THE DISCREPANCY OF SHORTEST PATHS*

GREG BODWIN †, CHENGYUAN DENG ‡, JIE GAO §, GARY HOPPENWORTH ¶, JALAJ UPADHYAY ‖, AND CHEN WANG #

Abstract. The *hereditary discrepancy* of a set system is a quantitative measure of the pseudo-random properties of the system. Roughly speaking, hereditary discrepancy measures how well one can 2-color the elements of the system so that each set contains approximately the same number of elements of each color. Hereditary discrepancy has numerous applications in computational geometry, communication complexity and derandomization. More recently, the hereditary discrepancy of the set system of *shortest paths* has found applications in differential privacy [Chen et al. SODA 23].

The contribution of this paper is to improve the upper and lower bounds on the hereditary discrepancy of set systems of unique shortest paths in graphs. In particular, we show that any system of unique shortest paths in an undirected weighted graph has hereditary discrepancy $O(n^{1/4})$, and we construct lower bound examples demonstrating that this bound is tight up to polylog n factors. Our lower bounds hold even for planar graphs and bipartite graphs and improve a previous lower bound of $\Omega(n^{1/6})$ obtained by applying the trace bound of Chazelle and Lvov [SoCG'00] to a classical point-line system of Erdős.

As applications, we improve the lower bound on the additive error for differentially-private all pairs shortest distances from $\Omega(n^{1/6})$ [Chen et al. SODA 23] to $\widetilde{\Omega}(n^{1/4})$, and we improve the lower bound on additive error for the differentially-private all sets range queries problem to $\widetilde{\Omega}(n^{1/4})$, which is tight up to polylog n factors [Deng et al. WADS 23].

Key words. Discrepancy, Shortest path.

MSC codes. 68Q25, 68R10, 68U05

1. Introduction.

In graph algorithms, a fundamental problem is to efficiently compute the distance or shortest path information of a given input graph. Over the last decade or so, the community has increasingly sought a principled understanding of the *combinatorial structure* of shortest paths, with the goal to exploit this structure in algorithm design. That is, in various graph settings, we can ask:

*What notable structural properties hold for **shortest** path systems,*
*that do not necessarily hold for **arbitrary** path systems?*

The following are a few of the major successes of this line of work:

- An extremely popular strategy in the literature is to use *hitting sets*, in which we (often randomly) generate a set of nodes S and argue that it will hit the shortest path for every pair of nodes that are sufficiently far apart. Hitting sets rarely exploit any structure of shortest paths, as evidenced by the fact that most hitting set algorithms generalize immediately to arbitrary set systems. However, they have inspired a successful line of work into graphs of bounded *highway dimension* [1, 9, 12]; very roughly, these are graphs whose shortest paths admit unusually efficient hitting sets of a certain kind.

*Submitted to the editors DATE. A preliminary version with review of related work and proof details omitted appeared in the 51st International Colloquium on Automata, Languages, and Programming (ICALP 2024). https://doi.org/10.4230/LIPIcs.ICALP.2024.27

Funding: This work has been partially supported by NSF through AF 2153680, CCF-2118953, CCF-2208663, IIS-2229876, DMS-2220271, DMS-2311064, and CRCNS-2207440.

†University of Michigan (bodwin@umich.edu).

‡Rutgers University (cd751@rutgers.edu).

§Rutgers University (jg1555@rutgers.edu).

¶University of Michigan (garytho@umich.edu)

‖Rutgers University (jalaj.upadhyay@rutgers.edu).

#Texas A&M University (cwangjhw@tamu.edu).

- Shortest paths exhibit the notable structural property of *consistency*, i.e., any subpath of the shortest path is itself the shortest path. This fact is used throughout the literature on graph algorithms [13, 27, 28], including, e.g., in the classic Floyd-Warshall algorithm for All-Pairs Shortest Paths. A recent line of work has sought to characterize the additional structure exhibited by shortest path systems beyond consistency [2, 6, 13, 23, 25, 26, 27].
- Planar graphs have received special attention within this research program, and planar shortest path systems carry some notable additional structure. For example, it is known that planar shortest paths have unusually efficient tree coverings [8, 17], and that their shortest paths can be compressed into surprisingly small space [18, 19]. Shortest path algorithms also often benefit from more general structural facts about planar graphs, such as separator theorems [40, 43].

The main result of this paper is a new structural separation between shortest path systems and arbitrary path systems, expressed through the lens of *discrepancy theory*. We will come to formal definitions of discrepancy in just a moment, but at a high level, discrepancy has been described as a quantitative measure of the combinatorial pseudorandomness of a discrete system [24], and it has widespread applications in discrete and computational geometry, random sampling and derandomization, communication complexity, and much more[1]. We will show the following:

THEOREM 1.1 (Main Result, Informal). *The discrepancy of unique shortest path systems in weighted graphs is inherently smaller than the discrepancy of arbitrary path systems in graphs.*

This separation between unique shortest paths and arbitrary paths is due to the structural property of *consistency* of unique shortest path systems, which is well-studied in the literature [13, 27, 28].

Our results can be placed within a larger context of prior work in computational geometry. A classical topic in this area is to determine the discrepancy of incidence structures between points and geometric range spaces such as axis-parallel rectangles, half-spaces, lines, and curves (cf. [20, Section 1.5]). These results have been used to show lower bounds for geometric range searching [54, 62].

Indeed, systems of unique shortest paths in graphs capture some of the geometric range spaces studied in prior work. For instance, arrangements of straight lines in Euclidean space can be interpreted as systems of unique shortest paths in an associated graph, implying a relation between the discrepancies of these two set systems. This connection has recently found applications in the study of differential privacy on shortest path distance and range query algorithms [22, 29].

More generally, discrepancy on graphs has also found applications in proving tight lower bounds on answering cut queries on graphs [33, 49].

1.1. Formal Definitions of Discrepancy. We first define some notation that we use throughout the paper. We use the letter $\mathbb{R}$ to denote the set of real numbers and $\mathbb{N}$ to denote the set of natural numbers. For $n \in \mathbb{N}$, we use the notation $[n]$ to denote the set $\{1, \cdots, n\}$. For a real number $r \in \mathbb{R}$, we use $|r|$ to denote the absolute value.

For a vector v, we use the notation v_i to denote its i-th coordinate, and for a matrix A, we use the notation $A_{i,j}$ to denote its (i, j)-th entry. For a vector $v \in \mathbb{R}^n$

[1] We refer to the excellent textbooks of Alexander, Beck, and Chen [3], Chazelle [20], and Matoušek [52] for discussion and further applications.

and any positive number $p \in \mathbb{N}_{\geq 0}$, we use different vector norms:

$$\|v\|_p := (|v_1|^p + \cdots + |v_n|^p)^{1/p}$$

where $|\cdot|$ denote the absolute value. By continuity, $\|v\|_\infty := \max_{1\leq i\leq n} |v_i|$. For positive $p, q \in \mathbb{N}_{\geq 0}$ and any matrix $A \in \mathbb{R}^{m\times n}$,

$$\|A\|_{p\to q} = \max \frac{\|Ax\|_q}{\|x\|_p}.$$

We reserve the symbol G for graphs with vertex set V and edge set E. We reserve the symbol Π to denote a given path system. For a subset $U \subseteq V$, we use the notation $G[U]$ to denote the subgraph induced by the vertex set U.

Throughout the paper, we use the $\widetilde{O}$ and $\widetilde{\Omega}$ to hide poly-logarithmic factors in the input parameter, n.

We first collect the basic definitions needed to understand this paper.

Definition 1.2 (Edge and Vertex Incidence Matrices). *Given a graph $G = (V, E)$ and a set of paths Π in G, the associated* vertex incidence matrix *is given by $A \in \{0, 1\}^{|\Pi|\times|V|}$, where for each $v \in V$ and $\pi \in \Pi$ the corresponding entry is*

$$A_{\pi,v} = \begin{cases} 1 & \text{if } v \in \pi \\ 0 & \text{if } v \notin \pi. \end{cases}$$

The associated edge incidence matrix *is given by a binary matrix $A \in \{0, 1\}^{|\Pi|\times|E|}$, where for each $e \in E$ and $\pi \in \Pi$ the corresponding entry is*

$$A_{\pi,e} = \begin{cases} 1 & \text{if } e \in \pi \\ 0 & \text{if } e \notin \pi. \end{cases}$$

Definition 1.3 (Discrepancy and Hereditary Discrepancy). *Given a matrix $A \in \mathbb{R}^{m\times n}$, its* discrepancy *is the quantity*

$$\operatorname{disc}(A) = \min_{x\in\{1,-1\}^n} \|Ax\|_\infty.$$

Its hereditary discrepancy *is the maximum discrepancy of any submatrix A_Y obtained by keeping all rows but only a subset $Y \subseteq [n]$ of the columns; that is,*

$$\operatorname{herdisc}(A) = \max_{Y\subseteq[n]} \operatorname{disc}(A_Y)$$

For a system of paths Π in a graph G, we will write $\operatorname{disc}_v(\Pi)$ and $\operatorname{herdisc}_v(\Pi)$ to denote the discrepancy and the hereditary discrepancy of its vertex incidence matrix, and $\operatorname{disc}_e(\Pi)$ and $\operatorname{herdisc}_e(\Pi)$ to denote the discrepancy and the hereditary discrepancy of its edge incidence matrix.

For intuition, the vertex discrepancy of a system of paths Π can be equivalently understood as follows. Suppose that we color each node in G either red or blue, with the goal to balance the red and blue nodes on each path as evenly as possible. The discrepancy associated to that particular coloring is the quantity

$$\max_{\pi\in\Pi} \Big| \, |\{v \in \pi \mid v \text{ is colored red}\}| - |\{v \in \pi \mid v \text{ is colored blue}\}| \, \Big|.$$

The discrepancy of the system Π is the minimum possible discrepancy over all colorings. The hereditary discrepancy is the maximum discrepancy taken over all *induced path subsystems* Π' of Π; that is, Π' is obtained from Π by focusing only on the colors[2] of a subset of vertices $Y \subseteq V$. Edge discrepancy can be understood in a similar way, coloring edges rather than vertices.

Another related quantity is the ℓ_2-discrepancy and ℓ_2-hereditary discrepancy, which use ℓ_2 norm instead of ℓ_∞ norm in the definition.

DEFINITION 1.4 (ℓ_2-Discrepancy and ℓ_2-Hereditary Discrepancy). *Given a matrix $A \in \mathbb{R}^{m\times n}$, its* ℓ_2-discrepancy *is the quantity*

$$\mathrm{disc}_2(A) = \min_{x\in\{1,-1\}^n} \frac{1}{\sqrt{m}}\|Ax\|_2.$$

Its hereditary discrepancy *is the maximum discrepancy of any submatrix A_Y obtained by keeping all rows but only a subset $Y \subseteq [n]$ of the columns; that is,*

$$\mathrm{herdisc}_2(A) = \max_{Y\subseteq[n]} \mathrm{disc}_v(A_Y)$$

For a system of paths Π in a graph G, we will write $\mathrm{disc}_{v,2}(\Pi)$ and $\mathrm{herdisc}_{v,2}(\Pi)$ to denote the ℓ_2-discrepancy and the ℓ_2-hereditary discrepancy of its vertex incidence matrix, and $\mathrm{disc}_{e,2}(\Pi)$ and $\mathrm{herdisc}_{e,2}(\Pi)$ to denote the ℓ_2-discrepancy and the ℓ_2-hereditary discrepancy of its edge incidence matrix.

1.2. Our Results. Our main result is an upper and lower bound on the hereditary discrepancy of unique shortest path systems in weighted graphs, which match up to hidden polylog n factors.

THEOREM 1.5 (Main Result).

- **(Upper Bound)** *For any n-node undirected weighted graph G with a unique shortest path between each pair of nodes, there exists a polynomial-time algorithm that finds a coloring for the system of shortest paths Π such that:*

$$\mathrm{herdisc}_v(\Pi) \le \widetilde{O}(n^{1/4}) \qquad \textit{and} \qquad \mathrm{herdisc}_e(\Pi) \le \widetilde{O}(n^{1/4}).$$

- **(Lower Bound)** *There are examples of n-node undirected weighted graphs G with a unique shortest path between each pair of nodes in which this system of shortest paths Π has $\mathrm{herdisc}_v(\Pi) \ge \widetilde{\Omega}(n^{1/4})$ and $\mathrm{herdisc}_e(\Pi) \ge \widetilde{\Omega}(n^{1/4})$. In fact, in these lower bound examples we can take G to be planar or bipartite.*

The upper bound in Theorem 1.5 is constructive and algorithmic; that is, we provide an algorithm that colors vertices (edges, respectively) of the input graph to achieve vertex (edge, respectively) discrepancy $\widetilde{O}(n^{1/4})$ on its shortest paths (or on a given subsystem of its shortest paths). Notably, Theorem 1.5 should be contrasted with the fact that the maximum possible discrepancy of any *simple path system*[3] of polynomial size in a general graph is $\widetilde{\Theta}(n^{1/2})$. The upper bound of $\widetilde{O}(n^{1/2})$ follows by coloring the nodes randomly and applying standard Chernoff bounds. The lower bound is non-trivial and is proved in Appendix A – in fact, the lower bound on discrepancy (as well as hereditary discrepancy) for a grid graph for a polynomial

[2]In the coloring interpretation, hereditary discrepancy allows a different choice of coloring for each subsystem Π', rather than fixing a coloring for Π and considering the induced coloring on each Π'.

[3]A path system is *simple* if no individual path repeats nodes.

TABLE 1

Overview of vertex/edge (hereditary) discrepancy on general undirected graphs and special families of graph: tree, bipartite and planar graphs. Here n is the number of vertices of the graph and m is the number of edges. D is the graph diameter or the longest number of hops of paths considered.

		Tree	Bipartite	Planar	Undirected Graph
Vertex	disc	$\Theta(1)$	$\Theta(1)$	$O(n^{1/4})$	$\widetilde{\Theta}(n^{1/4})$
	herdisc	$\Theta(1)$	$\widetilde{\Theta}(n^{1/4})$	$\widetilde{\Theta}(n^{1/4})$	$\Omega(n^{1/6})$[21] $\rightarrow \widetilde{\Theta}(n^{1/4})$
Edge	disc	$\Theta(1)$	$\Theta(1)$	$O(n^{1/4})$	$O(n^{1/4})$
	herdisc	$\Theta(1)$	$\widetilde{\Theta}(n^{1/4})$	$\widetilde{\Theta}(n^{1/4})$	$\Omega(n^{1/6})$[21] $\rightarrow \widetilde{\Theta}(n^{1/4})$

TABLE 2

Overview of vertex/edge (hereditary) discrepancy on general directed graphs and special families of graph: tree, bipartite and planar graphs. Here n is the number of vertices of the graph and m is the number of edges. D is the graph diameter or the longest number of hops of paths considered.

		Tree	Bipartite	Planar	Directed Graph
Vertex	disc	$\Theta(1)$	$\Theta(1)$	$O(n^{1/4})$	$\widetilde{\Theta}(n^{1/4})$
	herdisc	$\Theta(1)$	$\widetilde{\Theta}(n^{1/4})$	$\widetilde{\Theta}(n^{1/4})$	$\widetilde{\Theta}(n^{1/4})$
Edge	disc	$\Theta(1)$	$\Theta(1)$	$O(n^{1/4})$	$\min\left\{O(m^{1/4}), \widetilde{O}(D^{1/2})\right\}$
	herdisc	$\Theta(1)$	$\widetilde{\Theta}(n^{1/4})$	$\widetilde{\Theta}(n^{1/4})$	$\widetilde{\Omega}(n^{1/4})$

number of *simple* paths can be $\Omega(\sqrt{n})$. Thus, Theorem 1.5 represents a concrete separation between unique shortest path systems and general path systems.

We refer to Table 1 for our results for undirected graphs and to Table 2 for our results for directed graphs. All the bounds on discrepancy and hereditary discrepancy hold for ℓ_2-discrepancy and ℓ_2-hereditary discrepancy. See Section 8 for details.

The main open question that we leave in this work is on the hereditary edge discrepancy of shortest paths in *directed* weighted graphs. We show the following:

THEOREM 1.6. *For any n-node, m-edge directed weighted graph G with a unique shortest path between each pair of nodes, the system of shortest paths Π satisfies*

$$\text{herdisc}_v(\Pi) \leq O(n^{1/4}) \qquad \textit{and} \qquad \text{herdisc}_e(\Pi) \leq O(m^{1/4}).$$

Lower bounds in the undirected setting immediately apply to the directed setting as well, and so this essentially closes the problem for directed hereditary *vertex* discrepancy. It is an interesting open problem whether the upper bound for directed hereditary *edge* discrepancy can be improved to $\widetilde{O}(n^{1/4})$ as well.

We also leave open whether our lower bound for hereditary discrepancy extends to (non-hereditary) edge discrepancy as well, and to (non-hereditary) vertex or edge discrepancy of planar graphs.

Applications to Differential Privacy. One application of our discrepancy lower bound on unique shortest paths is in differential privacy (DP) [31, 32]. An algorithm is *differentially private* if its output distributions are relatively close regardless of whether an individual's data is present in the data set. More formally, for two databases Y and Y' that are identical except for one data entry, a randomized algorithm $\mathcal{M}$ is (ε, δ)-differentially private if, for any measurable set A in the range of the algorithm $\mathcal{M}$,

$$\Pr[\mathcal{M}(Y) \in A] \leq e^{\varepsilon} \Pr[\mathcal{M}(Y') \in A] + \delta.$$

The topic of discrepancy of paths on a graph is related to two problems already studied in differential privacy: *All Pairs Shortest Distances* (APSD) [22, 36, 59] and *All Sets Range Queries* (ASRQ) [29]. In both of these problems, we assume that the graph topology is public. In the APSD problem, the edge weights are not publicly known. A query in APSD is a pair of vertices $(u, v) \in V \times V$ and the answer is the shortest distance between u and v. In contrast, in ASRQ problem, the edge weights are assumed to be known, and every edge also has a private *attribute*. Here, the range is defined by the shortest path between two vertices (based on publicly known edge weights). The answer to the query $(u, v) \in V \times V$ then is the sum of private attributes along the shortest path. In what follows, we give a high-level argument for the lower bound on DP-APSD problem; the lower bound of $\widetilde{\Omega}(n^{1/4})$ for the DP-ASRQ problem also follows nearly the same arguement.

