phase we find a fractional signing of the next block of vectors and then round it to obtain a $\{\pm 1\}$ signing.

To describe the two phases in round r , let us use $M \in \mathbb{R}^{d \times n}$ to denote the matrix whose n columns are the vectors being processed in round r , i.e. $v_{(r-1)n+1}, \dots, v_{rn}$. For any subset $I \subseteq [n]$, let us also write M_I to denote the $d \times n$ matrix where we keep the j^{th} column of M if it occurs in I and replace it with zeros otherwise. We also define $w := \sum_{t \leq (r-1)n} x_t v_t$ to be the discrepancy vector at the end of round $r-1$.

In the first phase, we find a fractional signing satisfying some properties that help in maintaining the invariants after round r .

13:12 Prefix Discrepancy, Smoothed Analysis, and Combinatorial Vector Balancing

► **Lemma 17** (Fractional signing). *For an arbitrary $\|w\|_\infty \leq \Delta$, with probability at least $1 - 1/\text{poly}(T)$ there is a fractional signing $x \in [-1, 1]^n$ satisfying:*

1. $Mx + w = 0$, and
2. $\|M_I x\|_\infty \leq 4\Delta$ for every prefix interval $I \subseteq [n]$, i.e., $I \in \{[k] \mid k \leq n\}$.

The conditions above are linear constraints, so the lemma above shows that the corresponding linear program is feasible with high probability. The natural approach to showing this using strong duality encounters some obstacles (elaborated on later), so our proof relies on a stronger linear program where one imposes the second condition in the lemma for a different family of intervals. Using this stronger linear program and a careful chaining argument, we can infer the existence of a feasible solution satisfying the conditions above with high probability.

The second phase rounds the fractional signing to a ± 1 signing using the following lemma, which is similar to the bit-by-bit rounding procedure used by Lovász, Spencer, and Vesztegombi to establish the connection between hereditary and linear discrepancy [17].

► **Lemma 18** (Block Rounding). *Let $x \in [-1, 1]^n$ be a fractional signing. There exists a signing $x^* \in \{-1, 1\}^n$ such that $\|M_I(x^* - x)\|_\infty \leq \Delta$ for any prefix interval $I \subseteq [n]$.*

We will prove Lemma 17 and Lemma 18 in Section 3.1.2 and Section 3.1.1, but first we show how they imply Theorem 2.

Proof of Theorem 2. Assuming that the two invariants are maintained at the end of round $r - 1$, we show that they hold after round r as well. For the first invariant, the discrepancy vector at the end of round r is $w + Mx^*$. Therefore, using that $\|w\|_\infty \leq \Delta$, the first condition of Lemma 17 and Lemma 18, we have

$$\|w + Mx^*\|_\infty = \|w + Mx + M(x^* - x)\|_\infty = \|M(x^* - x)\|_\infty \leq \Delta .$$

For the second invariant, for each prefix of the first $(r - 1)$ blocks the discrepancy remains the same, so let us consider a prefix of the form $w + M_I x^*$ where I is some prefix interval of $[n]$. In this case, using the second condition of Lemma 17, we can bound

$$\|w + M_I x^*\|_\infty \leq \|w\|_\infty + \|M_I(x^* - x)\|_\infty + \|M_I x\|_\infty \leq 6\Delta. \quad \blacktriangleleft$$

3.1.1 Proof of Lemma 18

The proof of Lemma 18 is a slight variant of the proof used in [17] to upper bound linear discrepancy by hereditary discrepancy. We give it here for completeness.

Proof of Lemma 18. We first apply the coordinate-wise map $y = (x + 1)/2 \in [0, 1]^n$ for the fractional signing $x \in [-1, 1]^n$. Our goal is to show that there exists $y^* \in \{0, 1\}^n$ such that for any prefix interval $I \subseteq [n]$, we have $\|M_I(y^* - y)\| \leq \Delta/2$. The lemma immediately follows by taking the inverse map $x^* = 2y^* - 1 \in \{-1, 1\}^n$.

Write the binary expansion of each y_i as $y_i = \sum_{j \geq 1} q_{i,j} 2^{-j}$, where $q_{i,j} \in \{0, 1\}$ denotes the j -th bit of y_i after the decimal point. Let k be a positive integer, and assume for now that all the y_i 's have at most k bits in their binary expansions (if the binary expansion of some y_i has $k_i < k$ bits, we let $q_{i,k_i+1} = \dots = q_{i,k} = 0$). We shall eventually send $k \rightarrow \infty$ so the above assumption becomes without loss of generality. Consider the k th bit $q_{i,k}$ of each y_i . Let $M^{(k)}$ be the sub-matrix of M restricted to those columns i for which $q_{i,k} = 1$. By applying Banaszczyk's result in Theorem 1 to the columns of the matrix $M^{(k)}$, there is a ± 1 signing $\chi^{(k)}$ of the columns of $M^{(k)}$ such that the prefix discrepancy of $M^{(k)}$ is at most $\Delta/2$.

Viewing $\chi^{(k)}$ as a vector in \mathbb{R}^n with entries $\{-1, 0, 1\}$ (where the 0 entries correspond to the columns not in $M^{(k)}$), we update the vector $y' = y + 2^{-k}\chi^{(k)}$. Now note that the k -th bit of each y'_i is 0, as $\chi_i^{(k)} \in \{-1, 1\}$ if $q_{i,k} = 1$ and 0 if $q_{i,k} = 0$. Moreover, for any prefix interval $I \subseteq [n]$, we have $\|M_I(y' - y)\|_\infty = 2^{-k} \cdot \|M_I\chi^{(k)}\| \leq 2^{-k} \cdot \Delta/2$.

