

000
001
002
003
004
005
006
007
008
009
010
011
012
013
014
015
016
017
018
019
020
021
022
023
024
025
026
027
028
029
030
031
032
033
034
035
036
037
038
039
040
041
042
043
044
045
046
047
048
049
050
051
052
053

DYNAMIC KERNEL GRAPH SPARSIFIERS

Anonymous authors

Paper under double-blind review

ABSTRACT

A geometric graph associated with a set of points $P = \{x_1, x_2, \dots, x_n\} \subset \mathbb{R}^d$ and a fixed kernel function $K : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$ is a complete graph on P such that the weight of edge (x_i, x_j) is $K(x_i, x_j)$. We present a fully-dynamic data structure that maintains a spectral sparsifier of a geometric graph under updates that change the locations of points in P one at a time. The update time of our data structure is $n^{o(1)}$ with high probability, and the initialization time is $n^{1+o(1)}$. Under certain assumption, our data structure can be made robust against adaptive adversaries, which makes our sparsifier applicable in iterative optimization algorithms. We further show that the Laplacian matrices corresponding to geometric graphs admit a randomized sketch for maintaining matrix-vector multiplication and projection in $n^{o(1)}$ time, under *sparse* updates to the query vectors, or under modification of points in P .

1 INTRODUCTION

Kernel methods are a fundamental tool in modern data analysis and machine learning, with extensive applications in computer science, from clustering, ranking and classification, to ridge regression, principal-component analysis and semi-supervised learning (von Luxburg, 2007; Ng et al., 2002; Zhu, 2005a;b; Liu et al., 2019). Given a set of n points in \mathbb{R}^d and a nonnegative function $K : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$, a kernel matrix has the form that the i, j -th entry in the matrix is $K(x_i, x_j)$.

Kernel matrices and kernel linear-systems naturally arise in modern machine learning and optimization tasks, from Kernel PCA and ridge regression (Alaoui & Mahoney, 2015; Avron et al., 2017a;b; Lee et al., 2020), to Gaussian-process regression (GPR) (Rasmussen & Nickisch, 2010), federated learning (Konečný et al., 2016), and the ‘state-space model’ (SSM) in deep learning (Gu et al., 2021b;a). In most of these applications, the underlying data points x_i are dynamically changing across iterations, either by nature or by design, and therefore computational efficiency of numerical linear-algebraic operations in this setting requires dynamic algorithms to maintain the kernel matrix under insertions and deletions of data points.

One motivating application of dynamic linear algebra on geometric graphs is dynamically maintaining a *spectral clustering* (Ng et al., 2002) of the kernel matrix of a weighted graph. In the static setting, a common approach for spectral clustering is *spectral sparsification* (Spielman & Srivastava, 2011; Ng et al., 2002), i.e., to run a spectral clustering algorithm on top of a spectral-sparsifier for the Laplacian matrix of the weighted graph. This approach, however, fails to extend to the dynamic setting, where a small fraction of data points are continually changing, since rebuilding the spectral-sparsifier is prohibitively expensive – Changing a single point $x_i \in P$ changes *an entire row* of $K(x_i, x_j)$.

Another motivation, arising in statistical physics and astronomy, is the N-body simulation problem (Trenti & Hut, 2008). The problem asks to efficiently simulate a dynamical system of particles, usually under the influence of physical forces, such as gravity. Let $P \subset \mathbb{R}^d$ denote a set of points. In our terminology, this setup corresponds to maintaining, for each $i \in [d]$, a graph G_i on the points in P , and letting C_g denote the gravitational constant, and m_x denote the mass of point $x \in P$. Hence, for any two points $u, v \in P$, the non-negative weight/kernel function of the edge (u, v) is defined as $K_i(u, v) := (\frac{C_g \cdot m_u \cdot m_v}{\|u-v\|_2^2}) \cdot (\frac{|v_i - u_i|}{\|u-v\|_2})$. Denoting the weighted adjacency matrix of G_i by A_i , computing the force between the points in the static setting corresponds to $A_i \mathbf{1}$. Once again, this approach (Trenti & Hut, 2008) fails to extend to the dynamic setting in which the n -bodies are slowly moving over time, since re-computing $K()$ would take $\Omega(n)$ time.

Finally, we mention an application to *semi-supervised learning* tasks, where the goal is to extend a partial function, whose values are known only on a subset of the training data, to the entire domain, such that the weighted sum of differences over the set is minimized (Zhu, 2005b)¹. In the static setting of geometric kernels, this least-squares minimizer can be found by solving a (Laplacian) linear system on P . Extending this approach to the dynamic setting requires dynamic spectral sparsifiers.

As mentioned above, one of the main tools for fast linear algebra on geometric graphs is spectral sparsification. Alman, Chu, Schild and Song (Alman et al., 2020) presented a static algorithm for constructing an ϵ -spectral sparsifier on a geometric graph $H = H(G)$ such that $(1 - \epsilon)L_G \preceq L_H \preceq (1 + \epsilon)L_G$ in almost linear time, which avoids explicitly writing the underlying $n \times n$ dense Kernel matrix, and facilitates several basic linear-algebraic operations on geometric graphs, in $\tilde{O}(n \log^d n) \ll n^2$ time, when the dimension d is fixed. The main goal of this paper is to extend the toolbox of (Alman et al., 2020) to the *dynamic* setting, as motivated by the above applications. More formally:

Given a set of n points $P \subset \mathbb{R}^d$ and a function K , is there a dynamic algorithm that can update the spectral sparsifier of the geometric graph G on P and K in $n^{o(1)}$ time, where in each iteration the location of a point $x_i \in P$ is changed? Is it possible to maintain approximate matrix-vector queries w.r.t the Laplacian matrix of a dynamic geometric graph, and an approximate-inverse of the Laplacian, in $n^{o(1)}$ time?

Prior to this work, no nontrivial dynamization algorithms were known for geometric graphs for the above linear-algebra primitives. While fully-dynamic $(1 \pm \epsilon)$ -spectral edge sparsifiers for general graphs are known (Abraham et al., 2016) (in amortized $\text{poly}(\log(n), 1/\epsilon)$ update time), the setting of geometric graphs is fundamentally different, since as mentioned earlier in the introduction, each update of a point $x_i \in P$ results in an *entire row* update of K , i.e., $O(n)$ edges, making (Abraham et al., 2016) too slow.

The main technical contribution of this paper is a new *dynamic well separated pair decomposition* (WSPD, Definition B.15, (Fischer & Har-Peled, 2005; Alman et al., 2020)). The core of this data structure is a *smooth resampling* technique for efficiently maintaining a WSPD under point-location updates, with mild weight-increase to the sparsifier, by *reusing* randomness (in an adversarially-robust manner).

1.1 MODEL

Before we state our main results, let us formally state the dynamic model of geometric graphs we consider:

Definition 1.1 (Dynamic spectral sparsifier of geometric graph). *Given a set of points $P \subset \mathbb{R}^d$ and kernel function $K : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$. Let G denote the geometric graph on P with edge weight $w(x_i, x_j) := K(x_i, x_j)$. Let L_G be the Laplacian matrix of graph G . Let $\epsilon \in (0, 0.1)$ denote an accuracy parameter. We want to design a data structure that dynamically maintains a $(1 \pm \epsilon)$ -spectral sparsifier for G and supports the following operations:*

- **INITIALIZE**($P \subset \mathbb{R}^d, \epsilon \in (0, 0.1)$), this operation takes point set P and constructs a $(1 \pm \epsilon)$ -spectral sparsifier of L_G .
- **UPDATE**($i \in [n], z \in \mathbb{R}^d$), this operation takes a vector z as input, and to replace x_i (in point set P) by z , in the meanwhile, we want this update to be fast and the change in the spectral sparsifier to be small.

For the above problem, we focus our attention on kernel functions with a natural property called (C, L) -multiplicatively Lipschitz. For $C \geq 1$ and $L \geq 1$, we say a function is (C, L) -Lipschitz

¹Formally, we are given a function $f : P \rightarrow \mathbb{R}$ together with its value on some subset $X \subset P$. Then we aim to extend the function f to the whole set P , which can minimize $\sum_{u, v \in P, u \neq v} K(u, v)(f(u) - f(v))^2$

108 if that, for all $c \in [1/C, C]$, it holds that $\frac{1}{c^L} \leq \frac{f(cx)}{f(x)} \leq c^L$. Common (C, L) -Lipschitz kernels
 109 including piecewise exponential kernels, polynomial functions of distance (with non-negative coef-
 110 ficients), and rational function kernels (Alman et al., 2020).

111 We formally define the sketch task of matrix multiplication here.

112 **Definition 1.2** (Sketch of approximation to matrix multiplication). *Given a geometric graph G with respect to point set P and kernel function K , and an n -dimensional vector x , we want to maintain a low dimensional sketch of an approximation to the multiplication result $L_G x$, where an ϵ -approximation to multiplication result $L_G x$ is a vector b such that $\|b - L_G x\|_2 \leq \epsilon \cdot \|L_G\|_F \cdot \|x\|_2$.*

113 We also give the formal definition of the sketching approximation to Laplacian solving here.

114 **Definition 1.3** (Sketch of approximation to Laplacian solving). *Given a geometric graph G with respect to point set P and kernel function K , and an n -dimensional vector b , we want to maintain a low dimensional sketch of an approximation to the multiplication result $L_G^\dagger b$, i.e., a vector \tilde{z} such that $\|\tilde{z} - L_G^\dagger b\|_2 \leq \epsilon \cdot \|L_G^\dagger\|_F \cdot \|b\|_2$*

115 We here explain the necessity of maintaining a sketch of an approximation instead of the directly
 116 maintaining the multiplication result in the dynamic regime. Let the underlying geometry graph on
 117 n vertices be G and the vector be $v \in \mathbb{R}^n$. When a d -dimensional point is moved in the geometric
 118 graph, a column and a row are changed L_G . We can assume the first row and first column are
 119 changed with no loss of generality. When this happens, if the first entry of v is not 0, all entries will
 120 change in the multiplication result. Therefore, it takes at least $\Omega(n)$ time to update the multiplication
 121 result exactly. In order to spend subpolynomial time to maintain the multiplication result, we need
 122 to reduce the dimension of vectors. Therefore, we use a sketch matrix with $m = \text{poly log}(n)$ rows.

132 2 OUR RESULTS

133 Our first main result is a dynamic sparsifier for geometric graphs, with subpolynomial update time
 134 in the oblivious adversary model.

135 **Theorem 2.1** (Informal version of Theorem E.3). *Let K denote a (C, L) -multiplicative Lipschitz
 136 kernel function. For any given data point set $P \subset \mathbb{R}^d$ with size n , there is a randomized dynamic
 137 algorithm DYNAMICGEOSPAR that receives updates of locations of points in P one at a time, and
 138 maintains an almost linear spectral sparsifier in $n^{o(1)}$ time with probability $1 - 1/\text{poly}(n)$.*

139 By introducing additional assumptions regarding dimensions, we can generate outcomes for an ad-
 140 versarial setting.

141 **Theorem 2.2** (Informal version of Theorem I.5). *Let K denote a (C, L) -multiplicative Lipschitz
 142 kernel function. For any given data point set $P \subset \mathbb{R}^d$ with size n . Define $\alpha := \frac{\max_{x, y \in P} \|x - y\|_2}{\min_{x, y \in P} \|x - y\|_2}$. If
 143 $\alpha^d = O(\text{poly}(n))$, then there is a randomized dynamic algorithm that receives updates of locations
 144 of points in P one at a time, and maintains a almost linear spectral sparsifier in $n^{o(1)}$ time with
 145 probability $1 - 1/\text{poly}(n)$. It also supports adversarial updates.*

146 For the dynamic matrix-vector multiplication problem, we give an algorithm to maintain a sketch of
 147 the multiplication between the Laplacian matrix of a geometric graph and a given vector in subpoly-
 148 nomial time.

149 **Theorem 2.3** (Informal version of Theorem F.1). *Let G be a (C, L) -Lipschitz geometric graph on n
 150 points. Let v be a vector in \mathbb{R}^n . There exists an data structure MULTIPLY that maintains a vector \tilde{z}
 151 that is a low dimensional sketch of an approximation to the multiplication result $L_G \cdot v$. MULTIPLY
 152 supports the following operations: Part 1. UPDATEG(x_i, z): move a point from x_i to z and thus
 153 changing K_G and update the sketch. This takes $n^{o(1)}$ time. Part 2. UPDATEV(δ_v): change v to
 154 $v + \delta_v$ and update the sketch. This takes $n^{o(1)}$ time. Part 3. QUERY: return the up-to-date sketch.*

155 We also present a dynamic algorithm to maintain the sketch of the solution to a Laplacian system.

156 **Theorem 2.4** (Informal version of Theorem G.1). *Let G be a (C, L) -Lipschitz geometric graph on
 157 n points. Let b be a vector in \mathbb{R}^n . There exists an data structure SOLVE that maintains a vector \tilde{z}
 158 that is a low dimensional sketch of an ϵ -approximation to the multiplication result $L_G^\dagger \cdot b$. It supports*

162 the following operations: Part 1. $\text{UPDATEG}(x_i, z)$: move a point from x_i to z and thus changing
 163 K_G and update the sketch. This takes $n^{o(1)}$ time. Part 2. $\text{UPDATEB}(\delta_b)$: change b to $b + \delta_b$ and
 164 update the sketch. This takes $n^{o(1)}$ time. Part 3. QUERY : return the up-to-date sketch.
 165

166 **Roadmap.** We provide an overview of techniques used in Section 3. For all the formal proofs, we
 167 leave them to the Appendix. We conclude our work in Section 4.
 168

169 3 TECHNICAL OVERVIEW

171 3.1 NOTATIONS

173 For any two sets A, B , we use $A \triangle B$ to denote $(A \setminus B) \cup (B \setminus A)$. Given two symmetric matrices
 174 A, B , we say $A \preceq B$ if $\forall x, x^\top Ax \leq x^\top Bx$. For a vector x , we use $\|x\|_2$ to denote its ℓ_2 norm.
 175 For psd matrix A , we use A^\dagger to denote the pseudo inverse of A . For two point sets A, B , we denote
 176 the complete bipartite graph on A and B by $\text{Biclique}(A, B)$. We use $\mu_{n,i}$ to denote an elementary
 177 unit vector in \mathbb{R}^n with i -th entry 1 and others 0. We use $\mathcal{T}_{\text{mat}}(a, b, c)$ to denote the running time of
 178 computing the product of two matrices in the shape of $\mathbb{R}^{a \times b}$ and $\mathbb{R}^{b \times c}$ respectively. For a matrix
 179 $A \in \mathbb{R}^{m \times n}$, we use $\|A\|_F$ to denote its Frobenius norm, i.e., $\|A\|_F := (\sum_{i \in [m], j \in [n]} A_{i,j}^2)^{1/2}$. For
 180 a matrix, we use A^\dagger to denote the pseudo inverse of matrix A . For a vector x , and a psd matrix A ,
 181 we use $\|x\|_A := (x^\top Ax)^{1/2}$.
 182

183 3.2 FULLY DYNAMIC KERNEL SPARSIFICATION DATA STRUCTURE

185 A geometric graph w.r.t. kernel function $K : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ and points $x_1, \dots, x_n \in \mathbb{R}^d$ is a graph on
 186 x_1, \dots, x_n where the weight of the edge between x_i and x_j is $K(x_i, x_j)$. An update to a geometric
 187 graph occurs when the location of one of these points changes.

188 In a geometric graph, when an update occurs, the weights of $O(n)$ edges change. Therefore, directly
 189 applying the existing algorithms for dynamic spectral sparsifiers ((Abraham et al., 2016)) to update
 190 the geometric graph spectral sparsifier will take $\Omega(n)$ time per update. However by using the fact
 191 that the points are located in \mathbb{R}^d and exploiting the properties of the kernel function, we can achieve
 192 faster update.

193 Before presenting our dynamic data structure, we first give a high level idea of the static construction
 194 of the geometric spectral sparsifier, which is presented in (Alman et al., 2020).
 195

196 3.2.1 BUILDING BLOCKS OF THE SPARSIFIER

198 In order to construct a spectral sparsifier more efficiently, one can partition the graph into several
 199 subgraphs such that the edge weights on each subgraph are close. On each of these subgraphs, lever-
 200 age score sampling, which is introduced in (Spielman & Srivastava, 2011) and used for constructing
 201 sparsifiers, can be approximated by uniform sampling.

202 For a geometric graph built from a d -dimensional point set P , under the assumption that each edge
 203 weight is obtained from a (C, L) -Lipschitz kernel function (Definition B.1), each edge weight in
 204 the geometric graph is not distorted by a lot from the euclidean distance between the two points
 205 (Lemma B.22). Therefore, we can compute this partition efficiently by finding a well separated pair
 206 decomposition ((Callahan & Kosaraju, 1995), WSPD, Definition B.15) of the given point set.
 207

208 An s -WSPD of P is a collection of pairs (A_i, B_i) of subsets of P , such that for all $a \neq b \in P$, there
 209 exists a unique i satisfying $a \in A_i, b \in B_i$, and the distance between A_i and B_i (as point sets) is at
 210 least s times the diameters of A_i and B_i ((A_i, B_i) is a s -well separated (WS) pair). In this case, the
 211 distance between point sets A_i and B_i is a $(1 \pm 1/s)$ -multiplicative approximation of the distance
 212 between any point in A_i and any point in B_i .
 213

214 Each WS pair in the WSPD can be viewed as an unweighted biclique, where the two point sets
 215 are the two sides of the bipartite graph. On an unweighted biclique, uniform random sampling and
 216 leverage score sampling are equivalent. Therefore, a uniformly random sample of the biclique forms
 217 a spectral sparsifier of the biclique, and the union of the sampled edges from all bicliques form a
 218 spectral sparsifier of the geometric graph.
 219

216 However, the time needed for constructing a WSPD depends exponentially on the ambient dimension
 217 of the point set and thus WSPD cannot be computed efficiently when the dimension is high. To
 218 solve this problem, one can use the ultra low dimensional Johnson Lindenstrauss (JL) projection to
 219 project the point set down to $k = o(\log n)$ dimension, such that with high probability, the distance
 220 distortion (multiplicative difference between the distance between two points and the distance be-
 221 tween their low dimensional images) between any pair of points is at most $n^{C_{jl} \cdot (1/k)}$, where C_{jl} is a
 222 universal constant for the JL projection. This distortion becomes an overestimation of the leverage
 223 score in the resulting biclique, and can be compensated by sampling $n^{C_{jl} \cdot (1/k)}$ edges.

224 Then one can perform a 2-WSPD on the k -dimensional points. Since JL projection gives a bijection
 225 between the d -dimensional points and their k -dimensional images, a 2-WSPD of the k -dimensional
 226 point set gives us a $(2 \cdot n^{C_{jl} \cdot (1/k)})$ -WSPD of the d -dimensional point set P . The d -dimensional bicliques
 227 resulted from this $(2 \cdot n^{C_{jl} \cdot (1/k)})$ -WSPD of P is what we use to sample edges and construct the sparsifier.
 228

229 In summary, after receiving a set of points P , we use the ultra-low dimensional JL projection to
 230 project these points to a $k = o(\log n)$ dimensional space, run a 2-WSPD on the k -dimensional
 231 points, and then map the k -dimensional WSPD result back to the d -dimensional point set to obtain a
 232 $(2 \cdot n^{C_{jl} \cdot (1/k)})$ -WSPD of P . For each pair (A, B) in this d -dimensional WSPD, we randomly sample
 233 edges from $\text{Biclique}(A, B)$. The union of all sampled edges is a spectral sparsifier of the geometric
 234 graph on P .

235 3.2.2 DYNAMIC UPDATE OF THE GEOMETRIC SPECTRAL SPARSIFIER.

237 For a geometric graph build on a point set P , we want the above spectral sparsifier to be able to
 238 handle the following update² :

239 240 Point location change ($x_i \in P, z \in \mathbb{R}^d$): move the point from location x_i to location z . This is
 241 equivalent to removing point $x \in P$ and then adding z to P .

242 However, in order to update the geometric spectral sparsifier efficiently, there are a few barriers that
 243 we need to overcome.

244 **Updating WSPD** When the point set P changes, we want to update the 2-WSPD such that the
 245 number of WS pairs that are changed in the WSPD is small. (Fischer & Har-Peled, 2005) presented
 246 an algorithm to update the list of WS pairs, but it cannot be used directly in this situation, because the
 247 (Fischer & Har-Peled, 2005) algorithm is only able to return a list of WS pairs such that the singleton
 248 containing the inserted (or removed) point is one of the vertex subsets in these pairs. However, in
 249 order to use the up-to-date WSPD to update the sparsifier, we need to know not only the WS pairs
 250 (A, B) where A or B is a singleton consisting of the inserted (or removed) point, but also all other
 251 WS pairs (A, B) such that the inserted (or removed) point is in A or B . The (Fischer & Har-Peled,
 252 2005) algorithm is not able to do this.

253 Fortunately, one s -WSPD construction algorithm presented in (Har-Peled, 2011) has the property
 254 that each point x appears only in $s^{O(d)}O(\log \alpha)$ WS pairs. This allows us to find all WS pairs
 255 affected by a point location change in $2^{O(k)}O(\log \alpha)$ time, since we are maintaining a 2-WSPD and
 256 the dimension of the point set is k . We summarize this algorithm below. The detailed discussion of
 257 the WSPD update can be found in Section E.4.

