

Embeddings of Planar Quasimetrics into Directed ℓ_1 and Polylogarithmic Approximation for Directed Sparsest-Cut

Ken-ichi Kawarabayashi
National Institute of Informatics
Tokyo, Japan
Email: k_keniti@nii.ac.jp

Anastasios Sidiropoulos
Dept. of Computer Science
University of Illinois at Chicago
Chicago, USA
Email: sidiropo@uic.edu

Abstract—The multi-commodity flow-cut gap is a fundamental parameter that affects the performance of several divide & conquer algorithms, and has been extensively studied for various classes of undirected graphs. It has been shown by Linial, London and Rabinovich [20] and by Aumann and Rabani [5] that for general n -vertex graphs it is bounded by $O(\log n)$ and the Gupta-Newman-Rabinovich-Sinclair conjecture [13] asserts that it is $O(1)$ for any family of graphs that excludes some fixed minor.

We show that the multicommodity flow-cut gap on directed planar graphs is $O(\log^3 n)$. This is the first *sub-polynomial* bound for any family of directed graphs of super-constant treewidth. We remark that for general directed graphs, it has been shown by Chuzhoy and Khanna [11] that the gap is $\Omega(n^{1/7})$, even for directed acyclic graphs.

As a direct consequence of our result, we also obtain the first polynomial-time polylogarithmic-approximation algorithms for the Directed Non-Bipartite Sparsest-Cut, and the Directed Multicut problems for directed planar graphs, which extends the long-standing result for undirected planar graphs by Rao [22] (with a slightly weaker bound).

At the heart of our result we investigate low-distortion quasimetric embeddings into directed ℓ_1 . More precisely, we construct $O(\log^2 n)$ -Lipschitz quasipartitions for the shortest-path quasimetric spaces of planar digraphs, which generalize the notion of Lipschitz partitions from the theory of metric embeddings. This construction combines ideas from the theory of bi-Lipschitz embeddings, with tools from data structures on directed planar graphs.

Keywords—metric embeddings; approximation algorithms; sparsest cut; planar graphs; quasimetric spaces;

I. INTRODUCTION

The multi-commodity flow-cut gap has been proven instrumental in the design of routing and divide & conquer algorithms in graphs. Bounds on this parameter generalize the max-flow/min-cut theorem, and lead to deep connections between algorithm design, graph theory, and geometry [20], [5], [3].

While the flow-cut gap for several classes of undirected graphs has been studied extensively, the case of directed graphs is poorly understood. In this work we make progress towards overcoming this limitation by showing that the flow-cut gap on n -vertex planar digraphs is $O(\log^3 n)$. This is the first *sub-polynomial* bound on any family of digraphs of

super-constant treewidth (not directed treewidth but the abstract graphs ignoring directions are of constant treewidth). We contrast our result with the strong lower bound due to Chuzhoy and Khanna [11], who showed that for general directed graphs the gap is $\Omega(n^{1/7})$, even for directed acyclic graphs. Thus it is highly natural to deal with planar directed graphs. Note that there is a long standing conjecture that for undirected planar graphs the flow-cut gap is $O(1)$ (we will mention more details later). There is no progress for a long time. So it seems hard to improve our flow-cut gap to $O(1)$ (this improvement would imply the above mentioned conjecture).

In order to prove our main results, we investigate low-distortion metric embeddings into directed ℓ_1 . More precisely, our result is obtained by proving an equivalent result in the theory of bi-Lipschitz quasimetric embeddings, which are mappings that generalize the standard theory of bi-Lipschitz metric embeddings to the asymmetric case. The distortion bound of our quasimetric embedding implies the same bound for the standard LP relaxations of various cut problems on directed graphs. Therefore, as a direct consequence, we also obtain the first polynomial-time polylogarithmic-approximation algorithms for several cut problems on directed graphs, including Directed Non-Bipartite Sparsest-Cut, and Directed Multicut for directed planar graphs.

A. Multi-commodity flow-cut gaps for undirected graphs

A *multi-commodity flow* instance in an undirected graph G is defined by two non-negative functions: $c: E(G) \rightarrow \mathbb{R}$ and $d: V(G) \times V(G) \rightarrow \mathbb{R}$. We refer to c and d as the *capacity* and *demand* functions respectively. The *maximum concurrent flow* is the maximal value ε such that for every $u, v \in V(G)$, $\varepsilon \cdot d(u, v)$ can be simultaneously routed between u and v , without violating the edge capacities. We refer to this value as $\text{maxflow}(G, c, d)$.

For every $S \subseteq V(G)$, the *sparsity* of S is defined as follows:

$$\frac{\sum_{(u,v) \in E(G)} c(u,v) |\mathbf{1}_S(u) - \mathbf{1}_S(v)|}{\sum_{u,v \in V(G)} d(u,v) |\mathbf{1}_S(u) - \mathbf{1}_S(v)|},$$

where $\mathbf{1}_S: V(G) \rightarrow \{0, 1\}$ is the indicator for membership in S . The sparsity of a cut is a natural upper bound for $\max\{\text{flow}(G, c, d), \text{min-cut}(G, c, d)\}$. The *multi-commodity max-flow min-cut gap* for G , denoted by $\text{gap}(G)$, is the maximum gap between the value of the flow and the upper bound given by the sparsity formula, over all multi-commodity flow instances on G . The flow-cut gap on undirected graphs has been studied extensively, and several upper and lower bounds have been obtained for various graph classes. The gap is referred to as the *uniform* multi-commodity flow-cut gap for the special case where there is a unit demand between every pair of vertices. Leighton and Rao [19] showed that the uniform flow-cut gap is $\Theta(\log n)$ in undirected graphs. Subsequently Linial, London and Rabinovich [20] and Aumann and Rabani [5] showed that the non-uniform multi-commodity flow-cut gap for the Sparsest Cut problem with k demand pairs is upper bounded by $O(\log k)$.

The flow-cut gap immediately implies a polynomial-time approximation algorithm for Sparsest-Cut, with approximation ratio equal to the gap. For general graphs, improved approximation algorithms have been obtained via semidefinite programming relaxations. This approach, pioneered by Arora, Rao and Vazirani [4], leads to $O(\sqrt{\log k})$ -approximation for the uniform case, and has been extended to $O(\sqrt{\log k \log \log k})$ -approximation for the general case by Arora, Lee and Naor [3]. The latter approach relies upon embeddings of metric spaces of negative type into ℓ_1 .

Besides these there are various studies of the flow-cut gap for specific graph families. A central conjecture posed by Gupta, Newman, Rabinovich, and Sinclair in [13] asserts the following.

Conjecture 1 (GNRS Conjecture [13]). *The multi-commodity flow-cut gap on any family $\mathcal{F}$ of graphs is $O(1)$ if and only if $\mathcal{F}$ forbids some fixed minor.*

Arguably one of the most interesting cases of Conjecture 1, which is still open, is the case of planar graphs, which is often referred to as the planar embedding conjecture:

Conjecture 2 (Planar Embedding Conjecture). *The multi-commodity flow-cut gap on planar graphs is $O(1)$.*

Conjecture 2 has been verified for the case of series-parallel graphs [13], $O(1)$ -outerplanar graphs [8], $O(1)$ -pathwidth graphs [18], and for some special classes of planar metrics [26]. However, Conjecture 2 is still wide open, and the current best gap is $O(\sqrt{\log n})$ by Rao [22] from 1999. Since then, no improvement has been made. The best-known lower bound is 2 [17].

For graphs excluding any fixed minor the flow-cut gap is known to be $O(1)$ for uniform demands [15].

