Optimal Rounding for Sparsest Cut

Alan Chang
Washington University in St. Louis
St. Louis, Missouri, U.S.A
alanchang@wustl.edu

Assaf Naor
Princeton University
Princeton, New Jersey, U.S.A
naor@math.princeton.edu

Kevin Ren
Princeton University
Princeton, New Jersey, U.S.A
kevinren@princeton.edu

Abstract

We prove that the integrality gap of the Goemans–Linial semidefinite program for the Sparsest Cut problem (with general capacities and demands) on inputs of size $n \geqslant 2$ is $\Theta(\sqrt{\log n})$. We achieve this by establishing the following geometric/structural result. If $(\mathcal{M}, d)$ is an n-point metric space of negative type, then for every $\tau > 0$ there is a random subset $\mathcal{Z}$ of $\mathcal{M}$ such that for any pair of points $x, y \in \mathcal{M}$ with $d(x, y) \geqslant \tau$, the probability that both $x \in \mathcal{Z}$ and $d(y, \mathcal{Z}) \geqslant \beta\tau/\sqrt{1 + \log(|B(y, \kappa\beta\tau)|/|B(y, \beta\tau)|)}$ is $\Omega(1)$, where $0 < \beta < 1 < \kappa$ are universal constants. The proof relies on a refinement of the Arora–Rao–Vazirani rounding technique.

CCS Concepts

• **Theory of computation** → **Rounding techniques**; **Random projections and metric embeddings**; • **Mathematics of computing** → *Functional analysis*.

Keywords

Sparsest Cut problem, semidefinite programming, metric embeddings.

ACM Reference Format:

Alan Chang, Assaf Naor, and Kevin Ren. 2025. Optimal Rounding for Sparsest Cut. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing (STOC '25), June 23–27, 2025, Prague, Czechia.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3717823.3718285

1 Introduction

Given an integer $n \geqslant 2$, the Sparsest Cut problem (with general capacities and demands) on n vertices takes as input two n-by-n symmetric matrices with nonnegative entries $\mathsf{C} = (c_{ij}), \mathsf{D} = (d_{ij}) \in \mathsf{M}_n([0,\infty))$ and aims to evaluate (or estimate) in polynomial time the following quantity (where we denote $[n] = \{1, \dots, n\}$):

$$\mathrm{SparsestCut}(\mathsf{C},\mathsf{D}) \stackrel{\mathrm{def}}{=} \min_{\emptyset \neq S \subsetneq [n]} \frac{\sum_{i\in S}\sum_{j\in[n]\setminus S} c_{ij}}{\sum_{i\in S}\sum_{j\in[n]\setminus S} d_{ij}}. \qquad (1)$$

This is a famous algorithmic task of central importance and interest. It would be needlessly repetitive to recount here the rich history of work on this extensively studied topic and its deep connections to other areas, as multiple references contain thorough surveys. In particular, all of the background that is relevant to the discussion herein can be found in [4, 20, 69]; see also [40], and the most recent and up-to-date account can be found in [63].

STOC '25, Prague, Czechia

ACM ISBN 979-8-4007-1510-5/25/06
https://doi.org/10.1145/3717823.3718285

It suffices to recall at this juncture that in the mid-1990s, Goemans and Linial introduced a semidefinite program (SDP) that computes in polynomial time a number $\mathrm{SDP}_{\mathrm{GL}}(\mathsf{C},\mathsf{D}) \geqslant 0$ that satisfies $\mathrm{SDP}_{\mathrm{GL}}(\mathsf{C},\mathsf{D}) \leqslant \mathrm{SparsestCut}(\mathsf{C},\mathsf{D})$. Specifically, $\mathrm{SDP}_{\mathrm{GL}}(\mathsf{C},\mathsf{D})$ is defined to be the minimum of $\sum_{i=1}^n \sum_{j=1}^n c_{ij}\|v_i - v_j\|_2^2$ over all the vectors $v_1, \dots, v_n \in \mathbb{R}^n$ that satisfy the following constraints:

$$\sum_{i=1}^{n}\sum_{j=1}^{n} d_{ij}\|v_i - v_j\|_2^2 = 1, \quad \text{and} \quad \forall i,j,k \in [n], \quad \|v_i - v_j\|_2^2 \leqslant \|v_i - v_k\|_2^2 + \|v_k - v_j\|_2^2. \qquad (2)$$

The pertinent question is therefore to understand the growth rate as $n \to \infty$ of the integrality gap of the Goemans–Linial SDP for Sparsest Cut, which is defined to be the following quantity:

$$\sup_{\substack{\mathsf{C},\mathsf{D}\in\mathsf{M}_n([0,\infty))\\ \mathsf{C},\mathsf{D}\ \mathrm{symmeric}}} \frac{\mathrm{SparsestCut}(\mathsf{C},\mathsf{D})}{\mathrm{SDP}_{\mathrm{GL}}(\mathsf{C},\mathsf{D})}.$$

The algorithm that outputs $\mathrm{SDP}_{\mathrm{GL}}(\mathsf{C},\mathsf{D})$ is then guaranteed to estimate $\mathrm{SparsestCut}(\mathsf{C},\mathsf{D})$ within a factor that is at most this integrality gap. Our main algorithmic contribution is:

Theorem 1. *The n-vertex integrality gap of the Goemans–Linial SDP for Sparsest Cut is $\Theta(\sqrt{\log n})$.*

The $\Omega(\sqrt{\log n})$ lower bound of Theorem 1 is from [63]. Thus, our contribution herein is in terms of algorithm design rather than proving an impossibility result, i.e., we derive an improved (sharp) upper bound on the integrality gap of the Goemans–Linial SDP, and hence we provide the best-known polynomial time approximation algorithm for Sparsest Cut. The previously best-known upper bound in Theorem 1 is due to [4], which proved that the stated integrality gap is $O(\sqrt{\log n}\log\log n)$. As we will see, obtaining an optimal upper bound without any unbounded lower order factor whatsoever is more than a technical matter: Our proof of Theorem 1 introduces conceptual innovations of impact on both rounding of SDPs as well as on geometric questions (some of which are well-known and longstanding). The full version [19] will discuss those further applications, while in the present extended abstract we will focus only on the Sparsest Cut problem.

Let $v_1, \dots, v_n \in \mathbb{R}^n$ be a vector solution of the Goemans–Linial SDP. If we define $d(i,j) = \|v_i - v_j\|_2^2$ for every $i, j \in [n]$, then thanks to the second part of (2) we know that $([n], d)$ is a metric space, which, by design, has the property that the metric space $([n], \sqrt{d})$ is isometric to a subset of Euclidean space. Such metric spaces are commonly called metric spaces of negative type; see e.g. the monograph [24] or the survey [57] for more on this important and useful notion, including the reason for the nomenclature.

The (bi-Lipschitz) distortion of a finite metric space $(\mathcal{M}, d_{\mathcal{M}})$ in an infinite metric space $(\mathcal{N}, d_{\mathcal{N}})$, which following [50] is commonly denoted $\mathsf{c}_{(\mathcal{N},d_{\mathcal{N}})}(\mathcal{M},d_{\mathcal{M}})$ or simply $\mathsf{c}_{\mathcal{N}}(\mathcal{M})$ when the underlying metrics are clear from the context, is the infimum over those $D>0$ for which there are $f:\mathcal{M}\to\mathcal{N}$ and $s>0$ such that

$$\forall x,y\in\mathcal{M},\qquad sd_{\mathcal{M}}(x,y)\leqslant d_{\mathcal{N}}(f(x),f(y))\leqslant Dsd_{\mathcal{M}}(x,y).$$

If $p\geqslant 1$, then one commonly uses the following shorter notation $\mathsf{c}_p(\mathcal{M})=\mathsf{c}_{\ell_p}(\mathcal{M})$. The parameters $\mathsf{c}_2(\mathcal{M})$ and $\mathsf{c}_1(\mathcal{M})$ are naturally called, respectively, the Euclidean distortion of $\mathcal{M}$ and the L_1 distortion of $\mathcal{M}$.

The classical Fréchet embedding $\Phi_{(\mathcal{M},d_{\mathcal{M}})}$ of $\mathcal{M}$ into the set $\mathbb{R}^{2^{\mathcal{M}}\setminus\{\varnothing\}}$ of real-valued functions on the nonempty subsets of $\mathcal{M}$ is defined by setting for every $x\in\mathcal{M}$ and $\varnothing\neq\mathcal{Z}\subseteq\mathcal{M}$,

$$\Phi_{(\mathcal{M},d_{\mathcal{M}})}(x)(\mathcal{Z})\stackrel{\mathrm{def}}{=}d_{\mathcal{M}}(x,\mathcal{Z}).$$

Given $p,D\geqslant 1$, one says that $(\mathcal{M},d_{\mathcal{M}})$ embeds into L_p with distortion D via the Fréchet embedding if there exists a probability measure $\mathbb{P}$ on $2^{\mathcal{M}}\setminus\{\varnothing\}$ such that for every $x,y\in\mathcal{M}$ we have

$$\begin{aligned}d_{\mathcal{M}}(x,y)&\leqslant D\|\Phi_{(\mathcal{M},d_{\mathcal{M}})}(x)-\Phi_{(\mathcal{M},d_{\mathcal{M}})}(y)\|_{L_p(\mathbb{P})}\\&\stackrel{\mathrm{def}}{=}D\Big(\mathbb{E}_{\mathbb{P}}\big[|\Phi_{(\mathcal{M},d_{\mathcal{M}})}(x)-\Phi_{(\mathcal{M},d_{\mathcal{M}})}(y)|^p\big]\Big)^{\frac{1}{p}}.\end{aligned}\tag{3}$$

For example, the famous embedding of [16] is of this form. This terminology is consistent with what we recalled above because (3) implies that $\mathsf{c}_p(\mathcal{M})\leqslant D$. Indeed, observe that as $2^{\mathcal{M}}\setminus\{\varnothing\}$ is a finite set, $L_p(\mathbb{P})$ is isometric to a subspace of ℓ_p, and then contrast (3) with the trivial estimate

$$\begin{aligned}&\|\Phi_{(\mathcal{M},d_{\mathcal{M}})}(x)-\Phi_{(\mathcal{M},d_{\mathcal{M}})}(y)\|_{L_p(\mathbb{P})}\\&\qquad\leqslant\|\Phi_{(\mathcal{M},d_{\mathcal{M}})}(x)-\Phi_{(\mathcal{M},d_{\mathcal{M}})}(y)\|_{L_\infty(\mathbb{P})}\leqslant d_{\mathcal{M}}(x,y),\end{aligned}\tag{4}$$

where the first step of (4) holds as $\mathbb{P}$ is a probability measure and the second step of (4) is a straightforward consequence of the triangle inequality for $d_{\mathcal{M}}$. However, knowing that $(\mathcal{M},d_{\mathcal{M}})$ embeds into L_p with distortion D via the Fréchet embedding provides more information than mere embeddability into L_p, which is sometimes needed in applications (e.g. [26, 52]); furthermore, the former embedding may not be possible when it is known that the latter embedding does exist (e.g. [11, 55]).

