

Rounding via Low Dimensional Embeddings

Mark Braverman 

Department of Computer Science, Princeton University, NJ, USA

Dor Minzer 

Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA, USA

Abstract

A regular graph $G = (V, E)$ is an (ε, γ) small-set expander if for any set of vertices of fractional size at most ε , at least γ of the edges that are adjacent to it go outside. In this paper, we give a unified approach to several known complexity-theoretic results on small-set expanders. In particular, we show:

1. Max-Cut: we show that if a regular graph $G = (V, E)$ is an (ε, γ) small-set expander that contains a cut of fractional size at least $1 - \delta$, then one can find in G a cut of fractional size at least $1 - O\left(\frac{\delta}{\varepsilon\gamma^6}\right)$ in polynomial time.
2. Improved spectral partitioning, Cheeger's inequality and the parallel repetition theorem over small-set expanders. The general form of each one of these results involves square-root loss that comes from certain rounding procedure, and we show how this can be avoided over small set expanders.

Our main idea is to project a high dimensional vector solution into a low-dimensional space while roughly maintaining ℓ_2^2 distances, and then perform a pre-processing step using low-dimensional geometry and the properties of ℓ_2^2 distances over it. This pre-processing leverages the small-set expansion property of the graph to transform a vector valued solution to a different vector valued solution with additional structural properties, which give rise to more efficient integral-solution rounding schemes.

2012 ACM Subject Classification Theory of computation \rightarrow Problems, reductions and completeness; Theory of computation \rightarrow Discrete optimization

Keywords and phrases Parallel Repetition, Small Set Expanders, Semi-Definite Programs

Digital Object Identifier 10.4230/LIPIcs.ITCS.2023.26

Funding *Mark Braverman*: Research supported in part by the NSF Alan T. Waterman Award, Grant No. 1933331, a Packard Fellowship in Science and Engineering, and the Simons Collaboration on Algorithms and Geometry.

Dor Minzer: Supported by a Sloan Research Fellowship.

1 Introduction

Expander graphs are important combinatorial objects with myriad of applications throughout theoretical computer science (see [11]). One of the reasons for that are the numerous equivalent definitions of expander graphs, that offer different views on this basic object. Combinatorially, a d -regular graph $G = (V, E)$ is said to be a combinatorial expander if the *edge expansion* of any $S \subseteq V$ of size at most $n/2$ is at least an absolute constant.

► **Definition 1.** Let $G = (V \cup U, E)$ be an undirected regular graph. The edge expansion of $S \subseteq V$ is defined as $\Phi(S) = \frac{|\{(v,u) \mid v \in S, u \notin S\}|}{|\{(v,u) \mid v \in S\}|}$. In words, $\Phi(S)$ is the fraction of edges adjacent to S going outside it.

Spectrally, a graph is said to be an expander if, letting A_G be the normalized transition matrix of G , the second eigenvalue of A_G is at most $1 - \Omega(1)$. A basic result in spectral graph theory, called Cheeger's inequality, asserts that qualitatively, spectral expanders and combinatorial expanders are equivalent.



© Mark Braverman and Dor Minzer;
licensed under Creative Commons License CC-BY 4.0
14th Innovations in Theoretical Computer Science Conference (ITCS 2023).

Editor: Yael Tauman Kalai; Article No. 26; pp. 26:1–26:30



Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

26:2 Rounding via Low Dimensional Embeddings

The main object studied in this paper is small set expanders, which is a relaxation of expander graphs defined as follows.

► **Definition 2.** A graph $G = (V \cup U, E)$ is said to be (ε, γ) small-set expanding if for any $S \subseteq V$ of size at most $\varepsilon|V|$ we have that $\Phi(S) \geq \gamma$.

Upon first sight, small-set expanders may seem like a slight variation over the combinatorial definition of expander graphs. However this seemingly small change makes a big of difference, and indeed small-set expanders are more difficult to study. For instance, there is no general known equivalence between the spectrum of the adjacency matrix of a graph and the small-set expansion of the graph. This is partly the reason that small-set expanders are typically much harder to work with.

The main goal of this paper is to present the idea of “low-dimension embeddings” and show how to use it to give alternative, more direct proofs to several results concerning computational problems over small set expanders. In particular, we (1) recover the results of [23, 16] about strong parallel repetition for Unique-Games over small set expanders, (2) we show an improved Cheeger’s inequality as well as approximation algorithm for Max-Cut over small set expanders, recovering results of [14, 15].

Conceptually, we show that for several cases (such as the computational problems mentioned above), it is beneficial to consider an intermediate “low-dimensional” projection step which incurs less quantitative loss when compared to the usual integral rounding procedures. In a low dimension, one has additional tools in their disposal to manipulate a vector valued solution so that it is possible to perform a final integral rounding procedure with better performance. We believe that such ideas may be helpful towards getting improved analysis of other semi-definite programming relaxations of combinatorial optimization problems, and hope that the current work encourages research along this direction.

1.1 Our Results

In this section, we formally state the main results of this paper.

1.1.1 Solving Max-Cut on Small Set Expanders

Our first result addresses the Max Cut problem. We show that on small set expanders, one can get an algorithm achieving a better approximation ratio compared to the Geomans-Williamson algorithm [9].

► **Theorem 3.** *There exists an absolute constant $C > 0$, such that the following holds for all $\varepsilon, \delta, \gamma > 0$. There is an efficient algorithm for the following task:*

Input: a regular, (ε, γ) small-set expander $G = (V \cup U, E)$ containing a cut of fractional size $1 - \delta$.

Output: a bipartition $V = V_1 \cup V_2$ of the vertices that cuts at least $1 - C \frac{\delta}{\varepsilon \gamma^6}$ fraction of the edges.