258 The (Har-Peled, 2011) WSPD algorithm constructs a *compressed quadtree* associated with the point
 259 set P , and the WS pairs are pairs of nodes of the compressed quadtree. To summarize, a *quadtree*
 260 is a hierarchical partition of a k -dimensional hypercube enclosing P . It is obtained by recursively
 261 dividing the k -dimensional region into 2^k smaller regions called *cells* (divide equally along each
 262 axis), which can be further subdivided until each resulting cell contains only one point. The tree
 263 representing this hierarchy is called a quadtree and each cell in this hierarchy corresponds to a tree
 264 node of the quadtree. The cells containing only one point are the *leaf nodes* of the tree. In a quadtree,
 265 there can be a long chain of tree nodes that contain the same set of points. We replace this chain by
 266 the first and last nodes on the chain, and an edge between them. The resulted tree is the compressed

267 268 269 ²We assume that throughout the update, the aspect ratio of the point set, denoted by $\alpha = \frac{\max_{x, y \in P} \|x - y\|_2}{\min_{x, y \in P} \|x - y\|_2}$,
 does not change.

quadtree associated with P . The compressed quadtree has size $O(|P|)$, and supports the following operations in $O(\log n)$ time: (1) finding the leaf node that contains a given point x , or the parent node under which the leaf node containing x should be inserted if $x \notin P$, (2) inserting a leaf node containing a given point x , and (3) removing a leaf node containing a given point x .

The WSPD is a list of pairs of well separated compressed quadtree nodes. For efficient update, we let the WSPD data structure to be a container (of WS pairs) that supports looking up all WS pairs containing a tree node n for a given n in time linear in the size of output.

When a point location update occurs, suppose point x_i is moved to z . We can do the following to find all WS pairs that need to be updated.

- Use the compressed quad tree data structure to locate leaf nodes that contains x_i and z (since z is not in the point set before the update, we locate the parent node under which z should be inserted)
- Go from each of these leaf nodes to the root of the compressed quad tree, for each tree node n visited in this process, use the WSPD data structure to find all WS pairs containing n .
- Update all WS pairs found in the previous step and the compressed quadtree.

Algorithm 6 in Section E.4 is a detailed version of this WSPD update scheme.

Resampling from bicliques. After updating the WSPD, we want to generate a uniform sample of edges from the new biclique. We show that with high probability, this can be done in $n^{o(1)}$ time with high probability.

When a point location change happens and point x_i is moved to z , each pair (A, B) in the WSPD list will undergo one and only one of the following changes, Part 1. Remaining (A, B) . Part 2. Becoming $(A \setminus \{x_i\}, B)$ or $(A, B \setminus \{x_i\})$. Part 3. Becoming $(A \cup \{z\}, B)$ or $(A, B \cup \{z\})$. Part 4. Becoming $(A \setminus \{x_i\} \cup \{z\}, B)$, $(A, B \setminus \{x_i\} \cup \{z\})$, $(A \setminus \{x_i\}, B \cup \{z\})$ or $(A \cup \{z\}, B \setminus \{x_i\})$.

For each WS pair (A, B) that remains (A, B) , we do not need to do anything about it. For each WS pair (A, B) that is changed (A', B') , in order to maintain a spectral sparsifier of Biclique(A', B'), we need to find a new uniform sample from Biclique(A', B'). Simply drawing another uniform sample from $A' \times B'$ cannot be done fast enough when $|A' \times B'|$ is large and this resampling will cause a lot of edge weight changes in the final sparsifier, which is not optimal.

To overcome this barrier, suppose after an update, a WS pair (A, B) is changed to (A', B') . Since the size difference between A and A' and the size difference between B and B' are at most constant, the size of $(A' \times B') \cap (A \times B)$ is much larger than the size of $(A' \times B') \setminus (A \times B)$. Therefore, when we draw a uniform sample from $A' \times B'$, most of the edges in the sample should be drawn from $(A' \times B') \cap (A \times B)$. Since we already have a uniform sample E from $A \times B$, which contains a uniform sample from $(A' \times B') \cap (A \times B)$, we can reuse E in the following way:

Let $H = E \cap (A' \times B')$. For each edge that needs to be samples, we flip an unfair coin for which the probability of landing on head is $\frac{|(A' \times B') \cap (A \times B)|}{|A' \times B'|}$, and we do the following (See Figure 2 for a visual example):

- If the coin lands on head, we sample an edge from H without repetition;
- Otherwise we sample an edge from $(A' \times B') \setminus (A \times B)$ without repetition.

Algorithm 7 in Section E.5 is a detailed version of this resampling scheme. With properly set probability for the coin flip, doing the sampling this way generates a uniform sample of $A' \times B'$, and with high probability, the difference between the new sample and E is small.

However, in this process, although the difference between the new sample and E is small, we still need to flip a coin for each new sample point. When the sample size is big, this can be slow.

The running time of resampling can be improved by removing a small number of edges from E . Indeed, suppose we want to resample s edges from $A' \times B'$, the number of edges that need to be drawn from $(A' \times B') \setminus (A \times B)$ follows a Binomial distribution with parameters s and $\frac{|(A' \times B') \setminus (A \times B)|}{|A' \times B'|}$.

We have the following improved resampling algorithm:

324 Let $H = E \cap (A' \times B')$.
 325

326 • Generate a random number x under $\text{Binomial}(s, \frac{|(A' \times B') \setminus (A \times B)|}{|A' \times B'|})$.
 327
 328 • Remove $x + |H| - s$ pairs from H .
 329 • Sample x new edges uniformly from $(A' \times B') \setminus (A \times B)$ and add them to H .
 330

331 Since x has $n^{o(1)}$ expected value, with high probability (Markov inequality), x is $n^{o(1)}$, the difference
 332 between E and the new sample is $n^{o(1)}$, and the resampling process can be done in $n^{o(1)}$ time.
 333 We omitted the edge case where the size of H is less than $s - x$. Algorithm 8 in Section E.6 is a
 334 detailed version of this sublinear resampling scheme.
 335

336 **Dynamic update.** Combining the above, we can update the spectral sparsifier (see Section E.7 for
 337 details). When a point location update occurs, suppose point x_i is moved to z . We use the ultra
 338 low dimensional JL projection matrix to find the $O(k)$ -dimensional images of x_i and z . Then we
 339 update the $O(k)$ -dimensional WSPD. For each $O(k)$ -dimensional modified pair in the WSPD, we
 340 find the corresponding d -dimensional modified pairs, and resample edges from these d -dimensional
 341 modified pairs to update the spectral sparsifier. Since there are $2^{O(k)} \log \alpha$ modified pairs in each
 342 update and for each modified pair, with probability $1 - \delta$, the uniform sample can be updated in
 343 $O(\delta^{-1} \epsilon^{-2} n^{o(1)})$ time, the dynamic update can be completed in $O(\delta^{-1} \epsilon^{-2} n^{o(1)} \log \alpha)$ time per update.
 344

345 3.3 ADAPTIVE ADVERSARIAL UPDATES 346

347 The dynamic algorithm above is only able to handle oblivious updates. Recall the building blocks
 348 of the dynamic update algorithms. We compute the JL projection of the update points, update the
 349 WSPD for the low dimensional projections, and resample from the corresponding d -dimensional
 350 bicliques. Among these steps, the WSPD update algorithm is deterministic; the resampling algo-
 351 rithm uses fresh randomness for every round of updates. Therefore, the only building part that can
 352 be exploited by an adaptive adversary is the JL projection. Below in this overview, we explain how
 353 we achieve a JL distance estimation against adaptive adversaries. In this section, we provide an
 354 overview of techniques we use for adversarial analysis.
 355

356 3.3.1 ADVERSARIAL DISTANCE ESTIMATION

357 Let a random vector $V = (V_1, V_2, \dots, V_d) \in \mathbb{R}^d$ be sampled from Gaussian distribution and $U =$
 358 $\frac{1}{\|V\|}V$ be the normalized vector. Let vector $Z = (U_1, U_2, \dots, U_k) \in \mathbb{R}^k$ be the projection of U onto
 359 the first k components. From the properties of random variables sampled from Gaussian distribution,
 360 we can compute $\Pr[d(U_1^2 + \dots + U_k^2) \leq k\beta(U_1^2 + \dots + U_d^2)]$ via algebraic manipulations. Let
 361 $L = \|Z\|^2$. We show that when $\beta < 1$, we have $\Pr[L \leq \frac{\beta k}{d}] \leq \exp(\frac{k}{2}(1 - \beta + \ln \beta))$ and when
 362 $\beta \geq 1$, we have $\Pr[L \geq \frac{\beta k}{d}] \leq \exp(\frac{k}{2}(1 - \beta + \ln \beta))$. By carefully choosing $\beta = n^{-2c/k} < 1$, we
 363 can prove $\Pr[L \leq \frac{\beta k}{d}] \leq n^{-c}$. And when $\beta = n^{1/k}$, we can prove $\Pr[L \geq \frac{\beta k}{d}] \leq \exp(-\log^{1.9} n)$.
 364

365 With the above analysis in hand, we can prove that there exists a map $f : \mathbb{R}^d \rightarrow \mathbb{R}^k$ such that for
 366 each fixed points $u, v \in \mathbb{R}^d$, we have $\|u - v\|_2^2 \leq \|f(u) - f(v)\|_2^2 \leq \exp(c_0 \cdot \sqrt{\log n}) \|u - v\|_2^2$
 367 with high success probability. We design a ϵ_0 -net of $\{x \in \mathbb{R}^d \mid \|x\|_2 \leq 1\}$ denoted as N which
 368 contains $|N| \leq (10/\epsilon_0)^{O(\log n)}$ points (Here we assume $d = O(\log n)$). Then we prove that for all
 369 net points, the approximation guarantee still holds with high success probability via union bound.
 370 Finally, we want to generalize the distance estimation approximation guarantee to all points on the
 371 unit ball by quantizing the off-net point to its nearest on-net point. After rescaling the constant, we
 372 can obtain the same approximation guarantee with high probability.
 373

374 Given a set of data points $\{x_i\}_{i=1}^n$, and a sketching matrix $\Pi \in \mathbb{R}^{k \times d}$ defined in Definition H.1,
 375 we initialize a set of precomputed projected data points $\tilde{x}_i = \Pi \cdot x_i$. To answer the approximate
 376 distance between a query point and all points in the data structure, we compute the distance as
 377 $u_i = n^{1/k} \cdot \sqrt{d/k} \cdot \|\tilde{x}_i - \Pi q\|_2$ and prove it provides $\exp(\Theta(\sqrt{\log n}))$ -approximation guarantee
 378 against adversarially chosen queries. When we need to update the i -th data point with a new vector
 379 $z \in \mathbb{R}^d$, we update \tilde{x}_i with $\Pi \cdot x_i$.
 380

378
379

3.3.2 SPARSIFIER WITH ROBUSTNESS TO ADVERSARIAL UPDATES

380
381
382
383

With the estimation robust for adversarial query, we are able to get a spectral sparsifier which supports adversarial updates of points, by applying the data structure in the construction of sparsifier (Setting the sketching dimension to be $O(\sqrt{\log n})$). Here we provide overview of our design to make it possible.

384

385

Net argument. In order to make the distance estimation robust, one needs to argue that, for arbitrary point, it has high probability to have high precision. The data structure we use for distance estimation has a failure probability of n^{-c} , where c is a constant we can set to be small. We can build an ϵ -net N with size of $|N| = \text{poly}(n)$. Then by union bound over the net, the failure probability of distance estimation on the net is bounded by $n^{O(1)-c}$. Then by triangle inequality, we directly get the succeed probability guarantee for arbitrary point queries.

391

392

393

394

395

396

397

398

399

400

401

402

403

α and d induce the size of the net. From the discussion above, we note that, in order to make the ϵ -net sufficient for union bound, it must have the size of $\text{poly}(n)$. From another direction, we need to make that, all the points in the set are distinguishable in the nets, i.e., for two different points $A, B \in \mathbb{R}^d$, the closest points of the net to A and B are different. To make sure this, we must set the gap ϵ_1 of the net to be less than the minimum distance of the points in the set. Without loss of generality, we first make the assumption that, all the points are in the ℓ_2 unit ball of \mathbb{R}^d , i.e., the set $\{x \in \mathbb{R}^d \mid \|x\|_2 \leq 1\}$. Then by the definition of aspect ratio $\alpha := \frac{\max_{x, y \in P} d(x, y)}{\min_{x', y' \in P} d(x', y')}$, the minimum distance of the points in P is $1/\alpha$. Thus, when we set the gap $\epsilon_1 \leq C \cdot \alpha^{-1}$ for some constant C small enough, every pair of points $x, y \in P$ is distinguishable in the net. Then there are $O(\alpha^d)$ points in the net of the ℓ_2 unit ball in \mathbb{R}^d (See Figure 3).

404

405

406

407

408

409

Balancing the aspect ratio and dimension. By the above paragraph, we know the set size is $O(\alpha^d)$ to make the points distinguishable. Recall that, our distance estimation data structure has failure probability of n^{-c} . And in order to make the union bound sufficient for our net, we need to apply it over the $|N|^2$ pairs from N . That is, to make the total failure probability sufficient, we need to restrict $|N| = \text{poly}(n)$. And in the former paragraphs, we already know that $|N| = O(\alpha^d)$, thus we have the balancing constraint of the aspect ratio and dimension $\alpha^d = O(\text{poly}(n))$.

410

411

3.4 MAINTAINING A SKETCH OF AN APPROXIMATION TO LAPLACIAN MATRIX
MULTIPLICATION

412

413

414

Let M be an $n \times n$ matrix and x be a vector in \mathbb{R}^n . We say a vector b is an ϵ -approximation to Mx if $\|b - Mx\|_{M^\dagger} \leq \epsilon \|Mx\|_{M^\dagger}$. Note that $\|x\|_A := \sqrt{x^\top Ax}$. Let G be a graph and H be a ϵ -spectral sparsifier of G . By definition, this means $(1 - \epsilon)L_G \preceq L_H \preceq (1 + \epsilon)L_G$. Note that, if A is a symmetric PSD matrix and symmetric B is a matrix such that $(1 - \epsilon)A \preceq B \preceq (1 + \epsilon)A$, then we have $\|Bv - Av\|_{A^\dagger} \leq \epsilon \|Av\|_{A^\dagger}$ holds for all v . Then, we have: Let G be a graph on n vertices and H be a ϵ -spectral sparsifier of G . For any $v \in \mathbb{R}^n$, $L_H v$ is an ϵ -approximation of $L_G v$. Thus, to maintain a sketch of an ϵ -approximation of $L_G x$, it suffices to maintain a sketch of $L_H x$.

415

416

417

418

419

420

The high level idea is to combine the spectral sparsifier defined in Section E and a sketch matrix to compute a sketch of the multiplication result $L_H v$ and try to maintain this sketch when the graph and the vector change.

421

422

423

424

425

426

427

428

429

430

431

We here justify the decision of maintaining a sketch instead of the directly maintaining the multiplication result. Let the underlying geometry graph on n vertices be G and the vector be $v \in \mathbb{R}^n$. When a point is moved in the geometric graph, a column and a row are changed in L_G . We can assume the first row and first column are changed with no loss of generality. When this happens, if the first entry of v is not 0, all entries will change in the multiplication result. Therefore, it takes at least $\Omega(n)$ time to update the multiplication result. In order to spend subpolynomial time to maintain the multiplication result, we need to reduce the dimension of vectors. Therefore, we use a sketch matrix (with $m = \epsilon^{-2} \log(n/\delta)$ rows, see Lemma F.3 for details) to project vectors down to lower dimensions.

432 **Maintaining the multiplication result efficiently.** In order to speed up the update, we generate
 433 two independent sketches Φ and Ψ , and maintain a sketch of L_H , denoted by $\tilde{L}_H = \Phi L_H \Psi^\top$ and
 434 a sketch of v denoted by $\tilde{v} = \Psi v$. Since Φ and Ψ are generated independently, in expectation
 435 $\Phi L_H \Psi^\top \Psi v = \Phi L_H v$. We store this result as the sketch.

436 Our spectral sparsifier has the property that with high probability, each update to the geometric graph
 437 G incurs only a sparse changes in the sparsifier H , and this update can be computed efficiently.
 438 Therefore, when an update occurs to G , ΔL_H is sparse, so $\Phi \Delta L_H \Psi^\top$ can be computed efficiently.
 439 We use $\Phi \Delta L_H \Psi^\top$ to update the sketch. When a sparse update occurs to v , $\Psi \Delta v$ can be computed
 440 efficiently. Since \tilde{L}_H and $\Psi \Delta v$ are m -dimensional operator and vector, $\tilde{L}_H \Psi \Delta v$ can be computed
 441 efficiently. We use $\tilde{L}_H \Psi \Delta v$ to update the sketch.

442
 443 **3.5 MAINTAINING A SKETCH OF AN APPROXIMATION TO THE SOLUTION OF A LAPLACIAN
 444 SYSTEM**

445 We start with another folklore fact:

446 **Fact 3.1** (folklore). *If $(1-\epsilon)L_G \preceq L_H \preceq (1+\epsilon)L_G$, then we have $(1-2\epsilon)L_G^\dagger \preceq L_H^\dagger \preceq (1+2\epsilon)L_G^\dagger$.*

447 Let G be a graph on n vertices and H be a ϵ -spectral sparsifier of G . For any vector b , $L_H^\dagger b$ is an
 448 ϵ -approximation of $L_G^\dagger b$. Thus, to maintain a sketch of an ϵ -approximation of $L_G^\dagger x$, it suffices to
 449 maintain a sketch of $L_H^\dagger x$. The high level idea is again to combine the spectral sparsifier defined
 450 in Section E and a sketch matrix to compute a sketch of the multiplication result $L_H^\dagger v$ and try to
 451 maintain this sketch when the graph and the vector change.

452 **Caveat: using a different sketch.** When trying to maintain a sketch of a solution to $L_H x = b$,³
 453 the canonical way of doing this is to maintain \bar{x} such that $\Phi L_H \bar{x} = \Phi b$. However, here \bar{x} is still an
 454 n -dimensional vector and we want to maintain a sketch with lower dimension. Therefore, we apply
 455 another sketch Ψ to \bar{x} and maintain \tilde{x} such that $\Phi L_H \Psi^\top \tilde{x} = \Phi b$.

456 **Maintaining the inversion result efficiently.** We maintain a sketch of L_H , denoted by $\tilde{L}_H =$
 457 $\Phi L_H \Psi^\top$ and a sketch of b denoted by $\tilde{b} = \Phi b$. Since \tilde{L}_H is a m -dimensional operator, its pseudo-
 458 inverse can be computed efficiently in m^ω time, where ω is the matrix multiplication constant.
 459 We use \tilde{L}_H^\dagger to denote the pseudoinverse of \tilde{L}_H , and compute $\tilde{L}_H \cdot \tilde{b}$. We store this multiplication
 460 result as the sketch. Our spectral sparsifier has the property that with high probability, each update
 461 to the geometric graph G incurs only a sparse changes in the sparsifier H , and this update can be
 462 computed efficiently. Therefore, when an update occurs to G , ΔL_H is sparse, so $\Phi \Delta L_H \Psi^\top$ can be
 463 computed efficiently. We use $\Phi \Delta L_H \Psi^\top$ to update the \tilde{L}_H and recompute \tilde{L}_H^\dagger . We then update the
 464 sketch to $L_H^\dagger \cdot \tilde{b}$ with the updated L_H^\dagger . When a sparse update occurs to b , $\Phi \Delta b$ can be computed
 465 efficiently. Since \tilde{L}_H^\dagger and $\Phi \Delta b$ are m -dimensional operator and vector, $\tilde{L}_H^\dagger \Phi \Delta b$ can be computed
 466 efficiently. We use $\tilde{L}_H^\dagger \Psi \Delta b$ to update the sketch.

467
 468 **4 CONCLUSION**

469 In this work, we present dynamic algorithms for maintaining geometric graphs efficiently. Our main
 470 contributions include the introduction of the DYNAMICGEOSPAR data structure and techniques for
 471 handling adversarial queries and low-dimensional sketches with near-optimal initialization and up-
 472 date times, significantly improving existing methods. By combining spectral sparsification and
 473 Johnson-Lindenstrauss projections, we ensure efficient recomputation of graph structures with sparse
 474 changes. We prove that our data structure can dynamically maintain a $(1 \pm \epsilon)$ -spectral sparsifier with
 475 high probability, leverage JL projections to maintain low-dimensional sketches for efficient updates
 476 and queries, and design algorithms that are robust against adaptive adversarial queries. Our work
 477 has significant practical implications for real-time updates in geometric graphs.

478
 479 ³Although L_H is sparse, its pseudoinverse L_H^\dagger has no guarantee to be sparse (see bottom of page 113 in
 480 (Saad, 2003)).