Chen *et al.* [22] showed that DP-APSD can be formulated as a linear query problem. In this setting, we are given a vertex incidence matrix A of the $\binom{n}{2}$ shortest paths of a graph and a vector x of length n and asked to output Ax. They show that the hereditary discrepancy of the matrix A provides a lower bound on the ℓ_∞ error for any (ε, δ)-DP mechanism for this problem. With this argument, our new discrepancy lower bound immediately implies the following lower bounds.

THEOREM 1.7 (Informal version of Corollaries 7.1 and F.1). *The (ε, δ)-DP APSD problem and (ε, δ)-DP ASRQ problem require additive error at least $\widetilde{\Omega}(n^{1/4})$.*

The best known additive error bound for the DP-ASRQ problem is $\widetilde{O}(n^{1/4})$ [29], which, by Theorem 1.7, is tight up to a polylog(n) factor. Prior to this work, the only known lower bounds for DP-ASRQ and DP-APSD were from a point-line system with hereditary discrepancy of $\Omega(n^{1/6})$ [22]. The best known additive error upper bound for DP-APSD is $\widetilde{O}(n^{1/2})$ [22, 36]. Closing this gap of $n^{1/4}$ remains an interesting open problem.

1.3. Our Techniques.

We provide a brief overview of techniques used for our upper and lower bounds on discrepancy separately.

Upper Bound Techniques. A folklore structural property of unique shortest paths on undirected graphs is *consistency*. Formally, a system of paths Π is *consistent* if for any two paths π_1, π_2, their intersection $\pi_1 \cap \pi_2$ is a (possibly empty) contiguous subpath of each. It is well known that, for any undirected graph $G = (V, E, w)$ with unique shortest paths[4], its system of shortest paths Π is consistent. Our discrepancy upper bounds actually apply to *any* consistent system of paths – not just those that arise as unique shortest paths in an undirected graph. Notice that unique shortest paths on *general directed graphs* are not necessarily consistent, but indeed are on directed acyclic graphs (DAGs).

We use two different proof techniques to obtain discrepancy upper bounds. First, we consider the paths as a set system with vertices (for vertex discrepancy) or edges (for edge discrepancy) as the ground set and then apply a standard application of *primal shatter functions* (Definition 3.5), which bounds the number of subsets obtained

[4]In general, on a weighted undirected graph one can use random perturbation to ensure all shortest paths are unique.

by limiting to only s elements. For any family of consistent paths, as well as shortest paths on directed graphs (which are not necessarily consistent), the primal shatter function is upper bounded by $O(s^2)$. By a well known bound on discrepancy through the primal shatter function, this immediately gives us an upper bound of $O(n^{1/4})$ for vertex discrepancy and $O(m^{1/4})$ for edge discrepancy (since edge discrepancy is defined on a ground set of m edges in the graph G).

When the graph is dense, this upper bound on edge discrepancy deteriorates, becoming trivial when $m = \Theta(n^2)$. We thus present a second proof of $\widetilde{O}(n^{1/4})$ for *both* vertex and edge discrepancy for a family of consistent paths, which explicitly constructs a low-discrepancy coloring. This improves the bound for vertex discrepancy by polylogarithmic factors and edge discrepancy by polynomial factors. The main idea in this construction is to adapt the *path cover* technique, used in the recent breakthrough on shortcut sets [44]. That is, we start by finding a small base set of roughly $n^{1/2}$ node-disjoint shortest paths in the distance closure of the graph. These paths have the property that any other shortest path π in the graph contains at most $O(n^{1/2})$ nodes that are not in any paths in the base set. We then color *randomly*, as follows:

- For every node that is not contained in any path in the base set, we assign its color randomly. Thus, applying concentration bounds, the contribution of these nodes to the discrepancy of π will be bounded by $\pm\widetilde{O}(n^{1/4})$.
- For every path in the base set, we choose the color of the first node in the path at random, and then alternate colors along the path after that. Then we can argue that by consistency, the nodes in each base path randomly contribute $+1$ or -1 (or 0) to the discrepancy of π (see Figure 1 for a visualization). Since there are only $n^{1/2}$ paths in the base set, we may again apply concentration bounds to argue that the contribution to discrepancy from these base paths will only be $\pm\widetilde{O}(n^{1/4})$.

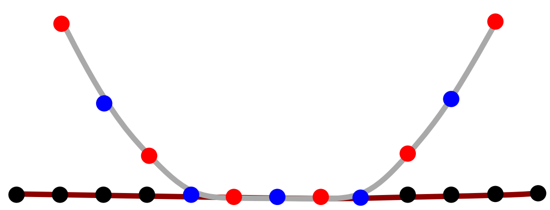

FIG. 1. *If we color the nodes of a unique shortest path with alternating colors, then its nodes will contribute discrepancy* 0, +1, *or* −1 *to all unique shortest paths that intersect it.*

Summing together these two parts, we obtain a bound of $\widetilde{O}(n^{1/4})$ on discrepancy with high probability. We can translate this to a bound on *hereditary* discrepancy using the fact that consistency is a hereditary property of path systems.

Lower Bound Techniques. Lower bounds on discrepancy are typically obtained by the trace bound (e.g., by Chazelle and Lvov [21]) on an explicit graph construction. The state-of-the-art lower bounds on the discrepancy of unique shortest paths were achieved using a point-line construction of Erdős [55], which had n points and n lines in $\mathbb{R}^2$ with each point staying on $\Theta(n^{1/3})$ lines and each line going through $\Theta(n^{1/3})$ points. This point-line system also implies tight lower bounds for the Szemerédi-Trotter theorem [61] and the discrepancy of arrangements of lines in the plane [21]. It can be associated with a graph that possesses useful properties derived from geometry. If edges in this graph are weighted by Euclidean distance, then the paths in the graph corresponding to straight lines are unique shortest paths by design. Two such shortest paths (along straight lines) only intersect at most once.

Probably not a coincidence, this point-line construction also provides lower bounds on the graph hopset problem. An (exact) hopset of a graph G with hop-bound β is a small set of additional edges H in the distance closure of G, such that every pair of nodes has a shortest path in $G \cup H$ containing at most β edges. Until recently, the state-of-the-art lower bound on the size of the hopset uses exactly the graph derived from the point-line incidence example. Recently, a construction by Bodwin and Hoppenworth [15] obtained stronger hopset lower bounds with a different geometric graph construction, which still took place in $\mathbb{R}^2$ but allowed shortest paths to have many vertices/edges in common. We show that this construction can be repurposed to derive a stronger lower bound of $\widetilde{\Omega}(n^{1/4})$ on vertex hereditary discrepancy by applying the trace bound of Larsen [46]. Combined with our upper bounds, this substantially improves our understanding of the discrepancy of unique shortest paths.

The above upper and lower bounds are for general graphs. Naturally, one can ask if we have better bounds for special families of graphs. We further show that the lower bounds remain the same for two interesting families: *planar graphs* and *bipartite graphs*. The lower bound construction mentioned above is not planar, and so this requires some additional work. A natural attempt is to restore planarity by adding vertices to the construction wherever two edges cross. However, this comes at a cost of an increase in the number of vertices and also with a potential danger of altering the shortest paths. We show that the number of crossings is not too much higher than n. Then, by carefully changing the weights of the edges and by exploiting the geometric properties of the construction, we show that the topology and incidence of shortest paths are not altered. For bipartite graphs, although the vertex discrepancy can be made very low – by coloring the vertices on one side $+1$ and vertices on the other side -1 – the hereditary discrepancy can be as high as the general graph setting. Specifically, we show a 2-lift of any graph G to a bipartite graph which essentially keeps the same hereditary discrepancy.

2. Related Work.

2.1. Discrepancy on graphs. In graph theory, the *discrepancy* of a graph introduced by Erdős [34] is defined as follows:

$$\max_{U \subset V} \left| e(G[U]) - p\binom{|U|}{2} \right|,$$

where $e(G[U])$ is the number of edges of the induced subgraph $G[U]$ on vertices $U \subseteq V$ and $p = \frac{|E|}{\binom{n}{2}}$ is the density of edges. If we consider a complete graph and randomly color each edge with probability p, the above definition of discrepancy quantifies the deviation of induced subgraphs of G from their expected size. Erdős and Spencer [35] showed that the graph discrepancy is $\Theta(n^{3/2})$ when $p = 1/2$. This definition and related definitions (e.g., positive discrepancy, dropping the absolute operator) have applications to quasi-randomness [24], graph cuts and edge expansion [4, 16, 57]. There is also study of multicolor discrepancy [38, 39] that we skip here.

Of particular relevance to our work, Balogh et al. [7] studied edge discrepancy (as defined in this paper) of (spanning) trees, paths and Hamilton cycles of a graph G. In particular, they showed that, for any labeling of edges of G, there is a path with discrepancy $\Omega(n)$, even when the graph is a grid. Prior to this, either probabilistic construction exhibiting such a lower bound was known [10, 30, 48] or an explicit construction of linear size non-planar graphs was known [5]. The construction for the planar graph in Balogh et al. [7] can be extended to coloring of vertices such that there is a path with vertex discrepancy $\Omega(n)$ when there are exponentially many paths and $\Omega(\sqrt{n})$ when there are polynomially many paths. (see Appendix A for details).

Discrepany of paths in directed graphs has also been studied. Reimer [58] showed that, if a directed graph has discrepancy $\Omega(n)$, then the graph must have $\Omega(n^2)$ edges. In the case when we do not allow antiparallel edges, Ben-Eliezer et al. [11] showed that there is a directed graph with $\Theta(n^2 \log^2(n))$ edges such that any mapping $\chi : E \to \{-1, 1\}$ will either have a path of length $\Omega(n)$ and all edges mapped to -1 or a path of length $\Omega(n \log(n))$ with all edges mapped to $+1$.

2.2. Connection with curve discrepancy. A classical topic in computational geometry is to study upper and lower bounds of the discrepancy/hereditary discrepancy of the incidence matrix of geometric objects and a set of points. For example, for a set of n points and n halfplanes in $\mathbb{R}^2$, the n by n incidence matrix (with rows corresponding to halfplanes and columns corresponding to points) has discrepancy of $\Omega(n^{1/4})$. For n points and n lines in the plane, the discrepancy of the incidence matrix is $\Omega(n^{1/6})$ [21]. In general the discrepancy of such incidence matrix is related to the 'complexity' of the geometric shapes. In our setting of a graph, the set of all pairs shortest paths defines a set system on the vertices. When the graph is planar, the shortest paths are essentially simple curves in the plane.

We would like to compare our results with discrepancy of curves and points in the geometric setting. Using the classification in Pach and Sharir [56], a family of simple curves have k *degree of freedom* and *multiplicity type* s if, for any k points, there are at most s curves passing through all of them, and any pair of curves intersect in at most s points. Lines in the plane have degree of 2 and multiplicity of 1. A set of curves with degree of 2 and multiplicity of 1 is called *pseudolines* – two pseudolines have at most one intersection. For n points and n lines, the discrepancy is upper bounded by $O(n^{1/6}(\log n)^{2/3})$ [20] which is nearly tight by a polylogarithmic factor. The proof uses the standard partial coloring argument with the Szeremédi-Trotter bound on point-line incidence – for any n points and m lines there are at most $O(m^{2/3}n^{2/3} + m + n)$ point-line incidences [61]. The Szeremédi-Trotter bound can be extended to a set of pseudolines [56, 60]. Therefore the same proof and upper bound hold for the discrepancy of pseudolines.

For a consistent set of shortest paths, two shortest paths will only intersect at a contiguous segment, which may have multiple vertices/points. Thus using the curve classification criterion, a consistent family of shortest paths in the plane has degree of 2 but multiplicity s that is possibly higher than a constant. In fact, our discrepancy lower bound construction in the planar graph setting uses a design with s possibly as high as $n^{1/2}$. This is the major difference of shortest paths in a planar graph with pseudolines, which allows the discrepancy of shortest paths to go beyond the pseudoline upper bound of $\widetilde{O}(n^{1/6})$[5].

3. Preliminaries. A *path system* is a pair $S = (V, \Pi)$ where V is a ground set of nodes and Π is a set of vertex sequences called *paths*. Each path may contain, at most, one instance of each node. We now formally define consistency, a structural

[5]Using the incidence upper bound for $k = 2$ and $s = n^{1/2}$ from [60] and partial coloring, one can obtain a discrepancy upper bound of $\widetilde{O}(n^{1/3})$ for our path construction. In contrast, we obtained a nearly tight bound of $\widetilde{\Theta}(n^{1/4})$.

property of unique shortest paths that will be useful.

DEFINITION 3.1. *A path system $S = (V, \Pi)$ is* consistent *if no two paths in S intersect, split apart, and then intersect again later. Formally:*

- *In the undirected setting, consistency means that for all $u, v \in V$ and all $\pi_1, \pi_2 \in \Pi$ such that $u, v \in \pi_1 \cap \pi_2$, we have that $\pi_1[u,v] = \pi_2[u,v]$, i.e., the intersection of π_1 and π_2 is a contiguous subpath (subsequence) of π_1 and π_2.*
- *In the directed setting, consistency means that for all $u, v \in V$ and all $\pi_1, \pi_2 \in \Pi$ such that u **precedes** v in both π_1 and π_2, we have that $\pi_1[u,v] = \pi_2[u,v]$.*

In every weighted graph for which all pairs shortest paths exist (i.e., no negative cycles), we can represent all-pairs shortest paths using a consistent path system. In particular, if all shortest paths are *unique*, then consistency is implied immediately.

We will investigate the combinatorial discrepancy of path systems (V, Π). Usually, we will assume that $|V| = n$ and $|\Pi|$ is polynomial in n. We define a vertex coloring $\chi : V \mapsto \{-1, 1\}$ and define the *discrepancy* of Π as

$$\text{disc}(\Pi) = \min_{\chi} \chi(\Pi), \quad \text{where } \chi(\Pi) = \max_{\pi \in \Pi} |\chi(\pi)|, \; \chi(\pi) = \sum_{v \in \pi} \chi(v).$$

Using a random coloring χ, we can guarantee that for all paths $\pi \in \Pi$ [20]:

$$|\chi(\pi)| \leq \sqrt{2|\pi| \ln(4|\Pi|)}.$$

This immediately provides a few observations.

OBSERVATION 3.2. *When Π is a set of paths with size polynomial in n, then* $\text{disc}(\Pi) = O(\sqrt{n \log n})$. *This bound is true even for paths that are possibly non-consistent.*

OBSERVATION 3.3. *When the longest path in Π has D vertices we have* $\text{disc}(\Pi) = O(\sqrt{D \log n})$. *Thus, for graphs that have a small diameter (e.g., small world graphs), the discrepancy of shortest paths is automatically small.*

Hereditary discrepancy is a more robust measure of the complexity of a path system (V, Π), defined as $\text{herdisc}(\Pi) = \max_{Y \subseteq V} \text{disc}(\Pi|_Y)$, where $\Pi|_Y$ is the collection of sets of the form $\pi \cap Y$ with $\pi \in \Pi$. Clearly, $\text{herdisc}(\Pi) \geq \text{disc}(\Pi)$. Sometimes the discrepancy of a set system may be small while the hereditary discrepancy is large [20]. Thus in the literature, we often talk about lower bounds on the hereditary discrepancy.

Now that we have defined vertex and edge (hereditary) discrepancy, one may wonder if there is an underlying relationship between vertex and edge (hereditary) discrepancy since they share the same bounds in most settings presented in Table 2. The following observation shows that vertex discrepancy bounds directly imply bounds on edge discrepancy.

OBSERVATION 3.4. *Denote by* $\text{disc}(n)$ *(and* $\text{herdisc}(n)$*) the maximum discrepancy (minimum hereditary discrepancy, respectively) of a consistent path system of a (undirected or directed) graph of n vertices. We have that*

1. *Let $g(x)$ be a non-decreasing function. If* $\text{herdisc}_v(n) \geq g(n)$, *then*
$$\text{herdisc}_e(n) \geq g(n/2).$$
2. *Let $f(x)$ be a non-decreasing function. If* $\text{disc}_v(n) \leq f(n)$, *then*
$$\text{disc}_e(m) \leq f(m).$$

Proof. We first show that if graph $G = (V, E, w)$ with the consistent path matrix A_v has hereditary discrepancy at least $g(n)$, we can obtain another graph $G' = (V', E', w')$ and matrix A_e as the (consistent path) edge incidence matrix with hereditary discrepancy at least $g(n/2)$. The construction is as follows.

(a) We first split each vertex $v \in V$ in G to two vertices $(v_{\mathsf{in}}, v_{\mathsf{out}})$ to obtain V'.
(b) For every $v \in V$, add a single edge $(v_{\mathsf{in}}, v_{\mathsf{out}})$ to E'.
(c) For any $v \in V$ and each edge $(u, v) \in E$ (with the fixed v), add edges $(u_{\mathsf{out}}, v_{\mathsf{in}})$ and $(u_{\mathsf{in}}, v_{\mathsf{out}})$ to E'.

