Project Summary — James R. Lee

Geometric spaces arise in computer science through a number of avenues. The most obvious of these occurs when the input data for a problem possesses an inherent geometric structure, e.g. the hop-distance between nodes in a network, or the similarity distance between two genomic sequences. Certainly these "geometries" may be far afield from the classical setting where one is concerned, for instance, with Euclidean spaces or smooth manifolds, and thus we require new techniques for the design of efficient algorithms on these spaces, and new models for the kind of geometries we should expect.

In the setting of combinatorial optimization, on the other hand, high-dimensional geometry often presents itself in a less straightforward manner. A basic example is the use of convex optimization (e.g. linear and semi-definite programming) in the solution—exact or approximate—to a variety of combinatorial problems. The most profound connections occur when a problem with an *a priori* purely combinatorial structure is shown to involve rich geometric phenomena. Furthermore, these connections arise in many of the most fundamental and important problems in optimization.

The proposed research addresses (1) the understanding of these connections and the development of new algorithmic techniques to exploit them, (2) the extent to which geometric obstacles are fundamentally relevant to the computational hardness of these problems, and (3) the use of geometric tools and intuition in understanding and manipulating large, dynamic systems (such as ad-hoc networks and scientific databases).

Intellectual merits. High-dimensional geometry has proven to be a unifying theme in producing the best efficient algorithms for a number of classical and fundamental problems. In particular, when a problem is NP-hard, we often seek solutions which are as close to optimal as possible, and in this setting of approximation, geometric optimization plays a central role. Furthermore, in recent years a stunning alignment has emerged between the geometric obstacles which prevent the design of better algorithms, and the complexity-theoretic indications which hint that perhaps no better algorithms exist. This suggests that the geometric approach is not only a useful algorithmic tool, but an essential one, lying at the heart of basic computational problems—vertex cover, finding maximum cuts, or 3-coloring the vertices of a graph—of which we still have a limited understanding.

In a different vein, given the massive size of data sets arising in a variety of domains, it becomes important to understand the structure of the data involved so that algorithms can exploit special features in order to avoid dismal worst-case complexities. Often geometric properties (e.g. the "intrinsic dimensionality" of a data set) provide an essential algorithmic foothold, and a basis for the design of light weight data structures.

Broader impacts. There is strong reason to believe that the algorithmic tools developed here will have broader application to fields such as vision, machine learning, and databases. The PI's previous and ongoing research, especially in the area of geometric search, has already fostered interactions with members of these communities, and prompted the inclusion of the PI's algorithms in software projects. We hope to expand these collaborations in the future, and to continue to develop algorithms which are useful in practical settings.

Secondly, the proposal includes the integration of research and teaching through graduate courses and the development of summer projects for undergraduates, as well as the use of courses to introduce new techniques to Ph.D. students outside of theory. We intend all resulting educational materials to be freely available and widely disseminated online, suitable for self-study or as a guide to similar courses taught elsewhere.

TABLE OF CONTENTS

For font size and page formatting specifications, see GPG section II.B.2.

Appendix Items:

*Proposers may select any numbering mechanism for the proposal. The entire proposal however, must be paginated. Complete both columns only if the proposal is numbered consecutively.

CAREER: Geometric phenomena in algorithms and complexity

James R. Lee

1 Introduction

Geometric spaces arise in computer science through a number of avenues. The most obvious of these occurs when the input data for a problem possesses an inherent geometric structure, e.g. the hop-distance between nodes in a network, or the similarity distance between two genomic sequences. Certainly these "geometries" may be far afield from the classical setting where one is concerned, for instance, with Euclidean spaces or smooth manifolds, and thus we require new techniques for the design of efficient algorithms on these spaces, and new models for the kind of geometries we should expect.

In the setting of combinatorial optimization, on the other hand, high-dimensional geometry often presents itself in a less straightforward manner. A basic example is the use of convex optimization (e.g. linear and semi-definite programming) in the solution—exact or approximate—to a variety of combinatorial problems. The most profound connections occur when a problem with an *a priori* purely combinatorial structure is shown to involve rich geometric phenomena. Furthermore, these connections arise in many of the most fundamental and important problems in optimization.

The proposed research addresses (1) the understanding of these connections and the development of new algorithmic techniques to exploit them, (2) the extent to which geometric obstacles are fundamentally relevant to the computational hardness of these problems, and (3) the use of geometric tools and intuition in understanding and manipulating large, dynamic systems (such as ad-hoc networks and scientific databases). We now summarize the basic high-level goals of the proposed research.

Basic goals.

- Develop new approximation algorithms for the central divide-and-conquer and clustering-type problems in graphs (e.g. Sparsest Cut). Further explore the relationship with semi-definite programs (SDPs), geometric embeddings, and graph flows. Examine the inherent limitations of algorithms based on these techniques.
- Understand the ultimate power of SDPs in general, and for specific classical problems—like minimum vertex cover, graph coloring, and satisfying linear equations modulo a prime—for which our understanding is still quite limited. Investigate the mysteriously strong relationship between the unique games conjecture and SDP-based approximation algorithms.
- Explore geometric algorithms for problems like near-neighbor search, clustering, and locality-sensitive hashing in non-Euclidean settings (e.g. general metric spaces, edit distances on strings, or negatively curved manifolds) and the extent to which various geometric restrictions appropriately model real-world data.
- Investigate rigorous notions of intrinsic dimension, dimensionality reduction, and the development of efficient algorithms and data structures for intrinsically low-dimensional data sets.

- Involve students of all levels in the proposed research, including experimental projects suitable for undergraduates (even those with non-theory interests), and cutting edge research collaboration with Ph.D. students.

- Develop courses and seminars which expose students to the latest developments in theoretical research, reinforce core concepts across all of theory, and introduce non-theory students to new algorithmic approaches that they can incorporate into their own research.

Institutional context. The University of Washington provides a superb setting for the research and teaching activities described in the proposal. First of all, they are extremely supportive and committed to my personal success, as well as the integration of my work with the larger mission of the department. The department is committing $5000 in matching funds to go along with an additional $5000 from the College of Engineering.

Secondly, the academic community in the department, the university, and the surrounding research centers (like Microsoft Research) provides a very strong setting for theoretical research. In particular, I expect fruitful collaborations with Paul Beame (on issues related to the power of lift-and-project SDP relaxations), Venkat Guruswami (on hardness of approximation and approximation algorithms), as well as colleagues in the nearby Microsoft Research theory group like Uriel Feige and Yuval Peres (with whom the PI has had considerable collaborations in the past). Furthermore, the computer science department at UW has outstanding computational biology and database groups with whom I expect to interact on a regular basis.

Proposal organization. In Sections 2 and 3, we give a detailed outline of the proposed research activities, their intellectual merits, and their place in the broader context of computer science. Although these descriptions paint a broad picture of the proposed work, as with any quality research program, new discoveries will often take us in unexpected directions. Section 4 outlines an educational plan and summarizes the broader impacts of this proposal. Finally, Section 5 details previous research of the PI done under NSF support.

2 High-dimensional geometry in combinatorial optimization

This section describes proposed research in combinatorial optimization and mathematical programming, chiefly the use of linear programs (LPs) and semi-definite programs (SDPs) in developing approximation algorithms for important NP-hard problems. The basic goals are to (1) design new algorithms for centrally important optimization problems, (2) to explore the limitations of our current algorithmic models, and (3) to further develop an understanding of whether these limitations are inherent in the computational complexity of these problems.

Approximation algorithms. When an optimization problem is NP-hard, one accepts that finding the optimal answer may be difficult, and instead settles for *approximately optimal* solutions, with the goal of efficiently finding as good a solution as possible.

We take the graph coloring problem as an example. Given a graph G, a *proper k-coloring of G* is an assignment of the numbers $\{1, 2, \ldots, k\}$ to the vertices of G so that no two adjacent vertices share the same color. Our goal is to color the vertices of G with as few colors as possible. Let $\mathsf{OPT}(G)$ be the smallest value of k for which G admits a proper k-coloring. For any algorithm A, we define $\mathsf{A}(G)$ to be the number of colors that A uses in properly coloring G. The goal in designing approximation algorithms is to produce an efficient (polynomial time) algorithm A for which $\frac{\mathsf{A}(G)}{\mathsf{OPT}(G)}$ is as small as possible. Often this *approximation ratio* will grow with n (the size of the input), in which case we express it as a function of n. If a problem is NP-hard, then we cannot achieve an approximation ratio of 1 unless P = NP. In some cases, it is possible to achieve a ratio of $1+\varepsilon$ for every $\varepsilon > 0$, in which case the problem is said to admit a *polynomial-time approximation scheme* (PTAS).

In the past 15 years, outstanding progress has been made in understanding the approximability of NP-hard problems, via both the production of new algorithms with improved performance guarantees, and new lower bounds based on probabilistically checkable proofs (PCPs) which show that even producing approximate solutions to a number of problems remains NP-hard. Despite this effort, there are still significant gaps in our understanding of fundamental optimization problems, and in the ultimate power of the current algorithmic paradigms.