Theorem 2. *For every $n\in\mathbb{N}$, the largest possible Euclidean distortion of an n-point metric space of negative type is $\Theta(\sqrt{\log n})$. Furthermore, the $O(\sqrt{\log n})$ upper bound here is achieved via the Fréchet embedding.*

Similarly to Theorem 1, our contribution to Theorem 2 is its sharp upper bound $O(\sqrt{\log n})$; the matching lower bound is due to [25]. This upper bound confirms a well-known conjecture that to the best of our knowledge was first posed explicitly in the published literature in [28, page 158].

The fact that the upper bound in Theorem 2 implies a $O(\sqrt{\log n})$ upper bound on the integrality gap of the Goemans–Linial SDP for Sparsest Cut (i.e, the new contribution of Theorem 1) is a standard result (going back at least to [28]); see [57, Lemma 4.5] for its detailed derivation (based on the duality argument in [54, Proposition 15.5.2], which is attributed there to unpublished work of Rabinovich).

We focused above on the algorithmic task of efficiently approximating the number $\mathrm{SparsestCut}(\mathsf{C},\mathsf{D})$, but the fact that our $O(\sqrt{\log n})$-embedding of any n-point metric space of negative type is actually (per Theorem 2) into ℓ_2 rather than "merely" into ℓ_1 implies formally that there is also an algorithm which outputs a subset $\varnothing\neq S\subsetneq[n]$ that is a near-minimizer of the right hand side of (1), up to the aforementioned $O(\sqrt{\log n})$ error tolerance. This deduction utilizes the observation from the seminal work [50] that optimal embeddings into a Hilbert space can themselves be found in polynomial time (unlike embeddings into L_1, see [24, 34]), as this task itself can be cast as a semidefinite program; the (standard) deduction of this assertion is worked out in e.g. [4, Section 5].

One can refine the above discussion for each $k\in\{2,\ldots,n\}$ to obtain a $O(\sqrt{\log k})$-factor approximation algorithm if the matrix D has support of size at most k, i.e., $|\{(i,j)\in[n]\times[n]:d_{ij}>0\}|\leqslant k$. Also this fact is a formal consequence of Theorem 2, as the stated distortion guarantee is via the Fréchet embedding. That embedding can be automatically extended to any super-space while remaining 1-Lipschitz, which is all that is needed in order to obtain the aforementioned $O(\sqrt{\log k})$-factor approximation guarantee. Again, the standard deduction of this assertion is worked out in [4], though note that [4, Section 5] incorporates a step that is irrelevant for our purposes because the embedding of [4] is not Fréchet, so [4] must justify why it could be extended.

It is worthwhile to summarize the above discussion as the following separate algorithmic statement:

Theorem 3. *There exists a polynomial time algorithm with the following property. Fix $n\in\mathbb{N}$ and $k\in\{2,\ldots,n\}$. Suppose that $\mathsf{C}=(c_{ij}),\mathsf{D}=(d_{ij})\in\mathsf{M}_n([0,\infty))$ are symmetric matrices such that*

$$|\{(i,j)\in[n]\times[n]:d_{ij}>0\}|\leqslant k.$$

Then, the aforementioned algorithm outputs a subset $\varnothing\neq S\subsetneq[n]$ that satisfies:[1]

$$\frac{\sum_{i\in S}\sum_{j\in[n]\setminus S}c_{ij}}{\sum_{i\in S}\sum_{j\in[n]\setminus S}d_{ij}}\lesssim\mathrm{SparsestCut}(\mathsf{C},\mathsf{D})\sqrt{\log k}.\tag{5}$$

2 Random zero sets with local growth guarantees

Our main geometric/structural contribution is Theorem 4 below. It demonstrates that for any given scale $\tau>0$, a metric space of negative type admits a distribution over random subsets that separates with constant probability any pair of points whose distance is at least τ, where the amount of separation improves if balls of radius proportional to τ (centered at one of those two points) grow slowly. Such "random zero sets with (super-Gaussian) local growth guarantees" can be viewed as suitable replacements for the random half spaces that are used ubiquitously in the Euclidean setting.

Theorem 4. *There are universal constants $0<\beta,\delta<1<\kappa<\infty$ such that if $(\mathcal{M},d)$ is a finite metric space of negative type, then for every $\tau>0$ there is a probability distribution $\mathbb{P}=\mathbb{P}^{\tau,\mathcal{M}}$ over $2^{\mathcal{M}}\setminus\{\varnothing\}$*

[1]We will use throughout the ensuing text the following (standard) conventions for asymptotic notation, in addition to the usual $O(\cdot),o(\cdot),\Omega(\cdot),\Theta(\cdot)$ notation. Given $a,b>0$, by writing $a\lesssim b$ or $b\gtrsim a$ we mean that $a\leqslant\kappa b$ for some universal constant $\kappa>0$, and $a\asymp b$ stands for $(a\lesssim b)\wedge(b\lesssim a)$. When we will need to allow for dependence on parameters, we will indicate it by subscripts. For example, in the presence of auxiliary objects q,U,ϕ, the notations $a\lesssim_{q,U,\phi}b$ and $a=O_{q,U,\phi}(b)$ mean that $a\leqslant\kappa(q,U,\phi)b$, where $\kappa(q,U,\phi)>0$ may depend only on q,U,ϕ, and similarly for the notations $a\gtrsim_{q,U,\phi}b$, $a\asymp_{q,U,\phi}b$, $a=\Omega_{q,U,\phi}(b)$, $a=\Theta_{q,U,\phi}(b)$.

that satisfies the following probabilistic estimate for every $x, y \in \mathcal{M}$ *for which* $d(x,y) \geqslant \tau$*:*

$$\mathbb{P}\left[\mathcal{Z} \subseteq \mathcal{M} : d(y,\mathcal{Z}) \geqslant \frac{\beta\tau}{\sqrt{1+\log\frac{|B(y,\kappa\beta\tau)|}{|B(y,\beta\tau)|}}} \quad \text{and} \quad x \in \mathcal{Z}\right] \geqslant \delta. \tag{6}$$

In the statement of Theorem 4 (and throughout what follows), balls in a metric space $(\mathcal{M}, d)$ are always closed balls, namely, $B(x,r) = \{y \in \mathcal{M} : d(x,y) \leqslant r\}$ for $x \in \mathcal{M}$ and $r \geqslant 0$. The distance of a point $y \in \mathcal{M}$ from a nonempty subset $\mathcal{Z}$ of $\mathcal{M}$ is $d(y,\mathcal{Z}) = \inf_{z\in\mathcal{Z}} d(y,z)$.

Theorem 4 settles a well-known problem that was considered by experts ever since the appearance of the Arora–Rao–Vazirani (ARV) rounding algorithm [6] and the measured descent embedding technique [42]. The question whether Theorem 4 holds was first posed in the literature in [2] (specifically, see the paragraph just before Section 3 there), though [2] expresses implicit doubt that Theorem 4 could be valid. The availability of Theorem 4 is indeed a somewhat surprising development as it leads to a change of perspective on (and a sharp improvement of) a major line of work that straddles algorithm design, metric geometry and functional analysis.

The upper bound in Theorem 1 is a consequence of Theorem 4 thanks to the measured descent embedding technique [43], which yields the following result:

Theorem 5 (special case of measured descent [43]). *Suppose that* $n \in \mathbb{N}$, $\alpha \geqslant \beta > 0$, *and* $0 < \varepsilon, \delta, \theta \leqslant 1$. *Let* $(\mathcal{M}, d)$ *be a metric space of size* n *with the property that for every* $\tau > 0$ *there is a probability measure* $\mathbb{P}^\tau$ *on the nonempty subsets of* $\mathcal{M}$ *such that for every* $x, y \in \mathcal{M}$ *that satisfy* $\tau \leqslant d(x,y) \leqslant (1+\theta)\tau$ *we have*

$$\mathbb{P}^\tau\left[\mathcal{Z} \subseteq \mathcal{M} :\ d(y,\mathcal{Z}) \geqslant \frac{\varepsilon\tau}{\sqrt{1+\log\frac{|B(y,\alpha\tau)|}{|B(y,\beta\tau)|}}} \quad \text{and} \quad x \in \mathcal{Z}\right] \geqslant \delta.$$

Then,

$$c_2(\mathcal{M}) = O_{\varepsilon,\delta,\theta,\alpha,\beta}\left(\sqrt{\log n}\right). \tag{7}$$

Furthermore, the bound (7) *on the Euclidean distortion of* $\mathcal{M}$ *is obtained via the Fréchet embedding.*

Even though those who are familiar with [43] will immediately recognize Theorem 5 as a special case of measured descent, Theorem 5 does not appear in [43] as a standalone statement. Instead, Theorem 5 follows from part of the proof of Lemma 1.8 in [43] (that lemma derives further facts that are relevant to the setup of [43], but not for Theorem 5). The full version [19] of the present extended abstract includes a self-contained proof of Theorem 5 which builds on the ideas of [43] while incorporating further enhancements so as to yield both a more general statement and the best dependence that we currently have of the implicit constant factor in (7) on the parameters $\varepsilon, \delta, \theta, \alpha, \beta$.

3 Overview of the proof of Theorem 4

The precursor to Theorem 4 appeared as [4, Theorem 3.1]. It states that there is a universal constant $c > 0$ such that for every n-point metric space $(\mathcal{M}, d)$ of negative type and any $\tau > 0$ there exists a probability measure $\mathbb{P} = \mathbb{P}^{\tau,\mathcal{M}}$ on $2^{\mathcal{M}} \smallsetminus \{\emptyset\}$ such that for all $x, y \in \mathcal{M}$ with $d(x,y) \geqslant \tau$,

$$\mathbb{P}\left[\mathcal{Z} \subseteq \mathcal{M} :\ d(y,\mathcal{Z}) \geqslant \frac{c\tau}{\sqrt{\log n}} \quad \text{and} \quad x \in \mathcal{Z}\right] = \Omega(1). \tag{8}$$

The main result of [2, 4] used (8) to prove the Euclidean distortion bound $c_2(\mathcal{M}) \lesssim \sqrt{\log n}\log\log n$, and [3] showed that the same bound on the Euclidean distortion of $\mathcal{M}$ can be obtained via the Fréchet embedding.

Evidently, conclusion (6) of Theorem 4 is quantitatively stronger than (8). However, its main contribution is a qualitative enhancement that allows one to exploit cancelation to derive the upper bounds of Theorem 1 and Theorem 2 without any unbounded lower order factors. The previously best-known upper bound [3] in the context of Theorem 2 had a redundant $\log\log n$ factor, and correspondingly conclusion (5) of Theorem 3 was obtained in [4] with the factor $\sqrt{\log k}$ replaced by $\sqrt{\log k}\log\log k$.