The path incident matrix A_e is defined as follows: for each path as a row a of A_v, construct a new path in G' by following the order of $u_{\mathsf{out}} \to v_{\mathsf{in}} \to v_{\mathsf{out}} \to w_{\mathsf{in}}$ for a $u \to v \to w$ sequence. For each row in A_v, we mark the used edges as 1 in A_e with the path constructed by the above process. Note that the new path system defined by A_e remains consistent: for any intersection between the two paths $P_1 \cap P_2 = (u_1, u_2, \cdots, u_\ell)$, the intersection remains a single path of $(u_{1,\mathsf{in}}, u_{1,\mathsf{out}}, \cdots, u_{\ell,\mathsf{in}}, u_{\ell,\mathsf{out}})$ in A_e.

Let Y be the columns that induces the $g(n)$ discrepancy on G, i.e.,

$$\min_{x \in \{-1,+1\}^{|Y|}} \|A_{v_Y} x\|_\infty = g(n).$$

Now, observe that, for each row in A_e, an edge $(v_{\mathsf{in}}, v_{\mathsf{out}})$ is marked as 1 if and only if v is marked as 1 in A_v. Therefore, we can take the new set Y' as the edges corresponding to Y, and there is

$$\min_{x \in \{-1,+1\}^{|Y'|}} \left\|A_{e_{Y'}} x\right\|_\infty = \min_{x \in \{-1,+1\}^{|Y|}} \|A_{v_Y} x\|_\infty = g(n).$$

Finally, since graph G' has $n' = 2n$ vertices, we have the hereditary discrepancy to be at least $g(n'/2)$, as desired.

We next show that the hereditary edge discrepancy of G is at most $f(m)$, which implies the discrepancy upper bound. For a graph with n vertices, m edges, and a path incident matrix A_e, suppose Y is the set of columns (edges) that attain the hereditary discrepancy. We can add a vertex v_e for each $e \in Y$ and construct a new path incident matrix A_v, which is a matrix with $|Y|$ rows. Concretely, for each row of A_v, we simply let vertices $v_e \in Y$ be 1 if the corresponding edge is used in A_e. By the consistency of A_e, the new path incident matrix also characterizes a consistent path system (we can think of the underlying graph as the complete graph on a vertex set of Y). Note that we can get $f(m)$ discrepancy for the path system characterized by A_e as there are at most m vertices in A_e. This implies a $f(m)$ hereditary edge discrepancy on the original path system, which in turn implies the desired discrepancy upper bound.

Finally, note that the argument remains valid when the graph is directed, which means the results hold for both undirected and directed graphs. ∎

We also use some technical tools from discrepancy theory and statistics.

3.1. Known Results in Discrepancy Theory. The first result that we discuss is the one that gives an upper bound on the discrepancy of a set system in terms of *primal shatter function.*

Definition 3.5 (Primal Shatter Function). *Let $(X, \mathcal{R})$ be a set system, i.e., X is a ground state and $S = \{S_1, S_2, \cdots, S_\ell\}$ with $S_i \subseteq X$ for all $1 \le i \le \ell$. Let s be a positive integer. The* primal shatter function*, denoted as $\pi_{\mathcal{R}}(s)$, is defined as*

$$\pi_{\mathcal{R}}(s) := \max_{A \subseteq X:\, |A| = s} |\{A \cap S \mid S \in \mathcal{R}\}|\,.$$

The following is a well-known result of the discrepancy theory.

Proposition 3.6 (Theorem 1.2 in Matousek [51]). *Given a set system $(X, \mathcal{R})$, the discrepancy of a range space $\mathcal{R}$ whose primal shatter function is bounded by $\pi_{\mathcal{R}}(x) = cx^d$, for some constant $c > 0$, $d > 1$, is*

$$O(n^{1/2-1/(2d)}),$$

where n is the size of the ground state, and $O(\cdot)$ hides the dependency on c and d.

For the lower bound on the hereditary discrepancy, one general result is the *trace bound* first shown by Chazelle and Lyov [21].

Lemma 3.7 (Trace Bound of Chazelle and Lyov [21]). *If A is an m by n incidence matrix and $M = A^\top A$, then there is an absolute constant $0 < c < 1$ such that*

$$\text{herdisc}(A) \geq \frac{1}{4} c^{n \cdot \text{tr}(M^2)/(\text{tr}(M))^2} \sqrt{\frac{\text{tr}(M)}{n}}.$$

Recently, this trace bound has been improved in the exponential term through a series of works culminating in the following bound in Larsen [46], which we also use in our work:

Lemma 3.8 (Trace Bound of Larsen [46]). *If A is an m by n incidence matrix and $M = A^\top A$, then*

$$\text{herdisc}(A) \geq \frac{(\text{tr}(M))^2}{8e\, \text{tr}(M^2) \cdot \min\{m, n\}} \sqrt{\frac{\text{tr}(M)}{\max\{m, n\}}}.$$

We give various interpretations of $\text{tr}(M)$ and $\text{tr}(M^2)$ in Lemma 3.8 that would be useful later on. Algebraically, $\text{tr}(M)$ is the sum of its eigenvalue while $\text{tr}(M^2)$ is the sum of the square of the eigenvalues. Combinatorially, $\text{tr}(M)$ is the number of ones in A and $\text{tr}(M^2)$ is the number of rectangles of all ones in A. Geometrically, $\text{tr}(M)$ is the count of point/region incidences, and $\text{tr}(M^2)$ is the number of pairs of points in all the pairwise intersections of regions. Finally, if A is the incidence matrix for the shortest path, $\text{tr}(M^2)$ is the number of length 4 cycles in the underlying graph. Based on the algebraic interpretation, it means that the trace bound is non-trivial whenever all the eigenvalues of A are fairly uniform. This can be seen by noticing that, if $\{\lambda_1, \cdots, \lambda_n\}$ are eigenvalues of A, then $\text{tr}(M)^2 = n \cos^2(\theta)\, \text{tr}(M)$, where θ is the angle between the vector $(\lambda_1, \cdots, \lambda_n)$ and the all one-vector.

Our lower bound construction requires a hard instance of a class of graphs. For that, we use the construction of Bodwin and Hoppenworth [15], whose key properties that we use are stated as the following lemma:

Lemma 3.9 (Lemma 1 of Bodwin and Hoppenworth [15]). *For any $p \in [1, n^2]$, there is an infinite family of n-node undirected weighted graphs $G = (V, E, w)$ and sets Π of p paths in G such that*

- *G has $\ell = \Theta\left(\frac{n}{\sqrt{p \log n}}\right)$ layers. Each path in Π starts in the first layer, ends in the last layer, and contains exactly one node in each layer.*
- *Each path in Π is the unique shortest path between its endpoints in G.*
- *For any two nodes $u, v \in V$, there are at most $\frac{\ell}{h(u,v)}$ paths in Π that contain both u and v, where $h(u, v)$ is the hop-distance (number of edges on the shortest path) between u and v in G and $1 \leq h(u, v) \leq \ell$.*
- *Each node $v \in V$ lies on at most $O\left(\frac{\ell p}{n}\right)$ distinct paths in Π.*

Concentration Inequalities. We use the following standard variants of the Chernoff-Hoeffding bound in our paper.

PROPOSITION 3.10 (Chernoff bound). *Let $X_1, \ldots, X_n$ be n independent random variables with support on $\{0,1\}$. Define $X := \sum_{i=1}^{n} X_i$. Then, for every $\delta > 0$, there is*

$$\Pr\left[X \geq (1+\delta) \cdot \mathbb{E}[X]\right] \leq \exp\left(-\frac{\delta^2}{\delta+2} \cdot \mathbb{E}[X]\right).$$

In particular, when $\delta \in (0,1]$, there is

$$\Pr[|X - \mathbb{E}[X]| > \delta \cdot \mathbb{E}[X]] \leq 2 \cdot \exp\left(-\frac{\delta^2 \mathbb{E}[X]}{3}\right).$$

PROPOSITION 3.11 (Additive Chernoff bound). *Let $X_1, \ldots, X_n$ be n independent random variables with support in $[0,1]$. Define $X := \sum_{i=1}^{n} X_i$. Then, for every $t > 0$,*

$$\Pr[|X - \mathbb{E}[X]| > t] \leq 2 \cdot \exp\left(-\frac{2t^2}{n}\right).$$

4. General Graphs: Upper Bound Existential Proof. This section collects the existential proof of the upper bounds on vertex- and edge-discrepancy for consistent path systems in (possibly) directed graphs. Our approach uses Proposition 3.6, which gives a discrepancy upper bound using the primal shatter function of a set system. This approach leads to the same upper bounds for undirected and directed graphs. In specifics, we show an upper bound of $O(n^{1/4})$ holds for vertex discrepancy, while the edge discrepancy is at most $O(m^{1/4})$. That is, we show an existential proof of Theorem 1.6 (on directed graphs) by this approach. Note that for undirected graphs, we have achieved better edge discrepancy bounds using explicit constructions (as shown in Subsection 5.3).

THEOREM 1.6. *For any n-node, m-edge directed weighted graph G with a unique shortest path between each pair of nodes, the system of shortest paths Π satisfies*

$$\text{herdisc}_v(\Pi) \leq O(n^{1/4}) \qquad \textit{and} \qquad \text{herdisc}_e(\Pi) \leq O(m^{1/4}).$$

Proof. We consider vertex discrepancy first. Let $S = (V, \Pi)$ be the path system containing all $|\Pi| = \binom{n}{2}$ unique shortest paths in G over vertex set V. We can interpret S as a set system (e.g., by ignoring the ordering of vertices in paths $\pi \in \Pi$).

We claim that the primal shatter function π_S of S is $\pi_S(x) = O(x^2)$. The intersection of any set $A \subseteq V$ of $|A| = x$ vertices with a path $\pi = \pi[s,t] \in \Pi$ is equal to $A \cap \pi[u,v]$ with u and v being the first and last vertex on path π in set A, respectively. Then we have

$$|\{A \cap \pi \mid \pi \in \Pi\}| \leq |A|^2 = x^2,$$

and $\pi_S(x) = O(x^2)$, as claimed. Since the size of the ground state is $|V| = n$, Proposition 3.6 implies that the (non-hereditary) vertex discrepancy of the incidence matrix for a family of consistent paths on an n-node graph is at most $O(n^{1/4})$. An upper bound of $O(m^{1/4})$ for edge discrepancy on m-edge graphs follows from Observation 3.4.

Finally, to show the upper bound on *hereditary discrepancy*, we observe that for any subset $U \subseteq V$, we can define the system $\Pi[U]$ of the paths in Π induced on the nodes in U. This path system $\Pi[U]$ is also consistent. Applying the above argument on $\Pi[U]$ therefore give us an $O(n^{1/4})$ upper bound for the discrepancy of $\Pi[U]$, implying our desired vertex hereditary discrepancy upper bound. A similar argument achieves an edge hereditary discrepancy upper bound of $O(m^{1/4})$. ∎

Notice that for a sparse graph (e.g., $m = O(n)$) this matches the bound on vertex discrepancy, but for a dense graph (e.g., $m = \Theta(n^2)$), the upper bound becomes $O(n^{1/2})$, which is no better than the upper bound by random coloring.

In Subsection 5.2, we present a vertex coloring achieving hereditary discrepancy of $\widetilde{O}(n^{1/4})$. Finally, we present an explicit edge coloring with the same hereditary discrepancy bound in Subsection 5.3. This is a significant improvement over $O(m^{1/4})$, especially for dense graphs.

5. Undirected Graphs: Lower Bound and Explicit Colorings. We now discuss the main result (Theorem 1.5). We first show in Subsection 5.1 a hereditary discrepancy lower bound of $\Omega(n^{1/4}/\sqrt{\log n})$ for both edge and vertex discrepancy in general undirected graphs. Then, in Subsection 5.2, we present a vertex coloring achieving hereditary discrepancy of $\widetilde{O}(n^{1/4})$. Finally, we present an explicit edge coloring with the same hereditary discrepancy bound in Subsection 5.3.

5.1. Lower Bound. As suggested by Observation 3.4, we focus on the vertex hereditary discrepancy. We then show that this theorem implies the same lower bound on (non-hereditary) vertex discrepancy as well.

THEOREM 5.1. *There are examples of n-vertex undirected weighted graphs G with a unique shortest path between each pair of vertices in which this system of shortest paths Π has*

$$\mathrm{herdisc}_v(\Pi) \geq \Omega\left(\frac{n^{1/4}}{\sqrt{\log(n)}}\right).$$

To obtain the lower bound, we employ the new graph construction by Bodwin and Hoppenworth [15], which shows that any exact hopset with $O(n)$ edges must have at least $\widetilde{\Omega}(n^{1/2})$ hop diameter. Despite seeming unrelated, this construction also sheds light on our problem. Another technique we use to show the hereditary discrepancy lower bound is the trace bound [46]. In the following proof section, we first summarize the construction related to our objective, then show the calculation using the trace bound that leads to our lower bound.

Proof. The key properties of the graph construction in Bodwin and Hoppenworth [15] that we need can be summarized in Lemma 3.9. We will make use of the shortest path vertex incidence matrix of the graph in Bodwin and Hoppenworth [15]. Recall that hereditary discrepancy considers the sub-incidence matrix induced by columns corresponding to a set of vertices. We select the set of vertices occurring in the paths in Π, and show it leads to hereditary discrepancy at least $\Omega(n^{1/4}/\sqrt{\log n})$. Specifically, take A as the incidence matrix so each row corresponds to one path in Π. A has dimension $p \times n$ where n is the number of vertices in G and the (i, j)-th entry of A is 1 is the vertex j is in the path i.

Now define $M = A^\top A$. Recall that $\mathrm{tr}(M)$ is the number of 1s in the matrix A. Since by construction, every path has length ℓ, we have $\mathrm{tr}(M) = p\ell$. Furthermore, let m_{ij} be the (i, j)-th element of matrix M, and observe that it is exactly the number of paths that contain vertices i and j. Note that $m_{ij} = m_{ji}$. Additionally, $\mathrm{tr}(M^2)$ is the number of length 4 closed walks in the bipartite graph representing the incidence matrix A. This implies that

$$\operatorname{tr}(M^2) = \sum_{j=1}^{p} \sum_{\substack{u,v\in P_j,\\ u\neq v}} m_{u,v} + \sum_{i=1}^{n} m_{ii}^2 = \sum_{j=1}^{p}\sum_{i=1}^{\ell} \sum_{\substack{u,v\in P_j,\\ h(u,v)=i}} m_{u,v} + n\cdot O\left(\frac{p\ell}{n}\right)^2$$

(5.1)

$$\leq \sum_{j=1}^{p}\sum_{i=1}^{\ell} \ell\cdot\frac{\ell}{i} + O\left(\frac{p^2\ell^2}{n}\right) \leq p\ell^2\log\ell + O\left(\frac{p^2\ell^2}{n}\right).$$

By setting $p = n\log n$, it follows that $\ell = \Theta\left(\frac{\sqrt{n}}{\log n}\right)$ and $\operatorname{tr}(M) = p\ell$. Further,

$$np\ell^2\log\ell \leq O\left(n\cdot n\log(n)\cdot\frac{n}{\log^2 n}\cdot\log(n)\right) = O(n^3) = O(p^2\ell^2).$$

By equation (5.1), we have $\operatorname{tr}(M^2) = O(p^2\ell^2/n)$. Using this and $\operatorname{tr}(M) = p\ell$ in the trace bound in Lemma 3.8 [46] gives us

$$\begin{aligned}\operatorname{herdisc}_v(\Pi) &\geq \frac{(\operatorname{tr}(M))^2}{8e\min\{p,n\}\cdot\operatorname{tr}(M^2)}\sqrt{\frac{\operatorname{tr}(M)}{\max\{p,n\}}}\\ &= \frac{(\operatorname{tr}(M))^2}{8en\cdot\operatorname{tr}(M^2)}\sqrt{\frac{\operatorname{tr}(M)}{p}} \geq \Omega\left(\frac{p^2\ell^2}{p^2\ell^2}\sqrt{\ell}\right)\\ &= \Omega(\sqrt{\ell}) = \Omega\left(\frac{n^{1/4}}{\sqrt{\log(n)}}\right)\end{aligned}$$

by setting the value of ℓ. The completes the proof of the lower bound. □

5.2. Vertex Discrepancy Upper Bound – Explicit Coloring. In this subsection, we will upper bound the discrepancy $\chi(\Pi)$ of a consistent path system (V,Π) with $|V| = n$ and $|\Pi| = \operatorname{poly}(n)$. This immediately implies an upper bound for the hereditary vertex discrepancy of unique shortest paths in undirected graphs.

THEOREM 5.2. *For a consistent path system $S = (V,\Pi)$ where $|V| = n$ and $|\Pi| =$ poly(n), there exists a labeling χ such that $\chi(\Pi) = O(n^{1/4}\log^{1/2}(n))$. Consequently, every n-vertex undirected graph has hereditary vertex discrepancy $O(n^{1/4}\log^{1/2}(n))$.*

Let $S = (V,\Pi)$ be a consistent path system with $|V| = n$ and $|\Pi| = \operatorname{poly}(n)$. As the first step towards constructing our labeling $\chi : V \mapsto \{-1,1\}$, we will construct a collection of paths Π' on V that will have a useful covering property, stated in Proposition 5.3, over the paths in Π.