2.1 Balanced separators and geometric embeddings; dividing & conquering

One of the fundamental techniques in designing combinatorial algorithms is that of "divide and conquer." In general, the "divide" step of any such approach must balance two competing goals. On the one hand, the division should be non-trivial, in the sense that the subproblems are indeed smaller (and thus easier to solve) than the original, and on the other hand, the subproblems should be independent enough that one doesn't need to exert too much effort in combining the subsolutions into a single global solution.

Balanced separators and sparsest cut. A basic combinatorial problem of this flavor can be stated as follows. Let $G = (V, E)$ be a graph. Consider partitioning V into two disjoint pieces $V = A \cup B$ such that $|A|, |B| \leq \frac{2}{3}|V|$, and such that $|E(A, B)|$, the number of edges running from A to B, is as small as possible. This is the balanced graph separator problem.

It is difficult to overstate the importance of this problem throughout computer science. Finding such separators is a fundamental step in a number of approximation algorithms for NP-hard problems; we refer to the papers [75, 2], the survey [93] and chapters 20–21 in the book [100]. Even in cases where this problem cannot be used directly as a subroutine, the techniques developed in its solution are often widely applicable (see, e.g. [48, 38]). Furthermore, variants of this problem are centrally important in vision and image processing [92], finite-element methods [95], and solving sparse linear systems [96]. The same notions (in the context of "graph expansion") appear in complexity theory [85, 84], network flows [75], and study of mixing times of Markov chains [94], to name a few.

An intimately relation notion is that of the *sparsest cut* of a graph. In addition to G, suppose that we have a subset $D \subseteq V \times V$ of pairs of vertices called *demands*. Then for a subset $S \subseteq V$ (which we think of as a cut in G), we define the *sparsity of* S as

$$\Phi_{G,D}(S) = \frac{|E(S, \bar{S})|}{|D(S, \bar{S})|},$$

where $|D(S, \bar{S})|$ represents the number of demand pairs $\{x, y\} \in D$ which cross the cut $(S, \bar{S})$.

The idea is that we should separate as many demand pairs as possible (representing the fact that we have "significantly" broken the graph into two pieces) while minimizing the number of edges cut (minimizing the interaction between the two sides). As one might expect, this kind of tradeoff leads to an NP-hard optimization problem. Thus we look for *approximately optimal solutions,* i.e. cuts $(S, \bar{S})$ for which

$$\Phi_{G,D}(S) \leq C \cdot \min_{S \subseteq V} \left[\Phi_{G,D}(S)\right].$$

When $D = V \times V$, i.e. every pair of vertices is a demand pair, this is referred to as a *uniform* instance of the Sparsest Cut problem. It is well-known that approximation algorithms for the uniform case lead directly to algorithms with similar guarantees for the balanced separator problem.

The relation to high-dimensional geometric embeddings. In [76], Linial, London, and Rabinovich (LLR) offered the following approach to finding sparse cuts (and, by proxy, balanced separators): Find a (high-dimensional) geometric representation of the graph G, i.e. an identification of the nodes of G with vectors in $\mathbb{R}^n$ such that (1) the nodes are fairly spread out and (2) on average, adjacent vertices in G are represented by vectors which are close together (in the ℓ_2 norm, say). Once the vectors are represented in this way, the LLR approach is to choose an appropriate $(n-1)$-dimensional hyperplane and to take the cut in G induced by the hyperplane, i.e. the sets of vertices corresponding to vectors on each side of the hyperplane becomes the cut $(S, \bar{S})$ in G. The hyperplane is usually chosen at random according to an appropriate distribution.

In order to produce the geometric representation (i.e. an identification of the vertices of G with vectors in $\mathbb{R}^n$), one first solves an appropriate linear program to arrive at a metric $d(x, y)$ on the vertices of G (i.e. a distance function satisfying the triangle inequality). Then the metric space (V, d) is *embedded* into $\mathbb{R}^n$ (equipped with an appropriate norm). In this scenario, an embedding is simply a map $f : V \to \mathbb{R}^n$. In order for the eventual hyperplane cut to be nearly optimal, the mapping f must preserve the structure of the metric space (V, d). To quantify this, we define the *distortion of* f as the smallest constant D for which

$$\frac{d(x,y)}{D} \leq \|f(x) - f(y)\|_2 \leq d(x,y) \qquad \forall x, y \in V. \tag{1}$$

In other words, a distortion D map preserves all pairwise distances up to a factor of D. The LLR analysis (see also [11]) shows that the approximation ratio of this algorithm is precisely the distortion D of the mapping f. Thus the problem of finding sparse cuts (and hence balanced separators) becomes intimately tied to the construction of good (low-distortion) geometric embeddings of finite metric spaces. If the p-norm is used in line (1) instead of the 2-norm, we call f an embedding into ℓ_p.

Bourgain's theorem [15] states that every n-point metric space admits an embedding into ℓ_2 with distortion $O(\log n)$, and this allowed [76, 11] to arrive at an $O(\log n)$-approximation to the Sparsest Cut problem, improving upon earlier work of Leighton and Rao [75] which gave a similar result for the uniform case. We now turn to our research goals in this area.

Approximation algorithms for the uniform Sparsest Cut problem.

The uniform case of Sparsest Cut (which includes the balanced separator problem) is one of the most fundamental open problems in approximation algorithms. A recent geometric breakthrough of Arora, Rao, and Vazirani (ARV) [9] allowed the authors to give a new analysis of the well-known Sparsest Cut SDP, yielding an $O(\sqrt{\log n})$-approximation for the uniform case. There is no known hardness of approximation result for this problem—it is possible that even a PTAS exists.

The approach used in [9] has a natural barrier (the hypercube) which prevents it from being pushed any further than $O(\sqrt{\log n})$ (although other improvements, made by the PI, were possible [66]). Thus significant new ideas are required in order to make further progress.

A possible first step is to show that if an ARV-type approach fails to do better than $O(\sqrt{\log n})$ when rounding the SDP, then the fractional SDP solution must contain large hypercubes with small distortion (e.g. dimension $\Omega(\log n)$). Possibly these hypercubes give better information on which integral cuts an algorithm should choose. Similar approaches were taken in the local theory of Banach spaces (see, e.g. [17]), and this may provide additional ideas.

It is also interesting to consider limitations on the power of the Sparsest Cut SDP. Recent work of Devanur, et. al. [32] gives a lower bound of $\Omega(\log \log n)$ on the integrality ratio, but this leaves an exponential gap between the known upper and lower bounds. Improving this gap seems to require new advances in fourier analysis of boolean functions, an important topic throughout hardness of approximation, learning, and complexity theory. In particular, when the total influence (see [16]) of a boolean function $f : \{0,1\}^d \to \{0,1\}$

exceeds $\Omega(\log d)$, we are no longer able to prove strong bounds on the structure of the function, as opposed to when the influence is $\ll \log d$ (see e.g. [79, 39]).

The general Sparsest Cut problem.

The best approximation ratio for the general case is $O(\sqrt{\log n}\log\log n)$, given by the PI, in joint work with Arora and Naor [8]. Our approach is based on the construction of better ℓ_2 embeddings of negative-type metrics (this is a natural kind of metric space that is produced by the Sparsest Cut SDP). The best-known lower bound on the embedding/SDP approach is $\Omega(\log\log n)$ [62].

A short-term goal is to improve this ratio to $O(\sqrt{\log n})$, which should be possible using enhancements of existing techniques (in a recent manuscript [67], the PI has already improved the bound to $O(\sqrt{\log n\log\log n})$). Longer term goals include determining the integrality gap of the SDP in the general case, and developing PCP-based hardness of approximation results which show that no better algorithm exists.

In recent work with Naor [73], the PI gave a new (super-constant) integrality gap for the general case based on a certain sub-Riemannian geometry on the 3-dimensional Heisenberg group. We believe that eventual analysis will allow this example to yield a gap of the form $\Omega(\log n)^{\delta}$ for some $\delta > 0$, moving the known upper and lower bounds to within polynomial factors. We hope that refined techniques from geometric group theory, e.g. the coarse differentiation theory of [35], will allow the PI to push the lower bound to $\Omega(\sqrt{\log n})$.

Unlike the uniform case, for which we have very limited tools in proving hardness of approximation (in particular, the construction of PCPs with strong expansion properties), it should be possible to prove NP-hardness of approximating the general Sparsest Cut problem (i.e., it is not necessary for the lower bound graphs to have strong expansion in this case). The PI intends to study this problem, using ideas from the previously mentioned integrality gap, along with recent hardness of approximation techniques of Chuzhoy and Khanna for directed cut problems [28].

The Sparsest Cut problem in planar graphs.

One of the most well-known open conjectures about geometric embeddings of finite metric spaces concerns the shortest-path metrics on planar graphs. In [43], it is conjectured that the shortest-path metric on any planar graph admits an embedding into ℓ_1 with $O(1)$ distortion (independent of the size of the graph). If true, this would yield an $O(1)$-approximation for Sparsest Cut in planar graphs, as well as providing a very strong approximate multi-commodity max-flow/min-cut theory [75, 76] for planar graphs.