486
487
ETHIC STATEMENT488
489
490
491
This paper does not involve human subjects, personally identifiable data, or sensitive applications.
We do not foresee direct ethical risks. We follow the ICLR Code of Ethics and affirm that all aspects
of this research comply with the principles of fairness, transparency, and integrity.492
493
REPRODUCIBILITY STATEMENT494
495
496
497
We ensure reproducibility of our theoretical results by including all formal assumptions, definitions,
and complete proofs in the appendix. The main text states each theorem clearly and refers to the
detailed proofs. No external data or software is required.498
499
REFERENCES500
501
502
503
504
Ittai Abraham, David Durfee, Ioannis Koutis, Sebastian Krininger, and Richard Peng. On fully
dynamic graph sparsifiers. In Irit Dinur (ed.), *IEEE 57th Annual Symposium on Foundations
of Computer Science, FOCS 2016, 9-11 October 2016, Hyatt Regency, New Brunswick, New
Jersey, USA*, pp. 335–344. IEEE Computer Society, 2016. doi: 10.1109/FOCS.2016.44. URL
<https://doi.org/10.1109/FOCS.2016.44>.505
506
507
Ahmed Alaoui and Michael W Mahoney. Fast randomized kernel ridge regression with statistical
guarantees. *Advances in Neural Information Processing Systems*, 28:775–783, 2015.508
509
Josh Alman and Zhao Song. Fast attention requires bounded entries. *arXiv preprint
arXiv:2302.13214*, 2023.510
511
512
513
Josh Alman, Timothy Chu, Aaron Schild, and Zhao Song. Algorithms and hardness for linear
algebra on geometric graphs. In *2020 IEEE 61st Annual Symposium on Foundations of Computer
Science (FOCS)*, pp. 541–552. IEEE, 2020.514
515
516
Haim Avron, Kenneth L Clarkson, and David P Woodruff. Sharper bounds for regularized data
fitting. *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Tech-
niques (Approx-Random)*, 2017a.517
518
519
520
Haim Avron, Michael Kapralov, Cameron Musco, Christopher Musco, Ameya Velingker, and Amir
Zandieh. Random fourier features for kernel ridge regression: Approximation bounds and statis-
tical guarantees. In *International Conference on Machine Learning*, pp. 253–262. PMLR, 2017b.521
522
523
Arturs Backurs, Moses Charikar, Piotr Indyk, and Paris Siminelakis. Efficient density evaluation
for smooth kernels. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science
(FOCS)*, pp. 615–626. IEEE, 2018.524
525
526
Rick Beatson and Leslie Greengard. A short course on fast multipole methods. *Wavelets, multilevel
methods and elliptic PDEs*, 1:1–37, 1997.527
528
529
530
Sabri Boughorbel, Jean-Philippe Tarel, François Fleuret, and Nozha Boujemaa. The gcs kernel
for svm-based image recognition. In *Artificial Neural Networks: Formal Models and Their
Applications-ICANN 2005: 15th International Conference, Warsaw, Poland, September 11-15,
2005. Proceedings, Part II 15*, pp. 595–600. Springer, 2005.531
532
Jan van den Brand, Binghui Peng, Zhao Song, and Omri Weinstein. Training (overparametrized)
neural networks in near-linear time. In *ITCS*. arXiv preprint arXiv:2006.11648, 2021.533
534
535
536
Jan van den Brand, Yu Gao, Arun Jambulapati, Yin Tat Lee, Yang P Liu, Richard Peng, and Aaron
Sidford. Faster maxflow via improved dynamic spectral vertex sparsifiers. In *Proceedings of the
54th Annual ACM SIGACT Symposium on Theory of Computing*, pp. 543–556, 2022.537
538
539
Paul B. Callahan and S. Rao Kosaraju. A decomposition of multidimensional point sets with appli-
cations to k -nearest-neighbors and n -body potential fields. *J. ACM*, 42(1):67–90, Jan 1995. ISSN
0004-5411. doi: 10.1145/200836.200853. URL <https://doi.org/10.1145/200836.200853>.

540 Moses Charikar and Paris Siminelakis. Hashing-based-estimators for kernel density in high dimen-
 541 sions. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pp.
 542 1032–1043. IEEE, 2017.

543 Michael B Cohen, Rasmus Kyng, Gary L Miller, Jakub W Pachocki, Richard Peng, Anup B Rao,
 544 and Shen Chen Xu. Solving sdd linear systems in nearly $m \log 1/2 n$ time. In *Proceedings of the*
 545 *forty-sixth annual ACM symposium on Theory of computing*, pp. 343–352, 2014.

546 Michael B Cohen, Jelani Nelson, and David P Woodruff. Optimal approximate matrix product in
 547 terms of stable rank. *arXiv preprint arXiv:1507.02268*, 2015.

548 Michael B Cohen, Yin Tat Lee, and Zhao Song. Solving linear programs in the current matrix
 549 multiplication time. In *STOC*, 2019.

550 Sanjoy Dasgupta and Anupam Gupta. An elementary proof of a theorem of johnson and linden-
 551 strauss. *Random Structures & Algorithms*, 22(1):60–65, 2003.

552 Yichuan Deng, Zhao Song, Lichen Zhang, and Ruizhe Zhang. Efficient algorithm for solving hy-
 553 perbolic programs. *arXiv preprint arXiv:2306.07587*, 2023.

554 David Durfee, Yu Gao, Gramoz Goranci, and Richard Peng. Fully dynamic spectral vertex sparsi-
 555 fiers and applications. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of*
 556 *Computing*, pp. 914–925, 2019.

557 John Fischer and Sariel Har-Peled. Dynamic well-separated pair decomposition made easy. In *17th*
 558 *Canadian Conference on Computational Geometry, CCCG 2005*, 2005.

559 François Fleuret, Hichem Sahbi, et al. Scale-invariance of support vector machines based on the
 560 triangular kernel. In *3rd International Workshop on Statistical and Computational Theories of*
 561 *Vision*, pp. 1–13. Citeseer, 2003.

562 Yu Gao, Yang P Liu, and Richard Peng. Fully dynamic electrical flows: Sparse maxflow faster
 563 than goldberg-rao. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science*
 564 (*FOCS*), pp. 516–527. IEEE, 2022.

565 Albert Gu, Karan Goel, and Christopher Ré. Efficiently modeling long sequences with struc-
 566 tured state spaces. *CoRR*, abs/2111.00396, 2021a. URL <https://arxiv.org/abs/2111.00396>.

567 Albert Gu, Isys Johnson, Karan Goel, Khaled Saab, Tri Dao, Atri Rudra, and Christopher Ré. Com-
 568 bining recurrent, convolutional, and continuous-time models with linear state-space layers, 2021b.
 569 URL <https://arxiv.org/abs/2110.13985>.

570 Yuzhou Gu and Zhao Song. A faster small treewidth sdp solver. *arXiv preprint arXiv:2211.06033*,
 571 2022.

572 Sariel Har-Peled. *Geometric approximation algorithms*. American Mathematical Soc., 2011. No.
 573 173.

574 Baihe Huang, Shunhua Jiang, Zhao Song, Runzhou Tao, and Ruizhe Zhang. Solving sdp faster: A
 575 robust ipm framework and efficient implementation, 2021.

576 Haotian Jiang, Yin Tat Lee, Zhao Song, and Lichen Zhang. Convex minimization with integer
 577 minima in $o(n^4)$ time. *arXiv preprint arXiv:2304.03426*, 2023.

578 Shunhua Jiang, Zhao Song, Omri Weinstein, and Hengjie Zhang. Faster dynamic matrix inverse for
 579 faster lps. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*
 580 (*STOC*), 2021.

581 Shunhua Jiang, Bento Natura, and Omri Weinstein. A faster interior-point method for sum-of-
 582 squares optimization, 2022.

583 William B Johnson and Joram Lindenstrauss. Extensions of lipschitz mappings into a hilbert space.
 584 *Contemporary mathematics*, 26(189–206):1, 1984.

594 Daniel M Kane and Jelani Nelson. Sparser johnson-lindenstrauss transforms. In *SODA*, pp. 1195.
 595 Society for Industrial and Applied Mathematics, 2012.

596

597 Jonathan A Kelner, Lorenzo Orecchia, Aaron Sidford, and Zeyuan Allen Zhu. A simple, combi-
 598 natorial algorithm for solving sdd systems in nearly-linear time. In *Proceedings of the forty-fifth*
 599 *annual ACM symposium on Theory of computing*, pp. 911–920, 2013.

600 Jakub Konečný, H Brendan McMahan, Felix X Yu, Peter Richtárik, Ananda Theertha Suresh, and
 601 Dave Bacon. Federated learning: Strategies for improving communication efficiency. *arXiv*
 602 *preprint arXiv:1610.05492*, 2016.

603

604 Ioannis Koutis, Gary L Miller, and Richard Peng. A nearly- $m \log n$ time solver for sdd linear
 605 systems. In *2011 IEEE 52nd Annual Symposium on Foundations of Computer Science*, pp. 590–
 606 598. IEEE, 2011.

607 Ioannis Koutis, Gary L Miller, and Richard Peng. Approaching optimality for solving sdd linear
 608 systems. *SIAM Journal on Computing*, 43(1):337–354, 2014.

609 Rasmus Kyng and Sushant Sachdeva. Approximate gaussian elimination for laplacians-fast, sparse,
 610 and simple. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*,
 611 pp. 573–582. IEEE, 2016.

612 Rasmus Kyng, Yin Tat Lee, Richard Peng, Sushant Sachdeva, and Daniel A Spielman. Sparsi-
 613 fied cholesky and multigrid solvers for connection laplacians. In *Proceedings of the forty-eighth*
 614 *annual ACM symposium on Theory of Computing*, pp. 842–850, 2016.

615 Jason D Lee, Ruoqi Shen, Zhao Song, Mengdi Wang, and Zheng Yu. Generalized leverage score
 616 sampling for neural networks. In *NeurIPS*, 2020.

617

618 Yin Tat Lee and Aaron Sidford. Efficient accelerated coordinate descent methods and faster al-
 619 gorithms for solving linear systems. In *2013 ieee 54th annual symposium on foundations of*
 620 *computer science*, pp. 147–156. IEEE, 2013.

621

622 Yin Tat Lee, Zhao Song, and Qiuyi Zhang. Solving empirical risk minimization in the current matrix
 623 multiplication time. In *Conference on Learning Theory*, pp. 2140–2157. PMLR, 2019.

624

625 Xuanqing Liu, Si Si, Xiaojin Zhu, Yang Li, and Cho-Jui Hsieh. A unified framework for data
 626 poisoning attack to graph-based semi-supervised learning. In *Advances in Neural Information*
 627 *Processing Systems (NeurIPS)*, 2019.

628

629 Per-Gunnar Martinsson. Encyclopedia entry on fast multipole methods. *University of Colorado at*
 630 *Boulder*, 1(5):9, 2012.

631

632 Charles A Micchelli. *Interpolation of scattered data: distance matrices and conditionally positive*
 633 *definite functions*. Springer, 1984.

634

635 Danupon Nanongkai, Thatchaphol Saranurak, and Christian Wulff-Nilsen. Dynamic minimum span-
 636 ning forest with subpolynomial worst-case update time. In *2017 IEEE 58th Annual Symposium*
 637 *on Foundations of Computer Science (FOCS)*, pp. 950–961. IEEE, 2017.

638

639 Jelani Nelson and Huy L Nguyêñ. Osnap: Faster numerical linear algebra algorithms via sparser
 640 subspace embeddings. In *FOCS*, pp. 117–126. IEEE, 2013.

641

642 Andrew Y Ng, Michael I Jordan, and Yair Weiss. On spectral clustering: Analysis and an algorithm.
 643 In *Advances in neural information processing systems (NeurIPS)*, pp. 849–856, 2002.

644

645 Lianke Qin, Zhao Song, Lichen Zhang, and Danyang Zhuo. An online and unified algorithm for
 646 projection matrix vector multiplication with application to empirical risk minimization. In *Inter-*
 647 *national Conference on Artificial Intelligence and Statistics*, pp. 101–156. PMLR, 2023.

648

649 Kent Quanrud. Spectral sparsification of metrics and kernels. In *Proceedings of the 2021 ACM-SIAM*
 650 *Symposium on Discrete Algorithms (SODA)*, pp. 1445–1464. SIAM, 2021.

651

652 Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. *Advances in*
 653 *neural information processing systems*, 20, 2007.

648 Carl Edward Rasmussen and Hannes Nickisch. Gaussian processes for machine learning (gpml)
 649 toolbox. *J. Mach. Learn. Res.*, 11:3011–3015, 2010. ISSN 1532-4435.
 650

651 Yousef Saad. *Iterative methods for sparse linear systems*. SIAM, 2003.
 652

653 Zhao Song, Shuo Yang, and Ruizhe Zhang. Does preprocessing help training over-parameterized
 654 neural networks? *Advances in Neural Information Processing Systems*, 34:22890–22904, 2021a.
 655

656 Zhao Song, Lichen Zhang, and Ruizhe Zhang. Training multi-layer over-parametrized neural net-
 657 work in subquadratic time. *arXiv preprint arXiv:2112.07628*, 2021b.
 658

659 César R Souza. Kernel functions for machine learning applications. *Creative commons attribution-
 660 noncommercial-share alike*, 3(29):1–1, 2010.
 661

662 Daniel A Spielman and Nikhil Srivastava. Graph sparsification by effective resistances. *SIAM
 663 Journal on Computing*, 40(6):1913–1926, 2011.
 664

665 Daniel A. Spielman and Shang-Hua Teng. Nearly-linear time algorithms for graph partitioning,
 666 graph sparsification, and solving linear systems. In *Proceedings of the Thirty-Sixth Annual ACM
 667 Symposium on Theory of Computing*, STOC ’04, pp. 81–90, New York, NY, USA, 2004. Asso-
 668 ciation for Computing Machinery. ISBN 1581138520. doi: 10.1145/1007352.1007372. URL
 669 <https://doi.org/10.1145/1007352.1007372>.
 670

671 Michele Trenti and Piet Hut. N-body simulations (gravitational). *Scholarpedia*, 3(5):3930, 2008.
 672

673 Ulrike von Luxburg. A tutorial on spectral clustering, 2007.
 674

675 David P Woodruff. Sketching as a tool for numerical linear algebra. *Foundations and Trends® in
 676 Theoretical Computer Science*, 10(1–2):1–157, 2014.
 677

678 Amir Zandieh, Insu Han, Majid Daliri, and Amin Karbasi. Kdeformer: Accelerating transformers
 679 via kernel density estimation. *arXiv preprint arXiv:2302.02451*, 2023.
 680

681 Xiaojin Zhu. *Semi-supervised learning with graphs*. PhD thesis, Carnegie Mellon University, lan-
 682 guage technologies institute, school of Computer Science, 2005a.
 683

684 Xiaojin Jerry Zhu. Semi-supervised learning literature survey. Technical report, University of
 685 Wisconsin-Madison Department of Computer Sciences, 2005b.
 686

687

688

689

690

691

692

693

694

695

696

697

698

699

700

701

702 Appendix

703
 704 **Roadmap.** We divide the appendix as follows. Section A provides related work of our paper.
 705 Section B gives the preliminary for our paper. Section C discusses the sketching techniques we use.
 706 Section D provide the full main algorithm. Section E gives the fully dynamic spectral sparsifier for
 707 geometric graphs. Section F gives our sketch data structure for matrix multiplication. Section G
 708 introduces the algorithm for solving Laplacian system. Section H introduces the distance estimation
 709 data structure supporting adversarial queries. Based on that, Section I gives our spectral sparsifier
 710 that is robust to adaptive adversary. Section J provides some figure to better explain this work.
 711 Section ?? discusses the limitation of this work. Section ?? provides a elaborate discussion about
 712 potential societal impact.

714 A RELATED WORK

715
 716 **Dynamic Sparsifier.** There has been some work focused on maintaining the dynamic sparsifier in
 717 a efficient time (Durfee et al., 2019). Their follow-up work (Gao et al., 2022; Brand et al., 2022)
 718 provides an algorithm for computing exact maximum flows on graphs with bounded integer edge
 719 capacities. Quanrud’s work on spectral sparsification of graphs with metrics and kernels (Quanrud,
 720 2021) provide efficient algorithm for constructing an sparsifier for the graphs.

721
 722 **Solvers of Laplacian System.** For a Laplacian linear system of a graph with m edges, it is a
 723 widely-studied problem (Spielman & Teng, 2004; Koutis et al., 2014; 2011; Kelner et al., 2013; Lee
 724 & Sidford, 2013; Cohen et al., 2014; Kyng et al., 2016; Kyng & Sachdeva, 2016). It has been shown
 725 that it can be solved in time $\tilde{O}(m \log(1/\epsilon))$. While the existing algorithm is very fast when the graph
 726 is sparse enough, we should focused on faster algorithms since our target graph might be dense.

727
 728 **Approximating the Kernel Density Function (KDF).** There are some works (Charikar &
 729 Siminelakis, 2017; Backurs et al., 2018) studying algorithms for kernel density function. (Charikar
 730 & Siminelakis, 2017) studied the Kernel Density Estimate (KDE) problem and they gave an effi-
 731 cient data structure such that, given a data set with a specific kernel function, it can approximates
 732 the kernel density of a query point in sublinear time. Later work (Backurs et al., 2018) presented
 733 a collection of algorithms for KDF approximating the “smooth” kernel functions. (Zandieh et al.,
 734 2023; Alman & Song, 2023) shows how to use kernel technique to compute attention matrix in large
 735 language models.

736
 737 **Kernel Functions.** Kernel method is a popular technique in data analysis and machine learning
 738 (Souza, 2010). The most popular and widely-used kernel functions are in the form of $K(x, y) =$
 739 $f(\|x - y\|_2)$. We list some of them here: the Gaussian kernel (Ng et al., 2002; Rahimi & Recht,
 740 2007), multiquadric kernel (Beatson & Greengard, 1997), circular kernel (Boughorbel et al., 2005),
 741 power kernel (Fleuret et al., 2003), log kernel (Beatson & Greengard, 1997; Martinsson, 2012) and
 742 inverse multiquadric kernel (Micchelli, 1984; Martinsson, 2012).

743
 744 **Dynamic Algorithms used in Optimization.** In addition, dynamic algorithms have been widely
 745 used in many of the optimization tasks. Usually, most optimization analysis are robust against
 746 noises and errors, such as linear programming (Cohen et al., 2019; Jiang et al., 2021), empirical risk
 747 minimization (Lee et al., 2019; Qin et al., 2023), semi-definite programming (Huang et al., 2021; Gu
 748 & Song, 2022), general programming (Deng et al., 2023), integral optimization (Jiang et al., 2023),
 749 training neural network (Brand et al., 2021; Song et al., 2021a;b), and sum of squares method (Jiang
 750 et al., 2022). Approximate solutions are sufficient to for these optimizations.

751 B PRELIMINARY

752 B.1 DEFINITIONS

753
 754 We define the (C, L) -Lipschitz function as follows:

756 **Definition B.1.** For $C \geq 1$ and $L \geq 1$, a function is (C, L) -Lipschitz if for all $c \in [1/C, C]$,

$$758 \quad 759 \quad 760 \quad \frac{1}{c^L} \leq \frac{f(cx)}{f(x)} \leq c^L.$$

761 We define the *Laplacian* of a graph:

762 **Definition B.2** (Laplacian of graph). Let $G = (V, E, w)$ be a connected weighted undirected graph
763 with n vertices and m edges, together with a positive weight function $w : E \rightarrow \mathbb{R}_+$. If we orient the
764 edges of G arbitrarily, we can write its Laplacian as

$$765 \quad 766 \quad L_G = A^\top W A,$$

767 where $A \in \mathbb{R}^{m \times n}$ is the signed edge-vertex incidence matrix, given by

$$768 \quad 769 \quad 770 \quad 771 \quad A(e, v) = \begin{cases} 1, & \text{if } v \text{ is the head of } e \\ -1 & \text{if } v \text{ is the tail of } e \\ 0 & \text{otherwise} \end{cases}$$

772 and $W \in \mathbb{R}_+^{m \times m}$ is the diagonal matrix such that $W(e, e) = w(e)$, for all $e \in E$. We use $\{a_e\}_{e \in E}$
773 to denote the row vectors of A .

774 It follows obviously that L_G is positive semidefinite since for any $x \in \mathbb{R}^n$,

$$775 \quad 776 \quad 777 \quad x^\top L_G x = x^\top A^\top W A x = \|W^{1/2} A x\|_2^2 \geq 0.$$

778 Since L_G is symmetric, we can diagonalize it and write

$$779 \quad 780 \quad 781 \quad 782 \quad L_G = \sum_{i=1}^{n-1} \lambda_i u_i u_i^\top,$$

783 where $\lambda_1, \dots, \lambda_{n-1}$ are the nonzero eigenvalues of L_G and u_1, \dots, u_{n-1} are the corresponding
784 orthonormal eigenvectors. The *Moore-Penrose Pseudoinverse* of L_G is

$$785 \quad 786 \quad 787 \quad 788 \quad L_G^\dagger := \sum_{i=1}^{n-1} \frac{1}{\lambda_i} u_i u_i^\top.$$

789 B.2 BASIC ALGEBRA

790 **Fact B.3** (Folklore). Let $\epsilon \in (0, 1/2)$. Given two positive semidefinite matrix $A \in \mathbb{R}^{n \times n}$ and
791 $B \in \mathbb{R}^{n \times n}$ such that

$$792 \quad (1 - \epsilon)A \preceq B \preceq (1 + \epsilon)A,$$

793 then we have:

- 794 • Part 1. $(1 + \epsilon)^{-1} A^\dagger \preceq B^\dagger \preceq (1 - \epsilon)^{-1} A^\dagger$.
- 795 • Part 2. $\|Bx - Ax\|_{A^\dagger} \leq \epsilon \|Ax\|_{A^\dagger}, \forall x \in \mathbb{R}^n$.

800 *Proof.* **Proof of Part 1.** The first statement follows from Fact B.4 directly.