Constructing path cover Π'. Initially, we let $\Pi' = \emptyset$. We define V' to be the set of all nodes in V belonging to a path in Π', i.e., $V' := \bigcup_{\pi'\in\Pi'}\pi'$. While $|\pi\setminus V'| \geq n^{1/2}$ for some $\pi\in\Pi$, find a (possibly non-contiguous) subpath of π of length $n^{1/2}$ that is vertex-disjoint from all paths in Π'. Formally, find a subpath $\pi'\subseteq\pi$ such that $|\pi'| = n^{1/2}$ and $\pi'\cap V' = \emptyset$. Add path π' to path cover Π' and update V'. Repeatedly add paths to path cover Π' in this manner until $|\pi\setminus V'| < n^{1/2}$ for all $\pi\in\Pi$.

PROPOSITION 5.3. *Path cover Π' satisfies the following properties:*
1. *for all $\pi\in\Pi'$, $|\pi| = n^{1/2}$,*
2. *the number of paths in Π' is $|\Pi'| \leq n^{1/2}$,*

3. (Disjointness Property) *The paths in* Π' *are pairwise vertex-disjoint,*
4. (Covering Property) *For all* $\pi \in \Pi$*, the number of nodes in* π *that do not lie in any path in path cover* Π' *is at most* $n^{1/2}$*. Formally, let* $V' = \cup_{\pi' \in \Pi'} \pi'$*. Then* $\forall \pi \in \Pi, |\pi \setminus V'| \leq n^{1/2}$,
5. (Consistency Property) *For all* $\pi \in \Pi$ *and* $\pi' \in \Pi'$*, the intersection* $\pi \cap \pi'$ *is a (possibly empty) contiguous subpath of* π'.[6]

Proof. Properties 1, 3, and 4 follow from the construction of Π'. Property 2 follows from Properties 1 and 3 and the fact that $|V| = n$. The Consistency Property of Π' is inherited from the consistency of path system S. Specifically, by the construction of Π', path $\pi' \in \Pi'$ is a subpath of a path $\pi'' \in \Pi$. Recall that by the consistency of path system S, the intersection $\pi \cap \pi''$ is a (possibly empty) contiguous subpath of π''. Then $\pi \cap \pi'$ is a contiguous subpath of π' since $\pi' \subseteq \pi''$. This concludes the proof. ∎

Constructing labeling χ**.** Let $\pi' = (v_1, \dots, v_k) \in \Pi'$ be a path in our path cover. We will label the nodes of π' using the following random process. With probability $1/2$ we define $\chi : \pi' \mapsto \{-1, 1\}$ to be

$$\chi(v_i) = \begin{cases} 1 & i \equiv 0 \mod 2 \text{ and } i \in [1, k] \\ -1 & i \equiv 1 \mod 2 \text{ and } i \in [1, k] \end{cases},$$

and with probability $1/2$ we define $\chi : \pi' \mapsto \{-1, 1\}$ to be

$$\chi(v_i) = \begin{cases} -1 & i \equiv 0 \mod 2 \text{ and } i \in [1, k] \\ 1 & i \equiv 1 \mod 2 \text{ and } i \in [1, k] \end{cases}.$$

The labels of consecutive nodes in π' alternate between 1 and -1, with vertex v_1 taking labels 1 and -1 with equal probability. Since the paths in path cover Π' are pairwise vertex-disjoint, the labeling χ is well-defined over $V' := \cup_{\pi' \in \Pi'} \pi'$. We choose random labeling for all nodes in $V \setminus V'$, i.e., we independently label each node $v \in V \setminus V'$ with $\chi(v) = -1$ with probability $1/2$ and $\chi(v) = 1$ with probability $1/2$. An illustration can be found in Figure 2.

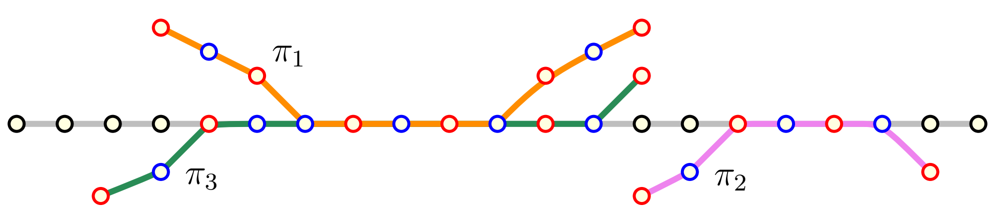

FIG. 2. *In this figure, paths* $\pi_1, \pi_2, \pi_3 \in \Pi'$ *from the path cover are intersecting a path* $\pi \in \Pi$*. Paths in the path cover are pairwise vertex-disjoint, and each path in the cover contributes discrepancy* 0, -1, *or* $+1$ *to* π.

Bounding the discrepancy $\chi(\Pi)$**.** Fix a path $\pi \in \Pi$. We will show that

$$\left| \sum_{v \in \pi} \chi(v) \right| = O(n^{1/4} \log^{1/2}(n))$$

with high probability. Theorem 5.2 will follow as $|\Pi| = \text{poly}(n)$.

[6]Note that it may not be true that $\pi \cap \pi'$ is a contiguous subpath of π.

Proposition 5.4. *For each path π' in path cover Π',*

$$\sum_{v\in\pi\cap\pi'}\chi(v)\in\{-1,0,1\}.$$

If $|\pi\cap\pi'|\equiv 0 \mod 2$, then $\sum_{v\in\pi\cap\pi'}\chi(v)=0$. Moreover,

$$\Pr\left[\sum_{v\in\pi\cap\pi'}\chi(v)=-1\right]=\Pr\left[\sum_{v\in\pi\cap\pi'}\chi(v)=1\right].$$

Proof. By the Consistency Property of Π' (as proven in Proposition 5.3), path $\pi\cap\pi'$ is a (possibly empty) contiguous subpath of π'. Then since consecutive nodes in π' alternate between -1 and 1, it follows that

$$\sum_{v\in\pi\cap\pi'}\chi(v)\in\{-1,0,1\}.$$

Now note that $\sum_{v\in\pi\cap\pi'}\chi(v)\neq 0$ iff $|\pi\cap\pi'|$ is odd. Moreover, the first vertex of $\pi\cap\pi'$ takes labels 1 and -1 with equal probability. This concludes the proof of Proposition 5.4. ∎

We are now ready to bound the discrepancy of π.

Proposition 5.5. *With high probability, $\chi(\pi)=O(n^{1/4}\log^{1/2}n)$.*

Proof. We partition the nodes of π into two sources of discrepancy that we will bound separately. Let $V':=\cup_{\pi'\in\Pi'}\pi'$.

Discrepancy of $\pi\cap V'$. For each path $\pi'\in\Pi'$, let $X_{\pi'}$ be the random variable defined as

$$X_{\pi'}:=\sum_{v\in\pi\cap\pi'}\chi(v).$$

We can restate the discrepancy of $\pi\cap V'$ as

$$\left|\sum_{v\in\pi\cap V'}\chi(v)\right|=\left|\sum_{\pi'\in\Pi'}X_{\pi'}\right|.$$

By Proposition 5.4, if $|\pi\cap\pi'|$ is even, then $X_{\pi'}=0$. Therefore, we may assume without any loss of generality that $|\pi\cap\pi'|$ is odd for all $\pi'\in\Pi'$. In this case, $\Pr[X_{\pi'}=-1]=\Pr[X_{\pi'}=1]=1/2$, implying that $\mathbb{E}[\sum_{\pi'\in\Pi'}X_{\pi'}]=0$. Then by Proposition 5.3 and the Chernoff bound, it follows that for any constant $c\geq 1$,

$$\Pr\left[\left|\sum_{\pi'\in\Pi'}X_{\pi'}\right|\geq c\cdot n^{1/4}\log^{1/2}(n)\right]\leq e^{-c^2\frac{n^{1/2}\log(n)}{2|\Pi'|}}\leq e^{-c^2/(2\cdot\log(n))}=n^{-c^2/2}.$$

Discrepancy of $\pi\setminus V'$. Note that by the Covering Property of the path cover (as proven in Proposition 5.3), $|\pi\setminus V'|\leq n^{1/2}$. Moreover, the nodes in $V\setminus V'$ are labeled independently at random, implying that $\mathbb{E}\left[\sum_{v\in\pi\setminus V'}\chi(v)\right]=0$. Then we may apply a Chernoff bound to argue that for any constant $c\geq 1$,

$$\Pr\left[\left|\sum_{v\in\pi\setminus V'}\chi(v)\right|\geq c\cdot n^{1/4}\log^{1/2}n\right]\leq \exp\{-c^2\frac{n^{1/2}\log(n)}{2|\pi\setminus V'|}\}$$
$$\leq e^{-c^2/(2\cdot\log(n))}=n^{-c^2/2}.$$

We have shown that the discrepancy of our labeling is $O(n^{1/4}\log^{1/2}(n))$ for $\pi \cap V'$ and $O(n^{1/4}\log^{1/2}(n))$ for $\pi \setminus V'$ with high probability. Therefore, with high probability, the total discrepancy of π is $O(n^{1/4}\log^{1/2}(n))$. This completes the proof of Proposition 5.5. ∎

Extending to hereditary discrepancy. Let A be the vertex incidence matrix of a path system $S = (V, \Pi)$ on n nodes, and let A_Y be the submatrix of A obtained by taking all of its rows but only a subset Y of its columns. Then there exists a subset $V_Y \subseteq V$ of the nodes in V such that A_Y is the vertex incidence matrix of the path system $S[V_Y]$ (path system S induced on V_Y). Moreover, if path system S is consistent, then $S[V_Y]$ is also consistent. Then we may apply our explicit vertex discrepancy upper bound to $S[V_Y]$. We conclude that the hereditary vertex discrepancy of S is $O(n^{1/4}\log^{1/2}(n))$.

5.3. Edge Discrepancy Upper Bound – Explicit Coloring. By Theorem 1.6, the edge discrepancy of the unique shortest paths of a (possibly directed) graph on m edges is $O(m^{1/4})$. However, in the case of undirected graphs and DAGs, we can improve the edge discrepancy to $O(n^{1/4}\log^{1/2}(n))$, where n is the number of vertices in the graph, by modifying the explicit construction for vertex discrepancy in Subsection 5.2. Our proof strategy will follow the same framework as the explicit construction for vertex discrepancy but with some added complications in the construction and analysis.

We first introduce some new notation that will be useful in this section. Given a path π and nodes $u, v \in \pi$, we denote by $u <_\pi v$ if u occurs before v on path π. Additionally, given a path system $S = (V, \Pi)$, we define the edge set $E \subseteq V \times V$ of the path system as the set of all pairs of nodes $u, v \in V$ that appear consecutively in some path in Π. Likewise, for any path π over the vertex set V, we define the edge set of π, $E(\pi) \subseteq \pi \times \pi$, as the set of all pairs of nodes $u, v \in \pi$ such that u, v appear consecutively in π and $(u, v) \in E$. Note that if path system S corresponds to paths in a graph G, then E will be precisely the edge set of G.

Recall that we wish to construct an edge labeling $\chi : E \mapsto \{-1, 1\}$ so that

$$\chi(\Pi) = \max_{\pi \in \Pi} \left| \sum_{e \in E(\pi)} \chi(e) \right|$$

is minimized. We will upper bound the discrepancy $\chi(\Pi)$ of consistent path systems such that $|V| = n$ and $|\Pi| = \text{poly}(n)$. This immediately implies an upper bound on the edge discrepancy of unique shortest paths in undirected graphs.

THEOREM 5.6. *For all consistent path systems $S = (V, \Pi)$ where $|V| = n$ and $|\Pi| = \text{poly}(n)$ with edge set E, there exists a labeling $\chi : E \to \{-1, 1\}$ such that*

$$\chi(\Pi) = O(n^{1/4}\log^{1/2}(n)).$$

Consequently, every n-vertex undirected graph has hereditary edge discrepancy $O(n^{1/4}\log^{1/2}(n))$.

Let $S = (V, \Pi)$ be a consistent path system with $|V| = n$ and $|\Pi| = \text{poly}(n)$. As the first step towards constructing our labeling $\chi : E \mapsto \{-1, 1\}$, we will construct a collection of paths Π' on V with a useful covering property over the paths in Π.

5.3.1. Constructing path cover Π'. Initially, we let $\Pi' = \emptyset$. We define V' to be the set of all nodes in V belonging to a path in Π', i.e.,

$$V' := \bigcup_{\pi' \in \Pi'} \pi'.$$

While there exists a path $\pi \in \Pi$ such that $|\pi \setminus V'| \geq n^{1/2}$, our goal is to find a (possibly non-contiguous) subpath of π of length $n^{1/2}$ that is vertex-disjoint from all paths in Π'. Specifically, let $\pi' \subseteq \pi$ be a (possibly non-contiguous) subpath of π containing exactly the first $n^{1/2}$ nodes in $\pi \setminus V'$. Add path π' to path cover Π' and update V'. Repeatedly add paths to path cover Π' in this manner until $|\pi \setminus V'| < n^{1/2}$ for all $\pi \in \Pi$.

Note that our path cover Π' is very similar to the path cover used in the explicit vertex discrepancy upper bound. Indeed, path cover Π' inherits all properties of the path cover defined in Subsection 5.2. The key difference here is that we require subpaths $\pi' \subseteq \pi$ in Π' to contain the *first* $n^{1/2}$ nodes in $\pi \setminus V'$. This will imply an additional property of our path cover, which we call the *No Repeats Property*.

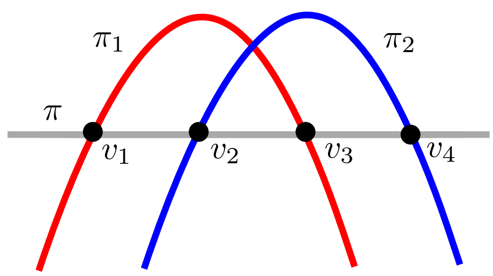

Fig. 3. *In this figure, paths $\pi_1, \pi_2 \in \Pi'$ are intersecting a path $\pi \in \Pi$. This arrangement of paths is forbidden by the No Repeats Property of Proposition 5.7.*

Proposition 5.7. *Path cover Π' satisfies all properties of Proposition 5.3, as well as the following additional properties:*

- (Edge Covering Property) *For all $\pi \in \Pi$, the number of edges in π that are not incident to any node lying in a path in path cover Π' is at most $n^{1/2}$. Formally, let $V' = \cup_{\pi' \in \Pi'} \pi'$. For all $\pi \in \Pi$,*

$$|\{(u, v) \in E(\pi) \mid u \notin V' \text{ and } v \notin V'\}| \leq n^{1/2},$$

- (No Repeats Property) *For all paths $\pi \in \Pi$, $\pi_1, \pi_2 \in \Pi'$, and nodes $v_1, v_2, v_3, v_4 \in \pi$ such that $v_1, v_3 \in \pi_1$ and $v_2, v_4 \in \pi_2$, the following ordering of the vertices in Π is impossible:*

$$v_1 <_\pi v_2 <_\pi v_3 <_\pi v_4,$$

where $x <_\pi y$ indicates that node x occurs in π before node y.

Proof. All properties from Proposition 5.3 follow from an identical argument as in the original proof. The Edge Covering Property follows immediately from the Covering Property of Proposition 5.3. What remains is to prove the No Repeats Property.

Suppose for the sake of contradiction that there exist paths $\pi \in \Pi$, $\pi_1, \pi_2 \in \Pi'$, and nodes $v_1, v_2, v_3, v_4 \in \pi$ such that $v_1, v_3 \in \pi_1$ and $v_2, v_4 \in \pi_2$, where $v_1 <_\pi v_2 <_\pi$

$v_3 <_\pi v_4$. We will assume that path π_1 was added to Π' before path π_2 (the case where π_2 was added to Π' first is symmetric). By the construction of Π', path $\pi_1 \in \Pi'$ is a (possibly non-contiguous) subpath of a path $\pi_1'' \in \Pi$ from which it is constructed. Additionally, by the consistency of the path system S, the intersection $\pi \cap \pi_1''$ is a contiguous subpath of π. Then $v_2 \in \pi \cap \pi_1''$, and specifically, $v_2 \in \pi_1''$.

We assumed that $v_2 \in \pi_2$, which implies that $v_2 \notin \pi_1$, since paths in Π' are pairwise vertex-disjoint. Since path π_1 was added to Π' before path π_2, this means that when π_1 was added to Π', node v_2 did not belong to any path in Π' (i.e., v_2 was not in V'). Recall that in our construction of Π', we constructed subpath $\pi_1 \subseteq \pi_1''$ so that it contained exactly the *first* $n^{1/2}$ nodes in $\pi_1'' \setminus V'$. However, $v_2 \notin \pi_1$, but $v_3 \in \pi_1$, and v_2 comes before v_3 in π_1''. This contradicts our construction of path π_1 in path cover Π'.

5.3.2. Constructing labeling χ. Let $\pi' \in \Pi'$ be a path of length k in our path cover. Let $e_1, \dots, e_k \in E(\pi')$ be the edges in π' listed in the order they appear in π'. Note that since π' is a possibly non-contiguous subpath of a path in Π, pairs of nodes $u, v \in V$ that appear consecutively in π do not necessarily correspond to edges in edge set E.

We will label the edges in $E(\pi')$ using the following random process. With probability 1/2 we define $\chi : E(\pi') \mapsto \{-1, 1\}$ to be

$$\chi(e_i) = \begin{cases} 1 & i \equiv 0 \mod 2 \text{ and } i \in [1, k] \\ -1 & i \equiv 1 \mod 2 \text{ and } i \in [1, k] \end{cases},$$

and with probability 1/2 we define $\chi : E(\pi') \mapsto \{-1, 1\}$ to be

$$\chi(e_i) = \begin{cases} -1 & i \equiv 0 \mod 2 \text{ and } i \in [1, k] \\ 1 & i \equiv 1 \mod 2 \text{ and } i \in [1, k] \end{cases}.$$

Note that the labels of consecutive edges e_i, e_{i+1} in π' alternate between 1 and -1, with edge e_1 taking labels 1 and -1 with equal probability.