1.1.2 Spectral Partitioning

Our second result is an improved quantitative for Cheeger’s inequality when the underlying graph is a small set expander. Recall that one side of Cheeger’s inequality states that if $G = (V, E)$ is a regular graph such that $\lambda_2(G) \leq 1 - \delta$, then one can find in polynomial time a set $S \subseteq V$ of size at most $|V|/2$ such that $\Phi(S) \leq \sqrt{2\delta}$. The other side of Cheeger’s inequality asserts that if $\lambda_2(G) = 1 - \delta$, then $\min_{|S| \leq |V|/2} \Phi(S) \geq \delta/2$, hence the upper

bound is tight up to the square root (and constant factors). We show that for small set expanders, one can considerably improve on the upper bound, and get a result which is much closer to the lower bound in Cheeger's inequality:

► **Theorem 4.** *There exists an absolute constant $C > 0$, such that if G is a (ε, γ) small set expander and $\lambda_2(G) \geq 1 - \delta$, then G contains a set of size at most $n/2$ such that $\Phi(S) \leq C \cdot \frac{\delta}{\gamma^3 \varepsilon^3}$. Furthermore, such set S can be found efficiently.*

1.1.3 Parallel Repetition

Our third result is about parallel repetition when applied on Unique-Games over small-set expanders. We begin by formally defining the Unique-Games problem.

► **Definition 5.** *A Unique-Games instance is composed of a graph $G = (V, E)$, a finite alphabet Σ and a collection of constraints $\{\phi_e\}_{e \in E}$, one for each edge. The constraint on an edge e is a 1-to-1 relation $\phi_e \subseteq \Sigma \times \Sigma$, specifying for each edge the tuples of labels that are considered satisfactory.*

Given a Unique-Games instance $\Psi = (G, \Sigma, \Phi = \{\phi_e\}_{e \in E})$, the value of Ψ , denoted by $\text{val}(\Psi)$, is the maximum fraction of constraints that may be satisfied by any assignment $A: V \rightarrow \Sigma$. A central conjecture in complexity theory, called the Unique-Games Conjecture [12], asserts for any $\varepsilon, \delta > 0$, there is $k \in \mathbb{N}$ such that given a Unique-Games instance Ψ with alphabet size k , it is NP-hard to distinguish between the case that $\text{val}(\Psi) \geq 1 - \varepsilon$ and the case that $\text{val}(\Psi) \leq \delta$. A well known approach to the Unique-Games Conjecture proceeds by first establishing a weak form of the conjecture, wherein δ is also close to 1, and then amplifying the gap via *parallel repetition*. Here, given a Unique-Games instance Ψ and a parameter $t \in \mathbb{N}$, the t -fold repeated game corresponds to the tensor product of the game is $\Psi^{\otimes t} = (G' = (V^t, E'), \Sigma^t, \Phi')$, where

$$E' = \{((u_1, \dots, u_t), (v_1, \dots, v_t)) \mid (u_i, v_i) \in E \forall i \in [t]\},$$

and $\Phi' = \{\phi_{e'}\}_{e' \in E'}$ where the constraint on edge e' corresponds to the AND of the t constraints on the t edges of the original graph G .

The parallel repetition theorem of Raz [20] and subsequent improvements [10, 4, 19] turn out not to be good enough; such results are only able to show that for a unique game Ψ , if $\text{val}(\Psi) \leq 1 - \varepsilon$, then the value of the t -fold repeated game is at most $(1 - \varepsilon^2)^{\Omega(t)}$. This square root often times does not matter in application, however in the context of Unique-Games it is crucial. Indeed, the weak forms of the Unique-Games conjecture that seem plausible may go to soundness which is as small as $\delta = 1 - \Theta(\sqrt{\varepsilon})$ (though current technology does not even close to this), thereby offering a quadratic difference at best. This makes the quadratic loss in parallel repetition unaffordable for the purposes of proving the Unique-Games Conjecture.

This raises the question: on which classes of graphs can one prove an improved parallel repetition theorem, surpassing the square barrier? Much effort had been devoted to study this question, starting with the tightness of the parallel repetition theorem [20, 13, 1, 8] and subsequently studying parallel repetition over special classes of graphs [3, 2, 21, 6, 23, 16, 17], most popularly on expanders and small-set expanders. It is worth noting that while it is known that the Unique-Games fails on such graph [2, 18], it is still possible that there is a regime of parameters for which a weak form of the Unique Games Conjecture holds, and for which a strong version of the parallel repetition theorem holds.

Using our technique, we are able to recover the result of [23, 16], and show that strong parallel repetition holds for unique games over small set expanders:

► **Theorem 6.** *There exists an absolute constant $c > 0$, such that the following holds for all $0 < \varepsilon, \gamma, \delta < 1$. If $\Psi = (G, \Sigma, \Phi)$ is a Unique-Games instance such that G is (ε, γ) small-set expanding, regular graph whose value is at most $1 - \delta$, then $\text{val}(\Psi^{\otimes t}) \leq (1 - \delta)^{c\varepsilon^3\gamma^3t}$.*

We remark that our proof applies to the more general class of projection games, but we state it only for Unique-Games for the purpose of this introduction.

1.2 Our Technique

The proofs of Theorems 3 and 4 follow a similar theme. In both cases, there are well known classical results that are very much similar to it except that they incur an additional square-root loss: Cheeger’s inequality and the Goemans Williamson algorithm [9]. The source for this quadratic loss comes from a rounding phase in their proofs, which transforms a vector-valued solution with an ℓ_2^2 distances guarantee to an integral valued solution with an ℓ_1 distances guarantee via rounding. Our proofs of Theorems 3 and 4 avoid this quadratic loss by incorporating a preprocessing phase and a “soft” rounding phase, which incur constant factors losses depending on the small set expansion parameters of the graphs, but ensures that the subsequent integral rounding phase would be almost lossless.

Low-dimensional embeddings

Our arguments also utilize the idea of low-dimensional embeddings. Here, one is given a collection of vectors (typically an SDP solution to some problem, such as Max-Cut or Unique-Games) which a-priori may lie in a high dimensional Euclidean space. Common wisdom, which manifests itself in standard rounding algorithms (such as Geomans-Williamson [9]), utilize such vectors to generate 1-dimensional Gaussians vectors with certain correlations. Our idea is to use higher (but constant) dimensional projections as a way to reduce the dimension of the vector while roughly preserving the quality of the solution. This allows us to apply ideas from low-dimensional geometry and construct a rounding scheme, which for small-set expanders, outperforms the standard rounding schemes for these problems.