We can now iterate this bit-rounding procedure to the bits $k-1, k-2, \dots, 1$ (by treating y' as the new vector y with only $k-1$ bits). This produces a $\{0, 1\}$ -vector $y^* = y + 2^{-k}\chi^{(k)} + 2^{-k+1}\chi^{(k-1)} + \dots + 2^{-1}\chi^{(1)}$ such that $\|M_I(y^* - y)\|_\infty \leq (2^{-k} + 2^{-k+1} + \dots + 2^{-1}) \cdot \Delta/2 \leq \Delta/2$ holds for any prefix interval $I \subseteq [n]$. As this rounding error does not depend on k , sending $k \rightarrow \infty$ completes the proof of the lemma. \blacktriangleleft

3.1.2 Proof of Lemma 17 via a Primal-dual Approach

As mentioned before, for our analysis it will be useful to consider a stronger linear program which imposes constraints on a different family of intervals other than just the prefix intervals of $[n]$. With this in mind, for a family \mathcal{I} of intervals of $[n]$, we consider the following linear programming feasibility problem:

$$\begin{aligned} \max \quad & 0 \\ \text{s.t.} \quad & (Mx)_i = -w_i && \forall i \in [d] \\ & -2\Delta \leq (M_I x)_i \leq 2\Delta && \forall I \in \mathcal{I}, i \in [d] \\ & -1 \leq x_j \leq 1 && \forall j \in [n] \end{aligned} \quad \text{LP}(\mathcal{I})$$

To understand when $\text{LP}(\mathcal{I})$ is feasible, we consider the corresponding dual linear program:

$$\begin{aligned} \min \quad & -y^\top w + 2\Delta \sum_{I \in \mathcal{I}} (\alpha_I^+ + \alpha_I^-)^\top \mathbf{1} + (\gamma^+ + \gamma^-)^\top \mathbf{1} \\ \text{s.t.} \quad & y^\top M + \sum_{I \in \mathcal{I}} (\alpha_I^+ - \alpha_I^-)^\top M_I + (\gamma^+ - \gamma^-)^\top = \mathbf{0} \\ & \alpha_I^+, \alpha_I^-, \gamma^+, \gamma^- \geq \mathbf{0} \quad \forall I \in \mathcal{I}. \end{aligned}$$

By strong duality, the primal is infeasible iff the dual objective is strictly negative. Next, we simplify the dual program and derive an analytical condition which will imply the feasibility of $\text{LP}(\mathcal{I})$. For this, we note that for any fixed value of the vector $\alpha_I := \alpha_I^+ - \alpha_I^-$, with α_I^+ and α_I^- non-negative, the least value of $(\alpha_I^+ + \alpha_I^-)^\top \mathbf{1}$ is $\|\alpha_I\|_1$. This holds similarly for $\gamma := \gamma^+ - \gamma^-$. So, the dual objective can be simplified as

$$-y^\top w + 2\Delta \sum_{I \in \mathcal{I}} \|\alpha_I\|_1 + \|\gamma\|_1 .$$

By the dual constraint $\|\gamma\|_1 = \|y^\top M + \sum_{I \in \mathcal{I}} \alpha_I^\top M_I\|_1$, the primal is *infeasible* if and only if there exist $y, \alpha_I \in \mathbb{R}^d$ for $I \in \mathcal{I}$ such that

$$-y^\top w + 2\Delta \sum_{I \in \mathcal{I}} \|\alpha_I\|_1 + \|y^\top M + \sum_{I \in \mathcal{I}} \alpha_I^\top M_I\|_1 < 0 .$$

Replacing α_I by $-\alpha_I$ and rearranging the above, we get that $\text{LP}(\mathcal{I})$ is feasible if for every $y, \{\alpha_I\}_{I \in \mathcal{I}} \in \mathbb{R}^d$ the following holds

$$\|y^\top M - \sum_{I \in \mathcal{I}} \alpha_I^\top M_I\|_1 \geq y^\top w - 2\Delta \sum_{I \in \mathcal{I}} \|\alpha_I\|_1 . \quad (2)$$

13:14 Prefix Discrepancy, Smoothed Analysis, and Combinatorial Vector Balancing

Using our assumption that $\|w\|_\infty \leq \Delta$, the above condition is satisfied if the left hand side in (2) is at least $\Delta(\|y\|_1 - 2\sum_{I \in \mathcal{I}} \|\alpha_I\|_1)$. We may further restrict our attention to the case when $\|y\|_1 \geq 2\sum_{I \in \mathcal{I}} \|\alpha_I\|_1$ since otherwise the inequality holds trivially. Furthermore, by rescaling, we may assume that $\|y\|_1 = d$. Thus, denoting by $K \subseteq \mathbb{R}^{d+|\mathcal{I}|d}$ the set of all vectors $(y, \{\alpha_I\}_{I \in \mathcal{I}})$ satisfying $\|y\|_1 = d$ and $\sum_{I \in \mathcal{I}} \|\alpha_I\|_1 \leq d/2$, we arrive at the following sufficient condition for feasibility.

► **Proposition 19.** *If $\|w\|_\infty \leq \Delta$ and for every $(y, \{\alpha_I\}_{I \in \mathcal{I}}) \in K$ we have*

$$\|y^\top M - \sum_{I \in \mathcal{I}} \alpha_I^\top M_I\|_1 \geq \Delta \cdot d, \quad (3)$$

then (2) holds as well, and thus $\text{LP}(\mathcal{I})$ is feasible.

Note that the left hand side of (3) is

$$\sum_{j=1}^n |(y - \sum_{I: j \in I} \alpha_I)^\top v_j| = \sum_{j=1}^n |z_j^\top v_j|, \quad (4)$$

defining $z_j := y - \sum_{I: j \in I} \alpha_I$ where $z_j \in \mathbb{R}^d$ satisfy $\|z_j\|_1 \geq \|y\|_1 - \sum_{i \in \mathcal{I}} \|\alpha_i\|_1 \geq d/2$ for every $j \in [n]$. If there is sufficient randomness in the vectors v_j , so that each term $|z_j^\top v_j| = 1/\text{poly}(d, \log T)$ with constant probability, by taking $n = \text{poly}(d, \log T)$, we can expect that (3) holds (recall that $\Delta = O(\sqrt{\log(dn)})$) with probability $1 - e^{-cn}$ for any given $(y, \{\alpha_I\}_{I \in \mathcal{I}}) \in K$.