Since the paths in path cover Π' are pairwise vertex-disjoint, the labeling χ is well-defined over

$$E' := \cup_{\pi' \in \Pi'} E(\pi'). \tag{5.2}$$

We take a random labeling for all edges in $E \setminus E'$, i.e., we independently label each edge $e \in E \setminus E'$ with $\chi(e) = -1$ with probability 1/2 and $\chi(e) = 1$ with probability 1/2.

5.3.3. Bounding the discrepancy ϕ. Fix a path $\pi := \pi[s, t] \in \Pi$. We will show that

$$\left| \sum_{e \in E(\pi)} \chi(e) \right| = O(n^{1/4} \log^{1/2}(n))$$

with high probability. This will complete the proof of Lemma 5.6 since $|\Pi| = \text{poly}(n)$. The proof of the following proposition follows from an argument identical to Proposition 5.4 and hence omitted.

PROPOSITION 5.8. *For each path π' in path cover Π',*

$$\sum_{e \in E(\pi) \cap E(\pi')} \chi(e) \in \{-1, 0, 1\}.$$

If $|E(\pi) \cap E(\pi')| \equiv 0 \mod 2$*, then* $\sum_{e \in E(\pi) \cap E(\pi')} \chi(e) = 0$*. Moreover,*

$$\Pr\left[\sum_{e \in E(\pi) \cap E(\pi')} \chi(e) = -1\right] = \Pr\left[\sum_{e \in E(\pi) \cap E(\pi')} \chi(e) = 1\right].$$

We are now ready to bound the edge discrepancy of π. Define

$$V' := \bigcup_{\pi' \in \Pi'} \pi' \quad \text{and} \quad E' := \bigcup_{\pi' \in \Pi'} E(\pi').$$

We partition the edges of the path π into three sources of discrepancy that we will bound separately. Specifically, using the definition of E' in equation (5.2), we split $E(\pi) \subseteq \pi \times \pi$ into the following sets E_1, E_2, E_3:

- $E_1 := E(\pi) \cap E'$,
- $E_2 := E(\pi) \cap ((V \setminus V') \times (V \setminus V'))$, and
- $E_3 := E(\pi) \setminus (E_1 \cup E_2)$.

Sets E_1 and E_2 roughly correspond to the two sources of discrepancy considered in the vertex discrepancy upper bound, while set E_3 corresponds to a new source of discrepancy requiring new arguments to bound. We begin with set E_1.

Proposition 5.9 (Discrepancy of E_1). *With high probability,* $\left|\sum_{e \in E_1} \chi(e)\right| = O(n^{1/4} \log^{1/2} n)$*.*

Proof. The proposition follows from an argument similar to Proposition 5.5. For each path $\pi' \in \Pi'$, let $X_{\pi'}$ be the random variable defined as

$$X_{\pi'} := \sum_{e \in E(\pi) \cap E(\pi')} \chi(e).$$

We can restate the discrepancy of $E_1 = E(\pi) \cap E'$ as

$$\left|\sum_{e \in E_1} \chi(e)\right| = \left|\sum_{\pi' \in \Pi'} X_{\pi'}\right|.$$

By Proposition 5.8, if $|E(\pi) \cap E(\pi')| \equiv 0 \mod 2$, then $X_{\pi'} = 0$, so without any loss of generality, we may assume that $|E(\pi) \cap E(\pi')|$ is odd for all $\pi' \in \Pi'$. In this case,

$$\Pr[X_{\pi'} = -1] = \Pr[X_{\pi'} = 1] = 1/2,$$

implying that $\mathbb{E}[\sum_{\pi' \in \Pi'} X_{\pi'}] = 0$. Then, by Proposition 5.7 and the Chernoff bound, it follows that for any constant $c \geq 1$,

$$\Pr\left[\left|\sum_{\pi' \in \Pi'} X_{\pi'}\right| \geq c \cdot n^{1/4} \log^{1/2}(n)\right] \leq e^{-c^2 \frac{n^{1/2} \log(n)}{2|\Pi'|}} \leq e^{-c^2/2 \cdot \log(n)} \leq n^{-c^2/2}.$$ ∎

We now bound the discrepancy of $E_2 = E(\pi) \cap ((V \setminus V') \times (V \setminus V'))$.

Proposition 5.10 (Discrepancy of E_2). *With high probability,* $\left|\sum_{e \in E_2} \chi(e)\right| = O(n^{1/4} \log^{1/2}(n))$*.*

Proof. The proposition follows from an argument similar to Proposition 5.5. Note that by the Edge Covering Property of the path cover (Proposition 5.7),

$$|E_2| = |\{(u,v) \in E(\pi) \mid u, v \notin V'\}| \leq n^{1/2}.$$

Moreover, the edges in $E \setminus E'$ are labeled independently at random, so we may apply a Chernoff bound to argue that for any constant $c \geq 1$,

$$\Pr\left[\left|\sum_{e \in E_2} \chi(e)\right| \geq c \cdot n^{1/4} \log^{1/2}(n)\right] \leq e^{-c^2 \frac{n^{1/2}\log(n)}{2|E_2|}} \leq e^{-c^2/2 \cdot \log(n)} \leq n^{-c^2/2}.$$

completing the proof. ∎

Finally, we upper bound the discrepancy of the remainder of the edges, $E_3 = E(\pi) \setminus (E_1 \cup E_2)$.

PROPOSITION 5.11 (Discrepancy of E_3). *With high probability,* $\left|\sum_{e \in E_3} \chi(e)\right| = O(n^{1/4} \log^{1/2}(n))$.

Proof. Let

$$k := |\{\pi' \in \Pi' \mid \pi \cap \pi' \neq \emptyset\}|$$

denote the number of paths in our path cover that intersect π. We define a function $f : \mathbb{Z}_{\geq 0} \mapsto \mathbb{Z}_{\geq 0}$ such that $f(\phi)$ equals the largest possible value of $|E_3|$ when $\phi = k$. Note that f is well-defined since $0 \leq |E_3| \leq |E|$. We will prove that $f(\phi) \leq 4\phi$, by recursively decomposing path π.

When $\phi = 1$, there is only one path $\pi' \in \Pi'$ that intersects π. Then the only edges in E_3 are of the form

$$E(\pi) \cap ((V' \times (V \setminus V')) \cup ((V \setminus V') \times V')) = E(\pi) \cap ((\pi' \times (V \setminus \pi')) \cup ((V \setminus \pi') \times \pi')).$$

By the Consistency Property of Proposition 5.7, path π' can intersect π and then split apart at most once. Then

$$f(1) = |E_3| = |E(\pi) \cap ((\pi' \times (V \setminus \pi')) \cup ((V \setminus \pi') \times \pi'))| \leq 2.$$

When $\phi > 1$, we will split our analysis into the two cases:

- **Case 1.** There exists paths $\pi'_1, \pi'_2 \in \Pi'$ and nodes $v_1, v_2, v_3 \in \pi$ such that $v_1, v_3 \in \pi'_1$ and $v_2 \in \pi'_2$ and $v_1 <_\pi v_2 <_\pi v_3$. In this case, we can assume without any loss of generality that $\pi[v_1, v_3] \cap \pi'_1 = \{v_1, v_3\}$ (e.g., by choosing v_1, v_3 so that this equality holds). Let x be the node immediately following v_1 in π, and let y be the node immediately preceding v_3 in π. Recall that s is the first node of π and t is the last node of π. It will be useful for the analysis to split π into three subpaths:

  $$\pi = \pi[s, v_1] \circ \pi[x, y] \circ \pi[v_3, t],$$

  where $\circ$ denotes the concatenation operation. Define

  $$\phi_1 := |\{\pi' \in \Pi' \mid \pi[x, y] \cap \pi' \neq \emptyset\}|$$
  $$\phi_2 := |\{\pi' \in \Pi' \mid (\pi[s, v_1] \circ \pi[v_3, t]) \cap \pi' \neq \emptyset\}|.$$

  We claim that $\phi_1 < \phi$, $\phi_2 < \phi$, and $\phi_1 + \phi_2 = \phi$. We will use these facts to establish a recurrence relation for f. By our assumption that $\pi[v_1, v_3] \cap \pi'_1 =$

$\{v_1, v_3\}$, it follows that $\pi[x, y] \cap \pi'_1 = \emptyset$, and so $\phi_1 < \phi$. Likewise, by the No Repeats Property of Proposition 5.7,

$$(\pi[s, v_1] \circ \pi[v_3, t]) \cap \pi'_2 = \emptyset.$$

Therefore, $\phi_2 < \phi$. Finally, observe that more generally, if there exists a path $\pi' \in \Pi'$ such that $\pi' \cap \pi[x, y] \neq \emptyset$ and $\pi' \cap (\pi[s, v_1] \circ \pi[v_3, t]) \neq \emptyset$, then the No Repeats Property of Proposition 5.7 is violated. We conclude that $\phi_1 + \phi_2 = \phi$.

Now $|E_3|$ can be upper bounded by the following inequality:

$$|E_3| \leq |E_3 \cap E(\pi[x, y])| + |E_3 \cap E(\pi[s, v_1] \circ \pi[v_3, t])| + 2.$$

Then using the observations about ϕ_1, ϕ_2, and ϕ in the previous paragraph, we obtain the following recurrence for f:

$$f(\phi) \leq f(\phi_1) + f(\phi_2) + 2 = f(i) + f(\phi - i) + 2,$$

where $0 < i < \phi$.

- **Case 2.** There exists a path $\pi' \in \Pi'$ and $v_1, v_2 \in \pi$ such that $\pi \cap \pi' = \pi[v_1, v_2] \cap V'$. Let x be the node immediately preceding v_1 in π, and let y be the node immediately following v_2 in π. Again, we split π into three subpaths:

$$\pi[s, t] = \pi[s, x] \circ \pi[v_1, v_2] \circ \pi[y, t].$$

Let

$$\phi_1 := |\{\pi' \in \Pi' \mid \pi[v_1, v_2] \cap \pi' \neq \emptyset\}|$$
$$\phi_2 := |\{\pi' \in \Pi' \mid (\pi[s, x] \circ \pi[y, t]) \cap \pi' \neq \emptyset\}|.$$

Our assumption in Case 2 follows that $\phi_1 = 1$ and $\phi_2 = \phi - 1$. Since $|E_3|$ can be upper bounded by the inequality

$$|E_3| \leq |E_3 \cap E(\pi[v_1, v_2])| + |E_3 \cap E(\pi[s, x] \circ \pi[y, t])| + 2,$$

we immediately obtain the recurrence

$$f(\phi) \leq f(\phi_1) + f(\phi_2) + 2 \leq f(1) + f(\phi - 1) + 2.$$

Taking our results from Case 1 and Case 2 together, we obtain the recurrence relation

$$f(\phi) \leq \begin{cases} \max\{f(i) + f(\phi - i) + 2,\ f(1) + f(\phi - 1) + 2\} & \phi > 1 \text{ and } 1 < i < \phi \\ 2 & \phi = 1 \end{cases}.$$

Applying this recurrence at most ϕ times, we find that

$$f(\phi) \leq \phi \cdot f(1) + 2\phi \leq 4\phi.$$

Finally, since $k \leq |\Pi'| \leq n^{1/2}$ and we defined f so that $f(k)$ equals the largest possible value of $|E_3|$, we conclude that

$$|E_3| \leq f(k) \leq f(n^{1/2}) = O(n^{1/2}).$$

Since the edges in $E_3 \subseteq E \setminus E'$ are labeled independently at random, we may apply a Chernoff bound as in Proposition 5.10 to argue that $\chi(E_3) = O(n^{1/4} \log^{1/2}(n))$ with high probability. ∎

We have shown that with high probability, the discrepancy of our edge labeling is $O(n^{1/4}\log^{1/2}(n))$ for E_1, E_2, and E_3, so we conclude that the total discrepancy of π is $O(n^{1/4}\log^{1/2}(n))$. A straightforward extension of this argument implies identical bounds for hereditary discrepancy. We defer this proof to the full version of our paper [14].

6. Planar Graphs. In this section, we will extend our hereditary vertex discrepancy lower bound for unique shortest paths in undirected graphs to the planar graph setting.

THEOREM 6.1. *There exists an n-vertex undirected planar graph with hereditary vertex discrepancy at least* $\Omega\left(\frac{n^{1/4}}{\log^2(n)}\right)$.

To prove this theorem, we will first give an abbreviated presentation of the graph construction in [15] that we used implicitly to obtain the $\Omega(n^{1/4}/\sqrt{\log n})$ hereditary vertex discrepancy lower bound in Theorem 5.1. Then we will describe a simple procedure to make this graph planar and argue that the shortest path structure of this planarized graph remains unchanged.

6.1. Graph Construction of Bodwin and Hoppenworth. Take n to be a large enough positive integer, and take $p = n\log n$. We will describe an n-node weighted undirected graph $G = (V, E, w)$ originally constructed in Bodwin and Hoppenworth [15].

Vertex Set V. We will use $\ell = \Theta\left(\frac{n^{1/2}}{\log n}\right)$ as a positive integer parameter for our construction. The graph G we create will consist of ℓ layers, denoted as $L_1, \ldots, L_\ell$. Each layer will have n/ℓ nodes, arranged from 1 to n/ℓ. Initially, we will assign a tuple label (i, j) to the jth node in the L_i layer. We will interpret the node labeled (i, j) as a point in $\mathbb{R}^2$ with integral coordinates. The vertex set V of graph G is made up of these n nodes distributed across ℓ layers.

Next we will randomize the node labels in V. For each layer L_i, where i ranges from 1 to ℓ, we randomly and uniformly pick a real number in the interval $(0, 1)$ and we call it ψ_i. After that, for each node in layer L_i of the graph G that is currently labeled (i, j), we relabel it as

$$\left(i, j + \sum_{k=1}^{j} \psi_k\right).$$

These new labels for the nodes in V are also treated as points in $\mathbb{R}^2$. We can imagine this process as adding a small epsilon of structured noise to the points corresponding to the nodes in the graph. The purpose of this noise is technical, but serves the purpose of achieving 'symmetry breaking' (see Section 2.4 of [15] for details).

Edge Set E. All edges will be between subsequent layers L_i, L_{i+1} within G. It will be helpful to think of the edges in G as directed from L_i to L_{i+1}, although in actuality G will be undirected. We represent the set of edges in G between layers L_i and L_{i+1} as E_i. For any edge $e = (v_1, v_2) \in E$, the edge e will be associated with the specific vector $\vec{u}_e := v_2 - v_1$. The 2nd coordinate of $\vec{u}_e$ will be labeled as u_e. Hence, for all e found in E, $\vec{u}_e$ is written as $(1, u_e)$.

For each $i \in [1, \ell - 1]$, let

$$C_i := \{(1, \psi_{i+1} + x) : x \in [0, n/\ell^2]\}.$$

We will refer to the vectors in C_i as *edge vectors*. For each $v \in L_i$ and edge vector $\vec{c} \in C_i$, if $v + \vec{c} \in V$, then add edge $(v, v + \vec{c})$ to E_i. After adding these edges to E_i, we will have that

$$C_i = \{\vec{u}_e \mid e \in E_i\}.$$

Finally, for each $e \in E$, if $\vec{u}_e = (1, u_e)$, then we assign edge e the weight $w(e) := u_e^2$. This completes the construction of our graph $G = (V, E, w)$.

PROPOSITION 6.2. *Consider the graph drawing of graph G where the nodes v in V are drawn as points at their associated coordinates in $\mathbb{R}^2$ and the edges (u, v) in E are drawn as straight-line segments from u to v. This graph drawing has $O(n \log^6 n)$ edge crossings.*

Proof. First note that if edges $e_1, e_2 \in E$ cross in our graph drawing of G, then edges e_1 and e_2 are between the same two layers of G (i.e., $e_1, e_2 \in L_i \times L_{i+1}$ for some $i \in [1, \ell - 1]$). Additionally, all edges between L_i and L_{i+1} are from the jth vertex in L_i to the $(j + k)$th vertex in L_{i+1}, where $j \in [1, n/\ell]$ and $k \in [0, \Theta(\log^2 n)]$.

Now fix an edge $(u, v) \in E \cap (L_i \times L_{i+1})$ for some $i \in [1, \ell - 1]$. If an edge $(u', v') \in E \cap (L_i \times L_{i+1})$ crosses (u, v), then $|u - u'| \leq \log^2 n$. Then there are at most $O(\log^2 n)$ nodes incident to edges that cross (u, v) in our drawing. Since each node in G has degree $O(\log^2 n)$, this implies that at most $O(\log^4 n)$ edges cross (u, v) in our drawing. Since $|E| = O(n \log^2 n)$, we conclude that our graph drawing has $O(n \log^6 n)$ edge crossings. ∎

Direction Vectors and Paths Π.. Our next step is to generate a set of unique shortest paths Π. The paths Π are identified by first constructing a set of vectors $D \subseteq \mathbb{R}^2$ called *direction vectors*, which are defined next.

Let $q = \Theta\left(\frac{\ell}{\log n}\right) = \Theta\left(\frac{n^{1/2}}{\log^2 n}\right)$ be an integer. We choose our set of direction vectors D to be

$$D := \left\{ \left(1,\ x + \frac{y}{q}\right) \quad \text{such that} \quad x \in \left[1, \frac{n}{4\ell^2} - 1\right] \text{ and } y \in [0, q] \right\}.$$

Note that adjacent direction vectors in D differ only by $1/q$ in their second coordinate. Each of our paths π in Π will have an associated direction vector $\vec{d} \in D$, and for all $i \in [1, \ell - 1]$, path π will take an edge vector in C_i that is closest to $\vec{d}$ in some sense.