The utility of Theorem 4 for removing altogether any unbounded lower order factor from [4] is obvious, as it is nothing more than a direct substitution of (6) into measured descent [43] (Theorem 5 herein). This possibility is discussed in [2], though with skepticism that the statement of Theorem 4 could be valid due to the nonlocal nature of the ARV rounding algorithm [7], which is the central input to (8). In contrast, a key feature of (6) is that the "performance" of the random zero set at the given scale τ for each given pair of points $x, y \in \mathcal{M}$ depends only the local "snapshot" $B(y,\kappa\beta\tau)$ of $\mathcal{M}$ near y at scale $\Theta(\tau)$, and moreover that it depends only on the extent to which the size of $B(y,\kappa\beta\tau)$ increases relative to the size of the proportionally smaller snapshot $B(y,\beta\tau)$ of $\mathcal{M}$.

In absence of such locality, [4] starts out with (8), from which point it does not make any further appeal to [7] and instead it proceeds by adapting measured descent in order to obtain the aforementioned distortion bound, i.e., while incurring a lower-order yet unbounded multiplicative loss. The reasoning that is used in [3] to show that this can, in fact, be achieved via the Fréchet embedding also takes (8) as a "black box" without any further appeal to [7] and proceeds by yet another enhancement of measured descent.

We do not see how to derive the aforementioned sharp embedding results via the above route, which is purely metric/analytic in contrast to the more structural nature of ARV. Instead, the present work "flips" the approach of [4] by leaving measured decent untouched (it can now be simply quoted as a "black box"), and enhancing, as we will next outline, the structural insights that are provided by the ARV framework.

After the appearance of the ARV algorithm in [6], simplifications, refinements, extensions, and reformulations of it were developed in multiple works, including notably its full journal version [7] and [1, 5, 21, 47, 61, 66]. All of those contributions were valuable to us in the process of developing the ensuing enhancement of ARV. Rothvoss' lecture notes [66] stand out here because they creatively redo the setup and reasoning in a natural, and, as it turns out, more flexible way. In particular, we introduce Definition 7 below of compatibility of a labelled graph with a mapping into $\mathbb{R}^n$, which arose from our efforts to understand the extent to which the proof in [66] can be strengthened.

3.1 Directional Euclidean sparsification of graphs

The above referenced versions of the ARV algorithm study (explicitly or implicitly) a natural way to "sparsify" certain (combinatorial) graphs, i.e., a procedure that removes in a meaningful way some of the edges of a given graph. Our proof of Theorem 4 investigates a more involved version of this procedure, applied to graphs that encode more geometric information than the proximity graphs that were used in those works. This section is devoted to explaining these ideas in suitable generality.

The aforementioned literature considers a finite metric space $(\mathcal{M}, d)$ of negative type and a scale $\tau > 0$, and studies the proximity graph $\mathsf{G} = (\mathcal{M}, E)$ whose vertex set is $\mathcal{M}$ and $\{x, y\} \in E$ if and only if $d(x, y) \leqslant \Delta$, where $\Delta = c\tau/\sqrt{\log|\mathcal{M}|}$ for some universal constant $c > 0$. In the present setting, we are led to consider a certain graph (specified below) whose vertex set is still $\mathcal{M}$, yet if a pair of points $x, y \in \mathcal{M}$ are joined by an edge, then this will encode information that combines both their proximity and the local growth rate of $\Theta(\tau)$-balls in $\mathcal{M}$ centered at x and y. The sparsification procedure in our context will involve a pairwise thresholding criterion (we will soon explain what this means) that is nonconstant, while previous works considered a fixed threshold that is independent of the given pair of points in $\mathcal{M}$.

Because the ensuing discussion needs to consider multiple metrics simultaneously, including multiple metrics on the same set (arising from both the original metric and the shortest-path metrics of graphs as above), it will be beneficial to use subscripts when denoting distances, balls, diameters. Thus, given a metric space $(\mathcal{M}, d_{\mathcal{M}})$, we will write $\mathrm{diam}_{\mathcal{M}}(A) = \sup_{a,b\in A} d_{\mathcal{M}}(a, b)$ and $d_{\mathcal{M}}(x, A) = \inf_{a\in A} d_{\mathcal{M}}(x, a)$ for, respectively, the $d_{\mathcal{M}}$-diameter of $\emptyset \neq A \subseteq \mathcal{M}$ and the $d_{\mathcal{M}}$-distance of $x \in \mathcal{M}$ from A. We will also write $B_{\mathcal{M}}(x, r) = \{y \in \mathcal{M} : d_{\mathcal{M}}(x, y) \leqslant r\}$ for the $d_{\mathcal{M}}$-ball centered at x of radius $r \geqslant 0$. Thus, given $n \in \mathbb{N}$, $p \geqslant 1$ and $x \in \mathbb{R}^n$, we will use the notation $B_{\ell_p^n}(x, r) = \{y \in \mathbb{R}^n : \|x - y\|_p \leqslant r\}$, where $\|z\|_p = (|z_1|^p + \ldots + |z_n|^p)^{1/p}$ for $z = (z_1, \ldots, z_n) \in \mathbb{R}^n$. The standard scalar product on $\mathbb{R}^n$ will be denoted $\langle\cdot,\cdot\rangle : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}$, i.e., $\langle z, w\rangle = z_1w_1 + \ldots + z_nw_n$ for $z = (z_1, \ldots, z_n), w = (w_1, \ldots, w_n) \in \mathbb{R}^n$. The standard Gaussian measure on $\mathbb{R}^n$ will be denoted (as usual) by γ_n, i.e., the density of γ_n at $z \in \mathbb{R}^n$ is proportional to $\exp(-\|z\|_2^2/2)$.

Throughout what follows, all graphs will be tacitly assumed to be finite and will be allowed to have self-loops. Given a (possibly disconnected) graph $\mathsf{G} = (V, E)$, denote the shortest-path/geodesic (extended) metric that it induces on its vertex set V by $d_{\mathsf{G}} : V \times V \to [0, \infty]$, under the natural convention that $d_{\mathsf{G}}(x, y) = \infty$ if and only if $(x, y) \in \Gamma \times \Gamma'$ for distinct connected components $\Gamma, \Gamma' \subseteq V$ of G. For $r \geqslant 0$ and $x \in V$, the corresponding ball in (V, d_{G}) will be denoted $B_{\mathsf{G}}(x, r) = \{y \in V : d_{\mathsf{G}}(x, y) \leqslant r\}$. In particular, $B_{\mathsf{G}}(x, 1) = \{x\} \cup \{y \in V : \{x, y\} \in E\}$, which is also denoted $N_{\mathsf{G}}(x)$, is the neighborhood in G of the vertex x.

When a graph $\mathsf{G} = (V, E)$ is accompanied by a mapping $f : V \to \mathbb{R}^n$, which we think of as a geometric representation of G, and an edge-labelling $\sigma : E \to \mathbb{R}$, which we think of as a thresholding criterion for determining which edges will be retained in the ensuing sparsification, for each vector $v \in \mathbb{R}^n$ we can consider the sub-graph of G that is obtained by deleting those $\{x, y\} \in E$ with $|\langle f(x) - f(y), v\rangle| \leqslant 4\sigma(\{x, y\})$:

Definition 6 (directional Euclidean sparsification). *Let $\mathsf{G} = (V, E)$ be a graph. If $n \in \mathbb{N}$ and $f : V \to \mathbb{R}^n$, then for every $\sigma : E \to \mathbb{R}$ and $v \in \mathbb{R}^n$ define*[2]

$$E(v; f, \sigma) \stackrel{\mathrm{def}}{=} \Big\{\{x, y\} \in E : |\langle f(x) - f(y), v\rangle| > 4\sigma(\{x, y\})\Big\}. \quad (9)$$

We thus obtain the following subgraph[3] *of G, which we call the Euclidean sparsification of G in direction v corresponding to its Euclidean representation f and the thresholding function σ:*

$$\mathsf{G}(v; f, \sigma) \stackrel{\mathrm{def}}{=} \Big(V, E(v; f, \sigma)\Big). \quad (10)$$

Understanding typical properties of $\mathsf{G}(v; f, \sigma)$ when v is chosen randomly according to the Gaussian measure γ_n is interesting in its own right. Here, we will investigate its matching number, which is also what [66] studies (as do other ARV-related works, usually implicitly), in the case when $\sigma \equiv 1/2$ is constant and $\mathsf{G} = (\mathcal{M}, E)$ is the above "vanilla" proximity graph on a finite metric space $(\mathcal{M}, d_{\mathcal{M}})$ of negative type, i.e., $\{x, y\} \in E$ if and only if $d_{\mathcal{M}}(x, y) \leqslant \Delta$ for some fixed $\Delta > 0$. Recall that the matching number $\nu(\mathsf{G})$ of a graph $\mathsf{G} = (V, E)$ is the maximum cardinality of a pairwise-disjoint collection of its edges. We will need only rudimentary properties of this basic combinatorial notion (that will be recalled when they will arise in proofs), which are covered in e.g. [51].

Since both the graphs $\mathsf{G} = (V, E)$ that we will investigate herein, and their labelings $\sigma : E \to \mathbb{R}$ that we will use as thresholds for Euclidean sparsification, are more complicated than the aforementioned special case, it is beneficial to first study the above setting abstractly. We arrived at the following definition, that will have an important role below, by examining the elegant proof in [66] with this goal in mind:

Definition 7 (compatibility of a graph with its Euclidean realization and edge labeling). *Fix $n \in \mathbb{N}$ and $C > 0$. Let $\mathsf{G} = (V, E)$ be a graph, $f : V \to \mathbb{R}^n$ and $\sigma : E \to [0, \infty)$. We say that G is C-compatible with f and σ if there exist $\Delta : V \to [0, \infty)$ and $K : V \to \mathbb{N}$ that have the following three properties:*

(1) For every $x \in V$ and $y \in B_{\mathsf{G}}(x, K(x) - 1)$, if $z \in V$ is such that $\{y, z\} \in E$, then

$$\Delta(x) \leqslant \sigma(\{y, z\}). \quad (11)$$

(2) For every $x \in V$ and every $y \in N_{\mathsf{G}}(x)$ we have

$$\int_{\mathbb{R}^n} \bigg(\max_{z\in B_{\mathsf{G}}(y,K(y))} \langle f(z) - f(y), v\rangle\bigg) \mathrm{d}\gamma_n(v) \leqslant K(x)\Delta(y). \quad (12)$$

(3) For every $x \in V$ we have

$$f\Big(B_{\mathsf{G}}(x, K(x))\Big) \subseteq B_{\ell_2^n}\Big(f(x), \frac{1}{C}\Delta(x)\Big). \quad (13)$$

[2] We decided to insert the factor 4 in the definition (9) of Euclidean sparsification for convenience only, as this yields a slight simplification of expressions in the ensuing reasoning. This choice is, of course, nothing more than a superficial normalization which could be removed if so desired by replacing throughout what follows the thresholding function σ by $\sigma/4$.

[3] Observe that even though graphs are allowed herein to have self-loops, if σ takes values in $[0, \infty)$, then the strict inequality in (9) implies that $\mathsf{G}(v; f, \sigma)$ will never have self-loops, i.e., $\{x, x\} \notin E(v; f, \sigma)$ for every $x \in V$.