B. Multi-commodity flow-cut gaps for directed graphs

For the case of directed graphs, the flow-cut gap is defined in terms of the Directed Non-Bipartite Sparsest Cut problem

which is an asymmetric variant of the Sparsest Cut problem, and is defined as follows. Let G be a directed graph and let $c: E(G) \rightarrow \mathbb{R}_{\geq 0}$ be a capacity function. Let $T = \{(s_1, t_1), (s_2, t_2), \dots, (s_k, t_k)\}$ be a set of terminal pairs, where each terminal pair (s_i, t_i) has a non-negative demand $\text{dem}(i)$. A cut in G is a subset of directed edges of $E(G)$. For a cut $S \subseteq E(G)$ in G , let I_S be the set of all indices $i \in \{1, 2, \dots, k\}$ such that all paths from s_i to t_i have at least one edge in S . Let $D(S) = \sum_{i \in I_S} \text{dem}(i)$ be the demand separated by S . Let $W(S) = \frac{C(S)}{D(S)}$ be the *sparsity* of S . The goal is to find a cut with minimum sparsity. The LP relaxation of this problem corresponds to the dual of the LP formulation of the directed maximum concurrent flow problem, and the integrality gap of this LP relaxation is the directed multi-commodity flow-cut gap. Hajiaghayi and Räcke [14] showed an upper bound of $O(\sqrt{n})$ for the flow-cut gap. This upper bound on the gap has been further improved by Gupta [12], and the current best approximation ratio is given by Agarwal, Alon and Charikar to $\tilde{O}(n^{11/23})$ in [2].

Besides these there are only a few studies of the flow-cut gap for specific (directed) graph families. For directed graphs whose abstract graphs are of treewidth t , it has been shown that the gap is at most $t \log^{O(1)} n$ by Mémoli, Sidiropoulos and Sridhar [21]. Salmasi, Sidiropoulos and Sridhar [25] have shown that the gap for the *uniform* case is $O(1)$ on directed graphs whose abstract graphs are series-parallel, and on digraphs whose abstract graphs are of bounded pathwidth, which implies $O(\log n)$ bounds for the gap on the non-uniform case.

On the lower bound side Saks *et al.* [24] showed that for general directed graphs the flow-cut gap is at least $k - \varepsilon$, for any constant $\varepsilon > 0$, and for any $k = O(\log n / \log \log n)$. The current best lower bound is given by Chuzhoy and Khanna who showed a $\tilde{\Omega}(n^{\frac{1}{7}})$ lower bound for the flow-cut gap in [11].

Our main result on multi-commodity flow-cut gaps.: In this paper we obtain the first poly-logarithmic upper bound for any family of digraphs whose abstract graph is of super-constant treewidth. Our main result on multi-commodity flow-cut gaps is as follows.

Theorem 1.1. *The uniform multi-commodity flow-cut gap on planar digraphs is $O(\log^2 n)$. The non-uniform multi-commodity flow-cut gap on planar digraphs is $O(\log^3 n)$.*

Given the strong lower bound by Chuzhoy and Khanna [11] for general directed graphs (even for directed acyclic graphs), it is highly natural to deal with planar digraphs.

This result can be considered as a first step to generalize Conjecture 2 for planar digraphs. Let us observe that if we could show $O(1)$ gap in Theorem 1.1, this would imply Conjecture 2. Moreover, this result extends the above mentioned long-standing result for undirected planar graphs by Rao [22] to planar digraphs (with a slightly weaker bound).

Theorem I.1 is a direct consequence of our main technical result, which uses the theory of Lipschitz quasipartitions, and is stated formally in Theorem I.4 (see [21], [25] for the reduction).

C. Cut problems of directed graphs

Better bounds on the flow-cut gap typically also imply better approximation ratios for solving cut problems. For the Directed Non-Bipartite Sparsest Cut problem the flow-cut gap upper bounds of [14] and [2] are also accompanied by $O(\sqrt{n})$ and $O(\tilde{n}^{11/23})$ polynomial time approximation algorithms respectively. Similarly for digraphs whose abstract graph are of treewidth t , a $t \log^{O(1)} n$ polynomial time approximation algorithm is also provided in [21].

Another closely related cut problem is the Directed Multicut problem which is defined as follows. Let G be a directed graph and let $c : E(G) \rightarrow \mathbb{R}_{\geq 0}$ be a capacity function. Let $T = \{(s_1, t_1), (s_2, t_2), \dots, (s_k, t_k)\}$ be a set of terminal pairs. A *cut* in G is a subset of $E(G)$. The *capacity* of a cut S is $c(S) = \sum_{e \in S} c(e)$. The goal is to find a cut separating all terminal pairs, minimizing the capacity of the cut. This problem is NP-hard. An $O(\sqrt{n \log n})$ approximation algorithm for Directed Multicut was presented by Cheriyan, Karloff and Rabani [9]. Subsequently an $\tilde{O}(n^{2/3}/OPT^{1/3})$ -approximation was given due to Kortsarts, Kortsarz and Nutov [16]. Finally [2] also gives an improved $\tilde{O}(n^{11/23})$ -approximation algorithm for this problem. Again for digraphs whose abstract graphs are of treewidth t a $t \log^{O(1)} n$ approximation algorithm was also shown in [21].

On the hardness side [10] demonstrated an $\Omega(\frac{\log n}{\log \log n})$ -hardness for the Directed Non-Bipartite Sparsest Cut problem and the Directed Multicut problem under the assumption that $\text{NP} \not\subseteq \text{DTIME}(n^{\log n^{O(1)}})$. This was further improved by them in a subsequent work [11] to obtain an $2^{\Omega(\log^{1-\varepsilon} n)}$ -hardness result for both problems for any constant $\varepsilon > 0$ assuming that $\text{NP} \subseteq \text{ZPP}$.

Our main results for these problems are the following theorems.

Theorem I.2. *There exists a polynomial-time $O(\log^2 n)$ -approximation algorithm for the Uniform Directed Sparsest Cut problem on planar digraphs. Moreover, there exists a polynomial-time $O(\log^3 n)$ -approximation algorithm for the General Directed Sparsest Cut problem on planar digraphs.*

Theorem I.3. *There exists a polynomial time $O(\log^2 n)$ -approximation algorithm for the Directed Multicut problem on planar digraphs.*

Both of the above results are direct consequences of our results on the theory of bi-Lipschitz quasimetric embeddings, which we describe in the next subsection. We refer the reader to [21], [25] for a detailed description of

the reduction from directed cut problems to bi-Lipschitz quasimetric embeddings.

D. Quasimetric spaces and embeddings

Random quasipartitions.: A quasimetric space is a pair (X, d) where X is a set of points and $d : X \times X \rightarrow \mathbb{R}_+ \cup \{+\infty\}$, that satisfies the following two conditions:

- (1) For all $x, y \in X$, $d(x, y) = 0$ iff $x = y$.
- (2) For all $x, y, z \in X$, $d(x, y) \leq d(x, z) + d(z, y)$.

The notion of random quasipartitions was introduced in [21]. A *quasipartition* of a quasimetric space (X, d) is a transitive reflexive relation on X . We remark that this definition is a generalization of the notion of a partition of a set of points, which can be viewed as a symmetric transitive binary relation, where two points are related if and only if they belong to the same cluster.

Let G be a digraph and let $M = (V(G), d_G)$ be its shortest-path quasimetric space. Let $F \subseteq E(G)$. Let

$$R = \{(u, v) \in V(G) \times V(G) : v \text{ is reachable from } u \text{ in } G \setminus F\}.$$

It is immediate to check that R is indeed a quasipartition. We say that R is *induced by the cutset* F .

Let $M = (X, d)$ be a quasimetric space. For any fixed $r \geq 0$, we say that a quasipartition Q of M is *r-bounded* if for every $x, y \in X$ with $(x, y) \in Q$, we have $d(x, y) \leq r$. For any $\beta > 0$, we say that a distribution over r -bounded quasipartitions of M , $\mathcal{D}$, is β -Lipschitz if for any $x, y \in X$, we have that

$$\Pr_{P \sim \mathcal{D}}[(x, y) \notin P] \leq \beta \frac{d(x, y)}{r}.$$

Given a distribution $\mathcal{D}$ over quasipartitions we sometimes use the term random quasipartition (with distribution $\mathcal{D}$) to refer to any quasipartition P sampled from $\mathcal{D}$. We consider the quasimetric space obtained from the shortest path distance of a directed graph. Mémoli, Sidiropoulos and Sridhar in [21] find an $O(1)$ -Lipschitz distribution over r -bounded quasipartitions of tree quasimetric spaces. They also prove the existence of a $O(t \log n)$ -Lipschitz distribution over r -bounded quasipartitions for any quasimetric that is obtained from a directed graph of treewidth t .