Paths Π.. We first define a set $S \subseteq L_1$ containing half of the nodes in the first layer L_1 of G:

$$S := \left\{ (1, j + \psi_1) \in L_1 \quad \text{such that} \quad j \in \left[1, \frac{n}{2\ell}\right] \right\}.$$

We will define a set of pairs of nodes P so that $P \subseteq S \times L_\ell$. For every node $s \in S$ and direction vector $\vec{d} \in D$, we will identify a pair of endpoints $(s, t) \in S \times L_\ell$ and a corresponding unique shortest path $\pi_{s,t}$ to add to Π.

Let $v_1 \in S$, and let $\vec{d} = (1, d) \in D$. The associated path π has start node v_1. We iteratively grow π, layer-by-layer, as follows. Suppose that currently $\pi = (v_1, \ldots, v_i)$, for $i < \ell$, with each $v_i \in L_i$. To determine the next node $v_{i+1} \in L_{i+1}$, let $E_i^{v_i} \subseteq E_i$ be the edges in E_i incident to v_i, and let

$$e_i := \text{argmin}_{e \in E_i^{v_i}} (|u_e - d|).$$

By definition, e_i is an edge whose first node is v_i; we define $v_{i+1} \in L_{i+1}$ to be the other node in e_i, and we append v_{i+1} to π. After this process terminates, we will have a path $\pi = (v_1, \ldots, v_\ell)$. Denote π as π_{v_1, v_ℓ} and add path π_{v_1, v_ℓ} to Π. Repeating for all $v_1 \in S$ and $\vec{d} \in D$ completes our construction of Π. Note that although we did not prove it, each path $\pi_{s,t} \in \Pi$ is a unique shortest $s \rightsquigarrow t$ path in G by Lemma 2 of Bodwin and Hoppenworth [15].

Lemma 1 of Bodwin and Hoppenworth [15] summarizes the key properties of G, Π that are needed to prove the hereditary vertex discrepancy lower bound for unique shortest paths in undirected graphs in Theorem 5.1. We restate this key lemma in Lemma 3.9 of Section 5.

To obtain a $\tilde{\Omega}(n^{1/4})$ lower bound for hereditary vertex discrepancy of unique shortest paths in *planar* graphs, we need to convert the graph G into a planar graph while ensuring that the unique shortest path structure of the graph remains unchanged.

6.2. Planarization of Graph G. In the previous subsection, we outlined the construction of the graph $G = (V, E, w)$ and set of paths Π from Bodwin and Hoppenworth [15]. This graph has an associated graph drawing with $\tilde{O}(n)$ edge crossings by Proposition 6.2. We will now 'planarize' graph G by embedding it within a larger planar graph G'. We will use the standard strategy of replacing each edge crossing in our graph drawing of G with a new vertex, causing each crossed edge to be subdivided into a path.

Planarization Procedure:.

1. We start with the current non-planar graph $G = (V, E, w)$ with the associated graph drawing described in Proposition 6.2.
2. For every edge crossing in the drawing of G, letting point $p \in \mathbb{R}^2$ be the location of the crossing, draw a vertical line in the plane through p. Add a new node to graph G at every point where this vertical line intersects the drawing of an edge. This step may blow up the number of nodes in the graph by quite a lot, but the resulting graph will be planar and layered.
3. We re-set all edge weights in the graph as follows. For each edge (u, v) in the graph, letting $p_u, p_v \in \mathbb{R}^2$ be the locations of nodes $u, v \in V$ in the drawing, we re-set the weight of edge (u, v) to be the squared Euclidean distance between p_u and p_v, i.e.,
$$w((u, v)) = \|p_u - p_v\|^2.$$
4. Finally, we remove excess nodes added to the graph in step 2. We perform the following operation for each node v of degree 2 in the resulting graph. Let (x, v) and (v, y) be the two edges incident to v. Add edge (x, y) to the graph and assign it weight $w((x, y)) = w((x, v)) + w((v, y))$. Remove node v and edges (x, v) and (v, y) from the graph. Note that the graph will remain planar after this operation.

Denote the planar graph resulting from this procedure as $G' = (V', E', w')$. The following proposition follows immediately from Proposition 6.2 and the planarization procedure.

PROPOSITION 6.3. *Graph G' is planar and has $O(n \log^6 n)$ nodes.*

Unique Shortest Paths in G'. Each edge $e = (u, v) \in E$ in graph G is the preimage of a $u \rightsquigarrow v$ path π_e in graph G' resulting from our planarization procedure. Likewise, each path $\pi \in \Pi$ is the preimage of a path π' in G' obtained by replacing each edge $e \in \pi$ with path π_e. Let the set Π' of paths in G' denote the image of the set of paths Π in G under our planarization procedure. As a final step towards proving Theorem 6.1, we need to argue that the unique shortest path structure of G is unchanged by our planarization procedure.

LEMMA 6.4. *Each path in Π' is the unique shortest path between its endpoints in G'.*

We now verify that graph G' and paths Π' have the unique shortest path property as stated in Lemma 6.4. We will require the following proposition about the construction of graph G' from [15] that we state without proof.

PROPOSITION 6.5 (c.f. Proposition 1 of [15]). *With probability* 1*, for every* $i \in [1, \ell-1]$ *and every direction vector* $\vec{d} = (1, d) \in D$*, there is a unique vector* $(1, c) \in C_i$ *that minimizes* $|c - d|$ *over all choices of* $(1, c) \in C_i$*.*

Additionally, our unique shortest paths argument will make use of the following technical proposition also proven in [15].

PROPOSITION 6.6 (c.f. Proposition 3 of [15]). *Let* $b, x_1, \ldots, x_k \in \mathbb{R}$*. Now consider* $\hat{x}_1, \ldots, \hat{x}_k$ *such that*

- $|\hat{x}_i - b| \leq |x_i - b|$ *for all* $i \in [1, k]$*, and*
- $\sum_{i=1}^k x_i = \sum_{i=1}^k \hat{x}_i$.

Then

$$\sum_{i=1}^{k} x_i^2 \geq \sum_{i=1}^{k} \hat{x}_i^2,$$

with equality only if $|\hat{x}_i - b| = |x_i - b|$ *for all* $i \in [1, k]$*.*

Using Propositions 6.5 and 6.6, we can now prove Lemma 6.4.

Proof of Lemma 6.4. As an immediate step toward proving Lemma 6.4, we will argue that we can make two assumptions about G' without loss of generality.

First, we may assume that G' is layered in the following sense: V' can be partitioned into k layers (for some $k > 0$) such that each path $\pi \in \Pi'$ begins in the first layer, ends in the last layer, and has exactly one node in each layer. Observe that after step 2 of the planarization procedure, graph G' is layered with respect to paths Π' in this sense. Moreover, step 4 of the planarization procedure does not change the structure of the set of paths Π'. Thus we can safely assume G' is layered with respect to paths Π'.

Second, we can assume, without loss of generality, that G' is a directed graph and that all edges in $L_i \times L_{i+1}$ in G' are directed from L_i to L_{i+1}. This assumption can be made using a blackbox reduction that is standard in the area (see Section 4.6 of 6.1 for details).

Fix an $s \leadsto t$ path $\pi' \in \Pi'$ in graph G', and let path $\pi \in \Pi$ in G be the associated preimage of π'. Let $(1, x) \in D$ be the direction vector associated with path π. Note that by Proposition 6.5, for each layer L_i, there is a unique vector $(1, c) \in C_i$ that minimizes $|c - d|$ over all choices of $(1, c) \in C_i$. By our construction of the paths in Π, path π will travel along an edge with edge vector $(1, c)$.

In graph G', there are additional layers between layers L_i and L_{i+1}, due to step 2 of our planarization procedure. If path π traveled along an edge with edge vector $(1, c)$ from L_i to L_{i+1} in G, then in each layer L' in G' between L_i and L_{i+1}, graph G' will take an edge vector (α, c), where $0 < \alpha \leq 1$. Moreover, again by Proposition 6.5, this edge vector (α, c) will be the unique edge vector from layer L' minimizing $|c - d|$.

Let ℓ' be the number of layers in G'. Let $\hat{x}_1, \ldots, \hat{x}_{\ell'-1} \in \mathbb{R}$ be real numbers such that the ith edge of π' has the corresponding vector $(\alpha_i, \hat{x}_i)$ for $i \in [1, \ell' - 1]$ and $0 < \alpha_i \leq 1$. Now consider an arbitrary $s \leadsto t$ path π^* in G, where $\pi^* \neq \pi'$. Since all edges in G are directed from L_i to L_{i+1}, it follows that π^* has $\ell' - 1$ edges. Let $x_1, \ldots, x_{\ell'-1} \in \mathbb{R}$ be real numbers such that the ith edge of π^* has the corresponding vector $(\alpha_i, x_i) \in C_i$ for $i \in [1, \ell - 1]$ and $0 < \alpha_i \leq 1$. Now observe that since π^* and

π' are both $s \rightsquigarrow t$ paths, it follows that

$$\sum_{i=1}^{\ell'-1} \hat{x}_i = \sum_{i=1}^{\ell'-1} x_i.$$

Additionally, by our construction of π', it follows that

$$|\hat{x}_i - x| \leq |x_i - x|$$

for all $i \in [1, \ell - 1]$. In particular, since $\pi^* \neq \pi'$, there must be some $j \in [1, \ell' - 1]$ such that $\hat{x}_j \neq x_j$, and so by Proposition 6.5, $|\hat{x}_j - x| < |x_j - x|$ with probability 1. Then by Proposition 6.6,

$$w(\pi') = \sum_{e \in \pi'} w(e) = \sum_{i=1}^{\ell-1} \hat{x}_i^2 < \sum_{i=1}^{\ell-1} x_i^2 = \sum_{e \in \pi^*} w(e) = w(\pi^*).$$

This implies that the path π' is a unique shortest $s \rightsquigarrow t$ path in G', as desired. ∎

Finishing the Proof.

LEMMA 6.7 (c.f. Lemma 1 of [15]). *There is an infinite family of $\Theta(n \log^6 n)$-node planar undirected weighted graphs $G' = (V', E', w')$ and sets Π' of $|\Pi'| = n \log n$ paths in G' with the following properties:*

- *Each path in Π' is the unique shortest path between its endpoints in G.*
- *Let G be the n-node undirected weighted graph and let Π be the set of $|\Pi| = n \log n$ paths described in Lemma 3.9 when $p = n \log n$. Then Π is an induced path subsystem of Π'.*

Proof. This follows immediately from Proposition 6.3, Lemma 6.4, and the above discussion about the set of paths Π' in G'. ∎

Let $N := \Theta(n \log^6 n)$ be the number of nodes in G'. By Lemma 6.7,

$$\text{herdisc}_v(\Pi') \geq \text{herdisc}_v(\Pi).$$

Likewise, by the proof of Theorem 5.1, $\text{herdisc}_v(\Pi) \geq \Omega(n^{1/4})$. We conclude that

$$\text{herdisc}_v(\Pi') \geq \text{herdisc}_v(\Pi) \geq \Omega\left(\frac{n^{1/4}}{\sqrt{\log n}}\right) = \Omega\left(\frac{N^{1/4}}{\log^2 N}\right).$$

7. Trees and Bipartite Graphs. For graphs with simple topology such as line, tree and bipartite graphs, both of the vertex and edge discrepancy are constant. However, a distinction can be observed on hereditary discrepancy for bipartite graphs. Formally, we have the following results.

LEMMA 7.1. *Let $T = (V, E, w)$ be a undirected tree graph, the hereditary discrepancy of the shortest path system induced by T is $\Theta(1)$.*

Proof. To start with, it is obvious that a lower bound of $\Omega(1)$ on both edge and vertex (hereditary) discrepancy always holds for any family of graphs. We therefore first focus on the $O(1)$ discrepancy upper bound for bipartite graphs ∎

LEMMA 7.2. *Let $G = (V, E, w)$ be a general bipartite graph, then it has $\Theta(1)$ discrepancy, but $\tilde{\Theta}(n^{1/4})$ hereditary discrepancy.*

Proof. To start with, it is obvious that a lower bound of $\Omega(1)$ on both edge and vertex (hereditary) discrepancy always holds for any family of graphs. We therefore first focus on the $O(1)$ discrepancy upper bound for bipartite graphs (including trees).

Analysis of discrepancy.. We start with the vertex discrepancy. For a bipartite graph $G = (L \cup R, E)$, a simple scheme achieves constant vertex discrepancy: assign coloring '+1' to every $v \in L$ and '−1' to every $u \in R$. Observe that every shortest path either has length of 1, or alternates between L and R, thus summing up assigned colors along the shortest path gives +1 vertex discrepancy at most 1. Finally, we apply Observation 3.4 to argue that the edge discrepancy is also $O(1)$.

Analysis of hereditary discrepancy.. We prove this statement by showing that we can reduce the hereditary discrepancy of bipartite graphs to general graphs by the 2-lift construction. Concretely, suppose we are given a path system that is characterized by $G = (V, E, w)$ and matrix A, such that the hereditary discrepancy is at least $f(n)$, and let the set of columns that attains the maximum hereditary discrepancy be Y. We will construct a new n'-vertex graph G' with a new matrix A', in which we have a set of columns Y' induces at least $f(n'/2)$ discrepancy. Such a graph is a valid instance of the family of the bipartite graphs, and an $\Omega(n^{1/4})$ hereditary discrepancy on G would imply an $\Omega(n'^{1/4})$ hereditary discrepancy on G'.

We now describe a detailed algorithm, Algorithm 7.1, for the 2-lift graph construction as follows. In the procedure, we slightly abuse the notation to interchangeably use the set with one element and the element itself, i.e., we use $\{a\}$ to denote a when the context is clear.

Algorithm 7.1 Construction 2-Lift: an algorithm to construct bipartite graph with high hereditary discrepancy.

Require: A consistent path system characterized by a general undirected graph $G = (V, E, w)$ and matrix A with hereditary discrepancy at least $f(n)$;

Ensure: A consistent path system characterized by bipartite graph graph $G' = (V', E', w')$ and matrix A' with hereditary discrepancy at least $f(n'/2)$;

Vertices V': each vertex $v \in V$, make two copies of vertices $v_L, v_R \in V'$.

Edges E': Maintain a "side indicator" $s \in \{L, R\}$, and initialize $s = L$.

for each vertex $v \in V$ with an arbitrary order: **do**

  Add *all* edges $(v_s, u_{\{L,R\} \setminus s})$ such that $u \in N(v)$.

  Delete v from $N(u)$ for all $u \in N(v)$.

  Switch the side indicator, i.e., $s \leftarrow \{L, R\} \setminus s$.

end for

Path system A': for each row a of A, starting from the the first vertex with 1, add 1 to the row of A' to the vertex whose degree is not 0.

Note that for any vertex $v \in V$, only one of (v_L, v_R) is used in the matrix A'. We now argue that A' is a valid collection of path systems. Note that for a single path P in A, we can always follow the vertices with non-zero degree, and connect the edges to a valid path in A'. Furthermore, two paths would conflict with each other only if there exists an edge that "shortcut" an even-sized path, i.e., both $(v_1, v_2, \cdots, v_{2k})$ and (v_1, v_{2k}) are in the path system. However, this would violate the consistency property of A. As such, all the rows in A' can find a valid path in G'.

Let Y be the columns that attains the $f(n)$ discrepancy on G, and we slightly abuse the notation to use $f(n)$ to denote both the indices of the columns in A and the vertex set $A \subseteq V$. Since we have an bijective mapping between the vertices in Y and the vertices we account for in G', we have the hereditary discrepancy to be at least $f(n) = f(n'/2)$, as desired. □

8. Bounds on ℓ_2-discrepancy. Consider a set of p paths on a graph G, the incidence matrix A of p rows and n columns has the (i,j)-th element to be 1 if the corresponding vertex v_j stays on the i-th path, and 0 otherwise. Then we consider the vertex ℓ_2-discrepancy $\mathrm{disc}_2(A)$ and hereditary discrepancy $\mathrm{herdisc}_2(A)$.

Since $\mathrm{disc}_2(A) \leq \mathrm{disc}(A)$ and $\mathrm{herdisc}_2(A) \leq \mathrm{herdisc}(A)$, the upper bounds on discrepancy and hereditary discrepany for path systems remain as upper bounds for ℓ_2 (hereditary) discrepancy, that is, $\tilde{O}(\sqrt{n})$ for general paths (that are not necessary shortest paths and not necessarily consistent), and $\tilde{O}(n^{1/4})$ for shortest paths or consistent paths.

For lower bound on ℓ_2 discrepancy, we first recall the following result in Larsen [46]:

LEMMA 8.1 ([46]). *For an $m \times n$ real matrix A, let $\lambda_1 \geq \lambda_2 \geq \cdots \lambda_n \geq 0$ denote the eigenvalues of $A^\top A$. For all positive integers $k \leq \min\{n,m\}$, we have*

$$\mathrm{herdisc}_2(A) \geq \frac{k}{e}\sqrt{\frac{\lambda_k}{8\pi mn}}.$$

In the proof of the trace bound, Chazelle and Lvov [21] have shown that for $k = \mathrm{tr}^2(M)/(2\,\mathrm{tr}\,M^2)$, we have $\lambda_k \geq \mathrm{tr}(M)/4\min\{n,m\}$. Setting it in Lemma 8.1, we have the same asymptotic trace bound for ℓ_2-hereditrary as in Lemma 3.8.

$$\mathrm{herdisc}_2(A) \geq \frac{\mathrm{tr}^2(M)}{2e\min\{m,n\}\,\mathrm{tr}\,M^2}\sqrt{\frac{\mathrm{tr}(M)}{32\pi\max\{m,n\}}}.$$

Using the same calculation as in the proof of Theorem 5.1, we get the equivalent bound on ℓ_2-hereditary discrepancy for the set Π of shortest paths in Lemma 3.9:

$$\mathrm{herdisc}_{v,2}(\Pi) \geq \Omega\left(\frac{n^{1/4}}{\sqrt{\log(n)}}\right).$$

Notice that the same lower bound argument generates a lower bound of $\Omega(n^{1/6})$ on ℓ_2 discrepancy for the Erdös point-line system.