The PI intends to resolve this conjecture either negatively or positively, and we believe that we now possess the tools required for this endeavor. In particular, a disproof would follow the multi-scale ℓ_1 lower bound techniques [26, 25] employed in our joint work with Naor [73]; the problem here is to construct an appropriate family of planar graphs for which a certain kind of discrete differentiation theory applies. In the PI's previous attempts to construct such examples, the Okamura-Seymour theorem [80] (which implies that every outerplanar graph embeds into ℓ_1 isometrically) has presented itself as a significant obstacle, and this may lead to techniques which yield a proof of the conjecture (the work of [27] on disjoint paths in planar graphs seems particularly relevant here).

Vertex separators and treewidth decompositions.

Up to this point, we have discussed only *edge separators* in graphs, but there are corresponding notions for *vertex separators*. In this case, the goal is to separate the graph into two large pieces by removing a small subset of the vertices. Finding small vertex separators is a key step in algorithms which construct approximately optimal *treewidth decompositions* of graphs (see [4, 14]), which are a fundamental component of many divide-and-conquer approaches to NP-hard problems (see, e.g. the discussion in [37]). In a formal sense, vertex separator problems are more difficult to solve than the corresponding edge versions (there is a simple reduction from the edge case to the vertex case).

In joint work with Feige and Hajiaghayi [37] (see also the prior work in Section 5), we show how a different kind of embedding (called *average distortion line embeddings*) can be used to produce approximation algorithms for vertex separator problems. In particular, we give an $O(\sqrt{\log n})$-approximation for the *uniform case* of vertex Sparsest Cut (leading to a similar approximation for balanced vertex separators, and approximate treewidth decompositions).

There are two main problems in this setting. First, we would like to extend the edge separator theory to the case of vertex separators, but using current techniques, we are only able to accomplish this partially. In particular, in [37], we show that the general vertex Sparsest Cut problem can be approximated within $O(\log n)$, but we cannot improve this (unlike the case of general edge Sparsest Cut). Improvements should come from a more refined study of average distortion line embeddings (in particular, an extension of the measured descent techniques [60] to this setting).

The second problem is to prove hardness of approximation for vertex separators—there is reason to believe that this could be significantly simpler than for the edge case, and could be a first step towards hardness results in the edge setting. To justify this assertion, one can observe that for the SDP relaxation for vertex Sparsest Cut, we are able to give a tight analysis of the integrality gap (*including* a matching lower bound), in drastic contrast to the edge case.

2.2 The ultimate power of SDPs and the unique games conjecture

Semi-definite programming has proven to be a powerful tool for the development of new algorithms, especially approximate algorithms for the optimization version of NP-hard constraint satisfaction problems like 3-SAT [102], Maximum Cut [41], Coloring 3-colorable graphs [49, 7], and the previously mentioned Sparsest Cut problem [9]. *But are SDPs the ultimate answer, or should we look elsewhere?*

We propose to study two sources (and their relationship with each other) which might provide a key to this question: (1) The power of systematically strengthened SDPs [78, 40] based on methods like the progressively stronger Lovasz-Schrijver lift-and-project relaxations [6, 3, 99, 19], and (2) the mysterious relationship between the Unique Games Conjecture (UGC) of Khot [51], SDP-based approximation ratios, and integrality gap "reductions" (see, e.g. [52, 79, 24, 54, 32, 53, 22, 87]). A key focus of our study will involve the use of the UGC as a basis for proving strong integrality gaps for multi-round lift-and-project relaxations, and the development of a systematic method for producing "bad" fractional solutions based on PCP reductions. The PI has recently initiated such work with Sanjeev Arora who will be on sabbatical next year at nearby Microsoft Research. Due to space limitations, we simply outline the some of the main problems that we intend our study to address.

1. **Vertex cover:** Given a graph $G = (V, E)$, find the smallest subset of vertices $S \subseteq V$ so that every edge is adjacent to at least one vertex of S.

   There is a classical factor 2 approximation known for this problem; assuming the UGC [53], this is tight, but the best NP-hardness result gives hardness of approximation within 1.36 [33], and the lower bounds for systematic SDP relaxations are still quite weak [6, 99].

2. **Coloring 3-colorable graphs:** Given a graph $G = (V, E)$ which is promised to have a proper 3-coloring, find a proper coloring with as few colors as possible.

   The best approximation algorithm colors such a graph with approximately $O(n^{0.2})$ colors [7]. It is NP-hard to 4-color a 3-colorable graph [44]; assuming an appropriate version of the UGC [46], the problem is hard to approximate within any constant factor.

3. The **uniform Sparsest Cut** problem, discussed in the preceding section.

3 Algorithms on metric spaces and intrinsic dimensionality

It is possible to model a diverse range of data sets by a metric space, i.e. a set X of points, together with a real-valued distance function $d(x, y)$ on pairs $x, y \in X$ which satisfies the triangle inequality $d(x, y) \leq d(x, z) + d(z, y)$ for all $x, y, z \in X$. The present section describes proposed research into the existence of efficient algorithms and data structures on metric spaces. One of the key obstacles this type of research must confront is the following: For a number of problems, e.g. nearest-neighbor search, it is trivially apparent that in the worst case, an algorithm will have to perform brute-force search on the data set in order to find an appropriate match. Thus we must impose some geometric restrictions on our metric space in order to pursue non-trivial algorithmic approaches. The question of which restrictions one is allowed to impose—and which can be harnessed for efficient algorithms—lies at the intersection of algorithm design and empirical study.

3.1 Intrinsic dimensionality

One fruitful approach to understanding the structure of data sets is by studying their *intrinsic dimensionality.* We know from classical computational geometry, than many problems can be solved very efficiently when the input is a point set in $\mathbb{R}^k$ for small values of k. But, of course, this is an unrealistic assumption to place on a variety of data sets. One might ask whether similar approaches might work if the input metric space is *intrinsically* low-dimensional, without actually being represented as a low-dimensional point set.

For this to work in a rigorous setting, we need a suitably general definition of dimension for general metric spaces which, at the same time, allows one to harness the property of intrinsic low-dimensionality for use by algorithms. In joint work with Gupta and Krauthgamer [42], the PI introduced the notion of the *doubling dimension* of a metric space, based on a concept of Assouad from geometric analysis [10]. We briefly recall this notion now.

Consider a finite metric space (X, d). For a subset $S \subseteq X$, we write $\operatorname{diam}(S) = \max_{x,y \in X} d(x, y)$ for the *diameter of* S. Now one defines the *doubling constant of* X as the least value λ_X such that every subset $S \subseteq X$ can be written as a union $S = S_1 \cup S_2 \cup \cdots \cup S_{\lambda_X}$ where $\operatorname{diam}(S_i) \leq \frac{1}{2} \operatorname{diam}(S)$ for every $1 \leq i \leq \lambda_X$, i.e. λ_X is the smallest number such that every subset can be covered by λ_X subsets of half the diameter. Finally, we define $\mathsf{dim}(X) = \log_2(\lambda_X)$ as the *doubling dimension of* X. This definition has a number of nice properties, including the fact that $\mathsf{dim}(X) \leq \mathsf{dim}(Y)$ when $X \subseteq Y$, and $\mathsf{dim}(\mathbb{R}^k, \ell_p) = \Theta(k)$ for any $p \geq 1$. In this sense, metric spaces of small doubling dimension generalize low-dimensional Euclidean spaces.

Example: Nearest-neighbor search. The $(1 + \varepsilon)$-NNS problem is defined as follows: Given a universe of points (X, d), and a subset $S \subseteq X$ of n points, we want to preprocess S so that given a query $q \in X$, we can very efficiently return a point $a \in S$ for which $d(q, a) \leq (1 + \varepsilon) \cdot d(q, a^*)$, where a^* is the closest point to q from S. In words, we want to return a $(1 + \varepsilon)$-approximate nearest neighbor to the query. In joint work with Krauthgamer [57, 58], the PI showed that one can design efficient, dynamic approximate nearest-neighbor search (NNS) data structures on spaces with small doubling dimension. This suggested that the doubling dimension captures a notion of algorithmic tractability. Subsequent work by the PI and others [45, 29, 97, 21] made outstanding progress on designing efficient algorithms for other problems on spaces of small doubling dimension. Additionally, modifications of our basic NNS algorithms were implemented in the context of learning [13], and in the software project isomap.

Research topics. We now discuss a selection of our proposed research in the area of intrinsic dimensionality.

A comprehensive methodology for transferring results from the Euclidean setting to spaces of low intrinsic dimension.

Presently, the algorithms for spaces of small doubling dimension (for, e.g. approximate NNS [45, 29, 57, 58], compact routing [21], spanners [20]) are similar to their counterparts in classical (Euclidean) computational geometry, but still seem to require highly non-trivial technical steps in order to be "adapted" to the more general case. We propose to find a general set of tools that sheds light on the similarity and differences between the two settings.