Our main result on Lipschitz quasipartitions is the following theorem. This is the main technical contribution of this paper. The fact that the quasipartition is efficiently samplable is needed in the algorithmic applications.

Theorem I.4. *Let $n \in \mathbb{N}$, and let G be an n -vertex planar digraph with non-negative edge lengths. Then for any $\Delta > 0$, there exists a Δ -bounded, $O(\log^2 n)$ -Lipschitz random quasipartition of the quasimetric space $(V(G), d_G)$. Moreover the quasipartition is samplable in polynomial time.*

Quasimetric embeddings.: Before stating our embedding results, we first need to introduce some notations and definitions. Let $M = (X, d)$ and $M' = (X', d')$ be quasimetric spaces. A mapping $f : X \rightarrow X'$ is called an *embedding of distortion* $c \geq 1$ if there exists some $\alpha > 0$, such that for all $x, y \in X$, we have $d(x, y) \leq \alpha \cdot d'(f(x), f(y)) \leq c \cdot d(x, y)$. We say that f is *isometric* when $c = 1$. We remark that this definition directly generalizes the standard notion of distortion of maps between metric spaces; however, here the ordering of any pair of points is important when dealing with distances.

Directed ℓ_1 (Charikar et al. [7]).: The directed ℓ_1 distance between two points x and y is given by $d_{\ell_1}(x, y) = \sum_i |x_i - y_i| + \sum_i |x_i| - \sum_i |y_i|$.

We can now state the main result on embeddability into directed ℓ_1 , which is an immediate consequence of Theorem I.4. For a detailed description of the reduction from the directed ℓ_1 embeddability to Lipschitz quasipartitions, we refer the reader to [21], [25].

Theorem I.5. *Let $n \in \mathbb{N}$, and let G be an n -vertex planar digraph with non-negative edge lengths. Then the quasimetric space $(V(G), d_G)$ admits an embedding into directed ℓ_1 with distortion $O(\log^3 n)$. Moreover the embedding is computable in polynomial time.*

E. High-level Overview of our Proof

As we have explained above, all of the results in this paper are direct consequences of a single result concerning random quasipartitions. Specifically, we show that for any planar digraph G , and for any “scale” $\Delta > 0$, there exists a Δ -bounded $O(\log^2 n)$ -Lipschitz quasipartition of the quasimetric space $(V(G), d_G)$.

We remark that it is known that the classical theory of Lipschitz partitions of metric spaces is known to not be sufficient for obtaining such a result (see [21], [25]). Thus, in order to handle the quasimetric case we combine ideas from the theory of Lipschitz partitions with methods developed within the literature of distance oracles for directed planar graphs, developed by Thorup [27].

Because we deal with directed graphs, some difficulties in using the distance oracle results appear in our proof. See Figure 1 for example. Moreover, we also have to consider some random partitioning scheme of directed graphs, which is similar to the one used by Klein, Plotkin and Rao [15] for undirected graphs.

We now explain what are the main ingredients used in the proof, and how they are combined in our main algorithm. The input to the algorithm consists of a directed planar graph G and $\Delta > 0$. The output is a random quasipartition of the quasimetric space $(V(G), d_G)$. The main steps of the algorithm are as follows. In order to simplify the exposition, we relax slightly the notation.

Step 1: Partitioning G into layers. It has been shown by Thorup [27] that any planar digraph can be decomposed into a sequence of vertex-disjoint “layers” $L_1, L_2, \dots$, such that each layer L_i has a directed rooted spanning tree T_i , such that at least one of the two following properties hold:

- (i) Every vertex in L_i can be reached from the root via a directed path of length at most Δ in T_i .
- (ii) Every vertex in L_i can reach the root via a directed path of length at most Δ in T_i .

We construct a randomized version of this decomposition scheme, which is similar to the random partitioning scheme used by Klein, Plotkin and Rao [15]. One important property is that during this random decomposition step, each edge $(u, v) \in E(G)$ is cut with probability at most $O(d_G(u, v)/\Delta)$. A key property is that, as in Thorup [27], every path of length at most Δ can be contained in at most three consecutive layers. This fact implies that in order to obtain a quasipartition for G , it is enough to compute a quasipartition for each layer, and then output their “common refinement” (which is formally defined by taking the union of the corresponding cutsets).

Step 2: Quasipartitioning each layer. It now remains to show how to compute a quasipartition for each layer. Fix some layer L . We now invoke another result of Thorup [27], who showed that there exists three shortest paths P_1, P_2, P_3 of length at most Δ in L , such that $V(P_1 \cup P_2 \cup P_3)$ is a balanced vertex separator of L . Of course, we cannot simply delete the paths P_1, P_2, P_3 , since this would cut some edges with probability 1 (since the choice of the paths is deterministic). Instead, we show how to delete some random set of edges $C \subseteq E(L)$, such that both of the following two conditions hold for the resulting digraph:

- (i) Each edge $(u, v) \in E(L)$ is cut with probability at most $O(d_L(u, v)/\Delta)$.
- (ii) If there exists a path R from u to v that intersects $P_1 \cup P_2 \cup P_3$, then it must be that $d_G(u, v) \leq \Delta$.

We refer to the above process as “quasipartitioning the neighborhood of the separator”; we give more details about this process in Step 3.

Equipped with this process, we easily obtain a quasipartition for L via recursion: we recursively run the algorithm, effectively quasipartitioning the neighborhoods of all separators of each connected component of $L \setminus V(P_1 \cup P_2 \cup P_3)$. The final output is the union of all cutsets computed in all recursive

calls. Since every edge can be contained in at most $O(\log n)$ recursive calls, it follows by the union bound that the probability of cutting any particular edge is at most $O(\log n)$ times the probability of cutting it in each quasipartitioning step, which we show below is at most $O(\log n \cdot d_L(u, v)/\Delta)$, which implies the result.

Step 3: Quasipartitioning the neighborhood of a separator. Because the separators involved in Step 2 consist of the union of at most three shortest paths, it can be shown that it is enough to handle the case of a single separator path P . An important difficulty here is that common intuition from the notion of a “neighborhood” in the case of undirected graphs, does not directly translate to the directed case that we are dealing with. Roughly speaking, the reason is that there are two different neighborhoods of P : one consisting of all vertices that are at distance at most Δ from some vertex in P , and one consisting of all vertices that are at distance at most Δ to some vertex in P (See Figure 1 for example). We obtain a quasipartitioning scheme that, roughly speaking, handles the “overlay” of these two neighborhoods. This process is inspired from the algorithm of Bartal [6] for computing Lipschitz partitions of undirected graphs. Bartal’s algorithm computes a random partition by growing balls of radii that are distributed according to some truncated exponential distribution. We adapt Bartal’s algorithm by growing random “quasiballs”, and showing that a similar analysis can be used to analyze their behavior.

This completes the high-level description of our approach.

Comparison to Klein-Plotkin-Rao and beyond. It has been shown by Klein, Plotkin and Rao [15], that for any undirected graph G that excludes some fixed minor, for any $\Delta > 0$, the shortest path metric $(V(G), d_G)$ admits a Δ -bounded, $O(1)$ -Lipschitz partition. Here, we obtain a slightly weaker bound (i.e. $O(\log^2 n)$ instead of $O(1)$), for the more general case of directed graphs, but our result holds only for planar digraphs. The reason is that the approach in [15] directly produces a forbidden minor, when the partitioning scheme fails to produce a bounded partition. This is, in fact, one of the very few known results for general minor-free graph families, that does not use the celebrated theory of graph minors of Robertson and Seymour [23]. In contrast, we use planarity in very specific ways in order to randomly “decompose” the input graph into more manageable pieces. More precisely, we only use the above mentioned path separator theorems by Thorup [27] for planar undirected graphs. Ittai and Gavioille [1] showed that the above planar result can be extended to minor-free undirected graphs. We leave it as an open problem to extend our approach to arbitrary families of minor-free directed graphs.