Similar, yet different ideas of low-dimensional projections have already made some appearances in the literature, in the context of higher order Cheeger inequalities [15] and in establishing improved bounds on the Grothendieck constant [5]. Below, we explain in more detail our proof strategy for our results.

1.2.1 Proof of Theorem 3

Our proof of Theorem 3 makes use of the Goemans-Williamson semi-definite program relaxation of Max-Cut [9]. Our algorithm proceeds by the following steps:

1. **Solving the SDP program:** We find unit vectors in \mathbb{R}^m that achieve value at least $1 - \delta$, where $m \leq n$. In the next two steps we pre-process these vectors to get a (different) collection of vectors that can be used in a rounding scheme with improved guarantees.
2. **Dimension reduction:** We show, via appropriate random projections, that we can reduce m to 3 as long as we are willing to take a multiplicative cost in the error. More precisely, we show we can find a solution consisting of unit vectors in \mathbb{R}^3 that achieve value at least $1 - C\delta$, for some absolute value $C > 0$. Denote them by $(x_v)_{v \in V}$.

We remark that it is also possible to reduce the dimension down to 2, and moreover that it simplifies subsequent steps. However, doing so may reduce the value of the resulting solution to $1 - C\delta \log(1/\delta)$, and we would like to avoid this logarithmic factor.

- 3. Soft rounding:** The main issue in the standard hyperplane rounding procedure in the Goemans-Williamson algorithm is the case that the vectors x_v are uniformly distributed on the sphere, satisfying that $\|x_u - x_v\|_2 = \Theta(\sqrt{\delta})$ for a typical edge (u, v) in the graph. In this step, we show that in the case of small set expanders, this cannot be the case. To be more precise, we show that such case can be avoided if we are willing to pay constant factors depending on the expansion parameters of the graph.

Towards this end, we show that given a low-dimension solution $(x_v)_{v \in V}$ from the previous step, one may produce a solution $(y_v)_{v \in V}$ with similar value, such that the set of vertices V may be divided into two sets, Z_1, Z_2 , where vectors corresponding to vertices in Z_2 occur in the union of at most $1/\varepsilon$ balls with small ℓ_2 radius, and for any edge (u, v) that has at least one of its endpoints in Z_1 must either have 0 or long ℓ_2 distance, i.e. $\|y_v - y_u\|_2 = 0$ or $\|y_v - y_u\|_2 \geq \varepsilon$. Thus, in a sense the “bad structure” that may occur for the hyperplane rounding procedure in the Goemans-Williamson algorithm can only occur inside Z_2 .

- 4. Rounding:** Finally, we show an improved analysis of the standard hyperplane rounding when performed on a collection of vectors with the structure from the previous step. Intuitively, as Z_2 is small, we can show via a union bound that with constant probability, a random hyperplane doesn’t cut any of the small balls in which vectors from Z_2 occur. Then, we sample a hyperplane conditioned on this event happening and show that a square loss does not occur for the rest of the edges (as the distance over these edges already has a “Boolean-type” behaviour of either being 0, or far from 0).

1.2.2 Proof of Theorem 4

To prove Theorem 4, we view the second eigenvector of G , call it x as an embedding of the graph into \mathbb{R} . Recall that the square-root loss in the standard proof of Cheeger’s inequality comes from an application of Cauchy-Schwarz, which bounds $\sum_{(u,v) \in E} |x_v| |x_u - x_v|$ by the square root of spectral gap of G . We wish to circumvent this loss, and for that we observe that if the ratios of entries in x were either 1 or bounded away from 1, i.e. outside the interval $[1 - \varepsilon, 1 + \varepsilon]$, then one may indeed proceed as

$$\begin{aligned} \sum_{(u,v) \in E} |x_v| |x_u - x_v| &= \sum_{(u,v) \in E} |x_v|^2 |x_u/x_v - 1| \leq \sum_{(u,v) \in E} |x_v|^2 \frac{|x_u/x_v - 1|^2}{\varepsilon} \\ &= \frac{1}{\varepsilon} \sum_{(u,v) \in E} |x_u - x_v|^2. \end{aligned}$$

Thus, our goal is to preprocess x so that we obtain this property, while at the same time not increasing $\sum_{(u,v) \in E} |x_u - x_v|^2$ by too much. Indeed, we show that a soft rounding strategy inspired by the above algorithms can be used in this case. We partition \mathbb{R} “dyadically” into intervals of the form $[(1 + \varepsilon)^i, (1 + \varepsilon)^{i+1})$. We show that inside each one of these intervals, there is a way to round each entry of x to one of the endpoints without incurring too much of a loss in $\sum_{(u,v) \in E} |x_u - x_v|^2$.

1.2.3 Proof of Theorem 6

The proof of Theorem 6 follows the information theoretic approach to parallel repetition [20, 19, 10, 4]. Our proof follows the outline of [4], except that in the rounding phase in that result – which uses Pinkser’s inequality to transform a KL-divergence closeness between

distributions to a statistical distance closeness guarantee between the same distributions – we perform a preprocessing step. Indeed, we show that one may appeal to our improved version of Cheeger’s inequality to change the distributions so that the KL-divergence between them does not change too much and yet the square root loss in Pinsker’s inequality does not occur. From there, the rest of our proof follows the outline of [4].

2 Solving Max-Cut on SSE’s: Proof of Theorem 3

In this section, we prove Theorem 3 following the outline given in the introduction.

2.1 The Goemans-Williamson Semi-definite Program Relaxation

Below is the standard semi-definite program relaxation of Max-Cut. Instead of thinking of the graph G , it will be more convenient for us to work with the graph $G' = (V', E')$, wherein each vertex has two copies, $V' = V \times \{-1, 1\}$, and the edges are $E' = \{(v, b), (u, b') \mid (v, u) \in E, b \neq b'\}$. We have the following claim.