I believe that a key a step in this process is to prove the existence of a PTAS for the traveling salesman problem (TSP) in metric spaces of small doubling dimension. Currently, only a quasi-PTAS (i.e. super-polynomial running time) is known [97], because Arora's framework [5] which provides a PTAS in the Euclidean case, cannot yet be completely carried over to the intrinsic dimension setting; Arora's more advanced techniques [5] seem to rely very delicately on properties of Euclidean geometry.

To restore some of the ambient Euclidean structure, we propose to study geometric algorithms on *conformal deformations* of Euclidean spaces. A *doubling measure* μ on $\mathbb{R}^k$, is one which satisfies $\mu(2B) \leq C \cdot \mu(B)$ for some constant C, and every ball[1] B in $\mathbb{R}^k$. Such a measure gives rise to a new distance function d_μ on $\mathbb{R}^k$ defined by

$$d_\mu(x, y) = \inf_\gamma \int_\gamma \mu^{1/k} \, ds$$

where γ is parameterized by arclength, and the infimum is over all rectifiable paths γ connecting x to y (informally, we are taking shortest paths, where the length of an infinitesimal segment is given by $\mu^{1/k}(B)$ where B is the smallest ball containing the segment). The main point is that, by a result of Semmes [88], every metric of small doubling dimension is close to a conformal deformation of some Euclidean space (and, in fact, one can assume a very strong property of the measure μ). Thus to design algorithms for general spaces of small intrinsic dimension, it suffices (essentially) to give appropriate algorithms for Euclidean deformations. It follows from an unpublished result of the PI, that there is an efficient algorithm which, given a metric space X, produces a Euclidean deformation that is close to X, and for which the dimension k of the embedding satisfies $k = O(\mathsf{dim}(X))$.

Non-linear dimension reduction in Euclidean spaces.

Even when we have a point set $X \subseteq \mathbb{R}^k$, it is useful to consider $\mathsf{dim}(X)$. If $k \gg \mathsf{dim}(X)$, then we would be foolish to run algorithms which depend on the embedding dimension k; rather, we should exploit the intrinsically low-dimensional structure which is being obscured by the high-dimensional representation. This is a scenario that occurs frequently when our data are measurements of a feature set taken from some process. For instance, consider a camera moving around an object, recording, say, 50 different features (e.g. hue, saturation, redness, shininess, etc.). While the camera might be mapping out some low-dimensional manifold in the feature space, the natural dimension of the recorded data will be large (50).

It is a natural and extremely interesting question whether there exists a map $f : X \to \mathbb{R}^{k'}$ which approximately preserves the structure of X, but for which $k' \approx \mathsf{dim}(X) \ll k$. We would like the map f to have low distortion, i.e. $\|f(x) - f(y)\|_2 \approx \|x - y\|_2$ for every $x, y \in X$. Whether such a map always exists (with distortion and dimension depending only on $\mathsf{dim}(X)$) was originally asked (in a different form) by Lang and Plaut [64], and independently by the PI with Gupta and Krauthgamer [42]. Although the PI and coauthors have worked on related questions before (see [56, 42]), previous techniques do not seem to make progress on this problem. In particular, it is not difficult to see that the map f must be non-linear (as opposed

[1]By $2B$, we mean the ball with the same center as B and twice the radius.

to dimension reduction techniques like principal component analysis and multi-dimensional scaling [30], or random projections [101]).

We now believe that the answer is negative, but proving this is an important first step in determining what kind of dimension reduction is possible. We believe that the lower bound lies in recent work of the PI with Naor [73] concerning sub-Riemannian geometry on 3-dimensional Heisenberg groups. Previously, we employed this geometry to provide a new counter-example to the Goemans-Linial conjecture for Sparsest Cut, but there is a natural embedding of this space into a (high-dimensional) Euclidean space that could provide a counter example for the dimension reduction problem as well. Future work would then seek weaker properties that could be preserved by a dimension reducing map, or stronger assumptions that would allow for a low-distortion embedding into a low-dimensional space.

Dimension reduction in ℓ_p spaces, $1 < p < 2$.

The famous Johnson-Lindenstrauss flattening lemma [47] tell us that given any n-point subset $X \subseteq \mathbb{R}^n$ (considered with the ℓ_2-norm), it is possible to construct a map $f : X \to \mathbb{R}^{O(\frac{\log n}{\varepsilon^2})}$ that preserves all pair-wise distances in X up to a $1 \pm \varepsilon$ factor. Brinkman and Charikar showed [18] that a similar result is impossible in the ℓ_1 norm, and the PI, in joint work with Naor [71] gave a very simple proof of this fact. On the other hand, dimension reduction for ℓ_p norms with $1 < p < 2$ is still an open question, which we intend to study. Since p-stable random variables are much better behaved when $p > 1$, there is hope that perhaps the case $p = 2$ might be extended. On the other hand, the PI, in joint work with Mendel and Naor [69] showed that the dimension reduction map cannot be linear (as in the case $p = 2$), so this would require a vastly different approach.

Weaker notions of intrinsic dimension.

Finally, there is hope that even weaker notions of intrinsic dimension might allow the construction of efficient algorithms. One possibility is the *Assouad-Nagata dimension* [65]—spaces which are low-dimensional in this sense can exhibit exponential volume growth (unlike the case for low doubling dimension), and this includes e.g. non-positively curved manifolds, a setting which could be important in a variety of applications; see, e.g. [90]. An intriguing question is whether every surface (i.e. 2-dimensional manifold) of bounded genus admits a very efficient approximate NNS data structure—such spaces occur often in point sets derived from 3-dimensional models.

3.2 Other non-classical geometries

We now consider two other types of non-classical geometries that we intend to study, and their relationship to problems arising in networking and computational biology.

Negatively curved spaces.

The classical negatively curved manifold is $\mathbb{H}^2$, i.e. the 2-dimensional hyperbolic space, but in general one can consider complete, simply-connected Riemannian manifolds with negative sectional curvature. In recent work with Krauthgamer [59], the PI considered how well one can do computational geometry on finite-dimensional spaces of negative curvature, yielding a number of results (e.g. efficient nearest-neighbor search, and a PTAS for TSP) which are known to hold in the (flat curvature) Euclidean setting. This branch of work has just been initiated, and we propose research that would result in a comprehensive toolbox for computational geometers working in this class of spaces.

Broader impact: Recent empirical work suggests that these spaces yield natural models in networking [91, 12] and vision [90]. To understand how, e.g. the internet might be negatively curved, consider that its hierarchical structure makes it roughly tree-like on large scales, but with local interconnections between groups of geographically close machines. This is perfectly modeled by negatively curved spaces.

Edit distances on strings.

Perhaps one of the poorest understood geometries involve distances which are computed iteratively, by a sequence of local operations. Consider the edit distances between two strings S and T: This is the minimum number of insert/delete character operations needed to convert S into T. Denoting this distance by d_{ED}, we recall that if S and T are DNA sequences, then $d_{\mathrm{ED}}(S,T)$ is often used to model the proximity between them (where edits are intended to model mutation and recombination).

Natural problems arise in this setting, e.g. doing nearest-neighbor search, where we are given a query string Q and we want to find the closest match from some database. Consider also the problem of distance-preserving hashes, i.e. whether we can compress our set of strings to much shorter hashes (say $S \to \tilde{S}$) such that we are able to recover (approximately) $d_{\mathrm{ED}}(S,T)$ from the hashes $\tilde{S}, \tilde{T}$. A well-known approach to both these problems is to embed the distance d_{ED} on n-character strings into ℓ_1 with small distortion (recall the definition of distortion from Section 2.1).

Unfortunately, the best upper and lower bounds for embeddings into ℓ_1 are quite disparate. It is known that the edit distance on $\{0,1\}^n$ embeds into ℓ_1 with distortion at most $2^{O(\sqrt{\log n}\log\log n)}$ [81], and there is a lower bound of $\Omega(\log n)$ [62]. The PI intends to work on both sides of this problem. Improved upper bounds will inevitably lead us to a better understanding of the edit distance properties, and improved lower bounds will likely involve advances in harmonic analysis of boolean functions, a topic very closely connected with hardness of approximation and PCPs, and with the uniform case of the Sparsest Cut problem discussed in Section 2.1.

4 Education plan and broader impacts

We now discuss the broader impacts of our proposed work, both through research contributions to other fields, and through the education of students at all levels.

Broader impacts in research. Large portions of this work are relevant to fields across computer science; as discussed in Section 2.1, the proposed work on graph separators is intimately related to disparate fields like vision and image processing, finite-element methods, solving sparse linear systems, Markov chains, and statistical physics.

But perhaps the most immediate impact will come from the research proposed in Section 3 concerning algorithms on metric spaces, intrinsic dimensionality, and geometric models for data sets. There is strong reason to believe that the algorithmic tools developed in these settings will have broad application to fields such as vision, machine learning, ad-hoc networks, databases, and computational biology. The PI's previous and ongoing research, especially in the area of geometric search, has already fostered interactions with members of these communities, and prompted the inclusion of the PI's algorithms in software projects; see, for instance, the cover trees project [13], which uses a modified version of the search algorithms proposed by the PI with Krauthgamer [57], and its use in speeding up (see [13]) the isomap software [98]. We hope to expand these collaborations in the future, and to continue to develop algorithms which are useful in practical settings.