Again if we drop the consistency property (or uniqueness of shortest paths), the ℓ_2 hereditary discrepancy can be much higher. Consider a graph G with two vertices s and t, together with $2n$ vertices u_i, v_i, $i \in [n]$. We connect s with u_1, v_1 and t with u_n, v_n. In addition, u_i, v_i are both connected to u_{i+1}, v_{i+1}, for $1 \leq i \leq n-1$. Now we can encode the $n \times n$ Hadanard matrix H by n paths from s to t. Each row of H corresponds to a path P_i. If the jth element is 1, we take u_j, otherwise, we take v_j. The incidence matrix A considering only vertices $v_1, v_2, \cdots v_n$ would be precisely $\frac{1}{2}(H+J)$, where J is an $n \times n$ matrix of all 1. It is known that $\|Ax\|_2^2 \geq n(n-1)/4$ [20]. Thus $\mathrm{herdisc}_2(A) = \Omega(\sqrt{n})$.

In summary, all bounds of hereditary discrepancy presented in the paper hold for ℓ_2 hereditary discrepancy for path systems on a graph.

9. Applications to Differential Privacy. In light of our new unique shortest path hereditary discrepancy lower bound result, significant progress can be made towards closing the gap in the error bounds for the problem of Differentially Private All Pairs Shortest Distances (APSD) [59, 22, 36]. Likewise, the problem of Differentially Private All Sets Range Query (ASRQ) [29] now has a tight error bound (up to logarithmic factors). We present the DP-APSD problem formally and show the proof of the new lower bound corresponding to Theorem 1.7. Details on the DP-ASRQ problem are deferred to Appendix D.

9.1. All Pairs Shortest Distances. Given a weighted undirected graph $G = (V, E, w)$ of size n, the private mechanism is supposed to output an n by n matrix $D' = \mathcal{M}(G)$ of approximate all pairs shortest paths distances in G, and the privacy guarantee is imposed on two sets of edge weights that are considered 'neighboring', i.e., with ℓ_1 difference at most 1. Our goal is to minimize the maximum additive error of any entry in the APSD matrix, i.e., the ℓ_∞ distance of $D' - D$ where D is the true APSD matrix. This line of work was initiated by [59], where an algorithm was proposed with $O(n)$ additive error. Recently, concurrent works [22, 36] breaks the linear barrier by presenting an upper bound of $\widetilde{O}(n^{1/2})$. Meanwhile, the only known lower bound is $\Omega(n^{1/6})$, due to [22], using a hereditary discrepancy lower bound based on the point-line system of [21]. With our improved hereditary discrepancy lower bound, we are able to show an $\Omega(n^{1/4}/\sqrt{\log n})$ lower bound on the additive error of the DP-APSD problem.

COROLLARY 9.1. *Given an n-node undirected graph, for any $\beta \in (0, 1)$ and $\varepsilon, \delta > 0$, no (ε, δ)-DP algorithm for APSD has additive error of $o(n^{1/4}/\sqrt{\log n})$ with probability $1 - \beta$.*

The connection between the APSD problem and the shortest paths hereditary discrepancy lower bound was shown in [22], which implies that simply plugging in the new exponent gives the result above. For the sake of completeness, we give the necessary definition to formally define the DP-APSD problem, and show the main arguments towards proving Corollary 9.1.

DEFINITION 9.2 (Neighboring weights [59]). *For a graph $G = (V, E)$, let $w, w' : E \to \mathbb{R}^{\geq 0}$ be two weight functions that map any $e \in E$ to a non-negative real number, we say w, w' are neighboring, denoted as $w \sim w'$ if $\sum_{e \in E} |w(e) - w'(e)| \leq 1$.*

DEFINITION 9.3 (Differentially Private APSD [59]). *Let $w, w' : E \to R^{\geq 0}$ be weight functions, and $\mathcal{A}$ be an algorithm taking a graph $G = (V, E)$ and w as input. The algorithm A is (ε, δ)-differentially private on G if for any neighboring weights $w \sim w'$ (See Definition 9.2) and all sets of possible output $\mathcal{C}$, we have: $\Pr[\mathcal{A}(G, w) \in \mathcal{C}] \leq e^\varepsilon \cdot \Pr[\mathcal{A}(G, w') \in \mathcal{C}] + \delta$.*

We say the private mechanism $\mathcal{A}$ is α-accurate if the ℓ_∞ norm of $|\mathcal{A}(G, w) - f(G, w)|$ is at most α, where f indicates the function returning the ground truth shortest distances.

Proof of Corollary 9.1. First, suppose $\mathbf{A} \in \mathbb{R}^{\binom{n}{2} \times n}$ is the shortest path vertex incidence matrix on the graph G. Previous work [22] has shown that the linear query problem on $\mathbf{A}$ can be reduced to the DP-APSD problem, formally stated as follows.

LEMMA 9.4 (Lemma 4.1 in [22]). *Let (V, Π) be a shortest path system with incidence matrix $\mathbf{A}$, if there exists an (ε, δ) DP algorithm that is α-accurate for the APSD problem with probability $1 - \beta$ on a graph of size $2|V|$, then there exists an (ε, δ) DP algorithm that is α-accurate for the $\mathbf{A}$-linear query problem with probability $1 - \beta$.*

Now all we need to show is that the $\mathbf{A}$-linear query problem has a lower bound of $\Omega(n^{1/4}/\sqrt{\log(n)})$. We note the following result by [54].

LEMMA 9.5. *For any $\beta \in (0, 1)$, there exists ε, δ such that for any $\mathbf{A}$, no (ε, δ)-DP algorithm is* $\mathrm{herdisc}(\mathbf{A})/2$*-accurate for the $\mathbf{A}$-query problem with probability $1 - \beta$.*

Combining Lemma 9.4 and 9.5, we find that additive error needed for the DP-APSD problem is at least the hereditary discrepancy of its vertex incidence matrix, implying Corollary 9.1. The lower bound for ASRQ also follows using the same argument (see Appendix D). ∎

10. Conclusion and Open Problems. This paper reported new bounds on the hereditary discrepancy of set systems of unique shortest paths in graphs. We leave several open questions:

1. An open problem is to improve our edge discrepancy upper bound in directed graphs. Standard techniques in discrepancy theory imply an upper bound of $\min\{O(m^{1/4}), \widetilde{O}(D^{1/2})\}$ for this problem, leaving a gap with our $\Omega(n^{1/4}/\sqrt{\log n})$ lower bound when $m = \omega(n)$. Unfortunately, we were not able to extend our low-discrepancy edge and vertex coloring arguments for undirected graphs to the directed setting, due to the pathological example in Figure 4.
2. Using our discrepancy lower bound, we gave an improved lower bound of $\widetilde{\Omega}(n^{1/4})$ on answering all pair shortest distance problem under the constraints of $(1, 1/n)$-differential privacy. In contrast, the best known upper bound remains $\widetilde{O}(n^{1/2})$ for the same problem. Closing this gap remains an interesting open question.

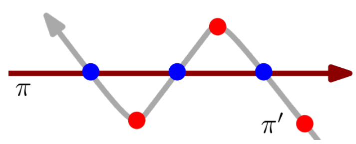

FIG. 4. *An example in directed graphs that demonstrates how coloring unique shortest paths with alternating colors can fail to imply low discrepancy.*

REFERENCES

[1] I. Abraham, A. Fiat, A. V. Goldberg, and R. F. Werneck, *Highway dimension, shortest paths, and provably efficient algorithms*, in Proceedings of the twenty-first annual ACM-SIAM symposium on Discrete Algorithms, SIAM, 2010, pp. 782–793.
[2] S. Akmal and N. Wein, *A local-to-global theorem for congested shortest paths*, in 31st Annual European Symposium on Algorithms (ESA 2023), 2023, pp. 8:1–8:17.
[3] J. R. Alexander, J. Beck, and W. W. Chen, *Geometric discrepancy theory and uniform distribution*, Handbook of discrete and computational geometry, (2004), p. 279.
[4] N. Alon, *On the edge-expansion of graphs*, Combinatorics, Probability and Computing, 6 (1997), pp. 145–152.
[5] N. Alon and F. R. Chung, *Explicit construction of linear sized tolerant networks*, Discrete Mathematics, 72 (1988), pp. 15–19.
[6] S. A. Amiri and J. Wargalla, *Disjoint shortest paths with congestion on DAGs*, arXiv preprint arXiv:2008.08368, (2020).
[7] J. Balogh, B. Csaba, Y. Jing, and A. Pluhár, *On the discrepancies of graphs*, Electronic Journal of Combinatorics, 27 (2020).
[8] L. Balzotti, *Non-crossing shortest paths are covered with exactly four forests*, arXiv preprint arXiv:2210.13036, (2022).
[9] H. Bast, D. Delling, A. Goldberg, M. Müller-Hannemann, T. Pajor, P. Sanders, D. Wagner, and R. F. Werneck, *Route planning in transportation networks*, Algorithm engineering: Selected results and surveys, (2016), pp. 19–80.
[10] J. Beck, *On size Ramsey number of paths, trees, and circuits. i*, Journal of Graph Theory, 7 (1983), pp. 115–129.
[11] I. Ben-Eliezer, M. Krivelevich, and B. Sudakov, *The size Ramsey number of a directed path*, Journal of Combinatorial Theory, Series B, 102 (2012), pp. 743–755.
[12] J. Blum, Y. Disser, A. E. Feldmann, S. Gupta, and A. Zych-Pawlewicz, *On sparse hitting sets: From fair vertex cover to highway dimension*, in 17th International Symposium on

Parameterized and Exact Computation (IPEC 2022), H. Dell and J. Nederlof, eds., vol. 249 of Leibniz International Proceedings in Informatics (LIPIcs), Dagstuhl, Germany, 2022, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, pp. 5:1–5:23.

[13] G. Bodwin, *On the structure of unique shortest paths in graphs*, in Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SIAM, 2019, pp. 2071–2089.

[14] G. Bodwin, C. Deng, J. Gao, G. Hoppenworth, J. Upadhyay, and C. Wang, *The discrepancy of shortest paths*, arXiv preprint arXiv:2401.15781, (2024).

[15] G. Bodwin and G. Hoppenworth, *Folklore sampling is optimal for exact hopsets: Confirming the $\sqrt{n}$ barrier*, in Proceedings of the 64th IEEE Symposium on Foundations of Computer Science (FOCS) 2023, 2023, pp. 701–720.

[16] B. Bollobás and A. D. Scott, *Discrepancy in graphs and hypergraphs*, More Sets, Graphs and Numbers: A Salute to Vera Sós and András Hajnal, (2006), pp. 33–56.

[17] H.-C. Chang, J. Conroy, H. Le, L. Milenkovic, and S. Solomon, *Covering planar metrics (and beyond): O(1) trees suffice*, in 2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS), 2023, pp. 2231–2261.

[18] H.-C. Chang, P. Gawrychowski, S. Mozes, and O. Weimann, *Near-optimal distance emulator for planar graphs*, in 26th Annual European Symposium on Algorithms (ESA 2018), 2018, pp. 16:1–16:17.

[19] H.-C. Chang, R. Krauthgamer, and Z. Tan, *Near-linear ε-emulators for planar graphs*, arXiv preprint arXiv:2206.10681, (2022).

[20] B. Chazelle, *The Discrepancy Method: Randomness and Complexity*, Cambridge University Press, 2000.

[21] B. Chazelle and A. Lvov, *A trace bound for the hereditary discrepancy*, in Proceedings of the sixteenth annual symposium on Computational geometry, SCG '00, New York, NY, USA, May 2000, Association for Computing Machinery, pp. 64–69.

[22] h. Chen, B. Ghazi, R. Kumar, P. Manurangsi, S. Narayanan, J. Nelson, and Y. Xu, *Differentially private all-pairs shortest path distances: Improved algorithms and lower bounds*, in Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2023, pp. 5040–5067.

[23] M. Chudnovsky, D. Cizma, and N. Linial, *The structure of metrizable graphs*, Discrete & Computational Geometry, (2024), pp. 1–20.

[24] F. R. K. Chung, R. L. Graham, and R. M. Wilson, *Quasi-random graphs*, Combinatorica, 9 (1989), pp. 345–362.

[25] D. Cizma and N. Linial, *Geodesic geometry on graphs*, Discrete & Computational Geometry, 68 (2022), pp. 298–347.

[26] D. Cizma and N. Linial, *Irreducible nonmetrizable path systems in graphs*, Journal of Graph Theory, 102 (2023), pp. 5–14.

[27] D. Coppersmith and M. Elkin, *Sparse sourcewise and pairwise distance preservers*, SIAM Journal on Discrete Mathematics, 20 (2006), pp. 463–501.

[28] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to algorithms*, MIT press, 2022.

[29] C. Deng, J. Gao, J. Upadhyay, and C. Wang, *Differentially private range query on shortest paths*, in Proceedings of the 18th Algorithms and Data Structures Symposium (WADS 2023), Springer, 2023, pp. 340–370.

[30] A. Dudek and P. Prałat, *An alternative proof of the linearity of the size-ramsey number of paths*, Combinatorics, Probability and Computing, 24 (2015), pp. 551–555.

[31] C. Dwork, F. McSherry, K. Nissim, and A. Smith, *Calibrating noise to sensitivity in private data analysis*, in Proceedings of the Third Conference on Theory of Cryptography, Springer, 2006, pp. 265–284.

[32] C. Dwork and A. Roth, *The algorithmic foundations of differential privacy*, Foundations and Trends® in Theoretical Computer Science, 9 (2014), pp. 211–407.

[33] M. Eliáš, M. Kapralov, J. Kulkarni, and Y. T. Lee, *Differentially private release of synthetic graphs*, in Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SIAM, 2020, pp. 560–578.

[34] P. Erdős, *On combinatorial questions connected with a theorem of Ramsey and van der Waerden*, Mat. Lapok, 14 (1963), pp. 29–37.

[35] P. Erdős and J. Spencer, *Imbalances in k-colorations*, Networks, 1 (1972), pp. 379–385.

[36] C. Fan, P. Li, and X. Li, *Breaking the linear error barrier in differentially private graph distance release*, arXiv preprint arXiv:2204.14247, (2022).

[37] H. Fichtenberger, M. Henzinger, and J. Upadhyay, *Constant matters: Fine-grained error bound on differentially private continual observation*, in International Conference on Machine Learning, PMLR, 2023, pp. 10072–10092.

[38] A. Freschi, J. Hyde, J. Lada, and A. Treglown, *A note on color-bias Hamilton cycles in dense graphs*, SIAM Journal on Discrete Mathematics, 35 (2021), pp. 970–975.
[39] L. Gishboliner, M. Krivelevich, and P. Michaeli, *Discrepancies of spanning trees and Hamilton cycles*, Journal of Combinatorial Theory, Series B, 154 (2022), pp. 262–291.
[40] S. Har-Peled, *A simple proof of the existence of a planar separator*, arXiv preprint arXiv:1105.0103, (2011).
[41] M. Henzinger, J. Upadhyay, and S. Upadhyay, *Almost tight error bounds on differentially private continual counting*, in Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), SIAM, 2023, pp. 5003–5039.
[42] M. Henzinger, J. Upadhyay, and S. Upadhyay, *A unifying framework for differentially private sums under continual observation*, in Symposium on Disrete Algorithms, 2024.
[43] M. R. Henzinger, P. Klein, S. Rao, and S. Subramanian, *Faster shortest-path algorithms for planar graphs*, Journal of Computer and System Sciences, 55 (1997), pp. 3–23.
[44] S. Kogan and M. Parter, *New diameter-reducing shortcuts and directed hopsets: Breaking the barrier*, in Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), SIAM, 2022, pp. 1326–1341.
[45] S. Kwapień and A. Pełczyński, *The main triangle projection in matrix spaces and its applications*, Studia Mathematica, 34 (1970), pp. 43–67.
[46] K. G. Larsen, *Constructive discrepancy minimization with hereditary L2 guarantees*, in 36th International Symposium on Theoretical Aspects of Computer Science (STACS 2019), R. Niedermeier and C. Paul, eds., vol. 126 of Leibniz International Proceedings in Informatics (LIPIcs), Dagstuhl, Germany, 2019, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, pp. 48:1–48:13.
[47] T. Lee, A. Shraibman, and R. Špalek, *A direct product theorem for discrepancy*, in 2008 23rd Annual IEEE Conference on Computational Complexity, IEEE, 2008, pp. 71–80.
[48] S. Letzter, *Path Ramsey number for random graphs*, Combinatorics, Probability and Computing, 25 (2016), pp. 612–622.
[49] J. Liu, J. Upadhyay, and Z. Zou, *Optimal bounds on private graph approximation*, in Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), SIAM, 2024, pp. 1019–1049.
[50] R. Mathias, *The Hadamard operator norm of a circulant and applications*, SIAM journal on matrix analysis and applications, 14 (1993), pp. 1152–1167.
[51] J. Matoušek, *Tight upper bounds for the discrepancy of half-spaces*, Discrete Comput. Geom., 13 (1995), pp. 593–601.
[52] J. Matoušek, *Geometric Discrepancy: An Illustrated Guide*, Springer Science & Business Media, May 1999.
[53] J. Matoušek, A. Nikolov, and K. Talwar, *Factorization norms and hereditary discrepancy*, International Mathematics Research Notices, 2020 (2020), pp. 751–780.
[54] S. Muthukrishnan and A. Nikolov, *Optimal private halfspace counting via discrepancy*, in Proceedings of the forty-fourth annual ACM symposium on Theory of computing, 2012, pp. 1285–1292.
[55] J. Pach and P. K. Agarwal, *Combinatorial Geometry*, John Wiley & Sons, Oct. 2011.
[56] J. Pach and M. Sharir, *On the number of incidences between points and curves*, Comb. Probab. Comput., 7 (1998), pp. 121–127.
[57] E. Räty, B. Sudakov, and I. Tomon, *Positive discrepancy, MaxCut, and eigenvalues of graphs*, arXiv preprint arXiv:2311.02070, (2023).
[58] D. Reimer, *The Ramsey size number of dipaths*, Discrete mathematics, 257 (2002), pp. 173–175.
[59] A. Sealfon, *Shortest paths and distances with differential privacy*, in Proceedings of the 35th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems, 2016, pp. 29–41.
[60] L. A. Székely, *Crossing numbers and hard Erdős problems in discrete geometry*, Comb. Probab. Comput., 6 (1997), pp. 353–358.
[61] E. Szemerédi and W. T. Trotter, Jr, *Extremal problems in discrete geometry*, Combinatorica, 3 (1983), pp. 381–392.
[62] C. D. Toth, J. O'Rourke, and J. E. Goodman, *Handbook of discrete and computational geometry*, CRC Press, 2017.