801 **Proof of Part 2.**

$$802 \quad 803 \quad \|Bx - Ax\|_{A^\dagger}^2 = x^\top (B - A) A^\dagger (B - A) x$$

$$804 \quad 805 \quad 806 \quad 807 \quad \epsilon^2 \cdot \|Ax\|_{A^\dagger}^2 = \epsilon^2 \cdot x^\top A A^\dagger A x$$

808 It is obvious that

$$809 \quad -\epsilon A \preceq B - A \preceq \epsilon A$$

810 Thus by Fact B.5, we have
 811

$$812 (B - A)A^\dagger(B - A) \preceq \epsilon^2 AA^\dagger A.$$

813 Thus we have for any $x \in \mathbb{R}^n$,
 814

$$815 \|Bx - Ax\|_{A^\dagger}^2 = x^\top (B - A)A^\dagger(B - A)x \\ 816 \leq \epsilon^2 \cdot x^\top (B - A)A^\dagger(B - A)x \\ 817 = \epsilon^2 \cdot \|Ax\|_{A^\dagger}^2.$$

818 Thus we complete the proof. \square
 819

820 **Fact B.4.** *If $A \preceq B$, then $B^\dagger \preceq A^\dagger$.*
 821

823 *Proof.* We denote the SVD of A and B by $A = U_A \Sigma_A V_A^\top$ and $B = U_B \Sigma_B V_B^\top$, then we have for
 824 any $x \in \mathbb{R}^n$,

$$825 x^\top (B^\dagger - A^\dagger)x = x^\top (V_B^\top \Sigma_B^{-1} U_B - V_A^\top \Sigma_A^{-1} U_A)x \\ 826 = x^\top V_B^\top \Sigma_B^{-1} U_B x - x^\top V_A^\top \Sigma_A^{-1} U_A x \\ 827 = (x^\top U_B^\top \Sigma_B V_B x)^{-1} - (x^\top U_A^\top \Sigma_A V_A x)^{-1} \\ 828 = (x^\top Bx)^{-1} - (x^\top Ax)^{-1} \\ 829 \geq 0,$$

832 where the last step follows from $A \preceq B$. Thus we complete the proof. \square
 833

834 **Fact B.5.** *Let A, C denote two psd matrices. Let B be a symmetric matrix. Suppose $-C \preceq B \preceq C$,
 835 then we have*

$$836 B A B \preceq C A C$$

838 B.3 JOHNSON-LINDERSTRAUSS TRANSFORM

840 B.3.1 ULTRA-LOW DIMENSION JL

842 **Lemma B.6** (Ultra-low Dimensional Projection (Johnson & Lindenstrauss, 1984; Dasgupta &
 843 Gupta, 2003)). *For $k = o(\log n)$, with high probability at least $1 - 1/\text{poly}(n)$ the maximum dis-
 844 tortion in pairwise distance obtained from projecting n points into k dimensions (with appropriate
 845 scaling) is at most $n^{O(1/k)}$, e.g.,*

$$846 \|x - y\|_2 \leq \|f(x) - f(y)\|_2 \leq n^{O(1/k)} \cdot \|x - y\|_2$$

848 where f is the projection from \mathbb{R}^d to \mathbb{R}^k .
 849

850 Throughout this paper, we use C_{jl} to denote the constant on the exponent, i.e. the distortion is
 851 bounded above by $n^{C_{\text{jl}} \cdot (1/k)}$.

852 **Lemma B.7** (Johnson & Lindenstrauss, 1984; Dasgupta & Gupta, 2003)). *Let $k < d$. Let
 853 V_1, V_2, \dots, V_d be d independent Gaussian $N(0, 1)$ random variables, $V = (V_1, V_2, \dots, V_d)$, and
 854 let $U = \frac{1}{\|V\|} V$. Let the vector $Z = (U_1, U_2, \dots, U_k) \in \mathbb{R}^k$ be the projection of U onto the first k
 855 components and let $L = \|Z\|^2$ be the square of the norm of Z . Then*

- 856 • Part 1. If $\beta < 1$, then

$$858 \Pr[L \leq \frac{\beta k}{d}] \leq \beta^{k/2} \cdot (1 + \frac{(1 - \beta)k}{(d - k)})^{(d - k)/2} \leq \exp(\frac{k}{2}(1 - \beta + \ln \beta))$$

- 861 • Part 2. If $\beta > 1$, then

$$863 \Pr[L \geq \frac{\beta k}{d}] \leq \beta^{k/2} \cdot (1 + \frac{(1 - \beta)k}{(d - k)})^{(d - k)/2} \leq \exp(\frac{k}{2}(1 - \beta + \ln \beta))$$

864 B.3.2 USEFUL LEMMAS ON JL
865

866 Using Lemma B.7, we can show that

867 **Lemma B.8.** Let $\beta = n^{1/k}$, then we have $\Pr[L \geq \frac{\beta k}{d}] \leq \exp(-\frac{1}{4}kn^{1/k}) \leq \exp(-\log^{1.9} n)$.
868869 *Proof.* We show that
870

871
$$\begin{aligned} \Pr[L \geq \frac{\beta k}{d}] &\leq \exp\left(\frac{k}{2}(1 - \beta + \ln \beta)\right) \\ 872 &\leq \exp\left(-\frac{k}{2}\frac{\beta}{2}\right) \\ 873 &= \exp\left(-\frac{k}{2}\frac{n^{1/k}}{2}\right) \\ 874 &\leq \exp(-\log^{1.9} n) \end{aligned}$$

875

876 where the first step follows from Lemma B.7, the second step follows from $\beta/2 \geq 1 + \ln \beta$, the third
877 step follows from $\beta = n^{1/k}$, and the last step follows from $n^{1/k} \geq \log^{1.9} n$. \square
878879 Using Lemma B.7 and choosing parameter β carefully, we can show that:
880881 **Lemma B.9.** Let $\beta = n^{-2c/k}/e = 2^{-2c(\log n)/k}/e < 1$, then we have $\Pr[L \leq \frac{\beta k}{d}] \leq n^{-c}$.
882883 *Proof.* We show that
884

885
$$\begin{aligned} \Pr[L \leq \frac{\beta k}{d}] &\leq \beta^{k/2} \cdot \left(1 + \frac{(1 - \beta)k}{(d - k)}\right)^{(d - k)/2} \\ 886 &= \beta^{k/2} \cdot \left(1 + \frac{(1 - \beta)k}{(d - k)}\right)^{\frac{(d - k)}{k(1 - \beta)} \cdot \frac{k(1 - \beta)}{2}} \\ 887 &\leq \beta^{k/2} \cdot e^{k/2} \\ 888 &\leq (\beta e)^{k/2} \\ 889 &\leq n^{-c} \end{aligned}$$

890 where the first step comes from Lemma B.7, the second step follows that $(d - k)/2 = \frac{(d - k)}{k(1 - \beta)} \cdot$
891 $\frac{k(1 - \beta)}{2}$, the third step follows $(1 + \frac{1}{a})^a = e$ and $\frac{k(1 - \beta)}{2} \leq \frac{k}{2}$, the fourth step simplifies the term, and
892 the last step follows from $\beta = n^{-2c/k}/e$.
893 \square 900
901 B.4 WELL SEPARATED PAIR DECOMPOSITION (WSPD)902 We assume that throughout the process, all points land in $[0, 1]^d$ and the *aspect ratio* of the point set
903 is at most α .
904905 **Definition B.10.** The aspect ratio (α) of a point set P is
906

907
$$\alpha := \frac{\max_{x, y \in P} d(x, y)}{\min_{x', y' \in P} d(x', y')}.$$

908

909 We state several standard definitions from literature (Callahan & Kosaraju, 1995).
910911 **Definition B.11** (Bounding rectangle). Let $P \subset \mathbb{R}^d$ be a set of points, we define the bounding
912 rectangle of P , denoted as $R(P)$, to be the smallest rectangle in \mathbb{R}^d such that encloses all points
913 in P , where “rectangle” means some cartesian product $[x_1, x'_1] \times [x_2, x'_2] \times \cdots \times [x_d, x'_d] \in \mathbb{R}^n$.
914 For all $i \in [d]$, We define the length of R in i -th dimension by $l_i(R) := x'_i - x_i$. We denote
915 $l_{\max}(R) := \max_{i \in [d]} l_i(R)$ and $l_{\min}(R) := \min_{i \in [d]} l_i(R)$. When $l_i(R)$ are all equal for $i \in [d]$,
916 we say R is a d -cube, and denote its length by $l(R)$. For any set of points $P \subseteq \mathbb{R}^d$, we denote
917 $l_i(P) = l_i(R(P))$.
918

918 **Definition B.12** (Well separated point sets). *Point sets P, Q are well separated with separation s if
 919 $R(P)$ and $R(Q)$ can be contained in two balls of radius r , and the distance between these two balls
 920 is at least $s \cdot r$, where we say s is the separation.*

921 **Definition B.13** (Interaction product). *The interaction product of point sets P, Q , denoted by $P \otimes Q$
 922 is defined as*

$$923 \quad P \otimes Q := \{\{p, q\} \mid p \in P, q \in Q, p \neq q\}$$

924 **Definition B.14** (Well separated realization). *Let $P, Q \subseteq \mathbb{R}^d$ be two sets of points. A well separated
 925 realization of $P \otimes Q$ is a set $\{(P_1, Q_1), \dots, (P_k, Q_k)\}$ such that*

- 927 1. $P_i \subset P, Q_i \subset Q$ for all $i \in [k]$.
- 928 2. $P_i \cap Q_i = \emptyset$ for all $i \in [k]$.
- 929 3. $P \otimes Q = \bigcup_{i=1}^k P_i \otimes Q_i$.
- 930 4. P_i and Q_i are well-separated.
- 931 5. $(P_i \otimes Q_i) \cap (P_j \otimes Q_j) = \emptyset$ for $i \neq j$.

932 Throughout the paper, we will mention that a set P is *associated with* a binary tree T . Here we
 933 mean the tree T has leaves labeled by a set containing only one point which is in P . All the non-leaf
 934 nodes are labeled by the union of the sets labeled with its subtree.

935 Given set $P \subseteq \mathbb{R}^d$, let T be a binary tree associated with P . For $A, B \subseteq P$, we say that a realization
 936 of $A \otimes B$ uses T if all the A_i and B_i in the realization are nodes in T .

937 **Definition B.15** (Well separated pair decomposition). *A well separated pair decomposition (WSPD)
 938 of a point set P is a structure consisting of a binary tree T associated with P and a well separated
 939 realization of $P \otimes P$ uses T .*

940 The result of (Callahan & Kosaraju, 1995; Har-Peled, 2011) states that for a point set P of n points,
 941 a well separated pair decomposition of P of $O(n)$ pairs can be computed in $O(n \log n)$ time. There
 942 are two steps of computing a well separated decomposition: (1) build compressed quad tree (de-
 943 fined below in Definition B.17) for the given point set; (2) find well separated pairs from the tree
 944 (Algorithm 1).

945 **Definition B.16** (Quad tree). *Given a point set $P \subset [0, 1]^d$, a tree structure \mathcal{T} can be constructed
 946 in the following way:*

- 947 • The root of \mathcal{T} is the region $[0, 1]^d$
- 948 • For each tree node $n \in \mathcal{T}$, we can obtain 2^d subregions by equally dividing n into two
 949 halves along each of the d axes. The children of n in \mathcal{T} are the subregions that contain
 950 points in P . n has at most 2^d children.
- 951 • The dividing stops when there is only one point in the cell.

952 In a quad tree, we define the **degree** of a tree node to be the number of children it has. There can be
 953 a lot of nodes in T that has degree 1. Particularly, there can be a path of degree one nodes. Every
 954 node on this path contain the same point set. To reduce the size of the quad tree, we compress these
 955 degree one paths.

956 **Definition B.17** (Compressed quad tree). *Given a quad tree T , for each a path of degree one nodes,
 957 we replace it with the first and last nodes on the path, with one edge between them. We call this
 958 resulting tree a compressed quad tree.*

959 **Lemma B.18** (Chapter 2 in (Har-Peled, 2011)). *The compressed quad tree data structure T has the
 960 following properties:*

- 961 • Given a point p , p exists in at most $O(\log \alpha)$ quad tree nodes.
- 962 • The height of the tree is $O(\min(n, \log \alpha))$.

963 *T supports the following operations:*

972 • $\text{QTFASTPL}(T, p)$ returns the leaf node containing p , or the parent node under which $\{p\}$
 973 should be inserted if p does not exist in T , in $O(\log n)$ time.
 974
 975 • $\text{QTINSERTP}(T, p)$ adds p to the T in $O(\log n)$ time.
 976
 977 • $\text{QTDELETEP}(T, p)$ removes p from T in $O(\log n)$ time.

978 **Lemma B.19** (Theorem 2.2.3 in (Har-Peled, 2011)). *Given a d -dimensional point set P of size n , a
 979 compressed quad tree of P can be constructed in $O(dn \log n)$ time.*

980
 981 **Algorithm 1** Finding well separated pairs

982 1: **procedure** $\text{WSPD}(Q, u, v)$ ▷ Q is a compressed quad tree, u, v are tree nodes on Q ,
 983 Lemma B.21
 984 2: **if** u and v are well separated **then**
 985 3: **return** (u, v)
 986 4: **else**
 987 5: **if** $l_{\max}(u) > l_{\max}(v)$ **then**
 988 6: Let u_1, \dots, u_m denote the children of u
 989 7: **return** $\bigcup_i \text{WSPD}(Q, u_i, v)$
 990 8: **else**
 991 9: Let v_1, \dots, v_m denote the children of v
 992 10: **return** $\bigcup_i \text{WSPD}(Q, u, v_i)$
 993 11: **end if**
 994 12: **end if**
 995 13: **end procedure**
 14:
 996 15: **procedure** $\text{COMPUTEWSPD}(Q)$ ▷ Q is a compressed quad tree
 997 16: $r \leftarrow \text{root of } Q$
 998 17: **return** $\text{WSPD}(Q, r, r)$
 999 18: **end procedure**

1000
 1001 **Theorem B.20** ((Callahan & Kosaraju, 1995)). *For point set $P \subseteq \mathbb{R}^d$ of size n and $s > 1$, a s -WSPD
 1002 of size $O(s^d n)$ can be found in $O(s^d n + n \log n)$ time and each point is in at most $2^{O(d)} \log \alpha$ pairs.*

1003
 1004 **Lemma B.21** ((Callahan & Kosaraju, 1995; Har-Peled, 2011)). *Given a compressed quad tree Q
 1005 of n points in \mathbb{R}^d , two nodes u, v in the tree Q . The procedure COMPUTEWSPD (Algorithm 1)
 1006 generates the WSPD from the tree Q , and runs in time*

$$O(2^d \cdot n \log n).$$

1008
 1009 B.5 PROPERTIES OF (C, L) -LIPSCHITZ FUNCTIONS

1010
 1011 **Lemma B.22** (Lemma 6.8 in (Alman et al., 2020)). *Let G be a graph and G' be another graph on
 1012 the same set of vertices with different edge weights satisfying*

1013
 1014 $\frac{1}{K} w_{G'}(e) \leq w_G(e) \leq K w_{G'}(e)$
 1015

1016 Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be a (C, L) -lipschitz kernel function (Definition B.1), for some $C < K$. Let $f(G)$ be
 1017 the graph obtained by switching each edge weight from $w(e)$ to $f(w(e))$. Then,

1018
 1019 $\frac{1}{K^{2L}} w_{G'}(e) \leq w_G(e) \leq K^{2L} w_{G'}(e).$
 1020

1021 B.6 LEVERAGE SCORE AND EFFECTIVE RESISTANCE

1022
 1023 **Definition B.23.** *Given a matrix $A \in \mathbb{R}^{m \times n}$, we define $\sigma \in \mathbb{R}^m$ to denote the leverage score of A ,
 1024 i.e.,*

1025 $\sigma_i = a_i^\top (A^\top A) a_i, \forall i \in [m].$

1026 Let (G, V, E) be an graph obtained by arbitrarily orienting the edges of an undirected graph, with n
 1027 points and m edges, together with a weight function $w : E \rightarrow \mathbb{R}_+$. We now describe the electrical
 1028 flows on the graph. We let vector $I_{\text{ext}} \in \mathbb{R}^n$ denote the currents injected at the vertices. Let
 1029 $I_{\text{edge}} \in \mathbb{R}^m$ denote currents induced in the edges (in the direction of orientation) and $V \in \mathbb{R}^n$
 1030 denotes the potentials induced at the vertices. Let A, W be defined as Definition B.2. By Kirchoff's
 1031 current law, the sum of the currents entering a vertex is equal to the amount injected at the vertex,
 1032 i.e.,

$$A^\top I_{\text{edge}} = I_{\text{ext}}.$$

1033 By Ohm's law, the current flow in an edge is equal to the potential difference across its ends times
 1034 its conductance, i.e.,
 1035

$$I_{\text{edge}} = W A V.$$

1036 Combining the above, we have that

$$I_{\text{ext}} = A^\top (W A V) = L_G V.$$

1037 If $I_{\text{ext}} \perp \text{Span}(1_n) = \ker(K)$, that is, the total amount of current injected is equal to the total amount
 1038 extracted, then we have that

$$V = L_G^\dagger I_{\text{ext}}.$$

1039 **Definition B.24** (Leverage score of a edge in a graph). *We define the effective resistance or leverage
 1040 score between two vertices u and v to be the potential difference between them when a unit current
 1041 is injected at one that extracted at the other.*

1042 **Lemma B.25** (Algebraic form of leverage score, (Spielman & Srivastava, 2011)). *Let (G, E, V) be
 1043 a graph described as above, for any edge $e \in E$, the leverage score (effective resistance) of e has
 1044 the following form*

$$R(e) = A L_G^\dagger A^\top (e, e),$$

1045 where the matrix A, L_G is defined as Definition B.2.

1046 *Proof.* We now derive an algebraic expression for the effective resistance in terms of L_G^\dagger . For a edge
 1047 $e \in E$, we use $R(e)$ to denote its effective resistance. To inject and extract a unit current across the
 1048 endpoints of an edge (u, v) , we set $I_{\text{ext}} = a_e^\top$, which is clearly orthogonal to 1_n . The potentials
 1049 induced by I_{ext} at the vertices are given by $V = L_G^\dagger a_e^\top$. To measure the potential difference across
 1050 $e = (u, v)$, we simply multiply by a_e on the left:

$$V_v - V_u = (\mu_{n,v} - \mu_{n,u})^\top V = a_e L_G^\dagger a_e^\top.$$

1051 It follows that, the effective resistance across e is given by $a_e L_G^\dagger a_e^\top$ and that the matrix $A L_G^\dagger A^\top$ has
 1052 its diagonal entries $A L_G^\dagger A^\top (e, e) = R(e)$. \square

1053 B.7 SPECTRAL SPARSIFIER

1054 Here we give the formal definition of spectral sparsifier of a graph:

1055 **Definition B.26** (Spectral sparsifier). *Given an arbitrary undirected graph G , let L_G denote the
 1056 Laplacian (Definition B.2) of G . We say H is a ϵ -spectral sparsifier of G if*

$$(1 - \epsilon) L_G \preceq L_H \preceq (1 + \epsilon) L_G.$$

1057 C SKETCHING TECHNIQUES

1058 C.1 DEFINITIONS

1059 We first introduce the formal definition of the sparse embedding matrix:

1060 **Definition C.1** (Sparse Embedding matrix (Nelson & Nguyn, 2013)). *Let $h : [n] \times [s] \rightarrow [b/s]$
 1061 be a random 2-wise independent hash function and $\sigma : [n] \times [s] \rightarrow \{+1, -1\}$ be a random 4-wise
 1062 independent hash function. Then $R \in \mathbb{R}^{b \times n}$ is a sparse embedding matrix with parameter s if we
 1063 set $R_{(j-1)b/s+h(i,j),i} = \sigma(i, j)/s$ for all $(i, j) \in [n] \times [s]$ and all other entries to zero.*

1080 We also define the JL-moment property:
 1081

1082 **Definition C.2** (JL-moment property, Definition 12 in (Woodruff, 2014)). *We say a distribution \mathcal{D} on matrices $S \in \mathbb{R}^{k \times d}$ has the (ϵ, δ, ℓ) -JL moment property if that for all $x \in \mathbb{R}^d$ with $\|x\|_2 = 1$, it holds that*
 1083

$$1084 \mathbb{E}_{S \sim \mathcal{D}} [|\|Sx\|_2^2 - 1|^\ell] \leq \epsilon^\ell \cdot \delta.$$

1085 We give the formal definition of approximating matrix product:
 1086

1087 **Definition C.3** (Approximating matrix product (AMP) (Kane & Nelson, 2012; Woodruff, 2014)).
 1088 *Let $\epsilon \in (0, 1)$ be a precision parameter. Let $\delta \in (0, 1)$ be the failure probability. Given any two
 1089 matrix A, B each with n rows, we say a randomized matrix $R \in \mathbb{R}^{b \times n}$ from a distribution Π satisfies
 1090 (ϵ, δ) -approximate matrix product of A and B if*
 1091

$$1092 \Pr_{R \sim \Pi} [\|A^\top R^\top RB - A^\top B\|_F > \epsilon \cdot \|A\|_F \cdot \|B\|_F] \leq \delta.$$

1093 C.2 USEFUL RESULTS OF SPARSE EMBEDDING MATRIX

1094 Here in this section, we introduce the following technical theorem from literature, which gives the
 1095 concentration property of the sparse embedding matrices.