4.1 Education plan

In addition to the research plan outlined in previous sections, the PI will be actively involved in education at both the undergraduate and graduate level, as well as the integration of research and teaching through summer projects for undergraduates, and new graduate courses intended to introduce algorithmic techniques to students in other fields. We intend all resulting educational materials to be freely available and widely disseminated online, suitable for self-study or as a guide to similar courses taught elsewhere.

Graduate teaching. The PI intends to develop a range of new graduate courses which expose students to the latest developments in theoretical research, reinforce core concepts across all of theory, and introduce non-theory students to new algorithmic approaches that they can incorporate into their own research. All graduate courses will produce a set of lecture notes (written in conjunction with the students); by allowing free access to these notes online, we hope to further refine the teaching of similar courses elsewhere.

1. Geometric embeddings and approximation algorithms.

   This course, which will be offered some time during the 2007–2008 academic year, is primarily intended for graduate students interested in doing theoretical work at the forefront of algorithms research. The use of geometric embeddings in combinatorial optimization has seen tremendous growth in the past ten years to which the PI has been a witness and contributor. Related courses have been taught at CMU (by Gupta and Ravi), at MIT (by Indyk), and at Stanford (by Roughgarden).

   There is no standard treatment of the material to be presented in this course, and we expect that the lecture notes produced will eventually form the basis for a comprehensive textbook. Furthermore, the PI will take great care to design problem sets that take students to the very edge of current research, as the field is very active, and there exists a number of accessible open problems.

2. Core course: Randomized algorithms & probabilistic analysis.

   This is a new core theory course (which counts as official preparation toward the Ph.D. qualifying exam) that I will be teaching in the Spring of 2007. We will cover the basic tools and techniques of probabilistic analysis (moments, deviation inequalities, the probabilistic method, tail bounds, martingales, and concentration of measure, random walks, etc.) with an emphasis on modern problems and applications in computer science. This is an incredibly useful source of tools and techniques for any computer scientist, and an absolutely essential collection for students intending to do theoretical research in algorithms and complexity.

In future years, we plan to develop a course that looks at the algorithmic issues arising in ad-hoc networks (e.g. peer-to-peer systems and sensor nets). There will be a heavy project component, and we will seek to pair theoretical and systems researchers together into teams to carry out the project. A similar project component in the graduate algorithms course at Berkeley (taught by Papadimitriou and Vazirani) enjoyed success, and we hope to continue this trend in the fields most suited to these types of collaboration.

Reading groups and seminars. The PI also intends to host yearly reading groups on issues at the forefront of theoretical research. Possible topics for the coming year include the use of additive number theory in recent complexity results (e.g. derandomization, extractor construction, and PCP constructions) or the study of heuristics that solve NP-hard problems in average-case settings, and relationships with statistical physics.

Undergraduate teaching. In my first year, I will be teaching the undergraduate "Discrete Structures" course. This is an established part of the curriculum, and does not appear to require major revisions. The basic aim of the course is to introduce undergraduate majors to methods of formal proof, logic, probability, and basic graph theory.

In future years, I intend to update the undergraduate "Introduction to Algorithms" course to include a strong infusion of current algorithmic ideas in addition to classical material—for instance, the inclusion of topics related to how search engines use link analysis to improve search quality. My hope is to give students a modern feel for the subject that is often absent in more dated treatments. The field of algorithm construction is alive and kicking; and more relevant than ever to the development of practical solutions to real-world problems. I feel that an algorithms course should incorporate this spirit.

Finally, the University of Washington has intense "capstone" project courses for advanced undergraduate students. Students take these courses in their senior year to synthesize things they have learned or to study special topics. We refer to the following discussion on undergraduate research as an example of the topics that I might encourage a student to work on.

Undergraduate research. The proposed budget includes support for an undergraduate summer research project. The design of successful short-term projects for undergraduate is a difficult task—our goal is for students to both learn the relevant theoretical ideas and get experience with high-level research, but also for them to accomplish something substantial by the end of the project.

An example of such a project, related to the proposed research on intrinsic dimensionality and geometric models (Section 3.1) is as follows. A student would implement code to examine the geometric properties of a variety of scientific data sets that come equipped with a distance function or similarity measure. This includes, for instance, the doubling dimension of these spaces, or the degree to which they exhibit a combinatorial version of negative curvature.

Based on the presence or absence of these features, we would work together to develop a comprehensive picture of the kind of natural processes that exhibit these geometric restrictions. Depending on the students inclination, the project could continue in a practical direction: Developing software that uses known algorithms to exploit the structure of the data, or a more theoretical route: Developing new algorithms to address insufficiencies in previous assumptions and approaches.

5 Results from prior NSF support

As a graduate student at U. C. Berkeley from 2002–2005, I was supported by an NSF Graduate Research Fellowship. During this period, I worked on four main projects.

1. Geometric embeddings related to cuts, flows, and graph separators.

This project was concerned with high-dimensional geometric representations of graphs, and the use of these representations (via semi-definite programming) to design new approximation algorithms for classical NP-hard problems. This required the use and development of a number of geometric, combinatorial, and algorithmic tools. A brief introduction to these and related topics is given in Section 2.1. We now describe the main algorithmic consequences of this line of work; each is accompanied by a short description of the research involved.

An $O(\sqrt{\log n}\log\log n)$-approximation algorithm for the general Sparsest Cut problem in graphs.

This final result is based on a joint paper of the PI with Arora and Naor [8]. In particular, we gave a nearly optimal answer to the well-known question of how well n-point metrics of negative type embed into a Euclidean space. This improves the previously known $O(\log n)^{3/4}$ approximation, due to Chawla, Gupta, and Räcke [23]. These papers, in turn, are based on a novel embedding technique called *measured descent*, devised by the PI in collaboration with Krauthgamer, Mendel, and Naor [60], and on an improved version of the Arora-Rao-Vazirani structure theorem given by the PI in [66], which we discuss next.

An optimal version of the geometric structure theorem of Arora-Rao-Vazirani [9].

In [9], Arora, Rao, and Vazirani (ARV) give an $O(\sqrt{\log n})$-approximation to the uniform Sparsest Cut problem which is based on a new geometric structure theorem concerning certain subsets of Euclidean space (roughly, it argues that certain high-dimensional configurations of vectors cannot exist). In [66], the PI presented an optimal version of this theorem based on an enhancement of the ARV induction; as discussed above, this had applications to Sparsest Cut, including a simplified $O(\sqrt{\log n})$-approximation for the uniform case. This also improved other applications of the ARV technique, with implications for Min 2CNF Deletion and directed cut problems [1] and for vertex cover [48].

An $O(\sqrt{\log n}\log\log n)$-approximation algorithm for the Minimum Linear Arrangement problem.

The Minimum Linear Arrangement (MLA) problem concerns the construction of linear orderings of the nodes of a graph so that, on average, two adjacent vertices appear close together in the ordering. In joint work with Feige [38], we considered a new SDP for MLA, and analyzed it using techniques of Rao and Richa [83], along with the ARV geometric structure theorem. This improved over the best previous bound of $O(\log n)$ [83].

New approximation algorithms for vertex separators and treewidth decompositions of graphs.

In joint work with Feige and Hajiaghayi [37], the PI showed that the usual embedding approach does not work for vertex separator problems (e.g. the vertex version of Sparsest Cut, discussed in Section 2.1). To get around this, we show that a stronger type of embedding (called an *average distortion line embedding*) can be used as an algorithmic primitive for these types of problems. We then devised a new SDP for finding balanced vertex separators in graphs, and used these stronger embeddings to compute a separator from the SDP solution. We were thus able to extend the ARV result to give an $O(\sqrt{\log n})$-approximation for balanced vertex separators as well.

The most striking application of our technique is to the construction of *tree decompositions of graphs* with small *treewidth*, a major component in the Robertson-Seymour graph minor theory [86], and an important tool in divide-and-conquer algorithms for a number of problems. For instance, combined with the bi-dimensionality theory of [31], our results give the first polynomial-time approximation schemes for problems like *minimum feedback vertex set* and *connected dominating set* in excluded-minor graphs (e.g. bounded genus graphs). This can be seen as the ultimate extension of classical divide-and-conquer algorithms on planar graphs (e.g. [77]) to more general families, and the missing ingredient was an $O(1)$-approximation for treewidth in such families, which our results provide.

Volume-respecting embeddings and graph bandwidth.

In Feige's seminal paper on graph bandwidth [36], he introduced a stronger notion of embedding which he termed a *volume-respecting embedding*. Using the measured descent technique, we gave optimal volume-respecting embeddings for general metric spaces [60] answering an open question of Feige. In [68], the PI constructed nearly-optimal volume-respecting embeddings for finite subsets of Euclidean space, answering the main open questions of Krauthgamer-Linial-Magen [61], and Dunagan and Vempala [34]. This leads to a slight improvement in the best-known algorithm for graph bandwidth.