▷ **Claim 7.** The following two assertions hold:

1. If G has a cut of fractional size $1 - \delta$, then G' has a cut of fractional size at least $1 - \delta$.
2. If G is an (ε, γ) small set expander, then G' is a $(\varepsilon/2, \gamma/2)$ small set expander.

Proof. The first item is obvious. For the second item, suppose $S' \subseteq V'$ is a set of fractional size at most $\varepsilon/2$, and let $S = \{v \mid \exists b \in \{-1, 1\}, (v, b) \in S'\}$. Then S has fractional size at most ε , and we may write

$$\Phi(S') \geq \Pr_{(v,b) \in S', (u,b') \text{ neighbour}} [(u, b') \notin S'] \geq \Pr_{(v,b) \in S', (u,b') \text{ neighbour}} [u \notin S].$$

Note that sampling $(v, b) \in S'$, the distribution of v may not be uniform over S , but the probability of each $v \in S$ is either $1/|S'|$ or $2/|S'|$, hence

$$\Pr_{(v,b) \in S', (u,b') \text{ neighbour}} [u \notin S] \geq \frac{1}{2} \Pr_{v \in S, u \text{ neighbour}} [u \notin S] = \frac{1}{2} \Phi(S) \geq \frac{\gamma}{2}. \quad \blacktriangleleft$$

We thus write the program below, which is the standard semi-definite program relaxation for Max-Cut for the graph G' :

$$\begin{aligned} \min \quad & \frac{1}{|E'|} \sum_{((u,b),(v,b')) \in E'} \|x_{u,b} - x_{v,b'}\|_2^2 \\ \text{subject to} \quad & \|x_{v,b}\|_2 = 1 & \forall (v,b) \in V', \\ & x_{(v,b)} \in \mathbb{R}^{2n} & \forall (v,b) \in V', \\ & x_{(v,-b)} = -x_{(v,b)} & \forall (v,b) \in V'. \end{aligned}$$

It is clear that if G' contains a cut of fractional size at least $1 - \delta$, then there a solution for the above program with value at most δ . Using the ellipsoid algorithm, we may efficiently find a collection of vectors $(x_v)_{v \in V}$ achieving value at most $\delta + \xi = \delta'$ where ξ decays with the runtime of the algorithm; we take $\xi = \delta$ and fix such solution henceforth.

2.2 Reducing the Dimension of the Semi-Definite Program Solution

The goal of this section is to prove the following lemma.

► **Lemma 8.** *There exists an absolute constant $C > 0$ such that the following holds. There is an efficient, randomized procedure that given a solution $(x_v)_{v \in V'}$ to SDP program to Max-Cut with value at most δ' consisting of vectors from \mathbb{R}^n , outputs a solution $(z_v)_{v \in V'}$ consisting of vectors from \mathbb{R}^3 , whose value is at most $C\delta'$.*

Proof. Let $g_1, g_2, g_3 \sim N(0, I_n)$ be independent multi-dimensional Gaussians. For each $v \in V$, define

$$y_v = (\langle x_v, g_1 \rangle, \langle x_v, g_2 \rangle, \langle x_v, g_3 \rangle), \quad z_v = \frac{y_v}{\|y_v\|_2}.$$

Clearly, $(z_v)_{v \in V'}$ is a solution to the SDP program, and we next analyze the objective value that it gets. Fix $u, v \in V'$, and let $\delta_{u,v} = \|x_u - x_v\|_2^2$. We will show that, $\mathbb{E}_{g_1, g_2, g_3} [\|z_v - z_u\|_2^2] = O(\delta_{u,v})$ from which the claim follows by linearity of expectation over all the edges of the graph.

Denote $z = \frac{y_u}{\|y_u\|_2}$, then

$$\begin{aligned} & \mathbb{E}_{g_1, g_2, g_3} [\|z_v - z_u\|_2^2] \\ & \leq 2 \mathbb{E}_{g_1, g_2, g_3} [\|z_v - z_u\|_2^2 \mathbf{1}_{\|y_v\|_2 \leq \|y_u\|_2}] \\ & \leq 4 \underbrace{\mathbb{E}_{g_1, g_2, g_3} [\|z_v - z\|_2^2 \mathbf{1}_{\|y_v\|_2 \leq \|y_u\|_2}]}_{(I)} + 4 \underbrace{\mathbb{E}_{g_1, g_2, g_3} [\|z_u - z\|_2^2 \mathbf{1}_{\|y_v\|_2 \leq \|y_u\|_2}]}_{(II)}, \end{aligned}$$

and we upper bound each expectation separately, each by $O(\delta_{u,v})$, thereby finishing the proof.

► **Claim 9.** $(II) \leq O(\delta_{u,v})$.

Proof. We write

$$\begin{aligned} \mathbb{E}_{g_1, g_2, g_3} [\|z_u - z\|_2^2 \mathbf{1}_{\|y_v\|_2 \leq \|y_u\|_2}] &= \mathbb{E}_{g_1, g_2, g_3} \left[\frac{\|y_u - y_v\|_2^2}{\|y_u\|_2^2} \mathbf{1}_{\|y_v\|_2 \leq \|y_u\|_2} \right] \\ &= \mathbb{E}_{g_1, g_2, g_3} \left[\frac{\|y_u - y_v\|_2^2}{\|y_u\|_2^2} \mathbf{1}_{\|y_v\|_2 \leq \|y_u\|_2} \mathbf{1}_{\|y_u\|_2 < \delta_{u,v}} \right] \\ &+ \mathbb{E}_{g_1, g_2, g_3} \left[\frac{\|y_u - y_v\|_2^2}{\|y_u\|_2^2} \mathbf{1}_{\|y_v\|_2 \leq \|y_u\|_2} \mathbf{1}_{\|y_u\|_2 \geq \delta_{u,v}} \right]. \end{aligned}$$