1096 **Lemma C.4** (Theorem 19 in (Kane & Nelson, 2012)⁴). *Let $\epsilon, \delta \in (0, 1)$ be two parameters. Let \mathcal{D} be a distribution over d columns that satisfies the (ϵ, δ, ℓ) -JL moment property for some $\ell \geq 2$. Then for two matrices A, B with n rows, it holds that*

$$1097 \Pr_{\Psi \sim \mathcal{D}} [\|A^\top \Psi^\top \Psi B - A^\top B\|_F > 3 \cdot \epsilon \cdot \|A\|_F \cdot \|B\|_F] \leq \delta.$$

1100 There is a result giving the JL-moment property of sparse embedding matrices in literature.

1101 **Lemma C.5** (Implicitly⁵ in (Cohen et al., 2015)). *The sparse embedding matrix (Definition C.1) with $m = O(\epsilon^{-2} \cdot \log(1/\delta))$ and $s = \Omega(\epsilon^{-1} \cdot \log(1/\delta))$ satisfies $(\epsilon, \delta, \log(1/\delta))$ -JL moment property.*

1102 Now we give the AMP property of the sparse embedding matrix.

1103 **Lemma C.6** (AMP of Sparse Embedding matrix). *Let $A \in \mathbb{R}^{n \times d_A}$ and $B \in \mathbb{R}^{n \times d_B}$ be two arbitrary matrices. Let $R \in \mathbb{R}^{m \times n}$ be a Sparse Embedding matrix as defined in Definition C.1 with $m = O(\epsilon^{-2} \cdot \log(1/\delta))$ and $s = \Omega(\epsilon^{-1} \cdot \log(1/\delta))$ non-zero entries of each column, then it satisfies (ϵ, δ) -AMP of A and B , and $A^\top R^\top RB$ can be computed in time*

$$1104 \quad s \cdot \text{nnz}(A) + s \cdot \text{nnz}(B) + \mathcal{T}_{\text{mat}}(d_A, m, d_B).$$

1105 *Proof.* By Lemma C.5, R satisfy $(\epsilon, \delta, \log(1/\delta))$ -JL moment property. Since $\log(1/\delta) > 2$ trivially holds, then by Lemma C.4, we proved the correctness of the lemma. It takes $s \cdot \text{nnz}(A)$ time to compute $A^\top R^\top$, $s \cdot \text{nnz}(B)$ time to compute RB , and $\mathcal{T}_{\text{mat}}(d_A, m, d_B)$ to compute $A^\top R^\top RB$. \square

1106 D ALGORITHM

1107 Here in this section, we give our main algorithm as follows. The main theorem of the algorithm and
 1108 the analysis with respect to the correctness and running time can be found in Appendix F.

1109 E FULLY DYNAMIC SPECTRAL SPARSIFIER FOR GEOMETRIC GRAPHS IN 1110 SUBLINEAR TIME

1111 A geometric graph w.r.t. kernel function $K : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ and points $x_1, \dots, x_n \in \mathbb{R}^d$ is a graph on
 1112 x_1, \dots, x_n where the weight of the edge between x_i and x_j is $K(x_i, x_j)$. An update to a geometric
 1113 graph occurs when the location of one of these points changes.

1114 ⁴For examples, see Theorem 17 in (Nelson & Nguyn, 2013) and Theorem 13 in (Woodruff, 2014)]

1115 ⁵See Remark 2 at page 9 of (Cohen et al., 2015)

1134 **Algorithm 2** Maintaining a sketch of an approximation to the solution to a Laplacian equation

1135 1: **data structure** SOLVE ▷ Theorem 2.4 and Theorem G.1

1136 2: **members**

1137 3: DYNAMICGEOSPAR dgs ▷ This is the sparsifier H

1138 4: $\Phi, \Psi \in \mathbb{R}^{m \times n}$: two independent sketching matrices

1139 5: $\tilde{L} \in \mathbb{R}^{m \times m}$ ▷ A sketch of L_H

1140 6: $\tilde{L}^\dagger \in \mathbb{R}^{m \times m}$ ▷ A sketch of L_H^\dagger

1141 7: $\tilde{b} \in \mathbb{R}^m$ ▷ A sketch of b

1142 8: $\tilde{z} \in \mathbb{R}^m$ ▷ A sketch of the multiplication result

1143 9: **EndMembers**

1144 10:

1145 11: **procedure** INIT($x_1, \dots, x_n \in \mathbb{R}^d, b \in \mathbb{R}^n$)

1146 12: Initialize Φ, Ψ

1147 13: dgs.INITIALIZE(x_1, \dots, x_n)

1148 14: $\tilde{b} \leftarrow \Phi b$

1149 15: $\tilde{L} \leftarrow \Phi \cdot \text{dgs.GETLAPLACIAN}() \cdot \Psi^\top$

1150 16: $\tilde{L}^\dagger \leftarrow \text{PSEUDOINVERSE}(\tilde{L})$

1151 17: $\tilde{z} \leftarrow \tilde{L}^\dagger \cdot \tilde{b}$

1152 18: **end procedure**

1153 19:

1154 20: **procedure** UPDATEG($x_i, z \in \mathbb{R}^d$)

1155 21: dgs.UPDATE(x_i, z)

1156 22: $\tilde{L} \leftarrow \tilde{L} + \Phi \cdot \text{dgs.GETDIFF}() \cdot \Psi^\top$

1157 23: $\tilde{L}_{\text{new}}^\dagger \leftarrow \text{PSEUDOINVERSE}(\tilde{L})$

1158 24: $\tilde{z} \leftarrow \tilde{L}_{\text{new}}^\dagger \cdot \tilde{b}$

1159 25: $L^\dagger \leftarrow \tilde{L}_{\text{new}}^\dagger$

1160 26: **end procedure**

1161 27:

1162 28: **procedure** UPDATEB($\Delta b \in \mathbb{R}^n$) ▷ Δb is sparse

1163 29: $\Delta \tilde{b} \leftarrow \Phi \cdot \Delta b$

1164 30: $\tilde{z} \leftarrow \tilde{z} + \tilde{L}^\dagger \cdot \Delta \tilde{b}$

1165 31: **end procedure**

1166 32:

1167 33: **procedure** QUERY

1168 34: **return** \tilde{z}

1169 35: **end procedure**

In a geometric graph, when an update occurs, the weights of $O(n)$ edges change. Therefore, directly applying the existing algorithms for dynamic spectral sparsifiers ((Abraham et al., 2016)) to update the geometric graph spectral sparsifier will take $\Omega(n)$ time per update. However by using the fact that the points are located in \mathbb{R}^d and exploiting the properties of the kernel function, we can achieve faster update.

Before presenting our dynamic data structure, we first have a high level idea of the static construction of the geometric spectral sparsifier, which is presented in (Alman et al., 2020).

Building Blocks of the Sparsifier. In order to construct a spectral sparsifier more efficiently, one can partition the graph into several subgraphs such that the edge weights on each subgraph are close. On each of these subgraphs, leverage score sampling, which is introduced in (Spielman & Srivastava, 2011) and used for constructing sparsifiers, can be approximated by uniform sampling.

For a geometric graph built from a d -dimensional point set P , under the assumption that each edge weight is obtained from a (C, L) -Lipschitz kernel function (Definition B.1), each edge weight in the geometric graph is not distorted by a lot from the euclidean distance between the two points (Lemma B.22). Therefore, we can compute this partition efficiently by finding a well separated pair decomposition (WSPD, Definition B.15) of the given point set.

A s -WSPD of P is a collection of well separated (WS) pairs (A_i, B_i) such that for all $a \neq b \in P$, there is a pair (A, B) satisfying $a \in A, b \in B$, and the distance between A and B is at least s times the diameters of A and B (A and B are s -well separated). Therefore, the distance between A and B is a $(1 + 1/s)$ -multiplicative approximation of the distance between any point in A and any point in B and each WS pair in the WSPD can be viewed as a unweighted biclique (complete bipartite graph). On an unweighted biclique, uniform random sampling and leverage score sampling are equivalent. Therefore, a uniformly random sample of the biclique forms a spectral sparsifier of the biclique, and union of the sampled edges from all bicliques form a spectral sparsifier of the geometric graph.

However, the time needed for constructing a WSPD is exponentially dependent on the ambient dimension of the point set and thus WSPD cannot be computed efficiently when the dimension is high. To solve this problem, one can use the ultra low dimensional Johnson Lindenstrauss (JL) projection to project the point set down to $k = o(\log n)$ dimension such that with high probability the distance distortion (multiplicative difference between the distance between two points and the distance between their low dimensional images) between any pair of points is at most $n^{C_{\text{JL}} \cdot (1/k)}$, where C_{JL} is a constant. This distortion becomes an overestimation of the leverage score in the resulting biclique, and can be compensated by sampling $n^{C_{\text{JL}} \cdot (1/k)}$ edges. Then one can perform a 2-WSPD on the k -dimensional points. Since JL projection gives a bijection between the d -dimensional points and their k -dimensional images, a 2-WSPD of the k -dimensional point set gives us a canonical $(2 \cdot n^{C_{\text{JL}} \cdot (1/k)})$ -WSPD of the d -dimensional point set P . This $(2 \cdot n^{C_{\text{JL}} \cdot (1/k)})$ -WSPD of P is what we use to construct the sparsifier.

Dynamic Update of the Geometric Spectral Sparsifier. We present the following way to update the above sparsifier. In order to do this, we need to update the ultra low dimensional JL projection, the WSPD and the sampled edges from each biclique. In order to update JL projection for $O(n)$ updates, we initialize the JL projection matrix with $O(n)$ points so that with high probability, the distortion is small for $O(n)$ updates.

To update the WSPD, we note that each point appears only in $O(\log \alpha)$ WS pairs and we can find all these pairs in $2^{O(k)} \log \alpha$ time (Section E.4).

To update the sampled edges, Algorithm 8 updates the old sample to a new one such that with high probability, the number of edges changed in the sample is at most $n^{o(1)}$ and this can be done in $n^{o(1)}$ time (Section E.6).

Combining the above, we can update the spectral sparsifier (Section E.7).

Below is the layout of this section. In Section E.1, we provide some definitions. In Section E.2, we define the members of our data structure. In Section E.3, we present the algorithm for initialization. In Section E.4 we state an algorithm to find modified pairs in WSPD when a point's location is changed. In Section E.5, we first propose a (slow) resampling algorithm takes $O(n)$ time to resample $n^{o(1)}$ edges. In Section E.6, we then explain how to improve the running time of (slow) resampling algorithm. In Section E.7, we prove the correctness of our update procedure. In Section E.8, we apply a black box reduction to our update algorithm to obtain a fully dynamic update algorithm.

E.1 DEFINITIONS

We define our problem as follows:

Definition E.1 (Restatement of Definition 1.1). *Given a set of points $P \subset \mathbb{R}^d$ and kernel function $K : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$. Let G denote the geometric graph that is corresponding to P with the (i, j) edge weight is $w_{i,j} := K(x_i, x_j)$. Let $L_{G,P}$ denote the Laplacian matrix of graph G . Let $\epsilon \in (0, 0.1)$ denote an accuracy parameter. The goal is to design a data structure that dynamically maintain a $(1 \pm \epsilon)$ -spectral sparsifier for G and supports the following operations:*

- **INITIALIZE**($P \subset \mathbb{R}^d, \epsilon \in (0, 0.1)$), this operation takes point set P and constructs a $(1 \pm \epsilon)$ -spectral sparsifier of L_G .
- **UPDATE**($i \in [n], z \in \mathbb{R}^d$), this operation takes a vector z as input, and to replace x_i (in point set P) by z , in the meanwhile, we want to spend a small amount of time and a small number of changes to spectral sparsifier so that

1242 **Definition E.2** (Restatement of Definition B.10). *Given a set of points $P = \{x_1, \dots, x_n\} \subset \mathbb{R}^d$.
1243 We define the aspect ratio α of P to be*

1244

$$\alpha := \frac{\max_{i,j} \|x_i - x_j\|_2}{\min_{i,j} \|x_i - x_j\|_2}.$$

1245

1246 The main result we want to prove in this section is

1247 **Theorem E.3** (Formal version of Theorem 2.1). *Let α be the aspect ratio of a d -dimensional point
1248 set P defined above. Let $k = o(\log n)$. There exists a data structure DYNAMICGOSPAR that
1249 maintains a ϵ -spectral sparsifier of size $O(n^{1+o(1)})$ for a (C, L) -Lipschitz geometric graph such
1250 that*

1251

- DYNAMICGOSPAR can be initialized in

1252

$$O(ndk + \epsilon^{-2} n^{1+o(L/k)} \log n \log \alpha)$$

1253

time.

1254

- DYNAMICGOSPAR can handle point location changes. For each change in point location,
1255 the spectral sparsifier can be updated in

1256

$$O(dk + 2^{O(k)} \epsilon^{-2} n^{o(1)} \log \alpha)$$

1257

time. With high probability, the number of edges changed in the sparsifier is at most

1258

$$\epsilon^{-2} 2^{O(k)} n^{o(1)} \log \alpha.$$

1259

1260 E.2 THE GEOMETRIC GRAPH SPECTRAL SPARSIFICATION DATA STRUCTURE

1261 In the following definition, we formally define the members we maintain in the data structure.

1262 **Definition E.4.** *In DYNAMICGOSPAR, we maintain the following objects:*

1263

- P : a set of points in \mathbb{R}^d
- \mathcal{H} : an $n^{1+o(1)}$ size ϵ -spectral sparsifier of the geometric graph generated by kernel K and
1264 points P
- Π : a JL projection matrix
- Q : the image of P after applying projection Π
- T : a quad tree of point set Q
- \mathcal{P} : a WSPD for point set Q obtained from P
- EDGES: a set of tuples (A_i, B_i, E_i) . E_i is a set of edges uniformly sampled from
1265 $\text{Biclique}(X_i, Y_i)$, where X_i and Y_i are the d -dimensional point sets corresponding to A_i
1266 and B_i respectively

1267 E.3 INITIALIZATION

1268 Here in this section, we assume that kernel function $K(x, y) = f(\|x - y\|_2^2)$ is (C, L) -Lipschitz
1269 (Definition B.1).

1270 **Lemma E.5.** *Let $\alpha \geq 0$ be defined as Definition B.10. INITIALIZE($P \subset \mathbb{R}^d, \epsilon \in (0, 0.1), \delta \in (0, 0.1), K$) (Algorithm 4) takes a d -dimensional point set P as inputs and runs in*

1271

$$O(ndk + \epsilon^{-2} n^{1+O(L/k)} 2^{O(k)} \log n \log \alpha)$$

1272

time, where k is the JL dimension, $k = o(\log n)$.

1273 *Proof.* The running time consists of the following parts:

1296

Algorithm 3 Data Structure

```

1: data structure DYNAMICGEOSPAR                                ▷ Theorem E.3
2: members                                                 ▷ Definition E.4
3:    $\mathcal{H}$                                                  ▷ An  $n^{1+o(1)}$  size sparsifier
4:    $P \subset \mathbb{R}^d$                                          ▷ A point set for the geometric graph
5:    $\Pi \in \mathbb{R}^{k \times d}$                                      ▷ Projection matrix
6:    $Q \subset \mathbb{R}^k$  ▷ A set of  $k = o(\log n)$  dimensional points obtained by applying  $\Pi$  to all points in
1303    $P$ 
7:    $T$                                                  ▷ A quad tree generated from  $P'$ 
1305    $\mathcal{P} = \{(A_i, B_i)\}_{i=1}^m$                                 ▷ A WSPD of  $P$  based on  $T$ 
1306    $\text{EDGES} = \{(A_i, B_i, E_i)\}_{i=1}^m$  ▷  $E_i$  is a set of edges sampled from biclique  $(X_i, Y_i)$ , where
1307      $X_i$  and  $Y_i$  are the  $d$ -dimensional point sets
10: end members
11: end data structure

```

1310

1311

Algorithm 4 DYNAMICGEOSPAR

```

1: data structure DYNAMICGEOSPAR                                ▷ Theorem E.3
2: procedure INITIALIZE( $P, \epsilon, \delta, K$ )                      ▷ Lemma E.5
3:    $P \leftarrow P$ 
4:    $\Pi \leftarrow$  a random  $(k \times d)$  JL-matrix                                ▷ Lemma B.6
5:    $Q \leftarrow \{\Pi \cdot p \mid p \in P\}$ 
6:    $T \leftarrow$  build a compressed quad tree for  $Q$                                 ▷ Lemma B.19
7:    $\mathcal{P} \leftarrow \text{WSPD}(T, \text{root}(T), \text{root}(T))$                                 ▷ Algorithm 1
8:    $\text{EDGES}, \mathcal{H} \leftarrow \text{INITSPARSIFIER}(\mathcal{P}, K, \epsilon, k)$                                 ▷ Algorithm 5
9: end procedure
10: end data structure

```

1322

1323
1324

- Line 4 and Line 5 takes time $O(ndk)$ to Generate the projection matrix and compute the projected sketch;

1325
1326

- By Lemma B.19, Line 6 takes time

1327

$$O(nk \log n);$$

1328

to build the quad tree.

1329

- By Lemma B.21, Line 7 takes time

1331

$$O(n \times 2^k \log n)$$

1332

to generate the WSPD.

1333

- By Lemma E.7, Line 8 takes time

1336

$$O(\epsilon^{-2} n^{1+O(L/k)} 2^{o(k)} \log n \log \alpha)$$

1337

to generate the sparsifier.

1338

Adding them together we have the total running time is

1340

$$O(ndk + \epsilon^{-2} n^{1+O(L/k)} 2^{o(k)} \log n \log \alpha).$$

1342

Thus we complete the proof. □

1343

1344

1345

We here state a trivial fact of sampling edges from a graph.

1346

1347

1348

1349

Fact E.6 (Random sample from a graph). *For any graph G and a positive integer $s \in \mathbb{Z}_+$, there exists a random algorithm $\text{RANDSAMPLE}(G, s)$ such that, it takes G and s as inputs, and outputs a set containing s edges which are uniformly sampled from G without replacement. This algorithm runs in time $O(s)$.*

Now we are able to introduce the initialization algorithm for the sparsifier.

Lemma E.7. *The procedure INITSPARSIFIER (Algorithm 5) takes $\mathcal{P}, K, \epsilon, k$ as input, where \mathcal{P} is a WSPD of the JL projection of point set P , K is a (C, L) -Lipschitz kernel function, $k = o(\log n)$ and ϵ is an error parameter; runs in time*

$$O(\epsilon^{-2} n^{1+O(L/k)} 2^{O(k)} \log n \log \alpha)$$

and outputs `EDGES`, \mathcal{H} , such that

- EDGES is the set of tuples such that for each $(A_i, B_i, E_i) \in \text{EDGES}$, E_i is a set of edges sampled from $\text{Biclique}(A_i, B_i)$.
- \mathcal{H} is a $(1 \pm \epsilon)$ -spectral sparsifier of the K -graph based on P
- the size of \mathcal{H} is size $O(\epsilon^{-2} n^{1+O(L/k)})$

Proof. We divide the proof into the following paragraphs.

Correctness We view each well separated pair as a biclique. Since \mathcal{P} is a 2-WSPD on a JL projection of P of distortion at most $n^{O(1/k)}$, by Lemma B.6, for any WS pair (A, B) and its corresponding d -dimensional pair (X, Y) , we have that

$$\frac{\max_{x \in X, y \in Y} \|x - y\|_2}{\min_{x \in X, y \in Y} \|x - y\|_2} \leq 2 \cdot n^{O(1/k)}.$$

By Lemma B.22, it holds that

$$\frac{\max_{x \in X, y \in Y} \mathsf{K}(\|x - y\|_2)}{\min_{x \in X, y \in Y} \mathsf{K}(\|x - y\|_2)} \leq 2 \cdot n^{O(L/k)}.$$

By seeing the biclique as an unweighted graph where all edge weights are equal to the smallest edge weight, one can achieve a overestimation of the leverage score of each edge. For each edge, the leverage score (Definition B.24) is overestimated by at most

$$O(n^{O(L/k)}(|X| + |Y|)/(|X||Y|)).$$

Therefore, by uniformly sampling

$$s \equiv O(\epsilon^{-2} n^{O(L/k)} \cdot (|X| + |Y|) \cdot \log(|X| + |Y|))$$

edges from $\text{Biclique}(X, Y)$ and normalize the edge weights by $|X||Y|/s$, we obtained a ϵ -spectral sparsifier of $\text{Biclique}(X, Y)$.

Since \mathcal{H} is the union of the sampled edges over all bicliques, \mathcal{H} is a ϵ -spectral sparsifier of K_G of at most $\epsilon^{-2}n^{1+O(L/k)}$ edges. EDGES stores the sampled edges from each biclique by definition.

Running time Since each vertex appears in at most $2^{O(k)} \log \alpha$ different WS pairs (Theorem B.20), the total time needed for sampling is at most

$$\epsilon^{-2}2^{O(k)} \cdot n^{1+O(L/k)} \log n \log \alpha.$$

Thus we complete the proof. \square

E.4 FIND MODIFIED PAIRS

Algorithm 6 Find modified pairs

```

1: data structure DYNAMICGEOSPAR
2: procedure FINDMODIFIEDPAIRS( $T, \mathcal{P}, p, p'$ ) ▷ Theorem E.3
3:    $l_p \leftarrow \text{QTFASTPL}(T, p)$  ▷ move  $p$  to  $p'$ , Lemma E.8
4:    $l_{p'} \leftarrow \text{QTFASTPL}(T, p')$  ▷ Lemma B.18
5:    $\mathcal{S} \leftarrow \emptyset$  ▷ Lemma B.18
6:    $\mathcal{P}^{\text{new}} \leftarrow \mathcal{P}$ 
7:   for  $n \in$  quad tree nodes on the path from  $l_p$  to the quad tree root do
8:      $P_n \leftarrow \text{WFINDPAIRS}(\mathcal{P}, n)$  ▷ Lemma B.18
9:     for every pair  $(A, B) \in P_n$  do
10:       if  $A' = \{p\}$  or  $B' = \{p\}$  then
11:          $A', B' \leftarrow \emptyset$ 
12:       else
13:         Remove  $p$  from  $(A, B)$  and obtain  $(A', B')$ 
14:       end if
15:        $\mathcal{S} \leftarrow \mathcal{S} \cup \{(A, B, A', B')\}$ 
16:     end for
17:   end for
18:   for  $n \in$  quad tree nodes on the path from  $l_{p'}$  to the quad tree root do
19:      $P_n \leftarrow \text{WFINDPAIRS}(\mathcal{P}, n)$  ▷ Lemma B.18
20:     for every pair  $(A, B) \in P_n$  do
21:       Add  $p'$  to  $(A, B)$  and obtain  $(A', B')$ 
22:        $\mathcal{S} \leftarrow \mathcal{S} \cup \{(A, B, A', B')\}$ 
23:     end for
24:   end for
25:   for  $(A, B, A', B') \in \mathcal{S}$  do
26:      $\mathcal{P}^{\text{new}} \leftarrow \text{replace } (A, B) \in \mathcal{P} \text{ with } (A', B').$ 
27:   end for
28:    $T^{\text{new}} \leftarrow \text{QTINSERTP}(\text{QTDELETEP}(T, p), p')$  ▷ Lemma B.18
29:   return  $\mathcal{S}, T^{\text{new}}, \mathcal{P}^{\text{new}}$ 
30: end procedure
31: end data structure

```

WSPD is stored as a list of pairs \mathcal{P} that supports:

- $\text{WFINDPAIRS}(\mathcal{P}, A)$, find all pairs (A, B) and $(B, A) \in \mathcal{P}$ time linear in the output size.