A new lower bound on the integrality ratio of the well-known SDP relaxation for Sparsest Cut.

The Goemans-Linial conjecture states that the integrality gap of a certain SDP for the Sparsest Cut problem is $O(1)$. In a breakthrough paper, Khot and Vishnoi [54] showed that the conjecture is false, and for n-vertex graphs, the gap is at least $\Omega(\log\log n)^{1/6}$. Their analysis was improved to $\Omega(\log\log n)$ by Krauthgamer and Rabani [62]. In joint work with Naor, the PI [73] presented a new super-constant integrality gap for the SDP which has a number of properties not present in the Khot-Vishnoi example. In particular, our lower bound construction is significantly simpler, and has the potential to yield an integrality gap of the form $\Omega(\log n)^{\delta}$ for some $\delta > 0$. This is a subject of ongoing research, and could eventually improve exponentially over the previous lower bounds.

2. Dimension reduction.

At its most basic level, dimension reduction studies the question of whether high-dimensional objects can be mapped into significantly lower dimensional spaces while still approximately preserving their fundamental properties. This is discussed in Section 3.1. We review our progress in this field.

Dimension reduction in ℓ_p norms.

The seminal result in this area is the Johnson-Lindenstrauss flattening lemma: Given any n-point subset $X \subseteq \mathbb{R}^n$ (considered with the ℓ_2-norm), it is possible to construct a map $f : X \to \mathbb{R}^{O(\frac{\log n}{\varepsilon^2})}$ that preserves all pair-wise distance in X up to a $1 \pm \varepsilon$ factor. Thus with only a small loss, the dimension of the point set can be reduced to from n to $O(\log n)$. The question of whether a similar result holds when the ℓ_2 norm is replaced by the ℓ_1 norm was a long-standing open problem.

Brinkman and Charikar [18] showed that such a result is impossible when distances are measured in the ℓ_1 norm, and in fact the dimension cannot be reduced past $n^{\Omega(1)}$. Their result involves a long and technical analysis of carefully constructed linear programs. In collaboration with Naor [71], we gave a very simple and intuitive geometric proof of the Brinkman-Charikar result. Our lower bound technique extends easily to other types of spaces, as shown in joint work with Mendel and Naor [69]. Additionally, we show that the lower bound construction used for these results cannot be extended to prove lower bounds for ℓ_p norms with $p \neq 1$.

Low-dimensional embeddings of general metric spaces.

As discussed in Section 3.1, the doubling dimension $\mathsf{dim}(X)$ of a metric space (X, d) is a useful notion of its intrinsic dimension. Given a space with small doubling dimension, it is natural to ask whether it admits a low-distortion embedding into a low-dimensional geometric space (e.g. $\mathbb{R}^k$ equipped with the ℓ_2 norm). Although this cannot hold in full generality (see [88, 89, 63]), Assouad [10] showed that a weaker version does hold: He proved that the space X equipped with the modified distance function $\sqrt{d(x, y)}$ admits an $O(1)$-distortion embedding into $\mathbb{R}^k$, where $k \leq 40^{\mathsf{dim}(X)}$.

In joint work with Gupta and Krauthgamer [42], the PI improves Assouad's results in two ways. First, we show that one can gain an exponential dependence in the dimension, i.e. we construct embeddings where $k = O(\mathsf{dim}(X))$. Secondly, we show that if X is a *tree metric*, i.e. the shortest-path metric on a weighted tree, then in fact (X, d) does embed into a low-dimensional space with low distortion (*without* modifying the distance function). These results show that spaces with low intrinsic dimension can be endowed with a certain kind of low-dimensional geometric representation.

Resolution of the Levin conjecture on intrinsic dimensionality of graphs.

Our work on intrinsic dimensionality and dimension reduction started with a conjecture posed by Levin, together with Linial, London, and Rabinovich in their seminal paper on geometric embeddings in computer science [76]. They propose a notion of dimension for graphs and ask whether the property of low dimensionality corresponds to a certain type of low-dimensional geometric representation. In joint work with Krauthgamer [56], we resolved the Levin conjecture and give optimal upper and lower bounds on the embedding dimension in terms of the intrinsic dimension.

3. Geometric algorithms on metric spaces.

This project concerns algorithms for geometric problems like nearest-neighbor search (described in Section 3.2) on special types of metric spaces. In classical computational geometry, one considers the Euclidean spaces $\mathbb{R}^k$ with $k = O(1)$, but for a number of application domains (e.g. ad-hoc networks or various scientific data sets), it is unreasonable to think of the data points as being represented in a geometric space.

Nearest-neighbor search algorithms in intrinsically low-dimensional spaces.

In joint work with Krauthgamer [57], we show that the approximate nearest-neighbor search problem can be solved efficiently on spaces with small doubling dimension. In particular, if (X, d) is an n-point metric space with $\mathsf{dim}(X) = O(1)$, we show how to build a data structure that answers $(1+\varepsilon)$-approximate NNS queries in time $O(\log \Delta) + (1/\varepsilon)^{O(1)}$, using only $O(n)$ space. Here Δ is the *aspect ratio* of the metric space X (i.e. the ratio of the largest to smallest distance in X). As a special case, this improves over the search structure of Karger and Ruhl [50]. In later work [58], we show that the query time can be reduced to $O(\log n)$ (*independent* of the aspect ratio Δ), but with larger storage requirements.

Computational geometry on negatively curved spaces.

Recent studies in networking and vision [91, 12, 90] suggest that there are a variety of interesting data sets that exhibit "negatively curved" properties, i.e. are naturally embeddable into negatively curved manifolds (like the classical d-dimensional hyperbolic space $\mathbb{H}^d$). Partially motivated by these studies, the PI, in joint work with Krauthgamer [59], studied whether classical computational problems which were known to be tractable on (flat) Euclidean spaces $\mathbb{R}^d$ might also have efficient algorithms on negatively curved spaces.

We give a novel set of algorithms and data structures for a variety of problems on these spaces. For instance, we show the existence of light-weight spanners and compact routing schemes; we demonstrate an efficient approximate NNS data structure; and we show how classical NP-hard problems like TSP can be approximated within $1+\varepsilon$ for any $\varepsilon > 0$ in polynomial time.

4. The mathematics of embeddings.

Sometimes in the development of new algorithms, pure mathematical questions arise that either (1) need to be solved in order to analyze a certain approach or (2) are independently interesting, but are particularly suited to attack by techniques from the computer science toolbox. We give two prominent examples of our work in these categories.

Extending smooth functions using randomized metric space partitions.

In joint work with Naor [72, 70], we showed how the theory of random partitions of metric spaces, developed in computer science, could be applied to yield new results in the study of extensions of Lipschitz maps taking values in Banach spaces. Consider the setup: (Y, d) is a metric space, $X \subseteq Y$, and Z is a Banach space (e.g. $Z = \ell_1$). Given a function $f : X \to Z$ satisfying $\|f(x) - f(y)\|_Z \leq d(x, y)$ for all $x, y \in X$, we would like to construct an extension $\tilde{f} : Y \to Z$ (so that $\tilde{f}(x) = f(x)$ for $x \in X$) which satisfies $\|\tilde{f}(x) - \tilde{f}(y)\|_Z \leq C \cdot d(x, y)$, where $C = C(Y)$ depends only on Y (but not on f). We show how $C(Y)$ relates to a geometric decomposability property of Y that arose in computer science research [55, 82], thereby extending and unifying a number of previous results.

Trees, high-dimensional random walks, and uniform convexity.

In joint work with Naor and Peres [74], we studied the conditions under which an infinite tree metric T embeds with finite distortion into an (infinite-dimensional) Euclidean space. We offer a complete characterization: Such an embedding exists if and only if T does not contain arbitrarily large (distorted) complete binary trees. A sharper quantitative version of this equivalence leads to the following algorithmic result: Given a finite tree metric T, it is possible to efficiently compute a nearly-optimal embedding of T into an ℓ_p space for any $1 \leq p < \infty$. Previously such a result was known only for $p = 1, 2$. A key point of the analysis concerns the asymptotic behavior of Markov chains defined on high-dimensional Euclidean spaces.

References

[1] A. Agarwal, M. Charikar, K. Makarychev, and Y. Makarychev. $O(\sqrt{\log n})$ approximation algorithms for Min UnCut, Min 2CNF deletion, and directed cut problems. In *STOC'05: Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, pages 573–581, New York, 2005. ACM.

[2] A. Agrawal, P. Klein, R. Ravi, and S. Rao. Approximation through multicommodity flow. In *31st Annual Symposium on Foundations of Computer Science*, pages 726–737. IEEE Computer Soc., Los Alamitos, CA, 1990.

[3] M. Alekhnovich, S. Arora, and I. Tourlakis. Towards strong nonapproximability results in the Lovasz-Schrijver hierarchy. In *STOC'05: Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, pages 294–303, New York, 2005. ACM.