Appendix A. Discrepancy Bounds for Paths Without Consistency. In a graph when we consider simple paths without the consistency requirement, there is a strong lower bound on both vertex and edge discrepancy. In particular, we have the following theorem.

THEOREM A.1. *There is a planar graph $G = (V, E)$ such that the following is true for any coloring $f : V \mapsto \{-1, 1\}$ of vertices V:*

1. *There is a family Π of simple paths with $|\Pi| = O(\exp(n))$ and vertex discrepancy of $\Omega(n)$.*
2. *There is a family Π of simple paths with $|\Pi| = O(n)$ and vertex discrepancy of $\Omega(\sqrt{n})$.*

The same claim holds true for edge discrepancy as well.

Proof. The first claim follows from Proposition 1.6 in [7], which says that the edge discrepancy of paths on a $k \times \ell$ grid graph is at least $\Omega(k\ell)$. To prove vertex discrepancy, we make two additional remarks about this construction. First, for our purpose, it is sufficient to consider only an $n \times 2$ grid graph G. The set of paths in the construction of [7] consists of all paths that start from the top left corner and bottom left corner going to the right and possibly taking a subset of the vertical edges in the grid graph. The number of paths is $O(2^n)$. Second, for vertex discrepancy, we define a companion graph G'. Specifically, for each grid edge e in G, we place a vertex v of G' on e. We connect two vertices in G' if and only if the corresponding edges in G share a common vertex. The graph G' is still planar. In additional, a path P in G maps to a corresponding path P' in G' where vertices on P' follow the same order of the corresponding edges on P. See Figure 5 for an example. Therefore the edge discrepancy in G and the vertex discrepancy of G' are the same.

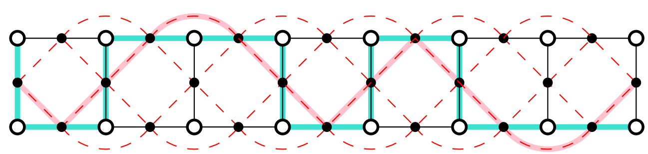

FIG. 5. *A $n \times 2$ grid graph (with vertices shown in hollow and edges in black) G with one path (in blue) starting from the top left corner go the right. The solid vertices and edges in dashed red define the companion graph G'. The corresponding path in G' is shown in pink.*

For the second claim, we take an $n \times n$ Hadamard matrix H with n as power of 2. The elements in H are $+1$ or -1. All the rows are pairwise orthogonal. For example, the Hadamard matrix with $n = 8$ is

$$H_8 = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & -1 & 1 & -1 & 1 & -1 & 1 & -1 \\ 1 & 1 & -1 & -1 & 1 & 1 & -1 & -1 \\ 1 & -1 & -1 & 1 & 1 & -1 & -1 & 1 \\ 1 & 1 & 1 & 1 & -1 & -1 & -1 & -1 \\ 1 & -1 & 1 & -1 & -1 & 1 & -1 & 1 \\ 1 & 1 & -1 & -1 & -1 & -1 & 1 & 1 \\ 1 & -1 & -1 & 1 & -1 & 1 & 1 & -1 \end{pmatrix}$$

It is known [20] that the matrix $A = \{a_{ij}\} = \frac{1}{2}(H + J)$ with J as an $n \times n$ matrix of all 1 has discrepancy at least $\Omega(\sqrt{n})$.

Now we try to embed the matrix A by paths on a $2 \times n$ grid graph G with a grid of two rows and n columns. Denote by e_j as the jth vertical edge in G, $1 \leq j \leq n$. For the ith row of A, we define set $X_i = \{e_j : a_{ij} = 1\}$. We then define two paths $P(X_i)$, $P'(X_i)$ on G.

- Path $P(X_i)$ starts from the top left corner of G going to the right and the vertical edges visited by $P(X_i)$ are precisely X_i.

- Path $P'(X_i)$ starts from the bottom left corner of G going to the right and, similar to $P(X_i)$, the vertical edges visited by $P(X_i)$ are precisely X_i.

Note that $P(X_i)$ and $P'(X_i)$ each contains edges X_i (see Figure 6 for an illustration). Also $P(X_i)$ and $P'(X_i)$ do not share any horizontal edges, and collectively cover all horizontal edges in G. In addition, we define two paths P and P' with P starting from the top left corner and visiting all the top horizontal edges and P' starting from the bottom left corner and visiting all bottom horizontal edges. We have $2n+2$ paths in total – each row i of the Hadamard matrix contributes 2 paths, and we additionally use P and P'.

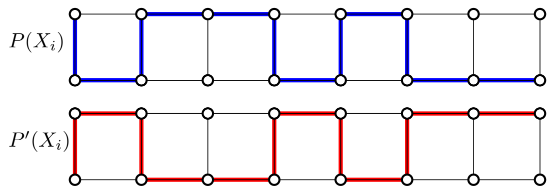

FIG. 6. *A $n \times 2$ grid graph G with two paths $P(X_i)$ and $P'(X_i)$. The ith row of matrix A, i.e., $A_i = (\frac{1}{2}(H+J))_i$, corresponds to $X_i = (1,1,0,1,1,1,0,0)$.*

For any $\{\pm 1\}$ coloring f on edges in the grid graph G, define D as the maximum absolute value of the sum of the colors of edges of each path, among all the $2n+2$ paths. We now argue that $D = \Omega(\sqrt{n})$.

First we define an n-dimensional vector $\chi \in \{+1,-1\}^n$ with $\chi_j = f(e_j)$, i.e., the color of the jth vertical edge in G. Since the discrepancy of matrix A is $\Omega(\sqrt{n})$, there must be one row vector a_i in A such that $a_i \cdot \chi \geq c\sqrt{n}$ for some constant c. In other words, the sum of the colors of the edges in X_i is $x = \sum_{e_j \in X_i} f(e_j)$ with $|x| \geq c\sqrt{n}$. Without loss of generality, we assume $x > 0$, which in turn implies $x \geq c\sqrt{n}$.

Now we define

$$z_1 := \sum_{e \in P(X_i) \setminus X_i} f(e) \quad \text{and} \quad z_2 := \sum_{e \in P'(X_i) \setminus X_i} f(e).$$

If $\max\{z_1, z_2\} \geq 0$, then we are done as the total sum of colors of either $P(X)$ or $P'(X)$ is at least $x \geq c\sqrt{n}$. Otherwise, suppose we have $z_1 < 0$ and $z_2 < 0$. We now have from path $P(X_i)$, $D \geq x + z_1$, and from path $P'(X_i)$, $D \geq x + z_2$. Further, consider paths P and P', the total sum of colors of all the horizontal edges is $z_1 + z_2$. Since both z_1 and z_2 are negative, there should be

$$2D \geq \sum_{e \in P} f(e) + \sum_{e \in P'} f(e) = -z_1 - z_2.$$

Summing up all three inequalities, we have $4D \geq 2x$. Thus $D \geq x/2 \geq c\sqrt{n}/2$. This finishes the proof for edge discrepancy, and the vertex discrepancy bound can be obtained using the same trick as in the proof of claim 1. ∎

The above theorem provides lower bounds for the discrepancy. Since hereditary discrepancy is at least as high as discrepancy, the lower bounds also hold for hereditary discrepancy.

Note that the paths used in the above theorem are 2-approximate shortest paths. The shortest path from top left corner to the top right corner of the $n \times 2$ grid is of length $n-1$ while all paths used are of length at most $2n-2$. The shortest path from top left corner to bottom right corner is of length n and all paths used in the construction are of length at most $2n-1$. Therefore when we relax from shortest paths to 2-approximate shortest paths the discrepancy bounds substantially go up. If we replace each horizontal edge by a chain of $\lceil 1/\varepsilon \rceil$ vertices, we can make these paths to be $(1+\varepsilon)$-approximate shortest paths for any $\varepsilon > 0$.

Grid graphs are a special family of planar graphs. Actually for grid graphs we can say a bit more on discrepancy of shortest paths. If we take *shortest paths* on an unweighted grid graph (without even requiring consistency property and there could be exponentially many shortest path between two vertices), the discrepancy is $O(1)$. Specifically, for vertex discrepancy on a $k \times \ell$ grid graph of left bottom corner at the origin and the top right corner at coordinate $(k-1, \ell-1)$, if we give a color of $+1$ to all vertices of coordinate (x, y) with even $x + y$ and a color of -1 to all other vertices, any shortest path visits a sequence of vertices with sum of coordinates alternating between even and odd values and thus has a total color of $O(1)$. For edge discrepancy, for all horizontal edges we give color $+1$ (-1) if the left endpoint is at an even (odd) x-coordinate and the right endpoint is at an odd (even) x-coordinate. We do the same for vertical edges. Again any shortest path has a 'staircase' shape and a total coloring of $O(1)$.

Appendix B. Relation Between Discrepancy and Hereditary Discrepancy in General Graphs.

In Theorem 5.1, we showed a construction of a path weighted graph G whose system of unique shortest paths Π satisfies $\text{herdisc}_v(\Pi) \geq \Omega(n^{1/4}/\sqrt{\log(n)})$. Here, we will observe that this result extends to discrepancy:

Theorem B.1. *There are examples of n-node undirected weighted graphs G with a unique shortest path between each pair of nodes in which this system of shortest paths Π has*

$$\text{disc}_v(\Pi) \geq \Omega\left(\frac{n^{1/4}}{\sqrt{\log(n)}}\right).$$

Proof. Let G be the graph from Theorem 5.1, and let Π be its system of unique shortest paths. By definition of hereditary discrepancy, there exists an induced path subsystem $\Pi' \subseteq \Pi$ with

$$\text{disc}_v(\Pi') \geq \Omega\left(\frac{n^{1/4}}{\sqrt{\log(n)}}\right).$$

Recall that, by induced subsystem, we mean that we may view the paths of Π as abstract sequences of nodes, and then Π' is obtained from Π by deleting zero or more nodes and deleting all occurrences of those nodes from the middle of paths. Thus the paths in Π' are not still paths in Π, but they are paths in a different graph G' on $n' \leq n$ nodes. It thus suffices to argue that all paths in Π' are

It thus suffices to argue that there is a graph G' on $n' \leq n$ nodes in which all paths in Π' are unique shortest paths. Indeed, this is well known, and is shown e.g. in [13] (c.f. Lemma 2.4.4 and 2.4.11). To sketch the proof: suppose that a node v is deleted from the initial system Π. Consider each path $\pi \in \Pi$ that contains v as an internal node, i.e., it has the form

$$\pi = (\ldots, u, v, x, \ldots).$$

When v is removed, the path now contains the nodes u, x consecutively, and so we must add (u, x) as a new edge to G' so that π is a path in G'. We judiciously set the edge weight to be $w(u,x) := w(u,v) + w(v,x)$. The weighted length of the path π does not change, and yet the distances of G' majorize those of G, which implies that π is still a unique shortest path in G'. Inducting this analysis over each deleted vertex leads to the desired claim. ∎

Appendix C. Application in Matrix Analysis. Hereditary discrepancy is intrinsically related to *factorization norm*, which has found applications in many ares of computer science, including but not limited to quantum channel capacity, communication complexity, etc. For any complex matrix $A \in \mathbb{C}^{m\times n}$, its factorization norm, denoted by $\gamma_2(A)$, is defined as the following optimization problem:

$$\gamma_2(A) = \min\left\{\|L\|_{2\to\infty}\|R\|_{1\to 2} : A = LR\right\}. \tag{C.1}$$

One can write (C.1) in a form of a semi-definite program (see Lee et al. [47]) and also show that the Slater point exists. In particular, the primal and dual program coincides. An interesting question in matrix analysis is to estimate the factorization norm of different class of matrices. In a series of work, various authors have computed tight bounds on the factorization norm of certain class of matrices:

- If A is a unitary matrix, then $\gamma_2(A) = 1$.
- If $A \in \mathbb{R}^{n\times n}$ is a positive semidefinite matrix with entries A_{ij}, then
$$\gamma_2(A) = \max_{1\le i\le n} A_{ii}.$$
- Mathias [50]: A satisfies the property that $\sqrt{A^\top A} \bullet \mathbb{I} = \sqrt{AA^\top} \bullet \mathbb{I} = \frac{\mathrm{tr}(A)}{n}\mathbb{I}$, then
$$\gamma_2(A) = \frac{\mathsf{Tr}(A^\top A)}{n}.$$
In particular, if $A \in \{0,1\}^n$ is a lower-triangular one matrix, then $\gamma_2(A) = \Theta(\log n)$ [37, 41, 45].
- If $A \in \mathbb{R}^{n\times n}$ is a lower-triangular Toeplitz matrix with entries decreasing either polynomially or exponentially, then $\gamma_2(A) = \Theta(1)$ [42].

Our tight bound on hereditary discrepancy for consistent path on graphs allows us to give tight bound on the factorization norm for the corresponding incidence matrix. In particular, we use the following result:

LEMMA C.1 (Matoušek et al. [53]). *For any real $m \times n$ matrix $A \in \mathbb{R}^{m\times n}$, there exists absolute constants $0 < c < C$ such that*

$$c\frac{\gamma_2(A)}{\sqrt{\log(m)}} \le \mathrm{herdisc}(A) \le C\gamma_2(A)\log(m).$$

Combining Lemma C.1 with our results, we have the following corollary:

COROLLARY C.2. *Let A_G be the incident matrix for unique shortest path system on an n vertices graph G. Then if G is bipartite, planar, or any general graph, then $\gamma_2(A_G) = \widetilde{\Theta}(n^{1/4})$.*

Appendix D. Differentially Private All Sets Range Queries. Here we introduce the DP-ASRQ problem and show its connection to the DP-APSD problem.

Given an undirected graph G, the problem of All Sets Range Queries (ASRQ) considers each edge associated with a certain attribute, and the range is the set of edges along a shortest path. Two type of queries are considered here: the *bottleneck* query returns the largest/smallest attribute in a range; while the *counting* query returns the summation of all attributes.

At a schematic level, the ASRQ problem with counting queries is very similar to the APSD problem: the graph topology is public and the edge attributes are considered private. Only that the shortest path structure is dictated by the edge weights to be protected in the APSD problem, however, irrelevant to the edge attributes in the ASRQ problem. This subtle difference has consequential caveat in the algorithm design: the graph topology can be used, for example, to construct an exact hopset first then apply perturbations to the edge attributes in the ASRQ problem; nevertheless, approach of this kind violates protecting the sensitive information in the APSD problem, since adversarial inference can be made on edge weights when the graph topology information is used. This observation also notes that the APSD problem is strictly harder than the ASRQ problem: recall that the best additive error upper bound of the APSD problem is still $\widetilde{O}(n^{1/2})$, while the other has $\widetilde{O}(n^{1/4})$ [29]. Therefore, plugging in our new hereditary lower bound essentially closes the gap for the ASRQ problem.

Corollary D.1 (Formal version of Theorem 1.7). *Given an n-node undirected graph, for any $\beta \in (0, 1)$ and any $\varepsilon, \delta > 0$, no (ε, δ)-DP algorithm for ASRQ has additive error of $o(n^{1/4})$ with probability $1 - \beta$.*

Definition D.2 (Differentially Private Range Queries). *Let $(\mathcal{R} = (X, \mathcal{S}), f)$ be a system of range queries and $w, w' : X \to \mathbb{R}^{\geq 0}$ be neighboring attribute functions. Furthermore, let $\mathcal{A}$ be an algorithm that takes $(\mathcal{R}, f, w)$ as input. Then $\mathcal{A}$ is (ε, δ)-differentially private on G if, for all pairs of neighboring attribute functions w, w' and all sets of possible outputs $\mathcal{C}$, we have $\Pr[\mathcal{A}(\mathcal{R}, f, w) \in \mathcal{C}] \leq e^{\varepsilon} \cdot \Pr[\mathcal{A}(\mathcal{R}, f, w') \in \mathcal{C}] + \delta$. If $\delta = 0$, we say $\mathcal{A}$ is ε-differentially private on G.*

To complete this section, we give the formal definition of the ASRQ problem above. The definition of neighboring attributes follows Definition 9.2. The lower bound proof of Corollary D.1 simply imitates the APSD problem, because the reduction from the linear query problem still holds despite the difference between two problems. The proof is omitted to avoid redundancy.