Definition 7 is scale-invariant in the following sense: if G is C-compatible with f and σ, then G is also C-compatible with λf and $\lambda\sigma$ for every $\lambda \geqslant 0$, as seen by considering $\lambda\Delta$. The only part of Definition 7 that involves the parameter C is the inclusion (13), which implies in particular that C-compatibility becomes a more stringent property as C grows. The precise role of the properties that Definition 7 requires from K, Δ will become apparent upon examining the details of how they are applied in the proofs (which appear in the full version [19] of the present extended abstract). In (very) broad strokes, their significance is that K, Δ take as input a single vertex, which we think of as a consistent choice of "local scales" at that vertex (for, respectively, the domain and range of f), yet they control the pairwise interactions σ through (11), and the oscillations of f through (12); both of these controls occur on (combinatorial) balls in G whose radius is determined by K.

Note the following feature of Definition 7: it stipulates the existence of Δ, K with the stated properties, but these are auxiliary objects that occur internally to the definition and are not part of the notion of G being C-compatible with f and σ. Thus, any statement about C-compatibility will not refer to K, Δ, i.e., they will only be used as tools within its proof. This is exemplified by the following theorem, which is central to our proof of Theorem 4. It asserts that C-compatibility for large $C > 0$ implies that the expected matching number $\nu(\mathrm{G}(v; f, \sigma))$ is small when v is distributed according to γ_n. This result is a generalization of the ARV reasoning as it was recast by Rothvoss, and its proof (which appears in the full version [19] of the present extended abstract), is an adaptation of the strategy that he introduced in [66, Section 6.1].

Theorem 8 (expected matching number of Euclidean sparsification). *Fix $C \geqslant 1$ and $n \in \mathbb{N}$. Suppose that $\mathrm{G} = (V, E)$ is a graph that is C-compatible with $f : V \to \mathbb{R}^n$ and $\sigma : E \to [0, \infty)$. Then,*[4]

$$\int_{\mathbb{R}^n} \nu\Big(\mathrm{G}(v; f, \sigma)\Big)\, \mathrm{d}\gamma_n(v) < 6e^{-\frac{1}{4}C^2}|V|. \tag{14}$$

The proof of Theorem 8 is the only part of our proof of Theorem 4 that elaborates on the central innovation of [7], namely, what is known today as the "ARV chaining argument" (through the approach to it that was devised in [66]). We will suitably implement this argument in the full version [19] of the present extended abstract to prove Theorem 8. That implementation involves work that is of a more technical nature, and the conceptually new contribution here is the mere introduction of Definition 7, which makes Theorem 8 possible.

3.2 From Theorem 8 to Theorem 4

The rest of the proof of Theorem 4 does not relate to ARV chaining. We will next describe the steps that remain in the derivation of Theorem 4 from Theorem 8. Proposition 9 below will be used for that purpose; it is taken from [22] though it is not a standalone statement there and its justification follows from the reasoning in [22, Section 3] (a strengthening of Proposition 9 is proved in the full version [19] of the present extended abstract). Like Theorem 8, Proposition 9 is about abstract combinatorics and probability: it does not relate to the assumption of Theorem 4 that $(\mathcal{M}, d_{\mathcal{M}})$ is a metric space, or that it is of negative type; these will occur only in subsequent stages of the reasoning.

Proposition 9. *There is a universal constant $\kappa > 1$ with the following property. Fix $n \in \mathbb{N}$ and $\Lambda > 0$. Let V be a finite set, $f : V \to \mathbb{R}^n$, and let ω be a probability measure on $V \times V$ satisfying*

$$\forall x, y \in V, \qquad \omega(x, y) > 0 \implies \|f(x) - f(y)\|_2 \geqslant \Lambda. \tag{15}$$

Then, for every $C \geqslant 1$ and $v \in \mathbb{R}^n$, there are $A(v) = A^{\omega,C}(v), B(v) = B^{\omega,C}(v) \subseteq V$ such that

$$\forall v \in \mathbb{R}^n,\ \forall (x, y) \in A(v) \times B(v), \quad |\langle v, f(x) - f(y)\rangle| > C\Lambda. \tag{16}$$

Furthermore, the γ_n-expected ω-measure of the product $A(v) \times B(v) \subseteq V \times V$ satisfies

$$\int_{\mathbb{R}^n} \omega\Big(A(v) \times B(v)\Big)\, \mathrm{d}\gamma_n(v) \gtrsim e^{-\kappa C^2}. \tag{17}$$

The following proposition is an especially important step in the proof of Theorem 4:

Proposition 10. *There is a universal constant $\zeta \geqslant 1$ with the following property. Let $(\mathcal{M}, d_{\mathcal{M}})$ be a finite metric space. For $\tau, C > 0$ define $\rho = \rho_{d_{\mathcal{M}},C,\tau} : \mathcal{M} \to [1, \infty)$ by*

$$\forall x \in \mathcal{M}, \qquad \rho(x) \stackrel{\mathrm{def}}{=} 1 + \frac{\zeta}{C}\sqrt{\log \frac{|B_{\mathcal{M}}(x, 19\tau)|}{|B_{\mathcal{M}}(x, \tau)|}}. \tag{18}$$

Then, there exists a mapping $q = q_{d_{\mathcal{M}},C,\tau} : \mathcal{M} \to \mathcal{M}$ satisfying $\max_{x \in \mathcal{M}} d_{\mathcal{M}}(q(x), x) \leqslant 7\tau$ with the following properties. Let $\mathrm{G} = \mathrm{G}_{d_{\mathcal{M}},C,\tau}$ be the graph whose vertex set is $\mathcal{M}$ and whose edge set $E = E_{d_{\mathcal{M}},C,\tau}$ is given by

$$\forall x, y \in \mathcal{M},\ \{x, y\} \in E \iff d_{\mathcal{M}}(x, y) \leqslant \frac{\tau}{\min\{\rho(x), \rho(y)\}}. \tag{19}$$

Then for any $n \in \mathbb{N}$ and any mapping $\varphi : \mathcal{M} \to \mathbb{R}^n$, the graph G is C-compatible with $\varphi \circ q : \mathcal{M} \to \mathbb{R}^n$ and $\sigma : E \to [0, \infty)$, where $\sigma = \sigma_{d_{\mathcal{M}},C,\varphi\circ q} : E \to [0, \infty)$ is defined by setting for every $\{x, y\} \in E$,[5]

$$\sigma(\{x, y\}) \stackrel{\mathrm{def}}{=} C\Bigg(\max_{\substack{a \in B_{\mathcal{M}}(x,2\tau)\cap B_{\mathcal{M}}(y,2\tau)\\ b \in B_{\mathcal{M}}(a,2\tau)}} \|\varphi \circ q(a) - \varphi \circ q(b)\|_2\Bigg). \tag{20}$$

Before proceeding to the rest of the steps of the proof of Theorem 4, will next discuss the significance of Proposition 10, whose proof, which appears in the full version [19], contains a key idea of the present work.

Proposition 10 is the first time that the assumption that $(\mathcal{M}, d_{\mathcal{M}})$ is a metric space is used in the proof of Theorem 4, but this proposition works for any metric space and we do not yet need to know that it is of negative type. Also, Proposition 10 introduces the type of graphs to which the general combinatorial/probabilistic statements of Theorem 8 and Proposition 9 will be applied: their vertices are points in $\mathcal{M}$ and their edges are given by (19). The relevance of Proposition 10 to Theorem 8 is evident, as Proposition 10 produces situations in which the C-compatibility assumption of Theorem 8

[4]Formally, for the integral in (14) to be defined one needs that the function $v \mapsto \nu(\mathrm{G}(v; f, \sigma))$ from $\mathbb{R}^n$ to $\mathbb{N} \cup \{0\}$ is measurable. Similar measurability requirements are needed in multiple other steps of the ensuing discussion, but they will be suppressed for the purpose of the present extended abstract; its full version [19] addresses such measurability issues (which are simple to ensure).

[5]This σ is well-defined: because $\rho \geqslant 1$ by (18), definition (19) implies that if $\{x, y\} \in E$, then $d_{\mathcal{M}}(x, y) \leqslant \tau$, so the maximum in (20) is over a nonempty subset of $\mathcal{M}$ (e.g., one can take there $a = x$ and $b = y$).

is satisfied. The link to Proposition 9 will be made later, in a subsequent step of the proof of Theorem 4.

The crucial geometric contribution of Proposition 10 is constructing the mapping $q:\mathcal{M}\to\mathcal{M}$, which should be viewed as a way to "compress" a given metric space: its image $q(\mathcal{M})\subseteq\mathcal{M}$ will typically be much smaller than $\mathcal{M}$, yet it will encode geometric properties of $\mathcal{M}$ that will be important for the purpose of working with ratios of sizes of balls as in (18), which is our main goal. In fact, Proposition 10 states that q is a universally compatible compression scheme in the sense that C-compatibility arises upon composition with q of any function whatsoever from $\mathcal{M}$ to $\mathbb{R}^n$ (for a suitable choice of edge-labelling σ).

Proposition 10 treats the specific function ρ in (18) because this is what is needed below, and also since this is what arises from the use of the Gaussian measure in requirement (12) of Definition 7. Nevertheless, the compression scheme that is presented in the full version [19] of the present extended abstract works for any edge set as in (19) when $\rho:\mathcal{M}\to[1,\infty)$ is arbitrary. The construction in the full version [19] is therefore more general than what is used for Proposition 10, and hence it could be useful for other purposes (perhaps in settings to which appropriate non-Gaussian or non-Euclidean versions of compatibility are pertinent). The idea is to consider and suitably analyse hierarchically nested (2τ)-nets in the level sets of ρ, arranged in increasing order.

Applying Proposition 10 to metrics of negative type, we obtain the following:

Proposition 11. *For every $r\geqslant 1$ there is $\beta=\beta(r)>0$ that has the following properties. Given $\tau,C>0$ and a finite metric space $(\mathcal{M},d_{\mathcal{M}})$ of negative type, let $\mathsf{G}_{d_{\mathcal{M}},rC,\beta\tau}=(\mathcal{M},E)$ be the graph from Proposition 10 with with (C,τ) replaced by $(rC,\beta\tau)$, i.e.,*

$$\forall x,y\in\mathcal{M},\ \{x,y\}\in E\iff d_{\mathcal{M}}(x,y)\leqslant\frac{\beta\tau}{\min\{\rho(x),\rho(y)\}},\tag{21}$$

where, with $\zeta\geqslant 1$ the universal constant from Proposition 10, we denote

$$\forall x\in\mathcal{M},\qquad \rho(x)\overset{\mathrm{def}}{=}1+\frac{\zeta}{rC}\sqrt{\log\frac{|B_{\mathcal{M}}(x,19\beta\tau)|}{|B_{\mathcal{M}}(x,\beta\tau)|}}.\tag{22}$$

Then, there are $f:\mathcal{M}\to\mathbb{R}^{|\mathcal{M}|}$, $\sigma:E\to[0,\infty)$, and $\Lambda>0$ that have the following properties:

- *G is (rC)-compatible with f and σ;*
- *Every $x,y\in\mathcal{M}$ satisfy*

$$\begin{aligned} d_{\mathcal{M}}(x,y)\geqslant\tau &\implies C\|f(x)-f(y)\|_2\geqslant\Lambda,\\ &\text{and}\\ \{x,y\}\in E &\implies 4\sigma(\{x,y\})\leqslant\Lambda.\end{aligned}\tag{23}$$