F. Organization

In Section II we introduce some definitions that will be used throughout the rest of the paper. In Section III we describe a general process for deleting random subsets of edges, that is analogous to Bartal’s Lipschitz partitioning scheme [6]. In Section IV we show how to apply this general process to the special case of quasipartitioning the “neighborhood” of a shortest path. This corresponds to the above Step 3. In Section V we introduce a randomized variant of Thorup’s decomposition scheme [27]; we also use (recursively) the path separator theorems from [27] in combination with the quasipartitioning scheme from Section IV, to obtain a quasipartitioning scheme for each part of the decomposition. This corresponds to the above Step 1. Finally, we put all the ingredients together in Section VI, where we present the main algorithm for sampling a Lipschitz quasipartition in any planar digraph.

II. DEFINITIONS AND PRELIMINARIES

We now introduce some notation that will be used throughout the paper. For any functions $f, g : \mathbb{N} \rightarrow \mathbb{N}$, we write $f(n) \lesssim g(n)$ whenever $f(n) = O(g(n))$.

Graphs and subgraphs. For any graph G and any $U \subset V(G)$, we denote by $G[U]$ the subgraph of G induced by U ; that is $V(G[U]) = U$, and $E(G[U]) = E(G) \cap (U \times U)$. For any $r > 0$, we write

$$N_G(U, r) = \{v \in V(G) : d_G(U, v) \leq r\},$$

where

$$d_G(U, v) := \min_{u \in U} d_G(u, v).$$

For any digraph G , any $u, v \in V(G)$, and any directed path P from u to v in G , we refer to u and v as the *head* and *tail* of P , respectively. We denote by $\text{len}(P)$ the length of P . For any $a, b \in V(P)$, we write $a \leq_P b$ if a precedes b in the traversal of P . We also write $b \geq_P a$ if $a \leq_P b$.

For any graph G , $X \subseteq V(G)$, and $\alpha \in [1, 0]$, we say that X is a α -balanced vertex separator of G if every connected component of $G \setminus X$ has at most $\alpha|V(G)|$ vertices.

For any digraph G , let $\overleftarrow{G}$ denote the digraph obtained by reversing the direction of every edge. For any path Q in G , let $\overleftarrow{Q}$ denote the path in $\overleftarrow{G}$ obtained by reversing the direction of every edge in Q .

For any digraph G , and any path P in G , and any $x, y \in V(P)$, such that either $x = y$, or x appears before y in P , we denote by $P[x, y]$ the subpath of P from x to y . Adapting standard set-theoretic notation, we also write $P[x, y) = P[x, y] \setminus \{y\}$, $P(x, y) = P[x, y] \setminus \{x\}$, and $P(x, y) = P[x, y] \setminus \{x, y\}$.

III. EXPONENTIAL RANDOM QUASIBALLS

In this section we analyze the behavior of a general process for constructing quasipartitions. The process involves computing a sequence of cutsets that are induced by quasiballs of exponentially distributed radius.

The process is completely analogous to the algorithm used by Bartal [6] to compute probabilistic partitions of metric spaces.

Let H be a digraph, with $|V(H)| = n$, and let $\Delta > 0$. We define a probability distribution $\mathcal{D}$ over cutsets of H , i.e. $2^{E(H)}$, as follows. First, we define a sequence of pairwise disjoint subsets $B_1, \dots, B_t \subseteq V(H)$. Let $I = [0, \Delta \ln n)$, and let $\mathcal{I}$ be the truncated exponential probability distribution on I with probability density function $p(x) = \frac{n}{n-1} \cdot \frac{1}{\Delta} \cdot e^{-x/\Delta}$. Initially all vertices in H are *unmarked*. Let $v_1 \in V(H)$ be chosen arbitrarily. Let $R_1 \sim \mathcal{I}$, and define

$$B_1 = \text{ball}_H(v_1, R_1).$$

We mark all vertices in B_1 . For any $i \geq 1$, we proceed inductively. If all vertices in H are marked, then we terminate the sequence at $t = i - 1$. Otherwise, we arbitrarily either terminate the sequence at $t = i - 1$, or we proceed to define B_i ; that is, this choice can depend on all previous random choices. We set v_i be to an arbitrary unmarked vertex. Let $R_i \sim \mathcal{I}$, and define

$$B_i = \text{ball}_H(v_i, R_i) \setminus (B_1 \cup \dots \cup B_{i-1}).$$

We mark all vertices in B_i . This completes the inductive construction of the sequence $B_1, \dots, B_t$. Finally, we set

$$F = \bigcup_{i=1}^t E(H) \cap (B_i \times (V(H) \setminus (B_1 \cup \dots \cup B_i))).$$

That is, for all $i \in \{1, \dots, t\}$, the cutset F contains all edges starting in B_i and ending at outside of $B_1 \cup \dots \cup B_i$. This completes the definition of the random cutset F . We say that the resulting probability distribution $\mathcal{D}$ is a Δ -shock.

We remark that the above process is quite general and allows for many different strategies for choosing the next vertex v_i , and for deciding whether to terminate the sequence before all vertices have been marked. In our main algorithm, we will apply this process by choosing the vertices $v_1, \dots, v_t$ from an inverse traversal of some shortest path (see Section IV).

We next analyze Δ -shocks.

Lemma III.1. *Let H be a digraph, let $\Delta > 0$, let $\mathcal{D}$ be a Δ -shock, and let $F \sim \mathcal{D}$. Then, for any $(u, v) \in E(H)$, we have*

$$\Pr[(u, v) \in F] \leq 2d_H(u, v)/\Delta.$$

The proof of Lemma III.1 is similar to the analysis of an analogous result for the case of metric cases from Bartal [6]. For the sake of completeness, the proof of Lemma III.1 is given in Section A.

IV. QUASIPARTITIONING THE NEIGHBORHOOD OF A DIRECTED PATH

Let H be a digraph, let $\Delta > 0$, and let Q be a shortest path in H from some $v \in V(H)$ to some $v' \in V(H)$, of length at

most Δ . We describe a procedure for computing a random quasipartition, given H and Q . Let $I = [0, \Delta \ln n)$, and let $\mathcal{I}$ be the truncated exponential probability distribution on I with probability density function $p(x) = \frac{n}{n-1} \cdot \frac{1}{\Delta} \cdot e^{-x/\Delta}$. The procedure consists of the following steps.

Procedure PATH-QUASIPARTITION(H, Q, Δ)

Step 1: Cutting balls away from Q . For any $i \in \mathbb{N}$, let $R_i \sim \mathcal{I}$. Let $a_1 = v'$,

$$B_1 = \text{ball}_H(a_1, R_1),$$

For any $i > 1$, let a_i be the first vertex not in $B_1 \cup \dots \cup B_{i-1}$ that we visit when traversing Q starting from v' . Let

$$B_i = \text{ball}_H(a_i, R_i) \setminus \left(\bigcup_{j=1}^{i-1} B_j \right).$$

For any $i \in \mathbb{N}$, let

$$F_i = E(H) \cap (B_i \times (V(H) \setminus B_i)).$$

See Figure 1 for an example.

Step 2: Cutting balls towards Q . For any $i \in \mathbb{N}$, let $R'_i \sim \mathcal{I}$. Let $a'_1 = v$,

$$B'_1 = \text{ball}_H(a'_1, R'_1),$$

For any $i > 1$, let a'_i be the first vertex not in $B'_1 \cup \dots \cup B'_{i-1}$ that we visit when traversing Q starting from v ; w.l.o.g. we may assume that $d_H(a'_i, a'_{i-1}) = R'_i$, by subdividing one edge of H . Let

$$B'_i = \text{ball}_H(a'_i, R'_i) \setminus \left(\bigcup_{j=1}^{i-1} B'_j \right).$$

For any $i \in \mathbb{N}$, let

$$F'_i = E(H) \cap ((V(H) \setminus B'_i) \times B'_i).$$

Step 3: Output. Let $F = \bigcup_{i \in \mathbb{N}} F_i$, $F' = \bigcup_{i \in \mathbb{N}} F'_i$. We output the cutset $C = F \cup F'$, and the quasipartition Z induced by C .