Given this, it is natural to attempt to take \mathcal{I} to be $\mathcal{I}_{\text{prefix}}$, the set of all prefixes of $[n]$, as required in our target Lemma 17. However, naively this requires us to perform a union bound over a set of vectors that is much larger than the e^{-cn} probability of the bad event. In particular, we need to do a union bound over all possible choices $(y, \{\alpha_I\}_{I \in \mathcal{I}}) \in K$. Since it suffices to discretize these vectors to have values which are integer multiples of $1/\text{poly}(n, d)$, given that $|\mathcal{I}_{\text{prefix}}| = n$, the total number of choices is $\text{poly}(n, d)^{nd} \gg e^{cn}$. It is unclear how to make this strategy work, so we work with a different family of intervals $\mathcal{I}_{\text{block}}$ where the constraints not only allow us to infer the target constraints of Lemma 17 but a careful chaining argument allows a more refined union bound.

Improved Bounds using a Block Decomposition $\mathcal{I}_{\text{block}}$. Let us partition $[n]$ into consecutive blocks of length $b = n^{0.1}$ (for convenience, we assume that b divides n). Consider the following family $\mathcal{I}_{\text{block}}$ of intervals of $[n]$ obtained by taking

- all prefix intervals of $[n]$ which end after an entire block, *i.e.*, intervals of the form $[ib]$ where $i \in [n/b]$, and
- all prefix intervals within each block, *i.e.*, intervals of the form $\{(i-1)b+1, \dots, (i-1)b+r\}$ where $i \in [b]$ and $r \in [n/b]$.

Note that any prefix interval $I \in \mathcal{I}_{\text{prefix}}$ can be expressed as the union of two intervals $I_1, I_2 \in \mathcal{I}_{\text{block}}$. Thus, if $\text{LP}(\mathcal{I}_{\text{block}})$ is feasible we get that $Mx + w = 0$ and also $\|M_I x\|_\infty \leq \|M_{I_1} x\|_\infty + \|M_{I_2} x\|_\infty \leq 4\Delta$, thus implying the constraints of Lemma 17.

To complete the proof, we show that $\text{LP}(\mathcal{I}_{\text{block}})$ is feasible w.h.p. by bounding the left hand side of (3). For this let us introduce some terminology: we say an interval $I \in \mathcal{I}_{\text{block}}$ is *long* if $|I| \geq b$ and *short* otherwise. Note that the size of the family $|\mathcal{I}_{\text{block}}| = \Theta(n)$ but the number of long intervals is $2n/b = o(n)$ and furthermore, any coordinate $j \in [n]$ is only contained in n/b short intervals.

To bound (3), we use a two-step union bound, best viewed as a chaining argument. In particular, we show the following lemma where we only consider $(y, \{\alpha_I\}) \in K$ as in Proposition 19.

▷ **Claim 20.** Fix any $(y, \{\alpha_I\}_{I \text{ long}})$. Then, for $n = \text{poly}(d, \log T)$ sufficiently large, the probability that (3) holds for **all** choices of $\{\alpha_I\}_{I \text{ short}}$ is at least $1 - e^{-c\epsilon^2 n}$ for a universal constant $c > 0$.

Note that there are $\Theta(n)$ short intervals and performing a naive union bound over all possible choices of $\{\alpha_I\}_{I \text{ short}}$ (discretized as before) is not enough to prove the above claim and our proof below will use a more refined argument. However, given the above claim, to show that (3) holds for all $(y, \{\alpha_I\}_{I \in \mathcal{I}_{\text{block}}}) \in K$, we only need to do a union bound over all the possible choices of $(y, \{\alpha_I\}_{I \text{ long}})$ where the number of long intervals is $O(n/b) = O(n^{0.9})$.

More precisely, discretizing each vector to integer multiples of $1/\text{poly}(n, d)$, the number of possible choices of $(y, \{\alpha_I\}_{I \text{ long}})$ is $(nd)^{O(n/b)} = e^{O(n^{0.91})}$ as $d \leq n$. Since $\epsilon = 1/\text{poly}(d, \log T)$, choosing $n = \text{poly}(d, \log T)$ to be a large enough polynomial, we get that $\text{LP}(\mathcal{I}_{\text{block}})$ is feasible with probability at least $1 - e^{-\Omega(n^{0.05})} = 1 - T^{-2}$.

To finish the proof, we now prove Claim 20.

Proof of Claim 20. For $j \in [n]$, let us define $z_{\text{long},j} := y - \sum_{I:j \in I, I \text{ long}} \alpha_I$ and $z_{\text{short},j} := \sum_{I:j \in I, I \text{ short}} \alpha_I$. For a fixed $(y, \{\alpha_I\}_{I \in \mathcal{I}})$, the contribution of coordinate $j \in [n]$ to the LHS of (3) is

$$\left| \left(y - \sum_{I:j \in I} \alpha_I \right)^\top v_j \right| = |(z_{\text{long},j} - z_{\text{short},j})^\top v_j| \geq \min \{0, |z_{\text{long},j}^\top v_j| - |z_{\text{short},j}^\top v_j|\}. \quad (5)$$

First we show that the contribution of $|z_{\text{long},j}^\top v_j|$ is typically large for most j . For this, we note that $\|z_{\text{long},j}\|_1 \geq \|y\|_1 - d/2 \geq d/2$, and so $\|z_{\text{long},j}\|_2 \geq \sqrt{d}/2$. By our assumption on the noise, we have $\|v_j\|_2 \leq 2$ and that $\text{Cov}(v_j) \geq \epsilon^2 I_d$. Thus, Proposition 15 implies that for any given $j \in [n]$,

$$\mathbb{P} \left[|z_{\text{long},j}^\top v_j| \geq \delta \right] \geq \frac{\epsilon^2}{10}, \quad (6)$$

where $\delta = \frac{1}{4} \cdot \epsilon \sqrt{d}$.