Proof. Define $\beta,\Lambda>0$ by

$$\begin{aligned}\beta=\beta(r)&\overset{\mathrm{def}}{=}\frac{1}{256r^2+14}\\ &\text{and}\\ \Lambda=\Lambda(r,C,\tau)&\overset{\mathrm{def}}{=}\frac{16rC\sqrt{\tau}}{\sqrt{256r^2+14}}.\end{aligned}\tag{24}$$

Let $q:\mathcal{M}\to\mathcal{M}$ be the function from Proposition 10, applied with (C,τ) replaced by $(rC,\beta\tau)$. Then,

$$\forall x\in\mathcal{M},\qquad d_{\mathcal{M}}(q(x),x)\leqslant 7\beta\tau.\tag{25}$$

Fix an isometric embedding $\varphi:\mathcal{M}\to\mathbb{R}^{|\mathcal{M}|}$ of $(\mathcal{M},\sqrt{d_{\mathcal{M}}})$ into $\ell_2^{|\mathcal{M}|}$ and let $f\overset{\mathrm{def}}{=}\varphi\circ q$. Thus,

$$\forall x,y\in\mathcal{M},\qquad \|f(x)-f(y)\|_2=\sqrt{d_{\mathcal{M}}(q(x),q(y))}.\tag{26}$$

Finally, let $\sigma:E\to[0,\infty)$ be the edge-labelling that Proposition 10 produces from the above choices, i.e., for all $\{x,y\}\in E$,

$$\sigma(\{x,y\})\overset{(20)\wedge(26)}{=}rC\Bigg(\max_{\substack{a\in B_{\mathcal{M}}(x,2\beta\tau)\cap B_{\mathcal{M}}(y,2\beta\tau)\\ b\in B_{\mathcal{M}}(a,2\beta\tau)}}\sqrt{d_{\mathcal{M}}(q(a),q(b))}\Bigg).\tag{27}$$

Thus, Proposition 10 ensures that G is rC-compatible with f and σ.

By the triangle inequality, for every $x,y\in\mathcal{M}$ with $d_{\mathcal{M}}(x,y)\geqslant\tau$ we have

$$\begin{aligned}C\|f(x)-f(y)\|_2 &\overset{(26)}{\geqslant} C\sqrt{\max\{d_{\mathcal{M}}(x,y)-d_{\mathcal{M}}(q(x),x)-d_{\mathcal{M}}(q(y),y),0\}}\\ &\overset{(25)}{\geqslant} C\sqrt{(1-14\beta)\tau}\overset{(24)}{=}\Lambda,\end{aligned}\tag{28}$$

which proves the first part of (23). The second part of (23) is proved similarly as follows: we can find $a\in B_{\mathcal{M}}(x,2\beta\tau)\cap B_{\mathcal{M}}(y,2\beta\tau)$ and $b\in B_{\mathcal{M}}(a,2\beta\tau)$ such that

$$\begin{aligned}\sigma(\{x,y\})&\overset{(27)}{\leqslant} rC\sqrt{d_{\mathcal{M}}(q(a),a)+d_{\mathcal{M}}(a,b)+d_{\mathcal{M}}(q(b),b)}\\ &\overset{(25)}{\leqslant} rC\sqrt{16\beta\tau}\overset{(24)}{=}\frac{\Lambda}{4}.\end{aligned}$$ □

The upshot of Proposition 11 is that it provides the compatibility of the graph that we care about with f and σ, which is the assumption of Theorem 8, but at the same time we see from (23) that σ is controlled by a constant Λ that satisfies the assumptions of Proposition 9. Thus, Proposition 11 will allow us to combine Theorem 8 and Proposition 9 to obtain the following theorem:

Theorem 12. *There are universal constants $0<\beta\leqslant 1\leqslant\alpha<\infty$ with the following properties. Fix a finite metric space $(\mathcal{M},d_{\mathcal{M}})$ of negative type. Suppose that $0<\tau\leqslant\mathrm{diam}(\mathcal{M})$ and that ω is a symmetric probability measure on $\mathcal{M}\times\mathcal{M}$ whose support is contained in the set $\{(x,y)\in\mathcal{M}\times\mathcal{M}:\ d_{\mathcal{M}}(x,y)\geqslant\tau\}$. In other words, $\omega(\mathcal{M}\times\mathcal{M})=1$, for every $x,y\in\mathcal{M}$ we have $\omega(x,y)=\omega(y,x)$, and $\omega(x,y)>0\implies d_{\mathcal{M}}(x,y)\geqslant\tau$. Then, for every $C\geqslant 1$ and $v\in\mathbb{R}^{|\mathcal{M}|}$ there exist nonempty subsets $A^*(v)=A^*_{\omega,C}(v)$, $B^*(v)=B^*_{\omega,C}(v)$ of $\mathcal{M}$ such that*

$$\begin{aligned}&\forall v\in\mathbb{R}^{|\mathcal{M}|},\ \forall (x,y)\in A^*(v)\times B^*(v),\\ &\qquad d_{\mathcal{M}}(x,y)>\frac{\beta\tau}{\min\{\rho(x),\rho(y)\}},\end{aligned}\tag{29}$$

where $\rho:\mathcal{M}\to[1,\infty)$ is defined by

$$\forall x\in\mathcal{M},\qquad \rho(x)\overset{\mathrm{def}}{=}1+\frac{1}{\alpha C}\sqrt{\log\frac{|B_{\mathcal{M}}(x,19\beta\tau)|}{|B_{\mathcal{M}}(x,\beta\tau)|}},\tag{30}$$

and furthermore, if $\kappa>1$ is the universal constant from Proposition 9, then

$$\int_{\mathbb{R}^n}\omega\Big(A^*(v)\times B^*(v)\Big)\,\mathrm{d}\gamma_n(v)\gtrsim e^{-\kappa C^2}.\tag{31}$$

The conclusion of Theorem 4 about random zero sets is a simple formal consequence of the fact that for every fixed probability measure ω as in Theorem 12, there exist random pairs of sets as in Theorem 12 that are separated per (29) and ω-large per (31). The reason for this is mainly duality (minimax theorem), together with a simple scale gluing argument; the details of this deduction appear in the full version [19] of the present extended abstract. The idea to insert into the ARV reasoning a "weighting" such as ω on pairs of points of $\mathcal{M}$ is a key insight of [21], where it was introduced in order to prove that any n-point metric space of negative type embeds into ℓ_2 with distortion $O((\log n)^{3/4})$; this idea played the same role in [2], as well as herein.

The proof that Theorem 12 follows from Theorem 8, Proposition 9, and Proposition 11, which, as we explained above, is all that remains to complete the proof of Theorem 4, appear in the full version][19] of the present extended abstract. We will next conclude this extended abstract by sketching in broad strokes the reason why this works.

Write $n = |\mathcal{M}|$. Start by applying Proposition 11 with $r = \zeta\alpha$, where ζ is the universal constant in (22), thus ensuring that (30) coincides with (22). Henceforth in this sketch, $\mathsf{G} = (\mathcal{M}, E)$ will stand for the graph from this application of Proposition 11, i.e., its edges are given by (21) where ρ is defined in (22). The first bullet point in the conclusion of Proposition 11 makes it possible to apply Theorem 8 to get that

$$\int_{\mathbb{R}^n} \nu\Big(\mathsf{G}(v; f, \sigma)\Big)\, \mathrm{d}\gamma_n(v) \lesssim e^{-\frac{r^2}{4}C^2} n = e^{-\frac{\zeta^2\alpha^2}{4}C^2} n. \qquad (32)$$

Think of the expectation estimate (32) as expressing the following structural information about the Euclidean sparsification $\mathsf{G}(v; f, \sigma)$ of G in a typical direction $v \in \mathbb{R}^n$: for such v the graph $\mathsf{G}(v; f, \sigma)$ is "clustered" in the sense that it cannot have a large collection of disjoint edges, and hence there is a small set of vertices (i.e., a small subset of $\mathcal{M}$) which is incident to all of the edges in $\mathsf{G}(v; f, \sigma)$. Even though this is not quite how we will use Theorem 8 in the full version [19], namely, we will actually use a similar statement about fractional matchings of $\mathsf{G}(v; f, \sigma)$ that is a simple formal consequence of Theorem 8, for the purpose of intuitively understanding within the present sketch the reason why the deduction of Theorem 12 works, it suffices to initially consider the above combinatorial implication of (32).

As Theorem 12 assumes that $\omega(x, y) > 0 \implies d_{\mathcal{M}}(x, y) \geqslant \tau$ for every $x, y \in \mathcal{M}$, the conclusion (23) of Proposition 11 implies the assumption (15) of Proposition 9 with f replaced by Cf. We may therefore proceed to apply Proposition 9 with f replaced by Cf to get subsets $A(v), B(v) \subseteq \mathcal{M}$ for each $v \in \mathbb{R}^n$ such that (17) holds and, by canceling C in (16) with f replaced by Cf,

$$\begin{aligned}&\forall v \in \mathbb{R}^n,\ \forall (x, y) \in A(v) \times B(v),\\ &\qquad \{x, y\} \in E \implies |\langle v, f(x) - f(y)\rangle| > \Lambda.\end{aligned} \qquad (33)$$

The next observation is crucial. Recalling Definition 6, by the second part of conclusion (23) of Proposition 11, it follows from (33) that for every $v \in \mathbb{R}^n$, if $(x, y) \in A(v) \times B(v)$ and $\{x, y\} \in E$, then $\{x, y\}$ is also an edge of $\mathsf{G}(v; f, \sigma)$. Equivalently, if $(x, y) \in A(v) \times B(v)$ yet $\{x, y\}$ is not an edge of $\mathsf{G}(v; f, \sigma)$, then necessarily $\{x, y\} \notin E$, i.e., by the definition (21) of E the desired inequality in (29) holds.

Per the above discussion, for typical $v \in \mathbb{R}^n$ there is a small subset of $\mathcal{M}$ that is incident to all of the edges in $\mathsf{G}(v; f, \sigma)$, so by removing it we get large subsets $A^*(v) \subseteq A(v)$ and $B^*(v) \subseteq B^*(v)$ such that (29) holds. The notion of "large" here must be interpreted as the size of $A^*(v) \times B^*(v)$ with respect to the given measure ω on $\mathcal{M} \times \mathcal{M}$, since the input to this reasoning (supplied by Proposition 9) is (17); this is why we will actually work with a weighted version of (32) for fractional matchings, but, as we stated above, it is a formal consequence of Theorem 8 that is quickly deduced in the full version [19]. Finally, for the above reasoning to succeed, the lower bound in (17) needs to dominate the upper bound in (32); this is why α is assumed to be a sufficiently large universal constant in Theorem 12 (specified in the full version [19]).

Roadmap to the full version

The full version [19] of the present extended abstract, whose title is "Random zero sets with local growth guarantees," is publicly available at https://arxiv.org/pdf/2410.21931. It provides complete details of the proofs, more general results, multiple additional applications, formulations of conjectures for further research, and proofs of impossibility statements. The purpose of the present section is to provide an extensive guide to the contents of this full version.