This completes the description of the procedure.

Portal assignment. The procedure PATH-QUASIPARTITION computes sets of vertices $A = \{a_i\}_i$, $A' = \{a'_i\}_i$. We refer to the vertices in $A \cup A'$ as *portals*. We define mappings

$$\pi_1 : V(Q) \rightarrow A,$$

and

$$\pi_2 : V(Q) \rightarrow A',$$

as follows. For any $a_i \in A$, and for any $x \in V(\overleftarrow{Q}[a_i, a_{i+1}])$, we set $\pi_1(x) = a_i$. For any $a'_i \in A'$, and for any $x \in V(Q[a_i, a_{i+1}])$, we set $\pi_2(x) = a_i$. Figure 2 depicts an example.




Figure 1. An example of Step 1 in procedure $\text{PATH-QUASIPARTITION}(H, Q, \Delta)$. Each B_i is depicted as a region with a distinct color. For clarity, only the regions $B_1, \dots, B_4$ are depicted. The sets $F_1, \dots, F_4$ are also depicted, each with a different color. We remark that Q is a shortest path from v to v' , but does not need to be a shortest path from v' to v . Consequently, the intersection of each B_i with Q does not need to consist of a single contiguous sub-path (e.g. B_3).




Figure 2. An example portal assignment.

Analysis.: We now analyze the quasipartition computed by procedure $\text{PATH-QUASIPARTITION}$.

Lemma IV.1. *Let H be a digraph, let Q be a shortest path in H of length at most $\Delta > 0$. Let C be a random cutset computed by Procedure $\text{PATH-QUASIPARTITION}(H, Q, \Delta)$, and let Z be the induced quasipartition. Then the following conditions hold:*

(1) *For any $(x, y) \in E(H)$,*

$$\Pr[(x, y) \notin Z] \lesssim \frac{d_H(x, y)}{\Delta}.$$

(2) *Let $x, y \in V(H)$. Suppose that there exists some path P from x to y in $H \setminus C$, such that P intersects Q . Then, $d_H(x, y) \lesssim \Delta \log n$.*

Proof: First we prove part (1). From Lemma III.1 it follows that for any $i \in \{1, 2\}$, during Step i , the probability of cutting (x, y) is at most $2d_H(x, y)/\Delta$. By taking a union bound over the two steps, the assertion follows.

It remains to prove part (2). Let a and b be the first and last vertices in $V(P) \cap V(Q)$ that we encounter along a traversal of P .

We argue that $d_H(x, \pi_2(a)) \lesssim \Delta \log n$. Recall that $C = F \cup F'$, where F is the cutset computed in Step 1, and F' is the cutset computed in Step 2. Let $P_1 = P[x, a]$. Since $P \subseteq H \setminus C$, it follows that $E(P_1) \cap F' = \emptyset$. By induction on $i \in \mathbb{N}$, we have that for all $i \geq 1$, there exists no path in $H \setminus F'$ from any vertex in $V(H) \setminus (B'_1 \cup \dots \cup B'_i)$ to any vertex in $B'_1 \cup \dots \cup B'_i$. Therefore, there exist $j, k \in \mathbb{N}$, with $j \leq k$, such that $x \in B'_j$ and $a \in B'_k$ (see Figure 3). Thus

$$\begin{aligned} d_H(x, \pi_2(a)) &\leq d_H(x, a_j) + d_H(a_j, \pi_2(a)) \\ &\quad (\text{triangle inequality}) \\ &\lesssim \Delta \log n + d_H(a_j, \pi_2(a)) \\ &\quad (x \in B'_j \subseteq \text{ball}_H(a_j, R_j)) \\ &\lesssim \Delta \log n + \Delta \\ &\quad (a_j \leq_Q a_k \leq_Q \pi_2(a) \text{ and } \text{len}(Q) \leq \Delta) \\ &\lesssim \Delta \log n, \end{aligned} \tag{1}$$

as required.

We next argue that $d_H(\pi_1(b), y) \lesssim \Delta \log n$. Arguing as in the paragraph above, we have that there exist $s, t \in \mathbb{N}$,




Figure 3. The configuration in the proof of Lemma IV.1. The path P_1 is not cut by F' .

with $s \geq t$, such that $b \in B_s$ and $y \in B_t$. Thus

$$\begin{aligned}
 d_H(\pi_1(b), y) &\leq d_H(\pi_1(b), a_t) + d_H(a_t, y) \\
 &\quad (\text{triangle inequality}) \\
 &\lesssim d_H(\pi_1(b), a_t) + \Delta \log n \\
 &\quad (y \in B_t \subseteq \text{ball}_H(a_t, R_t)) \\
 &\lesssim \Delta + \Delta \log n \\
 &\quad (\pi_1(b) \leq_Q a_s \leq_Q a_t \text{ and } \text{len}(Q) \leq \Delta) \\
 &\lesssim \Delta \log n, \tag{2}
 \end{aligned}$$

as required.

We next argue that $d_H(\pi_2(a), \pi_2(b)) \lesssim \Delta$. If $\pi_2(a) \leq_Q \pi_2(b)$, then $\pi_2(a)$ can reach $\pi_2(b)$ by following Q , and thus $d_H(\pi_2(a), \pi_2(b)) \leq \text{len}(Q) \leq \Delta$. Thus it remains to consider the case $\pi_2(a) \geq_Q \pi_2(b)$. Let $i, j \in \mathbb{N}$, such that $\pi_2(a) = a'_i$ and $\pi_2(b) = a'_j$. Since $a'_i \geq_Q a'_j$, it follows that $i \geq j$. If $i = j$, then $\pi_2(a) = \pi_2(b)$, and thus $d_H(\pi_2(a), \pi_2(b)) = 0$. Thus, it remains to consider the case $i > j$. By induction on the execution of Step 2, we have that for all $t \in \mathbb{N}$, any vertex in $V(H) \setminus (B'_1 \cup \dots \cup B'_t)$ cannot reach any vertex in $B'_1 \cup \dots \cup B'_t$ in $H \setminus F'$. Therefore, a cannot reach b , which is a contradiction. We have thus established that

$$d_H(\pi_2(a), \pi_2(b)) \leq \Delta. \tag{3}$$

Next, we argue that $d_H(\pi_2(b), \pi_1(b)) \leq \Delta$. By the choice of portals, we have that $\pi_2(b) \leq_Q b \leq_Q \pi_1(b)$. Since $\text{len}(Q) \leq \Delta$, this implies that

$$d_H(\pi_2(b), \pi_1(b)) \leq \text{len}(Q) \leq \Delta. \tag{4}$$

Combining (1), (2), (3) and (4) with the triangle inequality, we obtain $d_H(x, y) \leq d_H(x, \pi_2(a)) + d_H(\pi_2(a), \pi_2(b)) + d_H(\pi_2(b), \pi_1(b)) + d_H(\pi_1(b), y) \lesssim \Delta \log n + \Delta + \Delta + \Delta \log n \lesssim \Delta \log n$, which concludes the proof. ■

V. CRASHING WAVES

In this section, we refine some important concepts in [27]. This is one of the keys in our analysis in the next section.

We recall the following definition from [27]. A (t, Δ) -layered spanning tree T in a digraph H is a directed rooted spanning tree such that any root-to-leaf path in T is the




Figure 4. A sample from a (v, Δ) -wave.

concatenation of at most t shortest paths in H , each of length at most Δ . We say that H is (t, Δ) -layered if it contains such a spanning tree. We are interested in $(1, \Delta)$ -layered graphs and their properties. However, we use the slightly more general definition in order to maintain a compatible notation with [27].