To bound the contribution of $|z_{\text{short},j}^\top v_j|$, we define some terminology. Given a setting of $(y, \{\alpha_I\}_{I \in \mathcal{I}})$, we say that a coordinate $j \in [n]$ is *heavy* if it lies in at least one *short* interval I such that $\|\alpha_I\|_1 \geq \frac{\epsilon}{100b}$. We say that j is *light* if it is not heavy. Since $\sum_I \|\alpha_I\|_1 \leq d/2$, there are at most $\frac{50bd}{\epsilon}$ intervals I with $\|\alpha_I\|_1 \geq \frac{\epsilon}{100b}$. As every short interval contains at most b coordinates, it follows that at most $\frac{50b^2d}{\epsilon}$ coordinates are heavy. Since $b = n^{0.1}$ and $\epsilon = 1/\text{poly}(d, \log T)$, choosing $n = \text{poly}(d, 1/\epsilon) = \text{poly}(d, \log T)$ to be big enough, we can guarantee that at most $50bd/\epsilon \leq \epsilon^2 n/100$ coordinates $j \in [n]$ are heavy.

If j is light, then all short intervals I that contain j satisfy $\|\alpha_I\|_1 \leq \epsilon/(100b)$. Note that there are at most b short intervals that contain j . Thus, $\|z_{\text{short},j}\|_1 \leq \epsilon/100$ and so, using Cauchy-Schwarz, we have

$$|z_{\text{short},j}^\top v_j| \leq \|z_{\text{short},j}\|_2 \cdot \|v_j\|_2 \leq \sqrt{d} \|z_{\text{short},j}\|_1 \cdot 2 \leq \frac{\epsilon \sqrt{d}}{50} \leq \frac{\delta}{10}. \quad (7)$$

13:16 Prefix Discrepancy, Smoothed Analysis, and Combinatorial Vector Balancing

Now comes the simple but important observation: the value $z_{\text{long},j}$ only depends on $(y, \{\alpha_I\}_{I_{\text{long}}})$, but not on any α_I where I is a short interval. In particular, defining the event (which only depends on $(y, \{\alpha_I\}_{I_{\text{long}}})$)

$$\mathcal{E} := \mathcal{E}(y, \{\alpha_I\}_{I_{\text{long}}}) = \{ \text{at least } \epsilon^2 n/20 \text{ coordinates } j \in [n] \text{ satisfy } |z_{\text{long},j}^\top v_j| \geq \delta \} ,$$

we know from (6) and the multiplicative form of Chernoff bound that $\mathbb{P}[\mathcal{E}] \geq 1 - e^{-c\epsilon^2 n}$ for a universal constant $c > 0$.

When the event \mathcal{E} is true, we claim that (3) always holds. In particular, under the event \mathcal{E} , there are at least $\epsilon^2 n/20$ coordinates $j \in [n]$ with $|z_{\text{long},j}^\top v_j| \geq \delta$. Among these, at most $\epsilon^2 n/100$ coordinates j are heavy, so overall at least $\epsilon^2 n/25$ coordinates satisfy $|z_{\text{long},j}^\top v_j| \geq \delta$ and are light. Using (5) and (7), any such coordinate contributes at least $\delta/2$ to the left hand side of (5). It follows immediately that the left hand side of (3) is $\Omega(\epsilon^2 n \delta) = \Omega(n \cdot \epsilon^3 \sqrt{\delta})$. Since $\epsilon = 1/\text{poly}(d, \log T)$, taking $n = \text{poly}(d, 1/\epsilon, \log T) = \text{poly}(d, \log T)$ suffices to ensure that this is much bigger than Δd . This completes the proof of Claim 20. \blacktriangleleft

3.2 Proof of the lower bound

Proof of Theorem 16. The main idea is to partition the T vectors into T/n consecutive blocks M_i for $i \in [T/n]$ and show that with probability at least $1 - e^{-\sqrt{T}}$, one of the blocks is very close to the adversarial instance $w_1, \dots, w_n \in \mathbb{R}^d$ whose existence we assumed.

Since $\|w_t\|_2 = 1$ for each t , using Proposition 14, we have that for v_1, \dots, v_n i.i.d. sampled uniformly at random from \mathbb{S}^{d-1} , the following holds

$$\mathbb{P}[\|v_t - w_t\|_2 \leq 1/n, \forall t \in [n]] \geq \exp(-10nd \log n) .$$

Note that when $\|v_t - w_t\|_2 \leq 1/n$, the $d \times n$ matrix formed by the vectors v_1, \dots, v_n also has prefix discrepancy $\Omega(\sqrt{\log d + \log n})$. Thus, each block has prefix discrepancy $\Omega(\sqrt{\log d + \log n})$ with probability at least $\exp(-10nd \log n)$. Since $T/n = \exp(100nd \log n)$ blocks, the probability that none of the blocks has prefix discrepancy $\Omega(\sqrt{\log d + \log n})$ is upper bounded by

$$(1 - \exp(-10nd \log n))^{T/n} = \left(1 - (n/T)^{0.1}\right)^{T/n} \leq e^{-(T/n)^{0.9}} \leq e^{-\sqrt{T}} .$$

Whenever there exists a block M_i with prefix discrepancy at least $\Omega(\sqrt{\log d + \log n})$, then the prefix discrepancy for all the input vectors in blocks M_1, \dots, M_i is also $\Omega(\sqrt{\log d + \log n})$ as each prefix in M_i can be written as the difference between two prefixes of M_1, \dots, M_i . This finishes the proof of the theorem. \blacktriangleleft

4 Prefix Discrepancy for Trees and DAGs

In this section, we prove our results about prefix discrepancy for DAGs. Our first result, whose proof is given in Section 4.1, is Theorem 4 which gives an upper bound on prefix discrepancy for DAGs.

► **Theorem 4** (Prefix discrepancy for DAGs). *Given a T -node DAG G and input vectors $v_1, \dots, v_T \in \mathbb{B}_2^d$, we have $\text{DAGDisc}_G(v_1, \dots, v_T) \leq \text{HerDisc}(\text{prefix}(G)) \cdot O(\sqrt{\log d + \log T})$. Additionally, if $d = 1$, then $\text{DAGDisc}_G(v_1, \dots, v_T) = O(\text{HerDisc}(\text{prefix}(G)))$.*

Next in Section 4.2, we prove our lower bounds for prefix discrepancy of DAGs. We first show that the $\sqrt{\log T}$ in Theorem 4 cannot be improved in general, even for rooted-trees.