[4] E. Amir. Efficient approximation for triangulation of minimum treewidth. In *Proceedings of the 17th Annual Conference on Uncertainty in Artificial Intelligence*, pages 7–15. Morgan Kaufmann, 2001. Journal version titled "Approximation Algorithms for Treewidth" is available at the author's homepage.

[5] S. Arora. Polynomial time approximation schemes for Euclidean traveling salesman and other geometric problems. *J. ACM*, 45(5):753–782, 1998.

[6] S. Arora, B. Bellobas, and L. Lovász. Proving integrality gaps without knowing the linear program. In *Proceedings of the 43th Annual IEEE Conference on Foundations of Computer Science*, 2002.

[7] S. Arora, M. Charikar, and E. Chlamtac. New approximation guarantees for chromatic number. In *38th Annual Symposium on the Theory of Computing*, 2006.

[8] S. Arora, J. R. Lee, and A. Naor. Euclidean distortion and the Sparsest Cut. In *37th Annual Symposium on the Theory of Computing*, pages 553–562, 2005. Submitted to *Journal of the AMS*. Full version available at `http://www.cs.berkeley.edu/~jrl`.

[9] S. Arora, S. Rao, and U. Vazirani. Expander flows, geometric embeddings, and graph partitionings. In *36th Annual Symposium on the Theory of Computing*, pages 222–231, 2004.

[10] P. Assouad. Plongements lipschitziens dans $\mathbf{R}^n$. *Bull. Soc. Math. France*, 111(4):429–448, 1983.

[11] Y. Aumann and Y. Rabani. An $O(\log k)$ approximate min-cut max-flow theorem and approximation algorithm. *SIAM J. Comput.*, 27(1):291–301 (electronic), 1998.

[12] E. Begelfor and M. Werman. The world is not always flat or learning curved manifolds. Technical Report 65, Leibniz Center, The Hebrew University of Jerusalem, 2005.

[13] A. Beygelzimer, S. Kakade, and J. Langford. Cover trees for nearest neighbor. To appear, *International Conference on Machine Learning*, 2006.

[14] H. L. Bodlaender, J. R. Gilbert, H. Hafsteinsson, and T. Kloks. Approximating treewidth, pathwidth, frontsize, and shortest elimination tree. *J. Algorithms*, 18(2):238–255, 1995.

[15] J. Bourgain. On Lipschitz embedding of finite metric spaces in Hilbert space. *Israel J. Math.*, 52(1-2):46–52, 1985.

[16] J. Bourgain, J. Kahn, G. Kalai, Y. Katznelson, and N. Linial. The influence of variables in product spaces. *Israel J. Math.*, 77(1-2):55–64, 1992.

[17] J. Bourgain, V. Milman, and H. Wolfson. On type of metric spaces. *Trans. Amer. Math. Soc.*, 294(1):295–317, 1986.

[18] B. Brinkman and M. Charikar. On the impossibility of dimension reduction in ℓ_1. In *Proceedings of the 44th Annual IEEE Conference on Foundations of Computer Science*, pages 514–523, 2003.

[19] J. Buresh-Oppenheim, N. Galesi, S. Hoory, A. Magen, and T. Pitassi. Rank bounds and integrality gaps for cutting planes procedures. In *Proceedings of the 44th Annual IEEE Conference on Foundations of Computer Science*, 2003.

[20] H. T.-H. Chan and A. Gupta. Small hop-diameter sparse spanners for doubling metrics. In *17th Annual ACM-SIAM Symposium on Discrete Algorithms*, 2006.

[21] H. T.-H. Chan, A. Gupta, B. Maggs, and S. Zhou. On hierarchical routing in doubling metrics. In *Proceedings of the 16th Annual Symposium on Discrete Algorithms*, 2005.

[22] M. Charikar, K. Makarychev, and Y. Makarychev. Near-optimal algorithms for unique games. In *38th Annual Symposium on the Theory of Computing*, 2006.

[23] S. Chawla, A. Gupta, and H. Räcke. Embeddings of negative-type metrics and an improved approximation to generalized sparsest cut. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 102–111, Vancouver, 2005.

[24] S. Chawla, R. Krauthgamer, R. Kumar, Y. Rabani, and D. Sivakumar. On embeddability of negative type metrics into ℓ_1. In *Proceedings of the 20th Annual IEEE Conference on Computational Complexity*, pages 144–153, 2005.

[25] J. Cheeger and B. Kleiner. Differentiating Banach space valued functions on metric measure spaces. Preprint, 2006.

[26] J. Cheeger and B. Kleiner. Differentiating maps into L^1 and the geometry of BV functions. Preprint, 2006.

[27] C. Chekuri, S. Khanna, and F. Shepherd. Edge-disjoint paths in planar graphs with constant congestion. In *46th Annual Symposium on Foundations of Computer Science*. IEEE, 2005.

[28] J. Chuzhoy and S. Khanna. Hardness of cut problems in directed graphs. In *38th Annual Symposium on the Theory of Computing*, 2006.

[29] R. Cole and L. Gottlieb. Searching dynamic point sets in spaces with bounded doubling dimension. In *Proceedings of the 38th Annual Symposium on the Theory of Computing*, 2006.

[30] T. F. Cox and M. A. A. Cox. *Multidimensional Scaling*. Number 38 in Monographs on Statistics and Applied Probability. Chapman and Hall, 2nd edition, 2001.

[31] E. D. Demaine and M. Hajiaghayi. Bidimensionality: New connections between FPT algorithms and PTASs. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2005)*, Vancouver, Canada, January 2005.

[32] N. Devanur, S. Khot, R. Saket, and N. K. Vishnoi. Integrality gaps for Sparsest Cut and Minimum Linear Arrangement problems. In *38th Annual Symposium on the Theory of Computing*, 2006.

[33] I. Dinur and S. Safra. On the hardness of approximating minimum vertex cover. *Ann. of Math. (2)*, 162(1):439–485, 2005.

[34] J. Dunagan and S. Vempala. On Euclidean embeddings and bandwidth minimization. In *Randomization, approximation, and combinatorial optimization*, pages 229–240. Springer, 2001.

[35] A. Eskin, D. Fisher, and K. Whyte. Quasi-isometries and rigidity of solvable groups. Preprint, 2006.

[36] U. Feige. Approximating the bandwidth via volume respecting embeddings. *J. Comput. System Sci.*, 60(3):510–539, 2000.

[37] U. Feige, M. Hajiaghayi, and J. R. Lee. Improved approximation algorithms for minimum-weight vertex separators. In *STOC '05: Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 563–572, New York, NY, USA, 2005. ACM Press.

[38] U. Feige and J. R. Lee. An improved approximation ratio for the minimum linear arrangement problem. To appear, *Information Processing Letters*, 2006.

[39] E. Friedgut. Boolean functions with low average sensitivity depend on few coordinates. *Combinatorica*, 18(1):27–35, 1998.

[40] M. X. Goemans and L. Tunçel. When does the positive semidefiniteness constraint help in lifting procedures? *Math. Oper. Res.*, 26(4):796–815, 2001.

[41] M. X. Goemans and D. P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. ACM*, 42(6):1115–1145, Nov. 1995.

[42] A. Gupta, R. Krauthgamer, and J. R. Lee. Bounded geometries, fractals, and low-distortion embeddings. In *44th Symposium on Foundations of Computer Science*, pages 534–543, 2003.

[43] A. Gupta, I. Newman, Y. Rabinovich, and A. Sinclair. Cuts, trees and l_1-embeddings of graphs. In *40th Annual Symposium on Foundations of Computer Science*, pages 399–408. IEEE Computer Soc., Los Alamitos, CA, 1999.

[44] V. Guruswami and S. Khanna. On the hardness of 4-coloring a 3-colorable graph. *SIAM J. Discrete Math.*, 18(1):30–40 (electronic), 2004.

[45] S. Har-Peled and M. Mendel. Fast construction of nets in low dimensional metrics, and their applications. In *Proceedings of the 21st Annual Symposium on Computational Geometry*, 2005.

[46] E. M. I. Dinur and O. Regev. Conditional hardness for approximate coloring. In *38th Annual Symposium on the Theory of Computing*, 2006.

[47] W. B. Johnson and J. Lindenstrauss. Extensions of Lipschitz mappings into a Hilbert space. In *Conference in modern analysis and probability (New Haven, Conn., 1982)*, pages 189–206. Amer. Math. Soc., Providence, RI, 1984.

[48] G. Karakostas. A better approximation ratio for the vertex cover problem. In *32nd International Colloquium on Automata, Languages and Programming*, volume 3580 of *Lecture Notes in Computer Science*, pages 1043–1050. Springer, 2005.

[49] D. Karger, R. Motwani, and M. Sudan. Approximate graph coloring by semidefinite programming. *J. ACM*, 45(2):246–265, 1998.

[50] D. Karger and M. Ruhl. Finding nearest neighbors in growth-restricted metrics. In *34th Annual ACM Symposium on the Theory of Computing*, pages 63–66, 2002.