For the first expectation, as $\|y_v\|_2 \leq \|y_u\|_2$ we have $\frac{\|y_u - y_v\|_2^2}{\|y_u\|_2^2} \leq 4$, so the first expectation is bounded by

$$4 \mathbb{E}_{g_1, g_2, g_3} [\mathbf{1}_{\|y_v\|_2 \leq \|y_u\|_2} \mathbf{1}_{\|y_u\|_2 < \delta_{u,v}}] \leq 4 \mathbb{E}_{g_1, g_2, g_3} [\mathbf{1}_{\|y_u\|_2 < \delta_{u,v}}] \leq 4 \Pr_{G \sim N(0,1)} [|G| \leq \delta_{u,v}],$$

which is $O(\delta_{u,v})$. For the second expectation, we upper bound it by

$$\sum_{k=0}^{\infty} \mathbb{E}_{g_1, g_2, g_3} \left[\frac{\|y_u - y_v\|_2^2}{2^{2k} \delta_{u,v}^2} \mathbf{1}_{2^k \delta_{u,v} \leq \|y_u\|_2 < 2^{k+1} \delta_{u,v}} \right].$$

We bound the tail, that is the sum over $k \geq \log(1/\delta_{u,v})$, by

$$\begin{aligned} & \sum_{k=\lceil \log(1/\delta_{u,v}) \rceil}^{\infty} \mathbb{E}_{g_1, g_2, g_3} [\|y_u - y_v\|_2^2 \mathbf{1}_{2^k \delta_{u,v} \leq \|y_u\|_2 < 2^{k+1} \delta_{u,v}}] \leq \mathbb{E}_{g_1, g_2, g_3} [\|y_u - y_v\|_2^2] \\ & = O(\|x_u - x_v\|_2^2) \\ & = O(\delta_{u,v}). \end{aligned}$$

26:8 Rounding via Low Dimensional Embeddings

As for the sum on $k = 0, \dots, \lceil \log(1/\delta_{u,v}) \rceil - 1$, take $p \in \mathbb{N}$ sufficiently large ($p = 4$ will do), and let p' be the Hölder conjugate of p . Then

$$\begin{aligned} & \mathbb{E}_{g_1, g_2, g_3} \left[\frac{\|y_u - y_v\|_2^2}{2^{2k} \delta_{u,v}^2} \mathbf{1}_{2^k \delta_{u,v} \leq \|y_u\|_2 < 2^{k+1} \delta_{u,v}} \right] \\ & \leq \frac{1}{2^{2k} \delta_{u,v}^2} \left(\mathbb{E}_{g_1, g_2, g_3} \left[\|y_u - y_v\|_2^{2p} \right] \right)^{1/p} \left(\mathbb{E}_{g_1, g_2, g_3} \left[\mathbf{1}_{2^k \delta_{u,v} \leq \|y_u\|_2 < 2^{k+1} \delta_{u,v}} \right] \right)^{1/p'}. \end{aligned} \quad (1)$$

First, we have

$$\begin{aligned} \mathbb{E}_{g_1, g_2, g_3} \left[\|y_u - y_v\|_2^{2p} \right] &= \mathbb{E}_{g_1, g_2, g_3} \left[(G_1^2 + G_2^2 + G_3^2)^p \right] \leq 3^{p-1} \mathbb{E}_{g_1, g_2, g_3} \left[G_1^{2p} + G_2^{2p} + G_3^{2p} \right] \\ &= 3^p \mathbb{E}_{g_1, g_2, g_3} \left[G_1^{2p} \right], \end{aligned}$$

where $G_1 = \langle g_1, x_u - x_v \rangle$, $G_2 = \langle g_2, x_u - x_v \rangle$, $G_3 = \langle g_3, x_u - x_v \rangle$. As G_1 is a Gaussian random variable with mean 0 and variance $\delta_{u,v}$, we know that

$$\mathbb{E}_{g_1, g_2, g_3} \left[G_1^{2p} \right] = O_p(\delta_{u,v}^p).$$

Second,

$$\mathbb{E}_{g_1, g_2, g_3} \left[\mathbf{1}_{2^k \delta_{u,v} \leq \|y_u\|_2 < 2^{k+1} \delta_{u,v}} \right] \leq \Pr_{g_1, g_2, g_3} \left[|\langle g_1, x_u \rangle|, |\langle g_2, x_u \rangle|, |\langle g_3, x_u \rangle| \leq 2^{k+1} \delta_{u,v} \right],$$

and as g_1, g_2, g_3 are independent, the last probability is at most $2^{3(k+1)} \delta_{u,v}^3$.

Plugging everything into (1) we get that the sum on $k = 0, \dots, \lceil \log(1/\delta_{u,v}) \rceil - 1$ is at most

$$\sum_{k=0}^{\lceil \log(1/\delta_{u,v}) \rceil - 1} \frac{1}{2^{2k} \delta_{u,v}^2} O_p(\delta_{u,v}) \left(2^{3(k+1)} \delta_{u,v}^3 \right)^{1/p'} = \delta^{3/p'-1} O_p(1) \sum_{k=0}^{\lceil \log(1/\delta_{u,v}) \rceil - 1} 2^{(3/p'-2)k}.$$

For sufficiently large p , $3/p' - 2 > 0$, and we may bound the last sum by

$$O(2^{(3/p'-2)\lceil \log(1/\delta_{u,v}) \rceil}) = \delta_{u,v}^{2-3/p'},$$

and plugging that in yields the bound $O_p(\delta_{u,v})$. \triangleleft

\triangleright **Claim 10.** $(I) \leq O(\delta_{u,v})$.