A notable aspect of the full version [19] is that it demonstrates that the conclusions of Theorem 2 and Theorem 4 hold for a class of metric spaces that is much larger than those that are of negative type, so the applicability of the results that we obtain is significantly broader than what we presented herein, which suffices for the aforementioned implications to the Sparsest Cut problem. This is (a further generalization of) a realization of [61]. One might think that the fact that in the second constraint in (2) the power '2' coincides with the index '2' of the Euclidean norm $\|\cdot\|_2$ is of significance, but this is far from the truth and, in fact, the Euclidean distances could be raised there to any power, and furthermore much more dramatic deformations of the Euclidean metric could be allowed. By freeing oneself from adhering to the aforementioned coincidence it becomes easier to understand the geometric mechanism that leads to the stated results. Note that the use of the negative type condition (or the weaker requirement that we alluded to above) is confined to Proposition 11, which is the simplest step of the proof, while most of the reasoning occurs in a setting that has nothing to do with it.

Given $0 < s, \varepsilon < 1$, we say that a metric space $(\mathcal{M}, d_{\mathcal{M}})$ is (s, ε)-quasisymmetrically Hilbertian if there exists a one-to-one function f from $\mathcal{M}$ to a Hilbert space $(\mathbf{H}, \|\cdot\|_{\mathbf{H}})$ such that

$$\|f(x) - f(y)\|_{\mathbf{H}} \leqslant (1 - \varepsilon)\|f(x) - f(z)\|_{\mathbf{H}} \qquad (34)$$

for every $x, y, z \in \mathcal{M}$ satisfying $d_{\mathcal{M}}(x, y) \leqslant s d_{\mathcal{M}}(x, z)$. Call $\mathcal{M}$ quasisymetrically Hilbertian if there are $0 < s, \varepsilon < 1$ for which it is (s, ε)-quasisymmetrically Hilbertian. Observe that the requirement (34) is satisfied trivially for every metric space $\mathcal{M}$ when $\varepsilon = 0$ and any $s > 0$ (by assigning points in $\mathcal{M}$ to an orthonormal system). Thus, we are asking here for an embedding into a Hilbert space that is better—however slightly—than this trivial embedding, a requirement that holds when $\mathcal{M}$ is quasisymmetrically equivalent to a subset of a Hilbert space in the sense of Ahlfors–Beurling [13] and Tukia–Väisälä [70], which is a classical and extensively studied notion. The class of quasisymmetrically Hilbertian metric spaces encompasses disparate geometries, though not all metric spaces (e.g., any space that contains arbitrarily large expander graphs is not quasisymmetrically Hilbertian, and also ℓ_p for $p > 2$ is not quasisymmetrically Hilbertian [58]). Every metric space of negative type is $(s, 1-\sqrt{s})$-quasisymmetrically Hilbertian for any $0 < s < 1$. Furtheremore, the class of quasisymmetrically Hilbertian metric spaces encompasses those of finite Assouad–Nagata dimension [44], or more generally [62] metric spaces that admit a padded stochastic decomposition (thus, it includes, say, planar graphs [41] and doubling metric spaces [8, 31]), as well as $L_p(\mu)$ spaces for $1 \leqslant p \leqslant 2$ [18, 71], and the infinite dimensional Heisenberg group $\mathbb{H}^\infty$ [48].

One of the main embedding results in the full version [19], which includes as a special case the upper bound on the Euclidean distortion that is stated in Theorem 2, is that for every $0 < s, \varepsilon < 1$ and every integer $n \geqslant 2$, if an n-point metric space $(\mathcal{M}, d_\mathcal{M})$ is (s, ε)-quasisymmetrically Hilbertian, then

$$c_2(\mathcal{M}) \lesssim_{s,\varepsilon} \sqrt{\log n}. \tag{35}$$

In the influential work [36], Johnson and Lindenstrauss asked if $c_2(\mathcal{M}) \lesssim \sqrt{\log n}$ for any n-point metric space $(\mathcal{M}, d_\mathcal{M})$. If this were true, then it would have been a satisfactory analogue of John's theorem [35] that $c_2(\mathbf{X}) \leqslant \sqrt{\dim \mathbf{X}}$ for any finite-dimensional normed space $\mathbf{X}$, in accordance with the predictions of the Ribe program [10, 17, 27, 37, 59, 60, 64]. In [16, page 47] Bourgain reiterated the aforementioned question of Johnson–Lindentrauss, but famously answered it negatively (in the same work [16]); an asymptotically stronger (sharp) impossibility result was subsequently obtained (by another method) in [9, 50]. Hence, a positive answer to the Johnson–Lindenstrauss question would necessitate imposing restrictions on the metric space $(\mathcal{M}, d_\mathcal{M})$. The above stated result (35) from the full version shows that the answer is positive within the class of quasisymmetrically Hilbertian metric spaces.

The structural result that we stated above as Theorem 4 for metric spaces of negative type is shown in the full version [19] to hold whenever the metric space $(\mathcal{M}, d_\mathcal{M})$ is (s, ε)-quasisymmetrically Hilbertian for some $0 < s, \varepsilon < 1$, provided that one allows the parameter $\beta > 0$ that appears in the statement of Theorem 4 to depend on s, ε. In the full version [19] we derive the following consequence of this fact to a natural extremal question in the theory of concentration of measure, which was first broached in [61].

Gromov classically associated [30] a quantity called the observable diameter to any metric probability space $(\mathcal{M}, d_\mathcal{M}, \mu)$; see also [12, 46, 68] for an indication of the extensive literature on this notion. Prior to recalling the precise technical definition, we will next briefly and informally describe its intuitive meaning. The observable diameter quantifies the extent to which one could measure the size of the metric space using sufficiently smooth real-valued functions as observations, while accounting for possible observational errors by discarding at most a fixed fraction (with respect to the given probability measure μ) of the observations. Lévy's spherical isoperimetric theorem [49] implies that the ratio between the observable diameter of the Euclidean n-sphere to its actual (metric) diameter is of order $1/\sqrt{n}$. The full version [19] shows that the Euclidean n-sphere is extremal in this regard among all nondegenerate n-dimensional metric probability spaces that are quasisymmetrically Hilbertian (both "nondegenerate" and "n-dimensional" need to be suitably defined here; see below), namely, up to dimension-independent positive constant factors, the Euclidean n-sphere has the smallest possible observable diameter among all such spaces (this is new even for arbitrary Borel probability measures on $\mathbb{R}^n$).

Formally, for a metric space $(\mathcal{M}, d_\mathcal{M})$, a Borel probability measure μ on $\mathcal{M}$, and $\theta > 0$ the θ-observable diameter

$$\mathrm{ObsDiam}_\mu^\mathcal{M}(\theta)$$

of the metric measure space $(\mathcal{M}, d_\mathcal{M}, \mu)$ is defined [30] to be the supremum over all possible 1-Lipschitz functions $f : \mathcal{M} \to \mathbb{R}$ of the infimum of $\mathrm{diam}_\mathbb{R}(f(\mathcal{S})) = \sup\{|f(x) - f(y)| : x, y \in \mathcal{S}\}$ over all possible Borel subsets $\mathcal{S} \subseteq \mathcal{M}$ with $\mu(\mathcal{S}) \geqslant 1 - \theta$.

The Euclidean n-sphere is denoted $S^{n-1} = \{x \in \mathbb{R}^n : \|x\|_2 = 1\}$, and it will always be equipped with the metric that is inherited from ℓ_2^n and the normalized surface-area (probability) measure σ^{n-1}. Since ℓ_2^n is K_n-doubling for some $K_n \leqslant 5^n$ (e.g. [67, Section 2.2]), also $S^{n-1} \subseteq \ell_2^n$ is 5^n-doubling. Here we use the standard terminology (e.g. [23, 33, 45]) that the metric space $(\mathcal{M}, d_\mathcal{M})$ is K-doubling for some $K \in \mathbb{N}$ if for every $x \in \mathcal{M}$ and every $r > 0$ there are $x_1, \ldots, x_K \in \mathcal{M}$ such that $B_\mathcal{M}(x, 2r) \subseteq B_\mathcal{M}(x_1, r) \cup \cdots \cup B_\mathcal{M}(x_K, r)$, where $B_\mathcal{M}(x, \rho) = \{y \in \mathcal{M} : d_\mathcal{M}(x, y) \leqslant \rho\}$ denotes the closed $d_\mathcal{M}$-ball of radius $\rho \geqslant 0$ centered at $x \in \mathcal{M}$. If $(\mathcal{M}, d_\mathcal{M})$ is a metric space and μ is a Borel probability measure on $\mathcal{M}$, then we say that μ is normalized if the median of $d(x, y)$ is equal to 1 when $(x, y) \in \mathcal{M} \times \mathcal{M}$ is distributed according to $\mu \times \mu$. Note that σ^{n-1} is not normalized per this terminology, but one could normalize it by rescaling the metric by a factor that is bounded from above and from below by positive universal constants.

Using the above terminology, the full version [19] obtains the following theorem, which is the precise version of the extremal statement that we previously stated informally:

Theorem 13. *For every $0 < s, \varepsilon < 1$ and every $n \in \mathbb{N}$, if $(\mathcal{M}, d_\mathcal{M})$ is a metric space that is (s, ε)-quasisymmetrically Hilbertian and 5^n-doubling, then every normalized Borel probability measure μ on $\mathcal{M}$ satisfies*

$$\forall 0 < \theta \leqslant \theta_0, \qquad \mathrm{ObsDiam}_\mu^\mathcal{M}(\theta) \gtrsim_{s,\varepsilon} \mathrm{ObsDiam}_{\sigma^{n-1}}^{S^{n-1}}(\theta),$$

where $\theta_0 > 0$ is a universal constant.

We defer the discussion of further geometric results surrounding (the quasisymmetrically Hilbertian version of) Theorem 4 to the full version [19]. We will end by describing one more algorithmic application of Theorem 4 (more of its algorithmic ramifications appear in the full version [19]), which follows from the work of Makarychev, Makarychev and Vijayaraghavan [53] which beautifully relates the Sparsest Cut problem to perturbation resilience of the MaxCut problem (in the sense of Bilu–Linial [15]; see also the MaxCut-specific work [14], and the survey [52]).

The weighted MaxCut problem on n-vertices takes as its input an n-by-n symmetric matrix $\mathsf{W} = (w_{ij}) \in \mathsf{M}_n([0, \infty))$ whose entries are nonnegative, which one should think of as an edge-weighted graph whose vertex set is $[n]$, and aims to compute (or approximate) in polynomial time the following quantity:

$$\mathrm{MaxCut}(\mathsf{W}) \stackrel{\mathrm{def}}{=} \max_{\emptyset \neq S \subsetneq [n]} \sum_{i \in S} \sum_{j \in [n] \setminus S} w_{ij}. \tag{36}$$

Given $\gamma \geqslant 1$, a matrix W as above is said to be Bilu–Linial γ-stable (γ-stable, in short) for MaxCut if there is a unique nontrivial subset $\emptyset \neq S \subsetneq [n]$ such that $\mathrm{MaxCut}(\mathsf{W}) = \sum_{i \in S} \sum_{j \in [n] \setminus S} w_{ij}$, i.e., the right hand side of (36) has only one maximizer, and furthermore we have $\mathrm{MaxCut}(\mathsf{W}') = \sum_{i\in S}\sum_{j\in[n]\smallsetminus S} w'_{ij}$ for every symmetric matrix $\mathsf{W}' = (w'_{ij}) \in \mathsf{M}_n([0,\infty))$ that satisfies $w_{ij} \leqslant w'_{ij} \leqslant \gamma w_{ij}$ for every $i, j \in [n]$.