Definition V.1 (Wave). Let G be a digraph, $v \in V(G)$, $\Delta > 0$. We define a probability distribution, $\mathcal{D}$, supported over $2^{E(G)}$. Let $\tau_0, \dots, \tau_n \in [0, \Delta]$ be chosen uniformly and independently at random. Let

$$V_0 = \{u \in V(G) : d_G(v, u) \leq 2\Delta + \tau_0\},$$

and for any $i \geq 1$, let

$$V_i = \begin{cases} \{u \in V(G) : d_G(V_{i-1}, u) \leq 2\Delta + \tau_i\} & \text{if } i \text{ is even} \\ \{u \in V(G) : d_G(u, V_{i-1}) \leq 2\Delta + \tau_i\} & \text{if } i \text{ is odd} \end{cases}$$

For any $i \in \mathbb{N}$, let

$$E_i = \begin{cases} E(G) \cap (V_{i-1} \times V_i) & \text{if } i \text{ is odd} \\ E(G) \cap (V_i \times V_{i-1}) & \text{if } i \text{ is even} \end{cases}$$

Finally, let

$$E' = \bigcup_i E_i.$$

We define $\mathcal{D}$ to be the distribution of the random variable E' . We refer to $\mathcal{D}$ as a (v, Δ) -wave.

Figure 4 gives an example of the construction of a wave. From the above definition, we show the following lemma.

Lemma V.2. Let G be a digraph, $v \in V(G)$, $\Delta > 0$. Suppose that all edges in G have length at most Δ . Let $\mathcal{D}$ be a (v, Δ) -wave. Let E' be a subset of edges sampled from $\mathcal{D}$, i.e. $E' \sim \mathcal{D}$. Then the following properties hold:

(i) For any $(a, b) \in E(G)$, we have that

$$\Pr[(a, b) \in E'] \lesssim \frac{d_G(a, b)}{\Delta}.$$

(ii) The following holds with probability 1. Let P be any path in G of length at most Δ . There exists some $i \geq 0$, such that $V(P) \subseteq V_i \cup V_{i+1} \cup V_{i+2}$. Furthermore, P can be decomposed into consecutive subpaths $P = P_1 \circ P_2 \circ P_3$, such that for all

$i \in \{1, 2, 3\}$, there exists $j \in \{i, i+1, i+2\}$, such that $P_i \subseteq V_j$

(iii) The following holds with probability 1. For any $i \in \mathbb{N}$, contracting $\bigcup_{j < i} V_j$ into a single vertex, and deleting $\bigcup_{j > i} V_j$ results in a $(1, \Delta)$ -layered graph.

Proof: We first prove (i). Let $(a, b) \in E(G)$. Consider the process of sampling E' from $\mathcal{D}$. Let $U_0 = \{v\}$, and for any $i \geq 1$, let $U_i = V_0 \cup \dots \cup V_{i-1}$. Intuitively, we can think of the process of sampling E' as inductively constructing the sequence $\{V_i\}_{i \geq 0}$, and U_i is the set of vertices that have already been explored just before constructing V_i . In order to cut the edge (a, b) during the construction of V_i , it must be that

$$d_G(U_i, a) \leq 2\Delta + \tau_i < d_G(U_i, b). \quad (5)$$

This implies that

$$d_G(U_i, a) \leq 3\Delta. \quad (6)$$

By the definition of a wave we have that $N_G(U_i, 2\Delta) \subseteq U_{i+2}$. Combining with (6) we get

$$d_G(U_{i+2}, a) \leq \Delta. \quad (7)$$

However, (7) implies that, with probability 1, $\{a, b\} \subseteq U_{i+1}$, and thus (a, b) cannot be cut during the construction of V_{i+2} . Therefore, (a, b) can potentially be cut in at most two consecutive steps during the constructing of some V_i and V_{i+1} . Note that i is a random variable that depends on the choice of the random variables $\{\tau_i\}_{i \geq 0}$. In each one of these two steps, by triangle inequality, we have that (a, b) is cut when (5) holds, which happens with probability at most $\frac{d_G(U_i, b) - d_G(U_i, a)}{\Delta} < \frac{d_G(a, b)}{\Delta}$. By the union bound over these two steps, it follows that the probability that (a, b) is cut at any point during the construction of the wave is at most $2 \frac{d_G(a, b)}{\Delta}$, as required. This proves part (i).

Next, we prove part (ii). Let P be any path in G of length at most Δ . Let w and w' be the first and last vertices in P . Let i be the minimum integer such that $V(P) \cap V_i \neq \emptyset$. Let v be an arbitrary vertex in $V(P) \cap V_i$. Since P has length at most Δ , it follows that for any $v' \in V(P)$, either $d_G(v, v') \leq \Delta$, or $d_G(v', v) \leq \Delta$. In both cases, this implies that $v' \in V_i \cup V_{i+1} \cup V_{i+2}$. Since v' is arbitrary, we get $P \subseteq V_i \cup V_{i+1} \cup V_{i+2}$. It remains to show that P can be decomposed into consecutive subpaths $P = P_1 \circ P_2 \circ P_3$, such that for all $i \in \{1, 2, 3\}$, there exists $j \in \{i, i+1, i+2\}$, such that $P_i \subseteq V_j$. By the construction of a wave, it follows that for all t , we either cut all edges from V_t to V_{t+1} , or we cut all the edges from V_{t+1} to V_t . Thus, while traversing P , if we leave some set V_t , we can never return to V_t (see Figure 4). This implies that the intersection of P with each of the sets V_i, V_{i+1} , and V_{i+2} is either empty or a subpath of P , which concludes the proof of part (ii).

Finally, part (iii) follows by contracting U_i into a single vertex u_i , and setting the spanning tree T in the definition of layered graph to be the shortest path tree rooted at u_i . ■

We need the following result from [27].

Lemma V.3 (Thorup [27]). *Let $\Delta > 0$, and let H be a planar $(1, \Delta)$ -layered digraph. Then, there exist shortest paths P_1, P_2, P_3 in H , each having length at most Δ , and such that $V(P_1) \cup V(P_2) \cup V(P_3)$ is a $2/3$ -balanced vertex separator in H . Moreover, there exists a polynomial-time algorithm which given H and Δ outputs P_1, P_2 , and P_3 .*

Lemma V.4. *Let $\Delta > 0$. Let H be a planar $(1, \Delta)$ -layered digraph. Then there exists a Δ -bounded, $O(\log^2 n)$ -Lipschitz quasipartition of the quasimetric space $(V(H), d_H)$.*

Proof: We construct a random cutset C recursively. If H contains at most $c = O(1)$ vertices, then the result follows by letting C be the set of all edges in H of length at least Δ , with probability 1. Otherwise, we proceed as follows. By Lemma V.3, we compute shortest paths P_1, P_2, P_3 in H , each of length at most Δ , such that $V(P_1 \cup P_2 \cup P_3)$ is a $2/3$ -balanced vertex separator in H . For each $i \in \{1, 2, 3\}$, we run procedure `PATH-QUASIPARTITION`(H, P_i, Δ), and obtain a random cutset $C_i \subseteq E(H)$. We then recurse on each connected component of $H \setminus (P_1 \cup P_2 \cup P_3)$. Let $\mathcal{P}$ be the set of all paths computed by Lemma V.3 in all recursive calls.

Let also $\mathcal{C}$ be the set of all random cutsets computed during all recursive calls. We set

$$C = \bigcup_{C' \in \mathcal{C}} C',$$

and we output the quasipartition R induced by C . This completes the description of the algorithm.

We next show that R is $O(\log n)$ -Lipschitz. Let $(u, v) \in E(H)$. By construction, we have that (u, v) is contained in at most one subgraph at each level of the recursion. Since the depth of the recursion is $k = O(\log n)$, it follows that (u, v) is contained in at most k subgraphs $H_1, \dots, H_k$ of H , with

$$H = H_1 \supseteq H_2 \supseteq \dots \supseteq H_k.$$

For each $i \in \{1, \dots, k\}$ let $P_{i,1}, P_{i,2}, P_{i,3}$ be the shortest paths that the algorithm computes in H_i , such that $V(P_1 \cup P_2 \cup P_3)$ is a $2/3$ -balanced separator in H_i . For each $i \in \{1, \dots, k\}$ let C_i be the random cutset that the algorithm

computes in H_i . We have

$$\begin{aligned}
\Pr[(u, v) \notin R] &\leq \Pr[(u, v) \in C] \\
&= \Pr[(u, v) \in C_1 \cup \dots \cup C_k] \\
&\leq \sum_{i=1}^k \Pr[(u, v) \in C_i] \\
&\quad (\text{union bound}) \\
&\lesssim \sum_{i=1}^k \frac{d_{H_i}(u, v)}{\Delta} \\
&\quad (\text{part (1) of Lemma IV.1}) \\
&= \sum_{i=1}^k \frac{d_H(u, v)}{\Delta} \\
&\quad (\text{since } (u, v) \in E(H_1) \cap \dots \cap E(H_k)) \\
&\lesssim \log n \frac{d_H(u, v)}{\Delta}, \\
&\quad (k \lesssim \log n)
\end{aligned}$$

and thus R is $O(\log n)$ -Lipschitz.