Lemma E.8. *Given a compressed quad tree T of a $O(k)$ -dimensional point set P , a WSPD \mathcal{P} computed from T , a point $p \in P$ and another point p' , in the output of Algorithm 6, T^{new} is a quad tree T^{new} of $P \setminus \{p\} \cup \{p'\}$, \mathcal{P}^{new} is a WSPD of $P \setminus \{p\} \cup \{p'\}$ and \mathcal{S} is a collection of tuples (A, B, A', B') . \mathcal{P}^{new} can be obtained by doing the following:*

For all $(A, B, A', B') \in \mathcal{S}$, replace $(A, B) \in \mathcal{P}$ with (A', B') .

This can be done in $2^{O(k)} \log \alpha$ time.

1458 *Proof.* We divide the proof into the following parts.
 1459

1460 **Correctness** By Lemma B.18, we have that, the new generated tree T^{new} is a quad tree of $P \setminus \{p\} \cup$
 1461 $\{p'\}$.

1462 We now show that, after replacing $(A, B) \in \mathcal{P}$ with (A', B') for all (A, B, A', B') in \mathcal{S} in Line 26,
 1463 we get a WSPD of the updated point set.
 1464

1465 First in Line 3 and Line 3, we find the path from the root to the leaf node containing p and p' . Then
 1466 in the following two for-loops (Line 7 and Line 18), we iteratively visit the nodes on the paths. In
 1467 each iteration, we find the WS pairs related to the node by calling WFINDPAIRS. We record the
 1468 original sets and the updated sets. Then in Line 26, we replace the original pairs by the updated
 1469 pairs to get the up-to-date pair list.
 1470

1471 **Running time** By Lemma B.18, the two calls to QTFASTPL takes $O(\log n)$ time. For each of p
 1472 and p' , there are at most $2^{O(k)} \log n$ pairs that can contain p or p' . Therefore, the total running time
 1473 of WFINDPAIRS is $O(2^{O(k)}) \log n$, and there are $2^{O(k)} \log n$ tuples in \mathcal{S} . The number of times that
 1474 the loops on lines 9 and 20 are executed is at most $2^{O(k)} \log n$. Hence the total time complexity of
 1475 FINDMODIFIEDPAIRS is $2^{O(k)} \log n$.
 1476 \square
 1477

1478 E.5 LINEAR TIME RESAMPLING ALGORITHM

1479 Here in this section, we state our linear time resampling algorithm.
 1480

1481 **Algorithm 7** Linear Time Resampling Algorithm

1482 1: **procedure** RESAMPLE(E, A, B, A', B', s) ▷ Lemma E.9
 1483 2: $E \leftarrow E \cap (A' \times B')$
 1484 3: $\mathcal{R} \leftarrow \emptyset$
 1485 4: $q \leftarrow \frac{|(A \times B) \cap (A' \times B')|}{|A' \times B'|}$
 1486 5: **for** $j = 1 \rightarrow s$ **do**
 1487 6: Draw a random number x from $[0, 1]$
 1488 7: **if** $x \leq q$ **then**
 1489 8: **if** $E \setminus \mathcal{R} \neq \emptyset$ **then**
 1490 9: Sample one pair from E (without repetition) and add it to \mathcal{R}
 1491 10: **else** ▷ all points of E are sampled
 1492 11: Sample one pair from $((A \times B) \cap (A' \times B')) \setminus E$ (without repetition) and add it
 1493 to \mathcal{R}
 1494 12: **end if**
 1495 13: **else**
 1496 14: Sample one pair of points (a, b) from $(A' \times B') \setminus (A \times B)$ and add it to \mathcal{R}
 1497 15: **end if**
 1498 16: **end for**
 1499 17: **return** \mathcal{R}
 1500 18: **end procedure**
 1501

1502 **Lemma E.9** (Resample). *Let C_{j1} be the constant defined in Lemma B.6. Let V be a set, A, B be
 1503 subsets of V such that $A \cap B = \emptyset$, A', B' be two sets that are not necessarily subsets of V such that*

$$1504 \quad A' \cap B' = \emptyset \text{ and } |(A \times B) \Delta (A' \times B')| < o\left(\frac{|A' \times B'|}{|A'| + |B'|}\right).$$

1506 *Let $n = |V \cup A' \cup B'|$. Let E be a subset of $V \times V$.*

1508 *Let H be a graph on vertex set V , $A, B \subset V$ and $A \cap B = \emptyset$. Let A', B' be two other vertex sets
 1509 such that $A' \cap B' = \emptyset$ (A' and B' do not have to be subsets of V). If*

1510 • E is a uniform sample of size

$$1511 \quad \epsilon^{-2} n^{C_{j1} \cdot (L/k)} (|A| + |B|) \log(|A| + |B|)$$

1512 from $A \times B$.
 1513
 1514 • $s = \epsilon^{-2} n^{C_{jl} \cdot (L/k)} (|A'| + |B'|) \log(|A'| + |B'|)$
 1515 • $|s - |E|| = n^{o(1)}$
 1516

1517 then with high probability, RESAMPLE generates a uniform sample of size s from $A' \times B'$ in $n^{o(1)}$
 1518 time. Moreover, with probability at least $1 - \delta$, the size of difference between the new sample and E
 1519 is $n^{o(1)}$.
 1520

1521 *Proof.* To show that the sample is uniform, we can see this sampling process as follows: To draw s
 1522 samples from $A' \times B'$, the probability of each sample being drawn from $(A' \times B') \cap (A \times B)$ is

$$\frac{|(A' \times B') \cap (A \times B)|}{|A' \times B'|}.$$

1523
 1524
 1525
 1526 Therefore, for each sample, with this probability, we draw this sample from $(A' \times B') \cap (A \times B)$
 1527 (line 7) and sample from $(A' \times B') \setminus (A \times B)$ otherwise (line 10).

1528
 1529
 1530
 1531
 1532 Since E is a uniform sample from $A \times B$, $E \cap (A' \times B')$ is a uniform sample from $(A \times B) \cap (A' \times B')$
 1533 and any uniformly randomly chosen subset of it is also a uniform sample from $(A \times B) \cap (A' \times B')$.
 1534 Hence, to sample pairs from $(A \times B) \cap (A' \times B')$, we can sample from $E \cap (A \times B)$ first (line 9)
 1535 and sample from outside $E \cap (A \times B)$ when all pairs in $E \cap (A \times B)$ are sampled (line 11). The
 1536 resulting set is a uniform sample from $A' \times B'$.

1537 To see the size difference between E and \mathcal{R} , we note that since

$$|(A \times B) \Delta (A' \times B')| < o\left(\frac{|A' \times B'|}{|A'| + |B'|}\right),$$

1538 the probability

$$\begin{aligned} \frac{|(A \times B) \cap (A' \times B')|}{|A' \times B'|} &\geq 1 - \frac{|(A \times B) \Delta (A' \times B')|}{|A' \times B'|} \\ &\geq 1 - o\left(\frac{1}{|A'| + |B'|}\right) \end{aligned}$$

1539
 1540
 1541
 1542
 1543 Therefore, to draw $s = \epsilon^{-2} n^{C_{jl} \cdot (L/k)} (|A'| + |B'|) \log(|A'| + |B'|)$ samples from $A' \times B'$, the
 1544 expectation of number of samples drawn from $(A' \times B') \setminus (A \times B)$ is at most

$$\epsilon^{-2} n^{C_{jl} \cdot (L/k)} (|A'| + |B'|) \log(|A'| + |B'|) \cdot o\left(\frac{1}{|A'| + |B'|}\right) \leq \epsilon^{-2} n^{C_{jl} \cdot (L/k)} \log(|A'| + |B'|)$$

1545
 1546
 1547
 1548 By Markov inequality, with high probability $1 - \delta$, at most

$$\delta^{-1} \epsilon^{-2} n^{C_{jl} \cdot (L/k)} \log(|A'| + |B'|)$$

1549
 1550
 1551 pairs were drawn from $(A' \times B') \setminus (A \times B)$.

1552
 1553
 1554
 1555 Now we analyze the time complexity. The loop runs for $O(s)$ time and each sample can be done in
 1556 constant time. Therefore, the total time complexity is $O(s)$. By the third bullet point, this is in worst
 1557 case $O(n \log n)$ time. \square

1558 E.6 EFFICIENT SUBLINEAR TIME RESAMPLING ALGORITHM

1559
 1560
 1561
 1562 Algorithm 7 returns a set of pairs that is with high probability close to the input set E . However,
 1563 since it needs to sample all s pairs, the time complexity is bad. We modify it by trying to remove
 1564 samples from E instead of adding pairs from E to the new sample and obtain Algorithm 8.

1565
 1566 **Lemma E.10** (Fast resample). *Let C_{jl} be the constant defined in Lemma B.6. Let V be a set, A, B
 1567 be subsets of V such that $A \cap B = \emptyset$, A', B' be two sets that are not necessarily subsets of V such
 1568 that*

$$A' \cap B' = \emptyset \text{ and } |(A \times B) \Delta (A' \times B')| < o\left(\frac{|A' \times B'|}{|A'| + |B'|}\right).$$

1569
 1570
 1571 Let $n = |V \cup A' \cup B'|$. Let E be a subset of $V \times V$. If

1566

Algorithm 8 Efficient Sublinear Time Resampling Algorithm

```

1: procedure FASTRESAMPLE( $E, A, B, A', B', s$ ) ▷ Lemma E.10
2:    $\mathcal{R} \leftarrow \emptyset$ 
3:    $E \leftarrow E \cap (A' \times B')$ 
4:   Draw a random number  $x$  from Binomial( $s, \frac{|(A' \times B') \setminus (A \times B)|}{|A' \times B'|}$ )
5:   Add to  $\mathcal{R}$   $x$  points uniformly drawn from  $(A' \times B') \setminus (A \times B)$  (without repetition)
6:   if  $|E| > s - x$  then
7:     Draw  $x + |E| - s$  pairs from  $E$  uniformly randomly (without repetition)
8:     Add pairs in  $E$  to  $\mathcal{R}$  except these  $x + |E| - s$  pairs drawn on line 7
9:   else
10:    Add all pairs in  $E$  to  $\mathcal{R}$ 
11:    Add to  $\mathcal{R}$   $(s - x - |E|)$  points uniformly drawn from  $(A' \times B') \setminus (A \times B) \setminus E$  (without
repetition)
12:   end if
13:   return  $\mathcal{R}$ 
14: end procedure

```

1582

1583

- E is a uniform sample of size $\epsilon^{-2} n^{C_{\text{jl}} \cdot (L/k)} (|A| + |B|) \log(|A| + |B|)$ from $A \times B$.
- $s = \epsilon^{-2} n^{C_{\text{jl}} \cdot (L/k)} (|A'| + |B'|) \log(|A'| + |B'|)$
- $|E| < o(|A \times B|)$
- $s < o(|A' \times B'|)$
- $|s - |E|| = n^{o(1)}$

1592

1593

1594

then with high probability, FASTRESAMPLE generates a uniform sample of size s from $A' \times B'$ in $n^{o(1)}$ time. Moreover, with probability at least $1 - \delta$, the size of difference between the new sample and E is $n^{o(1)}$.

1595

1596

1597

Proof. To show that the sample is uniform, we can see this sampling process as follows: To draw s samples from $A' \times B'$, the probability of each sample being drawn from $(A' \times B') \setminus (A \times B)$ is

$$\frac{|(A' \times B') \setminus (A \times B)|}{|A' \times B'|}.$$

1600

1601

Therefore, the number of samples drawn from $(A' \times B') \setminus (A \times B)$ satisfies a binomial distribution with parameters

1602

1603

1604

$$s \text{ and } \frac{|(A' \times B') \setminus (A \times B)|}{|A' \times B'|}.$$

1605

1606

Let x be such a binomial random variable, we sample x pairs from $(A' \times B') \setminus (A \times B)$ and the rest from $(A' \times B') \cap (A \times B)$.

1607

1608

1609

1610

1611

Since E is a uniform sample from $A \times B$, $E \cap (A' \times B')$ is a uniform sample from $(A \times B) \cap (A' \times B')$, and any uniformly randomly chosen subset of it is also a uniform sample from $(A \times B) \cap (A' \times B')$. Hence, if $|E| > s - x$, we take $s - x$ pairs from E by discarding $|E| - s + x$ pairs in E (line 7 and 8). If $|E| \leq s - x$, we take all samples from E and add $s - x - |E|$ pairs from $(A' \times B') \cap (A \times B)$ (line 11).

1612

1613

Now we try to bound the difference between E and \mathcal{R} and the time complexity. Since x is drawn from a binomial distribution, we have that

1614

1615

1616

$$\mathbb{E}[x] = s \cdot \frac{|(A' \times B') \setminus (A \times B)|}{|A' \times B'|}.$$

1617

1618

By Markov inequality,

1619

$$\Pr\left[x > \frac{s}{\delta} \cdot \frac{|(A' \times B') \setminus (A \times B)|}{|A' \times B'|}\right] \leq \delta$$

1620 Since $|(A \times B) \Delta (A' \times B')| < o(\frac{|A' \times B'|}{|A'| + |B'|})$, with probability at least $1 - \delta$,
 1621
 1622 $x < \delta^{-1} \cdot \epsilon^{-2} \cdot n^{C_{j1} \cdot (L/k)} (|A'| + |B'|) \log(|A'| + |B'|) \cdot o(\frac{1}{|A'| + |B'|})$
 1623
 1624 $< \delta^{-1} \epsilon^{-2} n^{C_{j1} \cdot (L/k)} \log n$
 1625

1626 Therefore, drawing new samples (lines 7 and 8 or line 11) takes $O(x + |s - |E||)$ time. The difference
 1627 between E and the output sample set is also at most $O(x + |s - |E||) = n^{o(1)}$. With probability at
 1628 least $1 - \delta$, we have that

$$1629 \quad x \leq \delta^{-1} \cdot \epsilon^{-2} \cdot n^{o(1)}.$$

1630 Since $|s - |E|| \leq n^{o(1)}$, the overall time complexity and the difference between E and the output
 1631 set are at most $\delta^{-1} \cdot \epsilon^{-2} \cdot n^{o(1)}$.
 1632

1633 Thus we complete the proof. \square

1635 E.7 A DATA STRUCTURE THAT CAN HANDLE $O(n)$ UPDATES

1636 In this section, we combine the above algorithms and state our data structure.

1638 Algorithm 9 Data structure update

```

1: data structure DYNAMICGOSPAR
2: procedure UPDATE( $p \in \mathbb{R}^d, p' \in \mathbb{R}^d$ ) ▷ Theorem E.3
3:    $P^{\text{new}} \leftarrow P \cup p' \setminus p$  ▷ Lemma E.11
4:    $Q^{\text{new}} \leftarrow Q \cup (\Pi p') \setminus (\Pi p)$  ▷  $\Pi$  is a projection stored in memory and fixed over all the
iterations
5:    $\mathcal{S}, T^{\text{new}}, \mathcal{P}^{\text{new}} \leftarrow \text{FINDMODIFIEDPAIRS}(T, P, \Pi p, \Pi p')$  ▷ Algorithm 6
6:    $\mathcal{H}^{\text{new}} \leftarrow \mathcal{H}$ 
7:   for all  $(A, B, A', B') \in \mathcal{S}$  do
8:      $E \leftarrow \text{EDGES}(A, B)$ 
9:     Scale each edge in  $E$  by  $\epsilon^{-2} (n^{O(L/k)} (|A| + |B|) \log(|A| + |B|)) / |A||B|$ 
10:     $X, Y, X', Y' \leftarrow d$ -dimensional points corresponding to  $A, B, A', B'$ 
11:    if  $|(A \times B) \Delta (A' \times B')| < o(\frac{|A' \times B'|}{|A'| + |B'|})$  then
12:       $s \leftarrow \epsilon^{-2} n^{O(L/k)} (|X'| + |Y'|) \log(|X'| + |Y'|)$ 
13:       $E^{\text{new}} \leftarrow \text{FASTRESAMPLE}(E, X, Y, X', Y', s)$  ▷ Algorithm 8
14:      Scale each edge in  $E^{\text{new}}$  by  $|X'||Y'|/s$ 
15:    else
16:       $E^{\text{new}} \leftarrow$  all edges in  $\text{Biclique}(X', Y')$ 
17:    end if
18:     $\text{EDGES.UPDATE}(A, B, A', B', E, E^{\text{new}})$  ▷ Change  $(A, B, E)$  to  $(A', B', E^{\text{new}})$ 
19:     $\mathcal{H}^{\text{new}} \leftarrow \mathcal{H}^{\text{new}} \setminus E \cup E^{\text{new}}$ 
20:  end for
21:   $\mathcal{H} \leftarrow \mathcal{H}^{\text{new}}$ 
22:   $P \leftarrow P^{\text{new}}$ 
23:   $\mathcal{P} \leftarrow \mathcal{P}^{\text{new}}$ 
24:   $Q \leftarrow Q^{\text{new}}$ 
25:   $T \leftarrow T^{\text{new}}$ 
26: end procedure
27: end data structure

```

1667 **Lemma E.11.** *Given two points $p, p' \in \mathbb{R}^d$, with high probability $1 - \delta$, function UPDATE (Algorithm
 1668 9) can handle $O(n)$ updates to the geometric graph and can update the ϵ -spectral sparsifier in*

$$1669 \quad O(dk + \delta^{-1} \epsilon^{-2} n^{o(1)} \log \alpha)$$

1670 *time per update. Moreover, after each update, the number of edge weight that are changed in the
 1671 sparsifier is at most $\delta^{-1} \epsilon^{-2} n^{o(1)} \log \alpha$.*

1672 *Proof.* Similar to Lemma E.7, in order for the updated \mathcal{H} to be a spectral sparsifier of K_G , we need

- After removing Πp and adding $\Pi p'$, the resulting JL projection Q still has distortion at most $n^{1/k}$.
- The WSPD of Q is updated to a WSPD of $Q \setminus \{\Pi p\} \cup \{\Pi p'\}$
- For each WS pair (A', B') in the new WSPD, let X' and Y' be A' and B' 's corresponding d -dimensional point set respectively, we can obtain a uniform sample of

$$\epsilon^{-2}(n^{O(L/k)}(|X'| + |Y'|) \log(|X'| + |Y'|))$$

edges from $\text{Biclique}(X', Y')$.

For each of the above requirement, we divide the proof into the following paragraphs.

Bounding on the distortion of JL distance To show the first requirement, we note that by Lemma B.6, if the JL projection matrix is initialized with $O(n)$ points, after at most $O(n)$ updates, with high probability, the distance distortion between two points is still bounded above by $n^{O(1/k)}$.

Update to WSPD To show the second requirement, `FINDMODIFIEDPAIRS` returns a collections \mathcal{S} of pairs updates. By Lemma E.8, for each $(A, B, A', B') \in \mathcal{S}$, after replacing pair $(A, B) \in \mathcal{P}$ with pair (A', B') , we obtain an updated WSPD.

Sample size guarantee To show the third requirement, for each $(A, B, A', B') \in \mathcal{S}$, let X, Y, X', Y' be their corresponding d -dimensional point sets. We resample

$$\epsilon^{-2}n^{O(L/k)}(|X'| + |Y'|) \log(|X'| + |Y'|)$$

edges from $\text{Biclique}(X', Y')$ by updating the edges sampled from $\text{Biclique}(X, Y)$. To do this, we first multiply each edge weight in $\text{EDGES}(X, Y)$ by

$$\epsilon^{-2}(n^{O(L/k)}(|X| + |Y|) \log(|X| + |Y|)) / |X||Y|$$

so that each edge has the same weight in E and in $\text{biclique}(X', Y')$. Then we apply `FASTRESAMPLE`. Since

- $|(X \times Y) \triangle (X' \times Y')| \leq o(\frac{|X' \times Y'|}{|X'| + |Y'|})$ (line 11)
- E is a uniform sample from $X \times Y$ of size $\epsilon^{-2}n^{C_{\text{jl}} \cdot (L/k)}(|X| + |Y|) \log(|X| + |Y|)$ (line 8 and definition of EDGES)
- $s = \epsilon^{-2}n^{O(L/k)}(|X'| + |Y'|) \log(|X'| + |Y'|)$ (line 12)
- $|s - |E|| = O(\log n)$, because $||X| + |Y| - |X'| - |Y'||$ is at most 1.

by Lemma E.10, the new sample can be viewed as a uniform sample from $\text{biclique}(X', Y')$.