► **Theorem 5** (Prefix Discrepancy Lower Bound). *For any integer $T > 0$, there exists a T -node rooted tree \mathcal{T} and input vectors $v_1, \dots, v_T \in \mathbb{B}_2^2$ such that $\text{DAGDisc}_{\mathcal{T}}(v_1, \dots, v_T) = \Omega(\sqrt{\log T})$.*

One might wonder if our improvement to $\sqrt{\log d + \log \log T}$ in the smoothed setting (Theorem 2) could also be generalized to DAGs. We show that such an improvement is impossible, even for trees in the completely stochastic setting.

► **Theorem 6** (Prefix Discrepancy Lower Bound in the Random Setting). *Let the dimension $d = 2$. There exists a T -node rooted tree \mathcal{T} such that when unit vectors $v_t \in \mathbb{R}^d$ are drawn i.i.d. uniformly from \mathbb{S}^{d-1} , then w.h.p. every ± 1 -coloring satisfies $\text{DAGDisc}_{\mathcal{T}}(v_1, \dots, v_T) = \Omega(\sqrt{\log T / \log \log T})$.*

The proof of Theorem 5 is given in Section 4.2.1 and the proof of the lower bound in the random setting is given in Section 4.2.2.

4.1 Upper Bounds

The proof of Theorem 4 has two parts – first generalizing Banaszczyk’s result in Theorem 1 to prefixes of rooted trees (which can be naturally viewed as DAGs by orienting the edges from root to leaves), and second, reducing the case of an arbitrary DAG to a rooted tree.

For the first part, we prove the following for input vectors in \mathbb{R}^d .

► **Lemma 21** (Prefix Discrepancy for Trees). *Let \mathcal{T} be a given T node rooted tree and dimension $d \geq 2$. For any $v_1, \dots, v_T \in \mathbb{B}_2^d$, we have that $\text{DAGDisc}_{\mathcal{T}}(v_1, \dots, v_T) = O(\sqrt{\log dT})$. More generally, given any convex body $K \in \mathbb{R}^d$ with Gaussian measure $\mu_d(K) \geq 1 - 1/(2T)$, there exists a signing $x \in \{\pm 1\}^T$ such that for every $S \in \text{prefix}(\mathcal{T})$, we have that $\sum_{t \in S} x_t v_t \in O(1) \cdot K$.*

► **Remark 22.** For the $d = 1$ case, it is easy to see that for any rooted tree \mathcal{T} , we have

$$\text{HerDisc}(\text{prefix}(G)) \leq \Delta(\text{prefix}(G), 1) = \sup_{v_1, \dots, v_T \in [-1, 1]} \text{DAGDisc}_G(v_1, \dots, v_T) = 1, \quad (8)$$

by using a simple greedy procedure. Also, note that $\text{HerDisc}(\text{prefix}(\mathcal{T})) \geq 1$.

For the second part, let $G = ([T], E)$ be an arbitrary DAG and recall that $\text{prefix}(G)$ denotes the set of all paths in G starting from the topologically ordered root. We show that one can remove edges from G to obtain a tree $\mathcal{T} \subseteq G$ such that every path in $\text{prefix}(G)$ contains at most $O(\text{HerDisc}(\text{prefix}(G)))$ non-tree edges.

► **Lemma 8** (Reducing DAGs to Trees). *Given a DAG $G = ([T], E)$, there exists a (rooted) spanning tree $\mathcal{T} \subseteq G$ such that every directed path \mathcal{P} in G has at most $4 \cdot \text{HerDisc}(\text{prefix}(G))$ non-tree edges.*

The bound in Lemma 8 is the best possible one can hope for in terms of hereditary discrepancy. In particular, for any tree \mathcal{T} on the same vertex set $[T]$, there exists a directed path $\mathcal{P} \in \text{prefix}(G)$ such that \mathcal{P} has at least $\text{HerDisc}(\text{prefix}(G))/2$ non-tree edges. To see this, we assume, for the purpose of contradiction, that there is a tree \mathcal{T} with vertex set $[T]$ such that any path $\mathcal{P} \in \text{prefix}(G)$ contains strictly less than $\text{HerDisc}(\text{prefix}(G))/2$ non-tree edges. Then any path $\mathcal{P} \in \text{prefix}(G)$ can be written as the disjoint union of less than $\text{HerDisc}(\text{prefix}(G))/2$ paths whose edges lie entirely in \mathcal{T} . Now, restricted to any subset of $[T]$ as in Definition 9, since $\text{HerDisc}(\text{prefix}(\mathcal{T})) = 1$, there is a signing such that every path (not necessarily starting from the root) in \mathcal{T} has discrepancy at most 2. The discrepancy along any path $\mathcal{P} \in \text{prefix}(G)$ is thus strictly smaller than $\text{HerDisc}(\text{prefix}(G))$, contradicting the definition of hereditary discrepancy.

13:18 Prefix Discrepancy, Smoothed Analysis, and Combinatorial Vector Balancing

Therefore, Lemma 8 also gives a combinatorial characterization (up to a constant factor) of the hereditary discrepancy for $\text{prefix}(G)$ as follows.

► **Corollary 23** (Combinatorial Characterization of $\text{HerDisc}(\text{prefix}(G))$). *For any DAG $G = ([T], E)$, we have*

$$\text{HerDisc}(\text{prefix}(G)) = \Theta(1) \cdot \min_{\text{tree } \mathcal{T}} \max_{\mathcal{P} \in \text{prefix}(G)} |\mathcal{P} \cap (G - \mathcal{T})|,$$

where the minimum ranges over all trees \mathcal{T} on the vertex set $[T]$.

Theorem 4 then follows immediately by combining Lemma 8 with the prefix discrepancy bound for trees given in Lemma 21 and Equation (8).