It remains to show that R is $O(\Delta \log n)$ -bounded. Let $x, y \in V(H)$, such that $(x, y) \in R$. This implies that there exists some path L from x to y in $H \setminus C$. Let

$$X = \bigcup_{P \in \mathcal{P}} V(P).$$

If $V(L) \cap X = \emptyset$, then by the basis of the recursion we have that L can contain at most $c = O(1)$ edges, each of length less than Δ , and thus $d_H(x, y) = O(\Delta)$. Otherwise, $V(L) \cap X \neq \emptyset$. This implies that there exists some path $P \in \mathcal{P}$ such that L intersects P . By part (2) of Lemma IV.1 it follows that $d_H(x, y) \lesssim \Delta \log n$. Thus, in all cases we obtain that $d_H(x, y) \lesssim \Delta \log n$, which implies that R is $O(\Delta \log n)$ -bounded, concluding the proof. To obtain a Δ -bounded, $O(\log^2 n)$ -Lipschitz quasipartition, we just scale the quasi-metric space $(V(H), d_H)$ by a factor of $\Theta(1/\log n)$. ■

VI. LIPSCHITZ QUASIPARTITIONS FOR PLANAR DIGRAPHS

We now prove our main result on Lipschitz quasipartitions, which is the main technical ingredient of this paper. We begin by describing the algorithm for computing a quasipartition, and we then prove that it satisfies the desired properties.

The input to the algorithm is a directed planar graph G , and some parameter $\Delta > 0$. The algorithm proceeds in the following steps:

Step 1: Sampling from a wave. Pick an arbitrary $v \in V(G)$. Let $\mathcal{D}$ be a (v, Δ) -wave, and let Q' be a quasipartition sampled according to $\mathcal{D}$, induced by some cutset $E' \subseteq E(G)$; that is $Q' \sim \mathcal{D}$. Let $\{V_i\}_{i \in \mathbb{N}}$ be as in Definition V.1. For any i , let

G_i be a $(1, \Delta)$ -layered graph (with probability 1) obtained by part (iii) of Lemma V.2.

Step 2: Quasipartitioning each layer. By Lemma V.4 it follows that for each i , there exists a polynomial-time samplable $\Delta/3$ -bounded $O(\log^2 n)$ -Lipschitz quasipartition $\mathcal{D}_i$ of G_i . In polynomial time, we sample some $Q_i \sim \mathcal{D}_i$, induced by some cutset F_i .

Step 3: Output. Let $F = E' \cup (\bigcup_i F_i)$. The final output is the quasipartition R induced by the cutset F .

We are now ready to prove the main result of this Section.

Proof of Theorem I.4: Let F be the cutset computed by the algorithm described above, and let R be the induced quasipartition. Let also $\mathcal{D}$, E' , G_i , $\mathcal{D}_i$, F_i be as in the description of the algorithm.

We need to bound the probability that any ordered pair of vertices is not in R . First, consider some $(u, v) \in E(G)$. By part (i) of Lemma V.2 we have that

$$\Pr[(u, v) \in E'] \lesssim \frac{d_G(u, v)}{\Delta} \quad (8)$$

With probability 1, there exists at most one $i \in \mathbb{N}$, such that $(u, v) \in E(G_i)$. Conditioned on the event that $(u, v) \in E(G_i)$, by Lemma V.4 we have

$$\Pr[(u, v) \in F_i] \lesssim \log^2 n \frac{d_G(u, v)}{\Delta} \quad (9)$$

Combining (8), (9), and the fact that (u, v) can be in at most one graph G_i , it follows by the union bound that

$$\begin{aligned}
\Pr[(u, v) \in F] &= \Pr \left[(u, v) \in E' \cup \left(\bigcup_i F_i \right) \right] \\
&= \Pr \left[((u, v) \in E') \vee \left((u, v) \in \bigcup_i F_i \right) \right] \\
&\leq \Pr[(u, v) \in E'] + \Pr \left[(u, v) \in \bigcup_i F_i \right] \\
&\quad (\text{union bound}) \\
&\lesssim \log^2 n \frac{d_G(u, v)}{\Delta}. \quad (10)
\end{aligned}$$

We have thus established that the probability of cutting any particular edge is low. We next extend this bound to arbitrary ordered pairs of vertices. Let $u', v' \in V(G)$. If v' is not reachable from u' in G , the condition holds vacuously. Thus, we may assume w.l.o.g. that v' is reachable from u' in G . Let P be a shortest path from u' to v' in G . Let $P = w_1, \dots, w_t$, with $w_1 = u'$, $w_t = v'$. We have that if $(u', v') \notin R$, then at least one edge of P must be in F , since otherwise v' is reachable from u' in $G \setminus F$, which implies

$(u', v') \in R$. Thus, we have

$$\begin{aligned}
\Pr[(u', v') \notin R] &\leq \Pr[E(P) \cap F \neq \emptyset] \\
&= \Pr\left[\left(\bigcup_{i=1}^{t-1} (w_i, w_{i+1})\right) \cap F \neq \emptyset\right] \\
&= \Pr\left[\bigvee_{i=1}^{t-1} ((w_i, w_{i+1}) \cap F \neq \emptyset)\right] \\
&\leq \sum_{i=1}^{t-1} \Pr[(w_i, w_{i+1}) \cap F \neq \emptyset] \\
&\quad (\text{union bound}) \\
&\lesssim \log^2 n \sum_{i=1}^{t-1} \frac{d_G(w_i, w_{i+1})}{\Delta} \\
&\quad (\text{by (10)}) \\
&= \log^2 n \frac{d_G(u', v')}{\Delta} \\
&\quad (P \text{ is a } u'v' \text{ shortest path})
\end{aligned}$$

Thus, we have established that the probability that any ordered pair is not contained in the quasipartition is bounded, as required.

It remains to show that R is Δ -bounded. To that end, it suffices to show that for any $a, b \in V(G)$, if $d_G(a, b) > \Delta$, then $(a, b) \notin R$. In order to prove the latter property, it suffices to show that for any path Z from a to b , at least one edge of Z is in the cutset F ; that is, $E(Z) \cap F \neq \emptyset$. Let Z' be prefix of Z of length Δ (we may assume that Z' has length exactly Δ by subdividing one edge and inserting a vertex on Z at distance exactly Δ from u). By part (ii) of Lemma V.2 we have that, with probability 1, there exists some $i \in \mathbb{N}$, such that Z' can be decomposed into three subpaths $Z' = Z'_1 \circ Z'_2 \circ Z'_3$, such that each Z'_i is contained in some G_s . By setting Z'' to be the longest of the three paths Z'_1 , Z'_2 and Z'_3 , we conclude that there exists some subpath Z'' of Z' , of length strictly greater than $\Delta/3$, and there exists some $j \in \mathbb{N}$, such that $Z'' \subseteq G_j$. Since, by construction, $\mathcal{D}_j$ is $\Delta/3$ -bounded, it follows that $E(Z'') \cap F_i \neq \emptyset$. Since $F_i \subseteq F$, and $E(Z'') \subseteq E(Z)$, we get that $E(Z) \cap F \neq \emptyset$, which concludes the proof. ■