Makarychev, Makarychev and Vijayaraghavan studied [53] an SDP that can be evaluated (with $o(1)$ precision) in polynomial time and outputs a number $\mathrm{SDP}_{\mathsf{MMV}}(\mathsf{W}) \geqslant 0$ that satisfies

$$\mathrm{SDP}_{\mathsf{MMV}}(\mathsf{W}) \geqslant \mathrm{MaxCut}(\mathsf{W}). \tag{37}$$

The Makarychev–Makarychev–Vijayaraghavan SDP for MaxCut coincides with the Goemans–Williamson SDP for MaxCut [29] with the (by now standard) added squared-ℓ_2 triangle inequality constraints, except for the "twist" that those constraints are imposed on the symmetrized version of the vector solution, namely, if the SDP outputs unit vectors $v_1,\ldots,v_n \in S^{n-1}$, then they require that $\|u-v\|_2^2 \leqslant \|u-w\|_2^2 + \|w-v\|_2^2$ for every choice of three vectors $u, v, w \in \{v_1,\ldots,v_n,-v_1,\ldots,-v_n\}$ (see [1, 38] for earlier incarnations of this idea, along with demonstrations of its utility in other aspects of combinatorial optimization).

The Makarychev–Makarychev–Vijayaraghavan SDP for MaxCut is said to be integral on the input W (see [53, Definition 2.3]) if for every optimal vector solution $v_1,\ldots,v_n \in S^{n-1}$ of this SDP there is a unit vector $e \in S^{n-1}$ such that $v_i \in \{e,-e\}$ for every $i \in [n]$. When this occurs, every optimal solution corresponds to a valid cut, so thanks to (37) we have $\mathrm{SDP}_{\mathsf{MMV}}(\mathsf{W}) = \mathrm{MaxCut}(\mathsf{W})$, i.e., the Makarychev–Makarychev–Vijayaraghavan algorithm outputs an exact solution for MaxCut on the input W, in contrast to the fact that (under standard complexity hypotheses) one cannot hope to always obtain such an exact solution on general inputs [32, 39, 56, 65].

By substituting Theorem 2 into [53, Theorem 3.1] and [53, Theorem 5.2], we get the following evaluation of the growth rate of the critical $\gamma = \gamma(n)$ such that the Makarychev–Makarychev–Vijayaraghavan SDP for MaxCut is integral on every n-vertex instance of MaxCut that is γ-stable:

THEOREM 14. *There are universal constants $C > c > 0$ with the following property for any integer $n \geqslant 2$. If $\gamma \geqslant C\sqrt{\log n}$, then the Makarychev–Makarychev–Vijayaraghavan SDP relaxation of MaxCut is integral on any n-vertex input that is γ-stable for MaxCut. If $1 \leqslant \gamma \leqslant c\sqrt{\log n}$, then there is an n-vertex input that is γ-stable for MaxCut on which the Makarychev–Makarychev–Vijayaraghavan SDP relaxation of MaxCut is not integral.*

The connection between the Sparsest Cut problem and perturbation resilience of the MaxCut problem was discovered in [53], where Theorem 14 was obtained with the lower bound $\gamma \geqslant C\sqrt{\log n}$ replaced by $\gamma \geqslant C\sqrt{\log n}\log\log n$, as [53] appealed to [4]. Thus, the impact of Theorem 2 in this context is to obtain the sharp rate of growth of the threshold for stable integrality.

Acknowledgments

A. C. was supported by NSF grant DMS-2247233. A. N. was supported by NSF grant DMS-2054875, BSF grant 2018223, and the Simons Investigator award. K. R. was supported by an NSF GRFP fellowship.

References

[1] Amit Agarwal, Moses Charikar, Konstantin Makarychev, and Yury Makarychev. 2005. $O(\sqrt{\log n})$ 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*. ACM, New York, 573–581. doi:10.1145/1060590.1060675

[2] Sanjeev Arora, James R. Lee, and Assaf Naor. 2005. Euclidean distortion and the sparsest cut [extended abstract]. In *STOC'05: Proceedings of the 37th Annual ACM Symposium on Theory of Computing*. ACM, New York, 553–562. doi:10.1145/1060590.1060673

[3] Sanjeev Arora, James R. Lee, and Assaf Naor. 2007. Fréchet embeddings of negative type metrics. *Discrete Comput. Geom.* 38, 4 (2007), 726–739. doi:10.1007/s00454-007-9007-0

[4] Sanjeev Arora, James R. Lee, and Assaf Naor. 2008. Euclidean distortion and the sparsest cut. *J. Amer. Math. Soc.* 21, 1 (2008), 1–21. doi:10.1090/S0894-0347-07-00573-5

[5] Sanjeev Arora, James R. Lee, and Sushant Sachdeva. 2011. A Reformulation of the Arora–Rao–Vazirani Structure Theorem. (2011). Unpublished manuscript, available at https://arxiv.org/abs/1102.1456.

[6] Sanjeev Arora, Satish Rao, and Umesh Vazirani. 2004. Expander flows, geometric embeddings and graph partitioning. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing*. ACM, New York, 222–231. doi:10.1145/1007352.1007355

[7] Sanjeev Arora, Satish Rao, and Umesh Vazirani. 2009. Expander flows, geometric embeddings and graph partitioning. *Journal of the ACM (JACM)* 56, 2 (2009), 1–37. doi:10.1145/1502793.1502794

[8] Patrice Assouad. 1983. Plongements lipschitziens dans $\mathbf{R}^n$. *Bull. Soc. Math. France* 111, 4 (1983), 429–448. http://www.numdam.org/item?id=BSMF_1983__111__429_0

[9] Yonatan Aumann and Yuval Rabani. 1998. An $O(\log k)$ approximate min-cut max-flow theorem and approximation algorithm. *SIAM J. Comput.* 27, 1 (1998), 291–301. doi:10.1137/S0097539794285983

[10] Keith Ball. 2013. The Ribe programme. Number 352. Exp. No. 1047, viii, 147–159. Séminaire Bourbaki. Vol. 2011/2012. Exposés 1043–1058.

[11] Yair Bartal, Nathan Linial, Manor Mendel, and Assaf Naor. 2006. Limitations to Fréchet's metric embedding method. *Israel J. Math.* 151 (2006), 111–124. doi:10.1007/BF02777357

[12] Jerôme Bertrand and Max Fathi. 2022. Stability of eigenvalues and observable diameter in $\mathrm{RCD}(1,\infty)$ spaces. *J. Geom. Anal.* 32, 11 (2022), Paper No. 270, 38. doi:10.1007/s12220-022-00999-9

[13] A. Beurling and L. Ahlfors. 1956. The boundary correspondence under quasiconformal mappings. *Acta Math.* 96 (1956), 125–142. doi:10.1007/BF02392360

[14] Yonatan Bilu, Amit Daniely, Nati Linial, and Michael Saks. 2013. On the practically interesting instances of MAXCUT. In *30th International Symposium on Theoretical Aspects of Computer Science*. LIPIcs. Leibniz Int. Proc. Inform., Vol. 20. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 526–537.

[15] Yonatan Bilu and Nathan Linial. 2012. Are stable instances easy? *Combin. Probab. Comput.* 21, 5 (2012), 643–660. doi:10.1017/S0963548312000193

[16] J. Bourgain. 1985. On Lipschitz embedding of finite metric spaces in Hilbert space. *Israel J. Math.* 52, 1-2 (1985), 46–52. doi:10.1007/BF02776078

[17] J. Bourgain. 1986. The metrical interpretation of superreflexivity in Banach spaces. *Israel J. Math.* 56, 2 (1986), 222–230. doi:10.1007/BF02766125

[18] Jean Bretagnolle, Didier Dacunha-Castelle, and Jean-Louis Krivine. 1965/66. Lois stables et espaces L^p. *Ann. Inst. H. Poincaré Sect. B (N.S.)* 2 (1965/66), 231–259.

[19] Alan Chang, Assaf Naor, and Kevin Ren. 2024. Random zero sets with local growth guarantees. (2024). Preprint available at https://arxiv.org/pdf/2410.21931.

[20] Shuchi Chawla. 2008. Sparsest Cut. In *Encyclopedia of Algorithms*. Springer-Verlag US, 868–870. doi:10.1007/978-0-387-30162-4_388

[21] Shuchi Chawla, Anupam Gupta, and Harald Räcke. 2005. Embeddings of negative-type metrics and an improved approximation to generalized sparsest cut. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms*. ACM, New York, 102–111. doi:10.1145/1361192.1361199

[22] Shuchi Chawla, Anupam Gupta, and Harald Räcke. 2008. Embeddings of negative-type metrics and an improved approximation to generalized sparsest cut. *ACM Transactions on Algorithms (TALG)* 4, 2 (2008), 1–18. doi:10.1145/1361192.1361199

[23] Ronald R. Coifman and Guido Weiss. 1971. *Analyse harmonique non-commutative sur certains espaces homogènes*. Lecture Notes in Mathematics, Vol. Vol. 242. Springer-Verlag, Berlin-New York. v+160 pages. Étude de certaines intégrales singulières.

[24] Michel Marie Deza and Monique Laurent. 2010. *Geometry of cuts and metrics*. Algorithms and Combinatorics, Vol. 15. Springer, Heidelberg. xii+587 pages. doi:10.1007/978-3-642-04295-9 First softcover printing of the 1997 original [MR1460488].

[25] Per Enflo. 1969. On the nonexistence of uniform homeomorphisms between L_p-spaces. *Ark. Mat.* 8 (1969), 103–105. doi:10.1007/BF02589549

[26] Uriel Feige, Mohammadtaghi Hajiaghayi, and James R. Lee. 2008. Improved approximation algorithms for minimum weight vertex separators. *SIAM J. Comput.* 38, 2 (2008), 629–657. doi:10.1137/05064299X

[27] G. Godefroy. 2017. De Grothendieck à Naor: une promenade dans l'analyse métrique des espaces de Banach. *Gaz. Math.* 151 (2017), 13–24.

[28] Michel X. Goemans. 1997. Semidefinite programming in combinatorial optimization. Vol. 79. 143–161. doi:10.1016/S0025-5610(97)00051-8 Lectures on mathematical programming (ismp97) (Lausanne, 1997).