Proof of Theorem 4. By Lemma 8, there exists a (rooted) tree $\mathcal{T} \subseteq G$ with the same set of vertices such that for any directed path \mathcal{P} in G , the number of non-tree edges is $O(\text{HerDisc}(\text{prefix}(G)))$. We then use Lemma 21 to find a coloring $x \in \{-1, 1\}^T$ for the vectors $\{v_t\}_{t \in [T]}$ such that each path in \mathcal{T} has discrepancy at most $O(\sqrt{\log d + \log T})$. Now, any directed path \mathcal{P} of the DAG G can be decomposed into at most $O(\text{HerDisc}(\text{prefix}(G)))$ disjoint sub-paths in \mathcal{T} , and each of these sub-paths has discrepancy $O(\sqrt{\log d + \log T})$ under coloring x (since the discrepancy of any sub-path can be written as a difference of discrepancy of two prefixes). It thus follows that the discrepancy of \mathcal{P} under coloring x is at most $O(\text{HerDisc}(\text{prefix}(G)) \cdot \sqrt{\log d + \log T})$.

The second statement also follows immediately by using Eq. (8) instead of Lemma 21 above. This completes the proof of Theorem 4. ◀

Due to limited space, we refer interested readers to the full version at <https://arxiv.org/abs/2111.07049> for the proofs of Lemma 21 and Lemma 8.

4.2 Lower Bounds

For the proof of the lower bounds for prefixes in the case of rooted trees, we will need the following definition.

► **Definition 24** (Canonical Coloring for Vertices of a Complete Binary Tree). *Let $\mathcal{T} = ([T], E)$ be a complete binary tree with root r . For each node $t \in \mathcal{T}$, let $\mathcal{P}_t \subseteq [T]$ be the unique path from r to t . We define the canonical coloring for a vertex t as the coloring $y_t \in \{\pm 1\}^{\mathcal{P}_t \setminus \{t\}}$ which assigns a node $j \in \mathcal{P}_t \setminus \{t\}$ the color $x_t^+(j) = -1$ if \mathcal{P}_t goes to j 's left child, and $x_t^+(j) = +1$ otherwise.*

We say that any given coloring $x \in \{\pm 1\}^T$ agrees with the canonical coloring y_u for the vertex u if x assigns the same sign as y_u to each node of $\mathcal{P}_u \setminus \{u\}$. Furthermore, for any given coloring $x \in \{\pm 1\}^T$, there always exists a leaf node u in the tree such that x agrees with the canonical coloring y_u for the vertex u . To find such a leaf u , one simply follows the coloring x from the root to the leaves, moving to the left child if x colors the node -1 , and the right child otherwise.

4.2.1 Adversarial Lower Bound for Trees

In this subsection, we prove Theorem 5.

Proof of Theorem 5. Without loss of generality, we assume that $T = 2^h - 1$. Consider a T -node complete binary tree \mathcal{T} with depth h . We construct the vectors $v_t \in \mathbb{B}_2^2$ at each node $t \in \mathcal{T}$ inductively in the following way. We start by picking an arbitrary unit vector v_r

for the root r . Then for any node $t \in \mathcal{T}$, let j^- and j^+ be the left and right children of t . Let $d_t^+ = \sum_{i \in \mathcal{P}_t \setminus \{t\}} y_t(i)v_i + v_t$ and $d_t^- = \sum_{i \in \mathcal{P}_t \setminus \{t\}} y_t(i)v_i - v_t$ where y_t is the canonical coloring for vertex t . The vectors v_{j^-} and v_{j^+} are then picked as arbitrary unit vectors orthogonal to d_t^- and d_t^+ respectively.

Let $x \in \{\pm 1\}^T$ be a given coloring and u be the leaf node such that x agrees with the canonical coloring y_u for the vertex u . Then, on the path $\mathcal{P}_u \setminus \{u\}$, every vector is orthogonal to the previous discrepancy vector under coloring x , and thus the discrepancy vector along $\mathcal{P}_u \setminus \{u\}$ has ℓ_2 norm at least $\sqrt{h-1}$. Since the vectors are 2-dimensional, the ℓ_∞ norm of the discrepancy vector along $\mathcal{P}_u \setminus \{u\}$ is at least $\Omega(\sqrt{h}) = \Omega(\sqrt{\log T})$. This completes the proof of Theorem 5. \blacktriangleleft

4.2.2 Stochastic Lower Bound for Trees

In this subsection, we prove Theorem 6. The idea is to show that the lower bound example in Theorem 5 (with approximate orthogonality to the canonical discrepancy vectors) appears with high probability in an ℓ -ary tree, where $\ell = \text{poly}(\log T)$ to be specified later and the vectors at the vertices are sampled i.i.d. from the unit sphere in \mathbb{R}^2 .

Proof of Theorem 6. We consider a T -node complete ℓ -ary tree \mathcal{T}' with depth h and root r . Note that ℓ and h satisfies that $\sum_{i=0}^{h-1} \ell^i = \ell^{\Theta(h)} = T$. The vectors $v_t \in \mathbb{R}^2$ for each node t of the tree are picked i.i.d. from \mathbb{S}^1 . For each node $t \in \mathcal{T}$, let $\mathcal{P}_t \subseteq [T]$ be the unique path from r to t and let $d_t^+ = \sum_{i \in \mathcal{P}_t \setminus \{t\}} y_t(i)v_i + v_t$ and $d_t^- = \sum_{i \in \mathcal{P}_t \setminus \{t\}} y_t(i)v_i - v_t$ where y_t is the canonical coloring for vertex t .