Proof. We write

$$\begin{aligned} \mathbb{E}_{g_1, g_2, g_3} \left[\|z_v - z\|_2^2 \mathbf{1}_{\|y_v\|_2 \leq \|y_u\|_2} \right] &= \mathbb{E}_{g_1, g_2, g_3} \left[\left\| y_v \left(\frac{1}{\|y_v\|_2} - \frac{1}{\|y_u\|_2} \right) \right\|_2^2 \mathbf{1}_{\|y_v\|_2 \leq \|y_u\|_2} \right] \\ &= \mathbb{E}_{g_1, g_2, g_3} \left[\left(1 - \frac{\|y_v\|_2}{\|y_u\|_2} \right)^2 \mathbf{1}_{\|y_v\|_2 \leq \|y_u\|_2} \right] \\ &= \mathbb{E}_{g_1, g_2, g_3} \left[\left(\frac{\|y_u\|_2 - \|y_v\|_2}{\|y_u\|_2} \right)^2 \mathbf{1}_{\|y_v\|_2 \leq \|y_u\|_2} \right] \\ &\leq \mathbb{E}_{g_1, g_2, g_3} \left[\left(\frac{\|y_u - y_v\|_2}{\|y_u\|_2} \right)^2 \mathbf{1}_{\|y_v\|_2 \leq \|y_u\|_2} \right] \\ &= (I), \end{aligned}$$

hence the claim follows from the prior claim. We used the fact that $|||a| - |b|| \leq \|a - b\|$. \triangleleft

\blacktriangleleft

2.3 The Soft Rounding: the Cubicular Rounding

In this section, we pick off where the last section ended with a solution vector collection $(z_v)_{v \in V'}$ consisting of vectors from \mathbb{R}^3 that has value δ'' where $\delta'' \leq O(\delta') = O(\delta)$. Our goal is to construct a new vector solution $\{z''_v\}_{v \in V'}$ whose value is at most δ''' for $\delta''' = O(\delta''/\gamma)$ that satisfies the following property. The set of vectors $\{z''_v\}_{v \in V'}$ can be partitioned into two collections, Z_1 and Z_2 , such that $|Z_2| \leq 1/\varepsilon$, and for any edge (u, v) such that $z''_u \neq z''_v$ and $z''_u \in Z_1$ it holds that $\|z''_u - z''_v\|_2 \geq \Omega(\varepsilon)$. In Section 2.4 we show how to use this property to perform a more efficient analysis of the standard hyperplane rounding.

2.3.1 Partitioning the 3-dimensional sphere into cubes

Consider the 3 dimensional sphere $S^2 = \{z \in \mathbb{R}^3 \mid \|z\|_2 = 1\}$. We will use a simple construction of triangulation of the sphere (geodesic polyhedron), that we describe next. Starting with a standard cube in \mathbb{R}^3 , denoted it by \mathcal{T}_6 (that has 6 faces and 8 vertices), we may partition each face of it (which is a square) into 4 squares in the natural way (i.e. by adding the lines that connect the midpoints of opposing sides). We then project the 4 new points that were generated by each face into S^2 , and get a more refined regular polygon with square faces, with 4 times as many faces; denote it by \mathcal{T}_{24} . We will perform this operation several times to reach a sufficiently refined shape; we need the perimeter of each face to be small compared to ε , say at most $\frac{\varepsilon}{K}$ for sufficiently large absolute constant K . We note that at each iteration, the perimeter of the faces shrinks by factor 2, hence we may pick a power of 4, $t \geq \frac{K'}{\varepsilon^2}$ for K' depending only on K , and have that each phase of \mathcal{T}_t has perimeter at most $\frac{\varepsilon}{K}$. We summarize this discussion with the following standard fact.

► **Fact 11.** *For all $K > 0$ there is $K' = K'(K)$, such that for all $\varepsilon > 0$, taking t the smallest number of the form $6 \cdot 4^m$ that is larger than K'/ε^2 , there is a regular polygon $\mathcal{T}_t \subseteq \mathbb{R}^3$ that can be constructed in $O_\varepsilon(1)$ time such that:*

1. Each face of \mathcal{T}_t is a square.
2. Central symmetry: P is a face of \mathcal{T}_t if and only if $-P$ is a face of \mathcal{T}_t .
3. \mathcal{T}_t has t faces and $\frac{4}{3}t$ vertices.
4. The perimeter of each face is at most ε/K .

2.3.2 Division of the sphere

Using the triangulation \mathcal{T}_t from Fact 11, we may define a division of S^2 by projecting each face of \mathcal{T}_t into the sphere. This way, we get a partition $P_1 \cup \dots \cup P_t$ of S^2 . We also note that the perimeter of each of P_i is the same, and is at most order of the perimeter of a face from \mathcal{T}_t , i.e. at most $O(\varepsilon/K)$.

For each i define the neighbourhood of P_i by:

$$P_i^\uparrow = \{u \in S^2 \mid \exists p \in P_i, \|u - p\|_2 \leq \varepsilon\}.$$

For each subset of $A \subseteq S^2$, we define the mass of A as $\text{mass}(A) = |\{v \in V \mid z_v \in A\}|$.

► **Claim 12.** $\sum_{i=1}^t \text{mass}(P_i^\uparrow) = O(n)$.

Proof. Consider the bipartite graph $\tilde{G} = ([t] \times [t], \tilde{E})$ wherein we connect i to j if $P_i^\uparrow \cap P_j \neq \emptyset$. We note that the degree of each vertex in \tilde{G} is $O(1)$, and $P_i^\uparrow \subseteq \cup_{j:(i,j) \in \tilde{E}} P_j$. Thus,

$$\sum_{i=1}^t \text{mass}(P_i^\uparrow) \leq \sum_{i=1}^t \sum_{j:(i,j) \in \tilde{E}} \text{mass}(P_j) = \sum_{j=1}^t \text{deg}(j) \text{mass}(P_j) = O\left(\sum_{j=1}^t \text{mass}(P_j)\right) = O(n). \quad \blacktriangleleft$$

It follows that for all but at most $2/\varepsilon$ of the i 's we have that $\text{mass}(P_i \uparrow) \leq \varepsilon n/2$. We refer to such i 's as light parts of the partition. In the following section we show that one may change a vector z_v such that v is in P_i for a light i , to be a vector in the boundary of P_i , and only incur a constant factor loss in the objective value, effectively reducing these vectors to be 2-dimensional.