[29] Michel X. Goemans and David P. Williamson. 1994. .879-approximation algorithms for MAX CUT and MAX 2SAT. In *Proceedings of the Twenty-Sixth Annual ACM Symposium on Theory of Computing, 23-25 May 1994, Montréal, Québec, Canada*, Frank Thomson Leighton and Michael T. Goodrich (Eds.). ACM, 422–431. doi:10.1145/195058.195216

[30] Misha Gromov. 2007. *Metric structures for Riemannian and non-Riemannian spaces* (english ed.). Birkhäuser Boston, Inc., Boston, MA. xx+585 pages. Based on the 1981 French original, With appendices by M. Katz, P. Pansu and S. Semmes, Translated from the French by Sean Michael Bates.

[31] Anupam Gupta, Robert Krauthgamer, and James R. Lee. 2003. Bounded Geometries, Fractals, and Low-Distortion Embeddings. In *44th Symposium on Foundations of Computer Science (FOCS 2003), 11-14 October 2003, Cambridge, MA, USA, Proceedings*. IEEE Computer Society, 534–543. doi:10.1109/SFCS.2003.1238226

[32] Johan Håstad. 2001. Some optimal inapproximability results. *J. ACM* 48, 4 (2001), 798–859. doi:10.1145/502090.502098

[33] Juha Heinonen. 2001. *Lectures on analysis on metric spaces*. Springer-Verlag, New York. x+140 pages. doi:10.1007/978-1-4613-0131-8

[34] Piotr Indyk and Jirí Matoušek. 2004. Low-Distortion Embeddings of Finite Metric Spaces. In *Handbook of Discrete and Computational Geometry, Second Edition*, Jacob E. Goodman and Joseph O'Rourke (Eds.). Chapman and Hall/CRC, 177–196. doi:10.1201/9781420035315.CH8

[35] Fritz John. 1948. Extremum problems with inequalities as subsidiary conditions. In *Studies and Essays Presented to R. Courant on his 60th Birthday, January 8, 1948*. Interscience Publishers, New York, 187–204.

[36] William B. Johnson and Joram Lindenstrauss. 1984. Extensions of Lipschitz mappings into a Hilbert space. In *Conference in modern analysis and probability (New Haven, Conn., 1982)*. Contemp. Math., Vol. 26. Amer. Math. Soc., Providence, RI, 189–206. doi:10.1090/conm/026/737400

[37] Nigel J. Kalton. 2008. The nonlinear geometry of Banach spaces. *Rev. Mat. Complut.* 21, 1 (2008), 7–60. doi:10.5209/rev_REMA.2008.v21.n1.16426

[38] George Karakostas. 2005. A Better Approximation Ratio for the Vertex Cover Problem. In *Automata, Languages and Programming, 32nd International Colloquium, ICALP 2005, Lisbon, Portugal, July 11-15, 2005, Proceedings (Lecture Notes in Computer Science, Vol. 3580)*, Luís Caires, Giuseppe F. Italiano, Luís Monteiro, Catuscia Palamidessi, and Moti Yung (Eds.). Springer, 1043–1050. doi:10.1007/11523468_84

[39] Subhash Khot, Guy Kindler, Elchanan Mossel, and Ryan O'Donnell. 2007. Optimal inapproximability results for MAX-CUT and other 2-variable CSPs? *SIAM J. Comput.* 37, 1 (2007), 319–357. doi:10.1137/S0097539705447372

[40] Subhash A. Khot and Nisheeth K. Vishnoi. 2015. The unique games conjecture, integrability gap for cut problems and embeddability of negative-type metrics into ℓ_1. *J. ACM* 62, 1 (2015), Art. 8, 39. doi:10.1145/2629614

[41] Philip N. Klein, Serge A. Plotkin, and Satish Rao. 1993. Excluded minors, network decomposition, and multicommodity flow. In *Proceedings of the Twenty-Fifth Annual ACM Symposium on Theory of Computing, May 16-18, 1993, San Diego, CA, USA*, S. Rao Kosaraju, David S. Johnson, and Alok Aggarwal (Eds.). ACM, 682–690. doi:10.1145/167088.167261

[42] R. Krauthgamer, J.R. Lee, M. Mendel, and A. Naor. 2004. Measured descent: a new embedding method for finite metrics [extended abstract]. In *45th Annual IEEE Symposium on Foundations of Computer Science*. 434–443. doi:10.1109/FOCS.2004.41

[43] R. Krauthgamer, J. R. Lee, M. Mendel, and A. Naor. 2005. Measured descent: a new embedding method for finite metrics. *Geom. Funct. Anal.* 15, 4 (2005), 839–858. doi:10.1007/s00039-005-0527-6

[44] Urs Lang and Thilo Schlichenmaier. 2005. Nagata dimension, quasisymmetric embeddings, and Lipschitz extensions. *Int. Math. Res. Not.* 58 (2005), 3625–3655. doi:10.1155/IMRN.2005.3625

[45] D. G. Larman. 1967. A new theory of dimension. *Proc. London Math. Soc. (3)* 17 (1967), 178–192. doi:10.1112/plms/s3-17.1.178

[46] Michel Ledoux. 2001. *The concentration of measure phenomenon*. Mathematical Surveys and Monographs, Vol. 89. American Mathematical Society, Providence, RI. x+181 pages. doi:10.1090/surv/089

[47] James R. Lee. 2005. On distance scales, embeddings, and efficient relaxations of the cut cone. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2005, Vancouver, British Columbia, Canada, January 23-25, 2005*. SIAM, 92–101. http://dl.acm.org/citation.cfm?id=1070432.1070446

[48] James R. Lee and Assaf Naor. 2006. L_p metrics on the Heisenberg group and the Goemans–Linial conjecture. In *47th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2006), 21-24 October 2006, Berkeley, California, USA, Proceedings*. IEEE Computer Society, 99–108. doi:10.1109/FOCS.2006.47

[49] Paul Lévy. 1951. *Problèmes concrets d'analyse fonctionnelle. Avec un complément sur les fonctionnelles analytiques par F. Pellegrino*. Gauthier-Villars, Paris. xiv+484 pages. 2d ed.

[50] Nathan Linial, Eran London, and Yuri Rabinovich. 1995. The geometry of graphs and some of its algorithmic applications. *Combinatorica* 15, 2 (1995), 215–245. doi:10.1007/BF01200757

[51] László Lovász and Michael D. Plummer. 2009. *Matching theory*. AMS Chelsea Publishing, Providence, RI. xxxiv+554 pages. doi:10.1090/chel/367 Corrected reprint of the 1986 original [MR0859549].

[52] Konstantin Makarychev and Yury Makarychev. 2016. Metric extension operators, vertex sparsifiers and Lipschitz extendability. *Israel J. Math.* 212, 2 (2016), 913–959. doi:10.1007/s11856-016-1315-8

[53] Konstantin Makarychev, Yury Makarychev, and Aravindan Vijayaraghavan. 2014. Bilu-Linial stable instances of max cut and minimum multiway cut. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*. ACM, New York, 890–906. doi:10.1137/1.9781611973402.67

[54] Jiří Matoušek. 2002. *Lectures on discrete geometry*. Graduate Texts in Mathematics, Vol. 212. Springer-Verlag, New York. xvi+481 pages. doi:10.1007/978-1-4613-0039-7

[55] Jiří Matoušek and Yuri Rabinovich. 2001. On dominated l_1 metrics. *Israel J. Math.* 123 (2001), 285–301. doi:10.1007/BF02784132

[56] Elchanan Mossel, Ryan O'Donnell, and Krzysztof Oleszkiewicz. 2010. Noise stability of functions with low influences: invariance and optimality. *Ann. of Math. (2)* 171, 1 (2010), 295–341. doi:10.4007/annals.2010.171.295

[57] Assaf Naor. 2010. L_1 embeddings of the Heisenberg group and fast estimation of graph isoperimetry. In *Proceedings of the International Congress of Mathematicians. Volume III*. Hindustan Book Agency, New Delhi, 1549–1575. doi:10.1142/9789814324359_0110

[58] Assaf Naor. 2012. An application of metric cotype to quasisymmetric embeddings. In *Metric and differential geometry*. Progr. Math., Vol. 297. Birkhäuser/Springer, Basel, 175–178. doi:10.1007/978-3-0348-0257-4_7

[59] Assaf Naor. 2012. An introduction to the Ribe program. *Jpn. J. Math.* 7, 2 (2012), 167–233. doi:10.1007/s11537-012-1222-7

[60] Assaf Naor. 2018. Metric dimension reduction: a snapshot of the Ribe program. In *Proceedings of the International Congress of Mathematicians—Rio de Janeiro 2018. Vol. I. Plenary lectures*. World Sci. Publ., Hackensack, NJ, 759–837.

[61] Assaf Naor, Yuval Rabani, and Alistair Sinclair. 2005. Quasisymmetric embeddings, the observable diameter, and expansion properties of graphs. *Journal of Functional Analysis* 227, 2 (2005), 273–303. doi:10.1016/j.jfa.2005.04.003

[62] Assaf Naor and Lior Silberman. 2011. Poincaré inequalities, embeddings, and wild groups. *Compos. Math.* 147, 5 (2011), 1546–1572. doi:10.1112/S0010437X11005343

[63] Assaf Naor and Robert Young. 2018. Vertical perimeter versus horizontal perimeter. *Annals of Mathematics* 188, 1 (2018), 171–279. doi:10.4007/annals.2018.188.1.4

[64] Mikhail I. Ostrovskii. 2013. *Metric embeddings*. De Gruyter Studies in Mathematics, Vol. 49. De Gruyter, Berlin. xii+372 pages. doi:10.1515/9783110264012 Bilipschitz and coarse embeddings into Banach spaces.

[65] Christos H. Papadimitriou and Mihalis Yannakakis. 1991. Optimization, approximation, and complexity classes. *J. Comput. System Sci.* 43, 3 (1991), 425–440. doi:10.1016/0022-0000(91)90023-X

[66] Thomas Rothvoss. 2016. Lecture notes on the ARV algorithm for sparsest cut. (2016). doi:10.48550/arXiv.1607.00854 Unpublished lecture notes.

[67] Stephen Semmes. 2001. *Some novel types of fractal geometry*. The Clarendon Press, Oxford University Press, New York. x+164 pages.

[68] Takashi Shioya. 2016. *Metric measure geometry*. IRMA Lectures in Mathematics and Theoretical Physics, Vol. 25. EMS Publishing House, Zürich. xi+182 pages. doi:10.4171/158 Gromov's theory of convergence and concentration of metrics and measures.

[69] David B. Shmoys. 1996. Cut problems and their application to divide and conquer. In *Approximation algorithms for NP-hard problems*, Dorit S. Hochbaum (Ed.). Boston, MA: PWS Publishing, Chapter 5, 192–235.

[70] P. Tukia and J. Väisälä. 1980. Quasisymmetric embeddings of metric spaces. *Ann. Acad. Sci. Fenn. Ser. A I Math.* 5, 1 (1980), 97–114. doi:10.5186/aasfm.1980.0531

[71] J. H. Wells and L. R. Williams. 1975. *Embeddings and extensions in analysis*. Ergebnisse der Mathematik und ihrer Grenzgebiete [Results in Mathematics and Related Areas], Vol. Band 84. Springer-Verlag, New York-Heidelberg. vii+108 pages.

Received 2024-11-04; accepted 2025-02-01