Our goal is to show that w.h.p. there exists a complete binary subtree $\mathcal{T}' \subseteq \mathcal{T}$ with depth h such that for any non-leaf node $t \in \mathcal{T}'$, its left (resp. right) child j^- (resp. j^+) and negative (resp. positive) canonical discrepancy vector d_t^- (resp. d_t^+) satisfies $|\langle v_{j^-}, d_t^- \rangle| \leq 1/4$ (resp. $|\langle v_{j^+}, d_t^+ \rangle| \leq 1/4$). If the above event occurs, then the discrepancy of the prefixes of the tree must be $\Omega(\sqrt{h})$. To see this, take any coloring $x \in \{-1, 1\}^{|\mathcal{T}'|}$, and consider the leaf $u \in \mathcal{T}'$ such that x agrees with the canonical coloring for the leaf u . Since v_{j^-} (resp. v_{j^+}) is almost orthogonal to d_t^- (resp. d_t^+) for any internal node t on the path to u , the square of the ℓ_2 -discrepancy for the coloring x along the root-leaf path \mathcal{P}_u increases by at least $\Omega(1)$ after each node (except maybe for the leaf u). As the dimension $d = 2$, the ℓ_∞ -discrepancy for the coloring x along the root-leaf path \mathcal{P}_u is also $\Omega(\sqrt{h})$.

It thus remains to show that such a complete binary tree \mathcal{T}' exists w.h.p. We find such a tree \mathcal{T}' inductively and start by adding the root r to \mathcal{T}' . For any node t that was just added to \mathcal{T}' , we hope to find left and right children j^- and j^+ among the first and second $\ell/2$ children of t such that both $|\langle v_{j^-}, d_t^- \rangle| \leq 1/4$ and $|\langle v_{j^+}, d_t^+ \rangle| \leq 1/4$ hold. For any child j of t , since $\|d_t^\pm\|_2 \leq \sqrt{h}$, we have

$$\mathbb{P}_{v_j \sim \mathbb{S}^1} \left[|\langle v_j, d_t^\pm \rangle| \leq 1/4 \right] \geq \frac{1}{10\sqrt{h}} .$$

Thus the probability that none of the first $\ell/2$ children satisfies this property is thus upper bounded by

$$(1 - 1/(10\sqrt{h}))^{\ell/2} \leq \exp(-\Theta(\ell/\sqrt{h})) .$$

Thus taking $\ell = C \cdot \sqrt{h} \cdot \log T$ for a sufficiently large constant C suffices to guarantee that the above probability is at most $1/T^2$. One could therefore afford to do a union bound over all the nodes in the process of building the complete binary tree \mathcal{T}' . To satisfy $\ell^{\Theta(h)} = T$, we need to set $h = \log T / \log \log T$. Thus the prefix discrepancy on \mathcal{T}' is thus at least $\Omega(\sqrt{h}) = \Omega(\sqrt{\log T / \log \log T})$. This completes the proof of Theorem 6. \blacktriangleleft

5 Combinatorial Vector Balancing

In this section we prove our result for combinatorial vector balancing – given a set family \mathcal{S} on the ground set $[T]$ and input vectors $v_1, \dots, v_T \in \mathbb{B}_2^d$, the goal is to find a signing $x \in \{\pm 1\}^T$ that minimizes

$$\text{CombDisc}_{\mathcal{S}}(v_1, \dots, v_T) := \min_{x \in \{-1, +1\}^T} \max_{S \in \mathcal{S}} \left\| \sum_{i \in S} x_i v_i \right\|_{\infty}.$$

► **Theorem 7** (Combinatorial Vector Balancing Upper Bound). *Let \mathcal{S} be a set system on the ground set $[T]$. Given input vectors $v_1, \dots, v_T \in \mathbb{B}_2^d$, we have that*

$$\text{CombDisc}_{\mathcal{S}}(v_1, \dots, v_T) \leq \text{HerDisc}(\mathcal{S}) \cdot \log(\min\{|\mathcal{S}|, T\}) \cdot \sqrt{\log d + \log |\mathcal{S}|}.$$

Moreover, the upper bound is constructive in the sense that there is a polynomial time algorithm to find such a coloring.

To prove Theorem 7, we bound $\text{CombDisc}_{\mathcal{S}}(v_1, \dots, v_T)$ in terms of $\gamma_2(A_{\mathcal{S}})$, where $A_{\mathcal{S}} \in \{0, 1\}^{|\mathcal{S}| \times T}$ is the incidence matrix of the set system \mathcal{S} (meaning that $(A_{\mathcal{S}})_i = 1$ iff $i \in S$). Theorem 7 then immediately follows by using this bound and relating γ_2 norm to hereditary discrepancy via Lemma 11.

► **Lemma 25.** *In the setting of Theorem 7, we have that $\text{CombDisc}_{\mathcal{S}}(v_1, \dots, v_T) \leq \gamma_2(A_{\mathcal{S}}) \cdot \sqrt{\log d + \log |\mathcal{S}|}$. Moreover, the upper bound is constructive in the sense that there is a polynomial time algorithm to find such a coloring.*

Proof of Lemma 25. Let the optimal factorization for $\gamma_2(A_{\mathcal{S}})$ be $A_{\mathcal{S}} = BC$ where $\gamma_2(A_{\mathcal{S}}) = r(B) \cdot c(C)$ with $r(B)$ being the maximum ℓ_2 norm among all rows of B and $c(C)$ being the maximum ℓ_2 norm among all columns of C .

Denote the j th coordinate of vector $v_t \in \mathbb{R}^d$ as $v_t(j)$. For each $j \in [d]$, we define a matrix $D_j \in \mathbb{R}^{|\mathcal{S}| \times T}$ as $D_j := \text{diag}(v_1(j), \dots, v_T(j))$. Note that

$$\text{CombDisc}_{\mathcal{S}}(v_1, \dots, v_T) = \min_{x \in \{\pm 1\}^T} \max_{S \in \mathcal{S}, j \in [d]} \left| \sum_{t \in S} x_t v_t(j) \right| = \text{Disc}(A_{\mathcal{S}}^D),$$

where the right hand side is the standard notion of discrepancy of a matrix⁶ and the matrix $A_{\mathcal{S}}^D \in \mathbb{R}^{|\mathcal{S}| \times T}$ is defined as

$$A_{\mathcal{S}}^D = \begin{pmatrix} A_{\mathcal{S}} D_1 \\ A_{\mathcal{S}} D_2 \\ \vdots \\ A_{\mathcal{S}} D_d \end{pmatrix}.$$

Now a factorization of $A_{\mathcal{S}}^D$ is given by

$$A_{\mathcal{S}}^D = \begin{pmatrix} B & & & \\ & B & & \\ & & \ddots & \\ & & & B \end{pmatrix} \cdot \begin{pmatrix} C D_1 \\ C D_2 \\ \vdots \\ C D_d \end{pmatrix}.$$

⁶ For a matrix $A \in \mathbb{R}^{d \times T}$, its discrepancy is defined as $\text{Disc}(A) = \min_{x \in \{\pm 1\}^T} \|Ax\|_{\infty}$.