REFERENCES

- [1] I. Abraham and C. Gavoille. Object location using path separators. In *Proceedings of the Twenty-Fifth Annual ACM Symposium on Principles of Distributed Computing(PODC)*, pages 188–197, 2006.
- [2] A. Agarwal, N. Alon, and M. S. Charikar. Improved approximation for directed cut problems. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 671–680. ACM, 2007.
- [3] S. Arora, J. Lee, and A. Naor. Euclidean distortion and the sparsest cut. *Journal of the American Mathematical Society*, 21(1):1–21, 2008.
- [4] S. Arora, S. Rao, and U. V. Vazirani. Expander flows, geometric embeddings and graph partitioning. *J. ACM*, 56(2):5:1–5:37, 2009.
- [5] Y. Aumann and Y. Rabani. An $o(\log k)$ approximate min-cut max-flow theorem and approximation algorithm. *SIAM Journal on Computing*, 27(1):291–301, 1998.
- [6] Y. Bartal. Probabilistic approximation of metric spaces and its algorithmic applications. In *Proceedings of 37th Conference on Foundations of Computer Science*, pages 184–193. IEEE, 1996.
- [7] M. Charikar, K. Makarychev, and Y. Makarychev. Directed metrics and directed graph partitioning problems. In *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pages 51–60. Society for Industrial and Applied Mathematics, 2006.
- [8] C. Chekuri, A. Gupta, I. Newman, Y. Rabinovich, and A. Sinclair. Embedding k -outerplanar graphs into ℓ_1 . *SIAM Journal on Discrete Mathematics*, 20(1):119–136, 2006.
- [9] J. Cheriyan, H. Karloff, and Y. Rabani. Approximating directed multicuts. *Combinatorica*, 25(3):251–269, 2005.
- [10] J. Chuzhoy and S. Khanna. Hardness of cut problems in directed graphs. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 527–536. ACM, 2006.
- [11] J. Chuzhoy and S. Khanna. Polynomial flow-cut gaps and hardness of directed cut problems. *Journal of the ACM (JACM)*, 56(2):6, 2009.
- [12] A. Gupta. Improved results for directed multicut. In *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pages 454–455. Society for Industrial and Applied Mathematics, 2003.
- [13] A. Gupta, I. Newman, Y. Rabinovich, and A. Sinclair. Cuts, trees and ℓ_1 -embeddings of graphs. *Combinatorica*, 24(2):233–269, 2004.
- [14] M. T. Hajiaghayi and H. Räcke. An-approximation algorithm for directed sparsest cut. *Information Processing Letters*, 97(4):156–160, 2006.
- [15] P. Klein, S. A. Plotkin, and S. Rao. Excluded minors, network decomposition, and multicommodity flow. In *Proceedings of the twenty-fifth annual ACM symposium on Theory of computing*, pages 682–690. ACM, 1993.
- [16] Y. Kortsarts, G. Kortsarz, and Z. Nutov. Greedy approximation algorithms for directed multicuts. *Networks*, 45(4):214–217, 2005.
- [17] J. R. Lee and P. Raghavendra. Coarse differentiation and multi-flows in planar graphs. *Discrete & Computational Geometry*, 43(2):346–362, 2010.
- [18] J. R. Lee and A. Sidiropoulos. On the geometry of graphs with a forbidden minor. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pages 245–254. ACM, 2009.

- [19] T. Leighton and S. Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *Journal of the ACM (JACM)*, 46(6):787–832, 1999.
- [20] N. Linial, E. London, and Y. Rabinovich. The geometry of graphs and some of its algorithmic applications. In *Foundations of Computer Science, 1994 Proceedings., 35th Annual Symposium on*, pages 577–591. IEEE, 1994.
- [21] F. Mémoli, A. Sidiropoulos, and V. Sridhar. Quasimetric embeddings and their applications. *arXiv preprint arXiv:1608.01396*, 2016.
- [22] S. Rao. Small distortion and volume preserving embeddings for planar and euclidean metrics. In *Proceedings of the Fifteenth Annual Symposium on Computational Geometry*, pages 300–306, 1999.
- [23] N. Robertson and P. D. Seymour. Graph minors I – XXIII, 1982 – 2010. *Journal of Combinatorial Theory, Series B from 1982 till 2010*.
- [24] M. Saks, A. Samorodnitsky, and L. Zosin. A lower bound on the integrality gap for minimum multicut in directed networks. *Combinatorica*, 24(3):525–530, 2004.
- [25] A. Salmasi, A. Sidiropoulos, and V. Sridhar. On constant multi-commodity flow-cut gaps for families of directed minor-free graphs. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 535–553, 2019.
- [26] A. Sidiropoulos. Non-positive curvature and the planar embedding conjecture. In *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*, pages 177–186. IEEE, 2013.
- [27] M. Thorup. Compact oracles for reachability and approximate distances in planar digraphs. *Journal of the ACM (JACM)*, 51(6):993–1024, 2004.

APPENDIX

Proof of Lemma III.1: The proof is a direct adaptation of the argument of Bartal [6], which holds for the case of metric spaces. Here, we observe that the same analysis holds for the directed variant we have defined.

Fix some $(u, v) \in E(H)$. We will derive an upper bound on the probability that (u, v) is included in the cutset F . Let $i \in \{1, \dots, t\}$. If $d_H(v_i, u) > d_H(v_i, v)$, then the probability that (u, v) is included at the cutset while computing X_i is 0. Therefore, we may assume w.l.o.g. that $d_H(v_i, u) \leq d_H(v_i, v)$. Following Bartal [6], we define the following events:

A_i : $u, v \in V(H) \setminus (B_{i-1} \cup \dots \cup B_i)$, where we define $B_0 = \emptyset$. That is, none of u and v have been included in any X_j , for any $j < i$.

M_i^I : $d_H(v_i, v) \leq R_i$, conditioned on A_i . That is, this is the event that both u and v are included in B_i , conditioned on A_i .

M_i^X : $d_H(v_i, u) \leq R_i < d_H(v_i, v)$, conditioned on A_i . That is, this is the event that (u, v) is cut when computing B_i .

M_i^N : $R_i < d_H(v_i, u)$, conditioned on A_i . That is, this is the event that none of the vertices u and v is included in X_i , conditioned on A_i .

N_i : The event that for all $j \in \{i+1, \dots, t\}$, we have $|\{u, v\} \cap B_j| \leq 1$, conditioned on A_i . That is, the vertices u and v are not both included in the same cluster B_j , for all $j > i$, conditioned on A_i .

From [6], we have

$$\begin{aligned} \Pr[M_i^X] &= \int_{d_H(v_i, u)}^{d_H(v_i, v)} p(x) dx \\ &= \frac{n}{n-1} \cdot \left(1 - e^{-\frac{d_H(v_i, v) - d_H(v_i, u)}{\Delta}}\right) e^{-d_H(v_i, v)/\Delta} \\ &\leq \frac{n}{n-1} \cdot \frac{d_H(u, v)}{\Delta} e^{-d_H(v_i, v)/\Delta}. \end{aligned}$$

Moreover

$$\begin{aligned} \Pr[M_i^N] &= \int_0^{d_H(v_i, u)} p(x) dx \\ &= \frac{n}{n-1} \cdot \left(1 - e^{-d_H(v_i, u)/\Delta}\right). \end{aligned}$$

Following [6], we now prove by induction on i that

$$\Pr[N_i] \leq \left(2 - \frac{i}{n-1}\right) \frac{d_H(u, v)}{\Delta}.$$

We have

$$\begin{aligned} \Pr[N_i] &= \Pr[M_i^X] + \Pr[M_i^N] \Pr[N_{i+1}] \\ &\leq \frac{n}{n-1} \cdot \frac{d_H(u, v)}{\Delta} e^{-d_H(v_i, u)/\Delta} \\ &\quad + \frac{n}{n-1} \left(1 - e^{-d_H(v_i, u)/\Delta}\right) \left(2 - \frac{i+1}{n}\right) \frac{d_H(u, v)}{\Delta} \\ &\leq \left(2 - \frac{i}{n-1}\right) \frac{d_H(u, v)}{\Delta}, \end{aligned}$$

which completes the induction. We have thus established that

$$\Pr[(u, v) \in F] \leq \Pr[N_0] \leq 2d_H(u, v)/\Delta,$$

which concludes the proof. ■