Similar to Lemma E.7, the edges uniformly sampled from $\text{Biclique}(X', Y')$ form a ϵ -spectral sparsifier of $\text{Biclique}(X', Y')$ after scaling each edge weight by

$$|X'| |Y'| / (\epsilon^{-2}(n^{O(L/k)}(|X'| + |Y'|) \log(|X'| + |Y'|))).$$

If the number of edges in $\text{Biclique}(X', Y')$ itself is

$$O\left(\epsilon^{-2}(n^{O(L/k)}(|X'| + |Y'|) \log(|X'| + |Y'|))\right),$$

we use all edges in the biclique without scaling. The union of all sampled edges remains a spectral sparsifier of K_G .

The projection can be updated in $O(dk)$ time, where d is the ambient dimension of the points and $k = o(\log n)$.

1728 By Lemma E.8, FINDMODIFIEDPAIRS takes $O(2^{O(k)} \log(\alpha))$ time and the returned collection \mathcal{S}
 1729 contains at most $O(2^{O(k)} \log(\alpha))$ changed pairs.
 1730

1731 By Lemma E.10, with high probability $1 - \delta$ resampling takes $\delta^{-1} \epsilon^{-2} n^{o(1)}$ time and the number of
 1732 new edges in the sample is $\delta^{-1} \epsilon^{-2} n^{o(1)}$.

1733 Therefore, with high probability, the total number of edge updates in \mathcal{H} is with high probability
 1734

$$\delta^{-1} \epsilon^{-2} 2^{O(k)} \log \alpha n^{o(1)} = \delta^{-1} \epsilon^{-2} n^{o(1)} \log \alpha,$$

1736 and the time needed to update the sparsifier is
 1737

$$O(dk + \epsilon^{-2} \delta^{-1} n^{o(1)} \log \alpha).$$

□

1741 E.8 A DATA STRUCTURE THAT CAN HANDLE FULLY DYNAMIC UPDATE

1742 By the limitation of the ultra low dimensional JL projection, when it needs to handle more than $O(n)$
 1743 projections, the $n^{C_{jl} \cdot (1/k)}$ distortion bound cannot be preserved with high probability. Therefore,
 1744 Lemma E.11 states that DYNAMICGEOSPAR can only handle $O(n)$ updates.
 1745

1746 This essentially gives us an online algorithm, with support of batch update. Under the setting of
 1747 online batch, the dynamic data structure \mathfrak{D} undergoes batch updates defined by these two parameters:
 1748 the number of batches, denoted by ζ , and the sensitivity parameter, denoted by w . \mathfrak{D} has one
 1749 initialization phase and ζ phases: an initialization phase and ζ update phases and in each update
 1750 phase, the data structure \mathfrak{D} receives updates for no more than w times.

1751 This algorithm is designed to maintain \mathfrak{D} under the update batches. The data structure is maintained
 1752 to exactly match the original graph after series of update batches. We define the *amortized randomized*
 1753 *update time* t to be the time such that, with every batch size less than w , the running time of
 1754 each update to data structure is no more than t . The goal of this section is to minimize the time t .
 1755 We first introduce the following useful lemma from literature, which introduces the framework of
 1756 the online-batch setting.

1757 **Lemma E.12** (Section 5, (Nanongkai et al., 2017)). *We define G to be a geometric graph, with
 1758 updates come in batches. Let $\zeta \in \mathbb{R}$ denote batch number. Let $w \in \mathbb{R}$ denote the sensitivity
 1759 parameter. Then there exists a data structure \mathfrak{D} with the batch number of ζ and sensitivity of w ,
 1760 which supports:*

- An initialization procedure which runs in time $t_{\text{initialize}}$;
- An update procedure which runs in time t_{update} .

1761 The two running time parameter $t_{\text{initialize}}$ and t_{update} are defined to be functions such that, they
 1762 send the maximum value of measures of the graph to non-negative numbers. For example, the upper
 1763 bounds of the edges.

1764 Then we have the result that, for any parameter ξ such that $\xi \leq \min\{\zeta, \log_6(w/2)\}$, there exists
 1765 a fully dynamic data structure consists of a size- $O(2^\xi)$ set of data structures \mathfrak{D} . It can initialize in
 1766 time $O(2^\xi \cdot t_{\text{initialize}})$. And it has update time of $O(4^\xi \cdot (t_{\text{initialize}}/w + w^{(1/\xi)} t_{\text{update}}))$ in the worst
 1767 case. When the data structure is updated every time, the update procedure can select one instance
 1768 from the set, which satisfies that

1. The selected instance of \mathfrak{D} matches the updated graph.
2. The selected instance of \mathfrak{D} has been updated for at most ξ times, and the size of the update
 batch every time is at most w .

1773 By Lemma E.5, the initialization time of DYNAMICGEOSPAR is
 1774

$$O(ndk + \epsilon^{-2} n^{1+o(1)}) \log \alpha.$$

1775 By Lemma E.11, the update time of DYNAMICGEOSPAR is
 1776

$$\delta^{-1} \epsilon^{-2} n^{o(1)} \log \alpha$$

per update and it can handle $O(n)$ batches of updates, each containing 1 update. Therefore, we can apply Lemma E.12 to DYNAMICGEOPAR with $w = 1$ and $\zeta = O(n)$. We obtain a fully dynamic update data structure as stated below.

Corollary E.13 (Corollary of Lemma E.5, E.11, and E.12). *There is a fully dynamic algorithm with initialization time $O(n^{1+(o(1))} \log \alpha)$ and update time $O(n^{o(1)})$.*

Corollary E.13 completes the proof of Theorem E.3.

F MAINTAINING A SKETCH OF AN APPROXIMATION TO MATRIX MULTIPLICATION

The goal of this section is to prove the following statement,

Theorem F.1 (Formal version of Theorem 2.3). *Let G be a (C, L) -lipschitz geometric graph on n points. Let v be a vector in \mathbb{R}^d . Let k denote the sketch size. There exists an data structure MULTIPLY that maintains a vector \tilde{z} that is a low dimensional sketch of an ϵ -approximation of the multiplication $L_G \cdot v$, where b is said to be an ϵ -approximation of $L_G x$ if*

$$\|b - L_G x\|_2 \leq \epsilon \|L_G\|_F \cdot \|x\|_2.$$

MULTIPLY supports the following operations:

- **UPDATEG**(x_i, z): move a point from x_i to z and thus changing K_G . This takes $dk + n^{o(1)} \log \alpha$ time, where α is the aspect ratio of the graph.
- **UPDATEV**(δ_v): change v to $v + \delta_v$. This takes $O(\log n)$ time.
- **QUERY()**: return the up-to-date sketch.

We divide the section into the following parts. Section F.1 gives the high level overview of the section. Section F.2 introduces the necessity of sketching. Section F.3 introduces our algorithms.

F.1 HIGH LEVEL OVERVIEW

The high level idea is to combine the spectral sparsifier defined in Section E and a sketch matrix to compute a sketch of the multiplication result and try to maintain this sketch when the graph and the vector change. We first revisit the definition of spectral sparsifiers. Let $G = (V, E)$ be a graph and $H = (V, E')$ be a ϵ -spectral sparsifier of G . Suppose $|V| = n$. By definition, this means

$$(1 - \epsilon)L_G \preceq L_H \preceq (1 + \epsilon)L_G$$

Lemma F.2. *Let G be a graph and H be a ϵ -spectral sparsifier of G . $L_H x$ is an ϵ -approximation of $L_G x$*

Proof. By Definition B.26, $(1 - \epsilon)L_G \preceq L_H \preceq (1 + \epsilon)L_G$. Applying Proposition B.3, we get

$$\|L_H x - L_G x\|_{L_G^\dagger} \leq \epsilon \|L_G x\|_{L_G^\dagger}$$

which means $L_H x$ is an ϵ -approximation of $L_G x$. \square

Thus, to maintain a sketch of an ϵ -approximation of $L_G x$, it suffices to maintain a sketch of $L_H x$.

F.2 NECESSITY OF SKETCHING

We here justify the decision of maintaining a sketch instead of the directly maintaining the multiplication result. Let the underlying geometry graph on n vertices be G and the vector be $v \in \mathbb{R}^n$. When a point is moved in the geometric graph, a column and a row are changed in L_G . Without loss of generality, we can assume the first row and first column are changed. When this happens, if the first entry of v is not 0, all entries will change in the multiplication result. Therefore, it takes at least $O(n)$ time to update the multiplication result. In order to spend subpolynomial time to maintain the multiplication result, we need to reduce the dimension of vectors. Therefore, we use a sketch matrix to project vectors down to lower dimensions.

1836 **Lemma F.3** (Johnson & Lindenstrauss, 1984)). Let $\epsilon \in (0, 0.1)$ denote an accuracy parameter. Let
 1837 $\delta \in (0, 0.1)$ denote a failure probability. Let $X = \{x_1, \dots, x_n\} \in \mathbb{R}^d$ denote a set of points. Let
 1838 $\Phi \in \mathbb{R}^{m \times n}$ denote a randomized sketching matrix that, if $m = O(\epsilon^{-2} \log(n/\delta))$, with probability
 1839 $1 - \delta$, we have: for all $x \in X$

$$(1 - \epsilon) \cdot \|x\|_2 \leq \|\Phi x\|_2 \leq (1 + \epsilon) \cdot \|x\|_2.$$

1840 We also have the following result by using the sparse embedding matrix:
 1841

1842 **Lemma F.4.** Let $\Psi \in \mathbb{R}^{m \times n}$ be a sparse embedding matrix (Definition C.1) with $m = O(\epsilon^{-2} \cdot$
 1843 $\log(1/\delta))$. Then for a vector $v \in \mathbb{R}^n$ and a matrix $L \in \mathbb{R}^{n \times n}$, we have $\|(L\Psi^\top \Psi v) - (Lv)\|_2 \leq$
 1844 $\epsilon \cdot \|v\|_2 \cdot \|L\|_F$, with probability at least $1 - \delta$.
 1845

1846 *Proof.* By Lemma C.5, we have Ψ satisfies the $(\epsilon, \delta, \log(1/\delta))$ -JL moment property. Since
 1847 $\log(1/\delta) \geq 2$ is trivial, then by Lemma C.4, and rescaling ϵ with a constant factor, we complete
 1848 the proof. \square
 1849

1850

F.3 ALGORITHMS

1851 **F.3.1 MODIFICATION TO DYNAMICGOSPAR**

1852 For the applications in Section F and G, we add a member DIFF and methods GETDIFF and GET-
 1853 LAPLACIAN to DYNAMICGOSPAR and change methods INIT and Update to initialize and update
 1854 DIFF (Algorithm 10).
 1855

1856 **Algorithm 10** Interfaces for getting ΔL_H after H changes

1857 1: **data structure** DYNAMICGOSPAR ▷ Lemma F.5
 1858 2: **members**
 1859 3: diff ▷ the difference in the Laplacian after the graph is updated
 1860 4: **EndMembers**
 1861 5:
 1862 6: **procedure** GETDIFF()
 1863 7: diffValue \leftarrow diff
 1864 8: diff $\leftarrow \{\}$
 1865 9: **return** diffValue
 1866 10: **end procedure**
 1867 11:
 1868 12: **procedure** GETLAPLACIAN()
 1869 13: **return** The Laplacian matrix of \mathcal{H}
 1870 14: **end procedure**
 1871 15:
 1872 16: **procedure** INITIALIZE(x_i, z) ▷ Content in Algorithm 4
 1873 17: ...
 1874 18: diff $\leftarrow \{\}$
 1875 19: **end procedure**
 1876 20: **procedure** UPDATE(P) ▷ Content in Algorithm 9
 1877 21: ...
 1878 22: **for** each EDGES update pair (E, E^{new}) **do**
 1879 23: **for** each e in $E \setminus E^{\text{new}}$ **do**
 1880 24: Add $-e$ to diff
 1881 25: **end for**
 1882 26: **for** each e in $E^{\text{new}} \setminus E$ **do**
 1883 27: Add e to diff
 1884 28: **end for**
 1885 29: **end for**
 1886 30: **end procedure**
 1887

1888 **Lemma F.5.** In data structure DYNAMICGOSPAR, suppose GETDIFF (Algorithm 10) is called
 1889 right after each UPDATE. The returned diff is a sparse matrix of size $O(n^{o(1)} \log \alpha)$.

1890 *Proof.* By Theorem E.3, in expectation each update introduces $O(n^{o(1)} \log \alpha)$ edge changes in the
 1891 sparsifier. Therefore, after m updates, there are at most $O(n^{o(1)} \log \alpha)$ entries in diff . \square
 1892

1893 **F.3.2 DYNAMIC SKETCH ALGORITHM**
 1894

1895 Here we propose the dynamic sketch algorithm as follows.
 1896

1897 **Algorithm 11** Maintaining a sketch of an approximation to multiplication

1898 1: **data structure** **MULTIPLY** ▷ Theorem F.1
 1899 2: **members**
 1900 3: **DYNAMICGEOSPAR** **dgs** ▷ This is the sparsifier H
 1901 4: $\Phi, \Psi \in \mathbb{R}^{m \times n}$: two independent sketching matrices
 1902 5: $\tilde{L} \in \mathbb{R}^{m \times m}$ ▷ A sketch of L_H
 1903 6: $\tilde{v} \in \mathbb{R}^m$ ▷ A sketch of v
 1904 7: $\tilde{z} \in \mathbb{R}^m$ ▷ A sketch of the multiplication result
 1905 8: **EndMembers**
 1906 9:
 1907 10: **procedure** **INIT**($x_1, \dots, x_n \in \mathbb{R}^d, v \in \mathbb{R}^n$)
 1908 11: Initialize Φ and Ψ
 1909 12: **dgs.INITIALIZE**(x_1, \dots, x_n)
 1910 13: $\tilde{v} \leftarrow \Psi v$
 1911 14: $\tilde{L} \leftarrow \Phi \cdot \text{dgs.GETLAPLACIAN}() \cdot \Psi^\top$
 1912 15: $\tilde{z} \leftarrow \tilde{L} \cdot \tilde{v}$
 1913 16: **end procedure**
 1914 17:
 1915 18: **procedure** **UPDATEG**($x_i, z \in \mathbb{R}^d$)
 1916 19: **dgs.UPDATE**(x_i, z)
 1917 20: $\Delta \tilde{L} \leftarrow \Phi \cdot \text{dgs.GETDIFF}() \cdot \Psi^\top$
 1918 21: $\tilde{z} \leftarrow \tilde{z} + \Delta \tilde{L} \cdot \tilde{v}$
 1919 22: $\tilde{L} \leftarrow \tilde{L} + \Delta \tilde{L}$
 1920 23: **end procedure**
 1921 24:
 1922 25: **procedure** **UPDATEV**($\Delta v \in \mathbb{R}^n$) ▷ Δv is sparse
 1923 26: $\Delta \tilde{v} \leftarrow \Psi \cdot \Delta v$
 1924 27: $\tilde{z} \leftarrow \tilde{z} + \tilde{L} \cdot \Delta \tilde{v}$
 1925 28: **end procedure**
 1926 29:
 1927 30: **procedure** **QUERY**
 1928 31: **return** \tilde{z}
 32: **end procedure**

1929 Here we give the correctness proof of Theorem F.1.
 1930

1931 *Proof of Theorem F.1.* We divide the proof into correctness proof and running time proof as follows.
 1932

1933 **Correctness** By Lemma F.2, $L_H x$ is an ϵ -approximation of $L_G x$. It suffices to show that **MULTI-**
 1934 **PLY** maintains a sketch of $L_H x$.
 1935

1936 In function **INIT**, a ϵ -spectral sparsifier H of G is initialized on line 12. On line 14, \tilde{L} is computed
 1937 as $\Phi L_H \Psi^\top$. Therefore, $\tilde{z} = \Phi L_H \Psi^\top \Psi v$. By Lemma F.4 we have that
 1938

$$\|L_H \Psi^\top \Psi v - L_H v\|_2 \leq \epsilon \cdot \|L_H\|_F \cdot \|v\|_2. \quad (1)$$

1939 And by Lemma F.3 one has
 1940

$$\|\tilde{z}\|_2 \in (1 \pm \epsilon) \cdot \|L_H \Psi^\top \Psi v\|_2. \quad (2)$$

1941 Then by Eq (1) and Eq (2) and rescaling ϵ we have that $\|\tilde{z} - L_H x\|_2 \leq \epsilon \cdot \|L_H\|_F \cdot \|v\|_2$.
 1942

1944 In function UPDATEG, the algorithms updates the spectral sparsifier (line 19) and obtains the differ-
 1945 ence in the Laplacian (line 20). Note that $\Delta\tilde{L} \cdot \tilde{v} = \Phi\Delta L_H \cdot \Psi^\top \cdot \Psi v$. Again, since in expectation
 1946 $\Psi^\top \Psi = I_{n \times n}$ and Φ and Ψ are chosen independently, in expectation $\Delta\tilde{L} \cdot \tilde{v} = \Phi\Delta L_H v$. Therefore,
 1947 $\tilde{z} + \Delta\tilde{L} \cdot \tilde{v} = \Phi(L_H + \Delta L_H)v$ is the updated sketch of $L_H x$.
 1948

1949 **Running time** By Theorem E.3, line 19 takes $O(dk + n^{o(1)} \log \alpha)$ time. By Lemma F.5, $\Delta\tilde{L}$ is
 1950 sparse with $\epsilon^{-2}n^{o(1)} \log \alpha$ non-zero entries. This implies line 20 takes $O(\epsilon^{-2}n^{o(1)} \log \alpha)$ time. So
 1951 the overall time complexity of UPDATEG is
 1952

$$dk + \epsilon^{-2}n^{o(1)} \log \alpha.$$

1953 In function UPDATEV, note that
 1954

$$\tilde{L} \cdot \Delta\tilde{v} = \Phi \cdot L_H \cdot \Psi^\top \cdot \Psi \Delta v.$$

1955 Again, since in expectation $\Psi^\top \Psi = I_{n \times n}$ and Φ and Ψ are chosen independently, in expectation
 1956

$$\tilde{L} \cdot \Delta\tilde{v} = \Phi L_H \Delta v.$$

1957 Therefore,

$$\tilde{z} + \tilde{L} \cdot \Delta\tilde{v} = \Phi L_H(v + \Delta v)$$

1958 is the updated sketch of $L_H x$.
 1959

1960 Since Δv is sparse, $\Psi \Delta v$ can be computed in $O(m)$ time, and \tilde{z} can also be updated in $O(m)$ time,
 1961 where $m = O(\epsilon^{-2} \log(n/\delta))$.
 1962

1963 Since \tilde{z} is always an up-to-date sketch of $L_H \cdot v$, QUERY always returns a sketch of an approximation
 1964 to $L_G x$ in constant time.
 1965

1966 Thus we complete the proof. □
 1967

1968 G MAINTAINING A SKETCH OF AN APPROXIMATION TO SOLVING 1969 LAPLACIAN SYSTEM

1970 In this section, we provide a data structure which maintains a sketch of an approximation to solving
 1971 Laplacian system. In other words, we prove the following theorem,
 1972

1973 **Theorem G.1** (Formal version of Theorem 2.4). *Let K_G be a (C, L) -lipschitz geometric graph on
 1974 n points. Let b be a vector in \mathbb{R}^d . There exists a data structure **SOLVE** that maintains a vector \tilde{z}
 1975 that is a low dimensional sketch of multiplication $L_G^\dagger \cdot b$, where \tilde{z} is said to be an ϵ -approximation
 1976 of $L_G^\dagger b$ if*
 1977

$$\|\tilde{z} - L_G^\dagger b\|_2 \leq \epsilon \cdot \|L_G^\dagger\|_F \cdot \|b\|_2.$$

1978 **SOLVE** supports the following operations:
 1979

- 1980 • UPDATEG(x_i, z): move a point from x_i to z and thus changing K_G . This takes $n^{o(1)}$ time.
 1981
- 1982 • UPDATEB(δ_b): change b to $b + \delta_b$. This takes $n^{o(1)}$ time.
 1983
- 1984 • QUERY(): return the up-to-date sketch. This takes $O(1)$ time.
 1985

1986 By Fact B.3, for any vector b , $L_H^\dagger b$ is a ϵ -spectral sparsifier of $L_G^\dagger b$. It suffices to maintain a sketch
 1987 of $L_H b$.
 1988

1989 When trying to maintain a sketch of a solution to $L_H x = b$, the classical way of doing this is to
 1990 maintain \bar{x} such that $\Phi L_H \bar{x} = \Phi b$. However, here \bar{x} is still an n -dimensional vector and we want
 1991 to maintain a sketch of lower dimension. Therefore, we apply another sketch Ψ to \bar{x} and maintain \tilde{x}
 1992 such that $\Phi L_H \Psi^\top \tilde{x} = \Phi b$.
 1993

1994 *Proof of Theorem G.1.* We divide the proof into the following paragraphs.
 1995

1998 **Analysis of INIT** In function INIT, a ϵ -spectral sparsifier H of G is initialized on line 13. On line
 1999 16, \tilde{L}^\dagger is computed as $(\Phi L_H \Psi^\top)^\dagger$. Therefore,
 2000

$$\tilde{z} = (\Phi L_H \Psi^\top)^\dagger \Phi b,$$

2003 which is a sketch of $L_H^\dagger x$. Thus, after INIT, $L_H x$, which is a sketch of an approximation to $L_G x$ is
 2004 stored in \tilde{z} .
 2005