Note that the maximum ℓ_2 -norm of the rows of the first matrix is the same as $r(B)$, and the maximum ℓ_2 -norm of the columns of the second matrix is at most $c(C)$, as $\|v_i\|_2 \leq 1$ for each vector $i \in [T]$. It follows that $\gamma_2(A_S^D) \leq r(B) \cdot c(C) = \gamma_2(A_S)$. The lemma then follows immediately from the second inequality in Lemma 11. \blacktriangleleft

References

- 1 Ryan Alweiss, Yang P. Liu, and Mehtaab Sawhney. Discrepancy minimization via a self-balancing walk. In *Proceedings of STOC*, 2021.
- 2 Wojciech Banaszczyk. Balancing vectors and Gaussian measures of n -dimensional convex bodies. *Random Struct. Algorithms*, 12(4):351–360, 1998.
- 3 Wojciech Banaszczyk. On series of signed vectors and their rearrangements. *Random Struct. Algorithms*, 40(3):301–316, 2012.
- 4 Nikhil Bansal, Haotian Jiang, Raghu Meka, Sahil Singla, and Makrand Sinha. Online discrepancy minimization for stochastic arrivals. In *Proceedings of SODA*, pages 2842–2861, 2021.
- 5 Nikhil Bansal, Haotian Jiang, Sahil Singla, and Makrand Sinha. Online vector balancing and geometric discrepancy. In *Proceedings of STOC*, pages 1139–1152, 2020.
- 6 Nikhil Bansal and Raghu Meka. On the discrepancy of random low degree set systems. In *Proceedings of SODA 2019*, pages 2557–2564, 2019.
- 7 Imre Bárány. A vector-sum theorem and its application to improving flow shop guarantees. *Math. Oper. Res.*, 6(3):445–452, 1981.
- 8 Imre Bárány and Victor S Grinberg. On some combinatorial questions in finite-dimensional spaces. *Linear Algebra and its Applications*, 41:1–9, 1981.
- 9 Kevin Buchin, Jiří Matoušek, Robin A. Moser, and Dömötör Pálvölgyi. Vectors in a box. *Math. Program.*, 135(1-2):323–335, 2012.
- 10 Bernard Chazelle. *The discrepancy method: randomness and complexity*. Cambridge University Press, 2001.
- 11 Sergej Chobanyan. Convergence as of rearranged random series in Banach space and associated inequalities. In *Probability in Banach Spaces*, 9, pages 3–29. Springer, 1994.
- 12 Friedrich Eisenbrand and Robert Weismantel. Proximity results and faster algorithms for integer programming using the steinitz lemma. In *Proceedings of SODA*, pages 808–816, 2018.
- 13 Esther Ezra and Shachar Lovett. On the Beck-Fiala conjecture for random set systems. *Random Struct. Algorithms*, 54(4):665–675, 2019.
- 14 Nika Haghtalab, Tim Roughgarden, and Abhishek Shetty. Smoothed analysis with adaptive adversaries. *CoRR*, abs/2102.08446, 2021.
- 15 Rebecca Hoberg and Thomas Rothvoss. A Fourier-Analytic Approach for the Discrepancy of Random Set Systems. In *Proceedings of SODA*, pages 2547–2556, 2019.
- 16 Klaus Jansen and Lars Rohwedder. On integer programming and convolution. In *Proceedings of ITCS*, pages 43:1–43:17, 2019.
- 17 László Lovász, Joel Spencer, and Katalin Vesztegombi. Discrepancy of set-systems and matrices. *European Journal of Combinatorics*, 7(2):151–160, 1986.
- 18 Jiří Matoušek, Aleksandar Nikolov, and Kunal Talwar. Factorization norms and hereditary discrepancy. *CoRR*, abs/1408.1376, 2014.
- 19 Jiří Matousek. *Geometric discrepancy: An illustrated guide*, volume 18. Springer Science & Business Media, 2009.
- 20 Daniele Micciancio and Panagiotis Voulgaris. Faster exponential time algorithms for the shortest vector problem. In *Proceedings of SODA*, pages 1468–1480, 2010.
- 21 Aleksandar Nikolov. Tighter bounds for the discrepancy of boxes and polytopes. *Mathematika*, 63(3):1091–1113, 2017.
- 22 Aditya Potukuchi. A spectral bound on hypergraph discrepancy. In *Proceedings of ICALP*, pages 93:1–93:14, 2020.

13:22 Prefix Discrepancy, Smoothed Analysis, and Combinatorial Vector Balancing

- 23 Tim Roughgarden, editor. *Beyond the Worst-Case Analysis of Algorithms*. Cambridge University Press, 2020.
- 24 Sergey Sevast'janov. On some geometric methods in scheduling theory: A survey. *Discret. Appl. Math.*, 55(1):59–82, 1994.
- 25 Joel Spencer. Balancing games. *J. Comb. Theory, Ser. B*, 23(1):68–74, 1977.
- 26 Joel Spencer. Balancing vectors in the max norm. *Combinatorica*, 6(1):55–65, 1986.
- 27 Daniel A Spielman and Shang-Hua Teng. Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. *Journal of the ACM (JACM)*, 51(3):385–463, 2004.
- 28 Ernst Steinitz. Bedingt konvergente reihen und konvexe systeme.(schluß.). *Journal für die reine und angewandte Mathematik*, 146:1–52, 1916.