2.3.3 Rounding to the Boundary

We now modify the solution $(z_u)_{u \in V}$ to $(z'_u)_{u \in V}$ as follows. For u , consider the i such that $z_u \in P_i$. If P_i is heavy we set $z'_u = z_u$, and otherwise we modify it as follows. Let p_i be the center of P_i , and consider the segment from p_i to z_u , and in particular its two intersection points with the boundary of P_i . We let z'_u be the intersection point closer to z_u among these 2.

► **Lemma 13.** *The collection of vectors z'_u defined about is a solution to the Max-Cut program, and*

$$\frac{1}{|E'|} \sum_{(u,v) \in E'} \|z'_u - z'_v\|_2^2 \leq \delta'''$$

for $\delta''' = O(\delta''/\gamma^3)$.

Proof. First, it is clear that all of the z_u 's are norm 1 vectors, and also that they respect the conditions $z_{u,-b} = -z_{u,b}$ as the part of a vector z and its negation have the same mass (as if P_i is the part of z , then $-P_i$ is the part of z).

Fix a light P_i and denote $\text{cost}_i = \sum_{\substack{u: z_u \in P_i \\ (u,v) \in E'}} \|z_u - z_v\|_2^2$. We argue that cost_i grows multiplicatively by at most factor $O(1/\gamma^3)$ due to the change of the procedure on the vectors inside P_i . This finishes the proof, as the only other effect of the procedure we need to make note of is that for each edge $e = (u,v) \in E'$ such that $u \in P_i, v \in P_j, \|z'_u - z'_v\|_2 = O(\|z_u - z_v\|_2)$, so each edge in cost_j increases by a constant factor at most once by the procedure on other P_i 's.

We now argue that cost_i grows multiplicatively by at most factor $O(1/\gamma^3)$, and assume towards contradiction this is not the case. For simplicity of notation, we drop the subscript i as we will only be concerned with P_i from now on. Let $B_0 = \{u \mid z_u \in P\}$, and let $W > 0$ be a large absolute constant to be determined. As a result of our soft rounding, the contribution of edges whose endpoints are both outside B_0 grows by at most factor $O(1)$, and the contribution of edges that have at least one of their endpoints in B_0 grows additively by at most $2d|B_0|\varepsilon^2$. Hence, if $\text{cost} \geq \frac{\gamma^3 \varepsilon^2}{50W^2} d|B_0|$ we would have been done, so we assume from now on that $\text{cost} < \frac{\gamma^3 \varepsilon^2}{50W^2} d|B_0|$.

Consider $\ell_1 = \frac{\gamma}{W}\varepsilon$ and let $B_1 = \{u \mid \exists p \in B_0, \|z_u - p\|_2 \leq \ell_1\}$. Also, define $\ell_2 = 2\ell_1$ and set $B_2 = \{u \mid \exists p \in B_0, \|z_u - p\|_2 \leq \ell_2\}$. As $B_1 \subseteq P^\uparrow$, $\text{mass}(B_1) \leq \varepsilon n$ so by small set expansion at least $\gamma d|B_1|$ of the edges touching B_1 escape it. If $|B_2 \setminus B_1| \leq \frac{\gamma}{2}|B_1|$, we would get that at least $\frac{\gamma}{2}d|B_1|$ of the edges touching B_1 go outside B_2 , hence

$$\text{cost} \geq \frac{\gamma}{2}d|B_1|(\ell_2 - \ell_1)^2 = \frac{\gamma \ell_1^2}{2}d|B_1| > \frac{\gamma^3 \varepsilon^2}{50W^2} d|B_0|$$

and contradiction. Hence $|B_2| \geq (1 + \gamma/2)|B_1|$. Next, define $\ell_3 = \ell_2 + \sqrt{\frac{|B_0|}{|B_2|}}\ell_1$ as well as $B_3 = \{u \mid \exists p \in B_0, \|z_u - p\|_2 \leq \ell_3\}$. Then given $\ell_2 \leq \varepsilon$, we have that $\text{mass}(B_2) \leq \varepsilon$ and similar argument to before gives that $|B_3| \geq (1 + \gamma/2)|B_2|$. Indeed, otherwise we get that at least $\frac{\gamma}{2}d|B_2|$ of the edges touching B_2 go outside B_3 , and then

$$\text{cost} \geq \frac{\gamma}{2} d |B_2| (\ell_3 - \ell_2)^2 = \frac{\gamma \ell_1^2}{2} d |B_0|,$$

and contradiction. This way, we continue iteratively, and once ℓ_k and B_k have been defined we take

$$\ell_{k+1} = \ell_k + \sqrt{\frac{|B_0|}{|B_k|}} \ell_1, \quad B_{k+1} = \{u \mid \exists p \in B_0, \|z_u - p\| \leq \ell_{k+1}\}$$

and provided that $\ell_k \leq \varepsilon$ we conclude that $|B_{k+1}| \geq (1 + \gamma/2) |B_k|$. Note that provided we have gotten to the $k + 1$ step, we have

$$\ell_k \leq \sum_{r=0}^{\infty} (1 + \gamma/2)^{-r/2} \ell_1 \leq O(\ell_1/\gamma) = O(\varepsilon/W) \leq \varepsilon$$

for sufficiently large absolute constant W , so we can ensure that the argument goes through indefinitely. This is a contradiction, as at some point the size of B_k would exceed εn . ◀

2.3.4 Rounding to the Corners

Next, we modify the vector-valued solution $(z'_u)_{u \in V}$ to a different vector-valued solution $(z''_u)_{u \in V}$ wherein (roughly speaking) vectors belonging to vertices on light P_i 's can only be in the corners of \mathcal{T}_i .

Towards this end, consider the boundary of the partition \mathcal{P} and a particular arc in it. This arc has 2 neighbouring cells, say P_i and P_j . If P_i and P_j are light, we look at that arc, mark its middle point $p_{i,j}$. We then round the each vector z'_u on that arc to the closer of the two endpoints of the arc to it (namely, according to the side of z'_u with respect to $p_{i,j}$).