2006 **Analysis of UPDATEG** In function UPDATEG, the algorithms updates the spectral sparsifier (line
 2007 21) and obtains the new Laplacian (line 22) and its pseudoinverse (line 23). Note that in line 24
 2008

$$\tilde{L}_{\text{new}}^\dagger \cdot \tilde{b} = (\Phi(L_H + \Delta L_H) \Psi^\top)^\dagger \Phi b$$

2010 This is a sketch of $(L_H + \Delta L_H)^\dagger b$.
 2011

2012 By Theorem E.3, line 21 takes $O(dk + n^{o(1)} \log \alpha)$ time. By Lemma F.5, $\Delta \tilde{L}$ is sparse with
 2013 $\epsilon^{-2} n^{o(1)} \log \alpha$ non-zero entries. This implies line 22 takes $O(\epsilon^{-2} n^{o(1)} \log \alpha)$ time. Since \tilde{L}^\dagger
 2014 is a $m \times m$ matrix and $m = O(\epsilon^{-2} \log(n/\delta))$, computing its pseudoinverse takes at most
 2015 $m^\omega = (\epsilon^{-2} \log(n/\delta))^\omega$ ⁶ time.
 2016

2017 So the overall time complexity of UPDATEG is
 2018

$$O(dk) + \epsilon^{-2} n^{o(1)} \log \alpha + O((\epsilon^{-2} \log(n/\delta))^\omega).$$

2019 **Analysis of UPDATEB** In function UPDATEB, note that
 2020

$$\tilde{L}^\dagger \cdot \Delta \tilde{b} = (\Phi \cdot L_H \cdot \Psi^\top)^\dagger \cdot \Phi \Delta b.$$

2024 Therefore, it holds that
 2025

$$\tilde{z} + \tilde{L}^\dagger \cdot \Delta \tilde{b} = (\Phi \cdot L_H \cdot \Psi^\top)^\dagger \cdot \Phi(b + \Delta b).$$

2026 This is a sketch of $L_H^\dagger(b + \Delta b)$.
 2027

2028 Since Δb is sparse, $\Phi \Delta b$ can be computed in $O(m)$ time, and \tilde{z} can also be updated in $O(m)$ time,
 2029 where $m = O(\epsilon^{-2} \log(n/\delta))$.
 2030

2031 **Analysis of QUERY** Since \tilde{z} is always an up-to-date sketch of $L_H \cdot v$, QUERY always returns a
 2032 sketch of an approximation to $L_G x$ in constant time.
 2033

2034 Thus we complete the proof. □
 2035

H DYNAMIC DATA STRUCTURE

2038 In this section, we describe our data structure in Algorithm 12 to solve the dynamic distance estimation
 2039 problem with robustness to adversarial queries. We need to initialize a sketch $\Pi \in \mathbb{R}^{k \times d}$ defined
 2040 in Definition H.1, where $k = \Theta(\sqrt{\log n})$, and use the ultra-low dimensional projection matrix to
 2041 maintain a set of projected points $\{\tilde{x}_i \in \mathbb{R}^k\}_{i=1}^n$. During QUERY, the data structure compute the
 2042 estimated distance between the query point $q \in \mathbb{R}^d$ and the data point x_i by $n^{1/k} \cdot \sqrt{d/k} \cdot \|\tilde{x}_i - \Pi q\|_2$.
 2043

H.1 MAIN RESULT

2045 In this section, we introduce our main results, we start with defining ultra-low dimensional JL matrix.
 2046

2047 **Definition H.1** (Ultra-Low Dimensional JL matrix). *Let $\Pi \in \mathbb{R}^{k \times d}$ denote a random JL matrix
 2048 where each entry is i.i.d. Gaussian.*
 2049

2050 Next, we present our main result in accuracy-efficiency trade-offs, which relates to the energy con-
 2051 sumption in practice.

⁶ ω is the matrix multiplication constant

Theorem H.2 (Main result). *Let $d = \Theta(\log n)$. Let $k = \Theta(\sqrt{\log n})$. There is a data structure (Algorithm 12) for the Online Approximate Dynamic Ultra-Low Dimensional Distance Estimation Problem with the following procedures:*

- **INIT**($\{x_1, x_2, \dots, x_n\} \subset \mathbb{R}^d, n \in \mathbb{N}_+, d \in \mathbb{N}_+, \epsilon \in (0, 1)\}$): *Given n data points $\{x_1, x_2, \dots, x_n\} \subset \mathbb{R}^d$, an accuracy parameter ϵ , and input dimension d and number of input points n as input, the data structure preprocesses in time $O(ndk)$.*
- **UPDATE**($z \in \mathbb{R}^d, i \in [n]$): *Given an update vector $z \in \mathbb{R}^d$ and index $i \in [n]$, the UPDA-TEX takes z and i as input and updates the data structure with the new i -th data point in $O(dk)$ time.*
- **QUERY**($q \in \mathbb{R}^d$): *Given a query point $q \in \mathbb{R}^d$, the QUERY operation takes q as input and approximately estimates the norm distances from q to all the data points $\{x_1, x_2, \dots, x_n\} \subset \mathbb{R}^d$ in time $O(nk)$ i.e. it outputs a vector $u \in \mathbb{R}^n$ such that:*

$$\forall i \in [n] \quad \|a - x_i\|_2 \leq y_i \leq n^{O(1/k)} \cdot \|a - x_i\|_2$$

with probability at least $1 - 1/\text{poly}(n)$, even for a sequence of adversarially chosen queries.

2106 H.2 TIME
21072108 In the section, we will provide lemmas for the time complexity of each operation in our data struc-
2109 ture.2110 **Lemma H.3** (INIT time). *There is a procedure INIT which takes a set of d -dimensional vectors
2111 $\{x_1, \dots, x_n\}$, a precision parameter $\epsilon \in (0, 0.1)$ and $d, n \in \mathbb{N}_+$ as input, and runs in $O(ndk)$ time.*
21122113 *Proof.* Storing every vector x_i takes $O(nd)$ time. Computing and storing \tilde{x}_i takes $O(n \times dk) =$
2114 $O(ndk)$ time. Thus procedure INIT runs in $O(ndk)$ time. \square
21152116 We prove the time complexity of UPDATE operation in the following lemma:
21172118 **Lemma H.4** (UPDATE time). *There is a procedure UPDATE which takes an index $i \in [n]$ and a
2119 d -dimensional vector z as input, and runs in $O(nk)$ time.*
21202121 *Proof.* Updating x_i takes $O(d)$ time. Update \tilde{x}_i takes $O(dk)$ time. Thus procedure UPDATE runs in
2122 $O(dk)$ time. \square
21232124 We prove the time complexity of QUERY operation in the following lemma:
21252126 **Lemma H.5** (QUERY time). *There is a procedure QUERY which takes a d -dimensional vector q as
2127 input, and runs in $O(nk)$ time.*
21282129 *Proof.* Computing Πq takes $O(dk)$ time. Computing all the u_i takes $O(nk)$ time. Thus procedure
2130 QUERY runs in $O(nk)$ time. \square
21312132 H.3 CORRECTNESS
2133

2134 In this section, we provide lemmas to prove the correctness of operations in our data structure.

2135 **Lemma H.6** (INIT correctness). *There is a procedure INIT which takes a set $\{x_1, \dots, x_n\}$ of d -
2136 dimensional vectors and a precision parameter $\epsilon \in (0, 0.1)$, and stores an adjoint vector \tilde{x}_i for each
2137 x_i .*
21382139 *Proof.* During INIT operation in Algorithm 12, the data structure stores a set of adjoint vectors
2140 $\tilde{x}_i \leftarrow \Pi \cdot x_i$ for $i \in [n]$. This completes the proof. \square
21412142 Then we prove the correctness of UPDATE operation in Lemma H.7.
21432144 **Lemma H.7** (UPDATE correctness). *There is a procedure UPDATE which takes an index $i \in [n]$ and
2145 a d -dimensional vector z , and uses z to replace the current x_i .*
21462147 *Proof.* During UPDATE operation in Algorithm 12, the data structure update the i -th adjoint vector
2148 \tilde{x}_i by $\Pi \cdot z$. This completes the proof. \square
21492150 We prove the correctness of QUERY operation in Lemma H.8.
21512152 **Lemma H.8** (QUERY correctness). *There is a procedure QUERY which takes a d -dimensional vector
2153 q as input, and output an n -dimensional vector u such that for each $i \in [n]$, $\|q - x_i\|_2 \leq u_i \leq$
2154 $n^{O(1/k)} \cdot \|q - x_i\|_2$ with probability $1 - 2/n$.*
21552156 *Proof.* The proof follows by Lemma H.9, Lemma I.2 and Lemma I.4. This completes the proof.
2157

2160 H.4 HIGH PROBABILITY
2161

2162 With Lemma B.9 and Lemma B.8 ready, we want to prove the following lemma:

2163 **Lemma H.9** (High probability for each point). *For any integer n , let $d = c_0 \log n$. Let k be a
2164 positive integer such that $k = \sqrt{\log n}$. Let f be a map $f : \mathbb{R}^d \rightarrow \mathbb{R}^k$. Let $\delta_1 = n^{-c}$ denote the
2165 failure probability where $c > 1$ is a large constant. Let $c_0 > 1$ denote some fixed constant. Then for
2166 each fixed points $u, v \in \mathbb{R}^d$, such that,*

2167
$$\|u - v\|_2^2 \leq \|f(u) - f(v)\|_2^2 \leq \exp(c_0 \cdot \sqrt{\log n}) \|u - v\|_2^2.$$

2168

2169 with probability $1 - \delta_1$.2170
2171 *Proof.* If $d \leq k$, the theorem is trivial. Else let $v', u' \in \mathbb{R}^k$ be the projection of point $v, u \in \mathbb{R}^d$ into
2172 \mathbb{R}^k . Then, setting $L = \|u' - v'\|_2^2$ and $\mu = \frac{k}{d} \|u - v\|_2^2$. We have that
2173

2174
$$\begin{aligned} \Pr[L \leq (n^{-2c/k}/e)\mu] &= \Pr[L \leq (n^{-2c/\sqrt{\log n}}/e)\mu] \\ 2175 &= \Pr[L \leq e^{-2c\sqrt{\log n}-1}\mu] \\ 2176 &\leq n^{-c} \end{aligned} \tag{3}$$

2177

2178 where the first step comes from $k = \sqrt{\log n}$, the second step comes from $n^{-2c/\sqrt{\log n}} =$
2179 $\exp(-2c\sqrt{\log n})$, and the third step comes from Lemma B.9.
2180

2181 By Lemma B.8, we have:

2182
$$\begin{aligned} \Pr[L \geq n^{1/k}\mu] &= \Pr[L \geq n^{1/\sqrt{\log n}}\mu] \\ 2183 &= \Pr[L \geq \exp(c_0\sqrt{\log n})\mu] \\ 2184 &\leq \exp(-\log^{1.9} n) \\ 2185 &\leq n^{-c} \end{aligned} \tag{4}$$

2186

2187 where the first step comes from the definition of $k = \sqrt{\log n}$, the second step follows that
2188 $n^{1/\sqrt{\log n}} = \exp(c_0\sqrt{\log n})$, the third step comes from Lemma B.8, and the fourth step follows
2189 that $\log^{1.9}(n)$ is bigger than any constant c .
21902191 Therefore, rescaling from Eq. (3) and Eq. (4) we have:
2192

2193
$$\|u - v\|_2^2 \leq \|f(u) - f(v)\|_2^2 \leq \exp(c_0 \cdot \sqrt{\log n}) \|u - v\|_2^2.$$

2194

2195 with probability $1 - \delta_1$ where $\delta_1 = n^{-c}$.
2196

□

2197 I SPARSIFIER IN ADVERSARIAL SETTING
21982199 In Section H, we get a dynamic distance estimation data structure with robustness to adversarial
2200 queries. Here in this section, we provide the analysis to generalize our spectral sparsifier to adver-
2201 sarial setting, including discussion on the aspect ratio α (Definition B.10).
2202

2203 I.1 DISTANCE ESTIMATION FOR ADVERSARIAL SPARSIFIER

2204 **Fact I.1.** *Let α be defined as Definition B.10. Let N denote a ϵ_1 -net on the ℓ_2 unit ball $\{x \in$
2205 $\mathbb{R}^d \mid \|x\|_2 \leq 1\}$, where $d = O(\log n)$ and $\epsilon_1 = O(\alpha^{-1})$. Then we have that $|N| \leq \alpha^{O(\log n)}$.*
22062207 **Lemma I.2.** *Let α be defined as Definition B.10. For any integer n , let $d = O(\log n)$, let $k =$
2208 $\sqrt{\log n}$. Let $f : \mathbb{R}^d \rightarrow \mathbb{R}^k$ be a map. If $\alpha = O(1)$, then for an ϵ_1 -net N with $|N| \leq \alpha^{O(\log n)}$, for
2209 all $u, v \in N$,*

2210
$$\|u - v\|_2^2 \leq \|f(u) - f(v)\|_2^2 \leq \exp(\sqrt{\log n}) \|u - v\|_2^2$$

2211

2212 with probability $1 - \delta_2$.
2213

2214 *Proof.* By Lemma H.9, we have that for any fix set V of n points in \mathbb{R}^d , there exists a map $f : \mathbb{R}^d \rightarrow$
 2215 \mathbb{R}^k such that for all $u, v \in V$,

$$2217 \quad \|u - v\|_2^2 \leq \|f(u) - f(v)\|_2^2 \leq \exp(\sqrt{\log n})\|u - v\|_2^2,$$

2218 with probability $1 - \delta_1$, where $\delta_1 = n^{-c}$.

2220 We apply the lemma on N , and by union bound over the points in N , we have that for all points
 2221 $u, v \in N$,

$$2222 \quad \|u - v\|_2^2 \leq \|f(u) - f(v)\|_2^2 \leq \exp(\sqrt{\log n})\|u - v\|_2^2,$$

2224 with probability $1 - \delta_2$, where it holds that

$$\begin{aligned} 2226 \quad \delta_2 &= \delta_1 \cdot |N|^2 \\ 2227 &\leq \delta_1 \cdot \alpha^{O(\log n)} \\ 2228 &= n^{-c} \cdot \alpha^{O(\log n)} \\ 2229 &\leq n^{O(1)-c}, \end{aligned}$$

2231 where the first step follows from union bound, the second step follows from $|N| \leq \alpha^{O(\log n)}$, the
 2232 third step follows from $\delta_1 = n^{-c}$ (Lemma H.9), and the last step follows from $\alpha = O(1)$.

2234 By choosing c as a constant large enough, we can get the low failure probability. \square

2235 **Corollary I.3** (Failure probability on α and d). *We have that, the failure probability of Lemma I.2
 2236 is bounded as long as $\alpha^d = O(\text{poly}(n))$.*

2238 **Lemma I.4** (Adversarial Distance Estimation of the Spectral Sparsifier). *Let $k = \sqrt{\log n}$ be the
 2239 JL dimension, $f : \mathbb{R}^d \rightarrow \mathbb{R}^k$ be a JL function. Let α be defined as Definition B.10. Then we have
 2240 that for all points u, v in the ℓ_2 unit ball, there exists a point pair u', v' which is the closest to u, v
 2241 respectively such that,*

$$2242 \quad \|u - v\|_2^2 \leq \|f(u') - f(v')\|_2^2 \leq \exp(\sqrt{\log n})\|u - v\|_2^2$$

2243 with probability $1 - n^{-c_1}$.

2246 *Proof.* By Lemma I.2, we have that for all $u', v' \in N$, it holds that

$$2248 \quad \|u' - v'\|_2^2 \leq \|f(u') - f(v')\|_2^2 \leq \exp(\sqrt{\log n})\|u' - v'\|_2^2 \quad (5)$$

2249 with probability at least $1 - \delta_2$. From now on, we condition on the above event happens. Then for
 2250 arbitrary $u, v \notin N$, there exists $u', v' \in N$ such that

$$2252 \quad \|u - u'\| \leq \epsilon_1 \text{ and } \|v - v'\| \leq \epsilon_1.$$

2253 Recall that we set $\epsilon_1 = O(\alpha^{-1})$ and now all points are in ℓ_2 ball, thus we have that $u' \neq v'$ for $u \neq v$.
 2254 Then by triangle inequality we have that

$$\begin{aligned} 2256 \quad \|u' - v'\|_2 &= \|u' - u + u - v + v - v'\|_2 \\ 2257 &\leq \|u' - u\|_2 + \|u - v\|_2 + \|v - v'\|_2 \\ 2258 &\leq \|u - v\|_2 + 2\epsilon_1 \\ 2259 &\leq (1 + O(1)) \cdot \|u - v\|_2, \end{aligned} \quad (6)$$

2261 where the last step follows by setting $\epsilon_1 = O(\|u - v\|_2) = O(\alpha^{-1})$. Similarly, we also have

$$2263 \quad \|u' - v'\|_2 \geq (1 - O(1)) \cdot \|u - v\|_2. \quad (7)$$

2264 By the linearity of f together with Eq.(5), (7) and (6), we have that

$$2266 \quad (1 - O(1)) \cdot \|u - v\|_2^2 \leq \|f(u') - f(v')\|_2^2 \leq (1 + O(1)) \cdot \exp(\sqrt{\log n})\|u - v\|_2^2.$$

2267 Rescaling it, we get the desired result. Thus we complete the proof. \square

2268
2269

I.2 SPARSIFIER IN ADVERSARIAL SETTING

2270

Here in this section, we provide our result of spectral sparsifier that can handle adversarial updates.

2271

Theorem I.5 (Sparsifier in adversarial setting, formal version of Theorem 2.2). *Let α be the aspect ratio of a d -dimensional point set P defined above. Let $k = O(\sqrt{\log n})$. If $\alpha^d = O(\text{poly}(n))$, then there exists a data structure DYNAMICGEOSPAR that maintains a ϵ -spectral sparsifier of size $O(n^{1+o(1)})$ for a (C, L) -Lipschitz geometric graph such that*

2272

- DYNAMICGEOSPAR can be initialized in

2273
2274

$$O(ndk + \epsilon^{-2}n^{1+o(L/k)} \log n \log \alpha)$$

2275

time.

2276

- DYNAMICGEOSPAR can handle adversarial point location changes. For each change in point location, the spectral sparsifier can be updated in

2277
2278

$$O(dk + 2^{O(k)}\epsilon^{-2}n^{o(1)} \log \alpha)$$

2279

time. With high probability, the number of edges changed in the sparsifier is at most

2280
2281

$$\epsilon^{-2}2^{O(k)}n^{o(1)} \log \alpha.$$

2282
2283
2284*Proof.* By Lemma I.4, the estimation data structure works for adversarial query points. Then the theorem follows by Lemma E.5, E.11, and E.12. \square

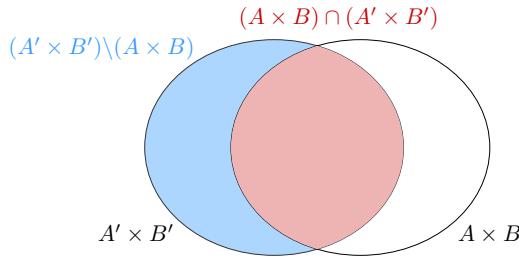
2285

J FIGURES

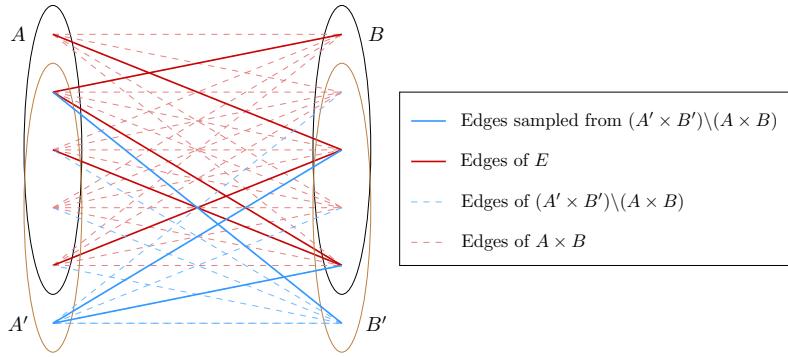
2286

We list our figures here.

2287

Figure 1: Division of the new biclique $(A' \times B')$: divided it into two parts (Blue part: $(A' \times B') \setminus (A \times B)$ and red part $(A \times B) \cap (A' \times B')$). And we sample from them respectively.

2288



2289

Figure 2: Resampling the biclique: E (The red edges) is uniformly sampled from $\text{Biclique}(A, B)$. After $A \times B$ becomes $A' \times B'$, we resample from $E \cap (A' \times B')$ with specific probabilities.

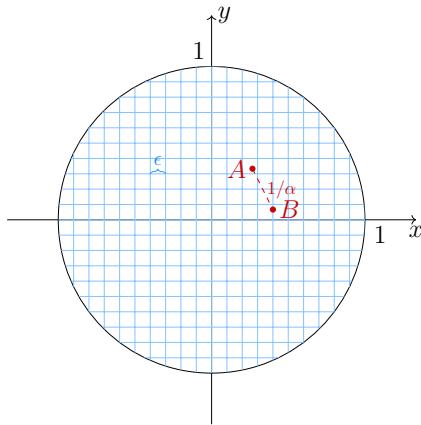


Figure 3: The net argument in our problem. Let $d = 2$. Here we restrict all the points to be in the ℓ_2 unit ball. By the definition of aspect ratio α , we know the minimum distance between two points is $1/\alpha$ (A and B in the figure). Thus, by setting $\epsilon = C \cdot \alpha^{-1}$ for some constant C small enough, every pair of points is distinguishable.

LLM USAGE DISCLOSURE

LLMs were used only to polish language, such as grammar and wording. These models did not contribute to idea creation or writing, and the authors take full responsibility for this paper's content.