[51] S. Khot. On the power of unique 2-prover 1-round games. In *34th Annual Symposium on the Theory of Computing*, 2002.

[52] S. Khot, G. Kindler, E. Mossel, and R. O'Donnell. Optimal inapproximability results for MAX-CUT and other 2-variable CSPs? In *46th Annual Symposium on Foundations of Computer Science*, 2005.

[53] S. Khot and O. Regev. Vertex cover might be hard to approximate to within 2-ε. In *18th IEEE Conference on Computational Complexity*, 2003.

[54] S. Khot and N. Vishnoi. The unique games conjecture, integrality gap for cut problems and embeddability of negative type metrics into ℓ_1. In *Proceedings of the 46th Annual IEEE Conference on Foundations of Computer Science*, 2005.

[55] P. N. Klein, S. A. Plotkin, and S. Rao. Excluded minors, network decomposition, and multicommodity flow. In *Proceedings of the 25th Annual ACM Symposium on Theory of Computing*, pages 682–690, 1993.

[56] R. Krauthgamer and J. R. Lee. The intrinsic dimensionality of graphs. In *Proceedings of the Thirty-Fifth Annual ACM Symposium on Theory of Computing*, pages 438–447 (electronic), New York, 2003. ACM. To appear, *Combinatorica*.

[57] R. Krauthgamer and J. R. Lee. Navigating nets: Simple algorithms for proximity search. In *Proceedings of the 15th annual ACM-SIAM Symposium on Discrete Algorithms*, 2004.

[58] R. Krauthgamer and J. R. Lee. The black-box complexity of nearest-neighbor search. *Theoret. Comput. Sci.*, 348(2-3):262–276, 2005.

[59] R. Krauthgamer and J. R. Lee. Algorithms on negatively curved spaces. In *47th Annual Symposium on Foundations of Computer Science*. IEEE, 2006. To appear.

[60] R. Krauthgamer, J. R. Lee, M. Mendel, and A. Naor. Measured descent: A new embedding method for finite metrics. *Geom. Funct. Anal.*, 15(4):839–858, 2005. Prelim. version in FOCS 2004.

[61] R. Krauthgamer, N. Linial, and A. Magen. Metric embeddings–beyond one-dimensional distortion. *Discrete Comput. Geom.*, 31(3):339–356, 2004.

[62] R. Krauthgamer and Y. Rabani. Improved lower bounds for embeddings into L_1. In *17th Annual ACM-SIAM Symposium on Discrete Algorithms*, 2006.

[63] T. J. Laakso. Plane with A_∞-weighted metric not bi-Lipschitz embeddable to $\mathbb{R}^N$. *Bull. London Math. Soc.*, 34(6):667–676, 2002.

[64] U. Lang and C. Plaut. Bilipschitz embeddings of metric spaces into space forms. *Geom. Dedicata*, 87(1-3):285–307, 2001.

[65] U. Lang and T. Schlichenmaier. Nagata dimension, quasisymmetric embeddings, and Lipschitz extensions. *Int. Math. Res. Not.*, (58):3625–3655, 2005.

[66] J. R. Lee. On distance scales, embeddings, and efficient relaxations of the cut cone. In *SODA '05: Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 92–101, Philadelphia, PA, USA, 2005. Society for Industrial and Applied Mathematics.

[67] J. R. Lee. Improved distortion estimates for negative-type metrics. Manuscript, 2006.

[68] J. R. Lee. Volume distortion for subsets of Euclidean spaces. In *Proceedings of the 22nd Annual Symposium on Computational Geometry*, 2006.

[69] J. R. Lee, M. Mendel, and A. Naor. Metric structures in L_1: dimension, snowflakes, and average distortion. *European J. Combin.*, 26(8):1180–1190, 2005.

[70] J. R. Lee and A. Naor. Absolute Lipschitz extendability. *C. R. Math. Acad. Sci. Paris*, 338(11):859–862, 2004.

[71] J. R. Lee and A. Naor. Embedding the diamond graph in L_p and dimension reduction in L_1. *Geom. Funct. Anal.*, 14(4):745–747, 2004.

[72] J. R. Lee and A. Naor. Extending Lipschitz functions via random metric partitions. *Invent. Math.*, 160(1):59–95, 2005.

[73] J. R. Lee and A. Naor. The Goemans-Linial conjecture and L_p metrics on the Heisenberg groups. In *47th Annual Symposium on Foundations of Computer Science*. IEEE, 2006. To appear.

[74] J. R. Lee, A. Naor, and Y. Peres. Trees and Markov convexity. In *17th Annual ACM-SIAM Symposium on Discrete Algorithms*, 2006.

[75] T. Leighton and S. Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *J. ACM*, 46(6):787–832, 1999.

[76] N. Linial, E. London, and Y. Rabinovich. The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15(2):215–245, 1995.

[77] R. J. Lipton and R. E. Tarjan. Applications of a planar separator theorem. *SIAM J. Comput.*, 9(3):615–627, 1980.

[78] L. Lovász and A. Schrijver. Cones of matrices and set-functions and 0-1 optimization. *SIAM J. Optim.*, 1(2):166–190, 1991.

[79] E. Mossel, R. O'Donnell, and K. Oleszkiewicz. Noise stability of functions with low in.uences invariance and optimality. In *46th Annual Symposium on Foundations of Computer Science*, pages 21–30, 2005.

[80] H. Okamura and P. D. Seymour. Multicommodity flows in planar graphs. *J. Combin. Theory Ser. B*, 31(1):75–81, 1981.

[81] R. Ostrovsky and Y. Rabani. Low distortion embeddings for edit distance. In *STOC'05: Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, pages 218–224, New York, 2005. ACM.

[82] S. Rao. Small distortion and volume preserving embeddings for planar and Euclidean metrics. In *Proceedings of the 15th Annual Symposium on Computational Geometry*, pages 300–306, New York, 1999.

[83] S. Rao and A. W. Richa. New approximation techniques for some linear ordering problems. *SIAM J. Comput.*, 34(2):388–404, 2004.

[84] O. Reingold. Undirected st-connectivity in log-space. In *37th Annual Symposium on the Theory of Computing*, pages 376–385, 2005.

[85] O. Reingold, S. Vadhan, and A. Wigderson. Entropy waves, the zig-zag graph product, and new constant-degree expanders. *Ann. of Math. (2)*, 155(1):157–187, 2002.

[86] N. Robertson and P. D. Seymour. Graph minors. II. Algorithmic aspects of tree-width. *J. Algorithms*, 7(3):309–322, 1986.

[87] A. Samorodnitsky and L. Trevisan. Gowers uniformity, influence of variables, and PCPs. In *38th Annual Symposium on the Theory of Computing*, 2006.

[88] S. Semmes. On the nonexistence of bi-Lipschitz parameterizations and geometric problems about A_∞-weights. *Rev. Mat. Iberoamericana*, 12(2):337–410, 1996.

[89] S. Semmes. Bilipschitz embeddings of metric spaces into Euclidean spaces. *Publ. Mat.*, 43(2):571–653, 1999.

[90] E. Sharon and D. Mumford. 2d-shape analysis using conformal mapping. In *CVPR (2)*, pages 350–357, 2004.

[91] Y. Shavitt and T. Tankel. On the curvature of the internet and its usage for overlay construction and distance estimation. In *Proceedings of the 23th Annual INFOCOM*, 2004.

[92] J. Shi and J. Malik. Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(8):888–905, 2000.

[93] D. B. Shmoys. Cut problems and their application to divide-and-conquer. In *Approximation Algorithms for NP-hard Problems, (D.S. Hochbaum, ed.)*, pages 192–235. PWS, 1997.

[94] A. Sinclair. Improved bounds for mixing rates of marked chains and multicommodity flow. In *LATIN*, pages 474–487, 1992.

[95] D. A. Spielman and S.-H. Teng. Spectral partitioning works: planar graphs and finite element meshes. In *37th Annual Symposium on Foundations of Computer Science (Burlington, VT, 1996)*, pages 96–105. IEEE Comput. Soc. Press, Los Alamitos, CA, 1996.

[96] D. A. Spielman and S.-H. Teng. Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing*, pages 81–90 (electronic), New York, 2004. ACM.

[97] K. Talwar. Bypassing the embedding: Approximation schemes and distance labeling schemes for growth restricted metrics. In *36th Annual ACM Symposium on Theory of Computing*, 2004.

[98] J. B. Tenenbaum, V. de Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000.

[99] I. Tourlakis. New lower bounds for vertex cover in the lovasz-schrijver hierarchy. In *21st IEEE Conference on Computational Complexity*, 2006.

[100] V. V. Vazirani. *Approximation algorithms*. Springer-Verlag, Berlin, 2001.

[101] S. S. Vempala. *The random projection method*. DIMACS Series in Discrete Mathematics and Theoretical Computer Science, 65. American Mathematical Society, Providence, RI, 2004. With a foreword by Christos H. Papadimitriou.

[102] U. Zwick. Approximation algorithms for constraint satisfaction problems involving at most three variables per constraint. In *11th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 201–210, 2006.