► **Lemma 14.** *The collection of vectors z''_u defined about is a solution to the Max-Cut program, and*

$$\frac{1}{|E'|} \sum_{(u,v) \in E'} \|z''_u - z''_v\|_2^2 \leq \delta'''$$

for $\delta''' = O(\delta'''/\gamma^3)$.

Proof. Consider a single arc on which we perform the operation, denote it by \mathcal{L} , and let

$$\text{cost}_{\mathcal{L}} = \sum_{(u,v) \in E: z'_u \in \mathcal{L}} \|z'_u - z'_v\|_2^2.$$

We argue that as a result of the above operation, $\text{cost}_{\mathcal{L}}$ increases by factor at most $O(1/\gamma^3)$. This quickly finishes the proof since the only other effect of the operation we need to take note of, is the effect of other arcs on $\text{cost}_{\mathcal{L}}$. For that, it suffices to observe that for each edge $(u,v) \in E$ such that $u \in \mathcal{L}$ and $v \in \mathcal{L}'$, the quantity $\|z'_u - z'_v\|_2^2$ increases by factor at most $O(1)$ due to the operation on \mathcal{L}' .

Let L be the length of the arc \mathcal{L} (noting it is of the order of ε), p be its midpoint and let $W > 0$ be a large absolute constant. Let $\ell_0 = \frac{\gamma L}{W}$, and consider

$$B_0 = \{u \mid z'_u \in \mathcal{L}, \|z'_u - p\| \leq \ell_0\}.$$

26:12 Rounding via Low Dimensional Embeddings

As a result of our operation, $\text{cost}_{\mathcal{L}}$ changes as follows: the contribution of edges that have both endpoints outside B_0 grows by factor $O(1)$ at most, whereas the contribution of edges that have an endpoint inside B_0 increases additively by at most $L^2 d |B_0|$. Hence, if $\text{cost}_{\mathcal{L}} \geq \frac{\gamma^3 L^2}{50W^2} d |B_0|$ we are done, so assume otherwise.

Let $\ell_1 = 2\ell_0$, and note that $m(B_0) \leq m(\mathcal{L}) \leq \varepsilon n$, so by small set expansion at least $\gamma d |B_0|$ of the edges touching B_0 go outside it. At most $d |B_1 \setminus B_0|$ of these edges go to B_1 , hence if $|B_1 \setminus B_0| \leq \gamma/2 |B_0|$ we get that

$$\text{cost}_{\mathcal{L}} \geq \frac{1}{2} \gamma d |B_0| (\ell_1 - \ell_0)^2 > \frac{\gamma^3 L^2}{50W^2} d |B_0|,$$

and contradiction. Thus, $|B_1 \setminus B_0| > \gamma/2 |B_0|$, and so $|B_1| \geq (1 + \gamma/2) |B_0|$. Continuing this way, iteratively define

$$\ell_{j+1} = \ell_j + \sqrt{\frac{|B_0|}{|B_j|}} \ell_1, \quad B_{j+1} = \{u \mid z'_u \in \mathcal{L}, \|z'_u - p\| \leq \ell_{j+1}\}.$$

If $\ell_{j+1} \leq L$ we can apply the small set expansion argument again to argue that $|B_{j+1}| \geq (1 + \gamma/2) |B_j|$. Thus, provided that $\ell_j \leq L$ we may continue with the argument indefinitely and keep on increasing B_j , which is clearly impossible, giving us the desired contradiction. We note that if W is large enough, then indeed we always have $\ell_j \leq L$, as we may bound

$$\ell_j \leq \sum_{r=0}^j \sqrt{\frac{|B_0|}{|B_r|}} \ell_1 \leq \sum_{r=0}^j (1 + \gamma/2)^{-r} \ell_1 = O(\ell_1 / \gamma) = O(L/W) < L$$

for large enough W . ◀

2.4 The Integral Rounding Procedure

We next describe the integral rounding procedure using the vector solution z''_u from Lemma 14. Let Heavy be the collection of all parts P_i that are not light, denote by Heavy' be the collection of all P_j 's adjacent to a face in Heavy , and let \mathcal{H} be the union of boundaries of faces in $\text{Heavy} \cup \text{Heavy}'$.

Consider choice of a hyperplane $H \subseteq \mathbb{R}^3$ uniformly at random.

▷ **Claim 15.** There exists $K > 0$, such that in the above set-up,

$$\Pr_H [\exists \mathcal{L} \in \mathcal{H} \text{ such that } H \cap \mathcal{L} \neq \emptyset] \leq 1/2.$$

Proof. By Claim 12, the number of heavy P_i 's is $O(1/\varepsilon)$, and as each one of them has $O(1)$ many arcs in its boundary, we deduce that $|\text{Heavy} \cup \text{Heavy}'| = O(1/\varepsilon)$. As each face has $O(1)$ arcs in the boundary, we conclude that $|\mathcal{H}| = O(1/\varepsilon)$. The length of each $\mathcal{L} \in \mathcal{H}$ is at most ε/K , so we get that the probability that H intersects a given arc is at most $O(\varepsilon/K)$. Thus, by the Union bound the probability that H intersects some arc in \mathcal{H} is at most $O(1/K) \leq 1/2$, provided that K is large enough. ◁

We now sample H conditioned on it not intersecting any face or arc from Heavy . We claim that, for any edge (u, v) such that H intersects the line between z''_u and z''_v , it must be the case that $\|z''_u - z''_v\|_2 \geq \Omega(\varepsilon)$. There are a few cases to consider:

1. Either z''_u or z''_v is not a corner of \mathcal{P} , say z''_u . In this case, z_u is in a face P_j that is either heavy or adjacent to a heavy P_i . In either case, to cross the line between z''_u and z''_v the hyperplane H must cross the boundary of P_j , but this does occur due to the conditioning on H .

2. Else, z''_u and z''_v are two corners of \mathcal{P} , so either they are the same corner, or at least $\Omega(\varepsilon)$ apart in ℓ_2 -distance.

We consider the cut $S = \{v \in V \mid \langle z''_{v,1}, x_H \rangle \geq 0\}$, and prove in the following lemma that the expected size of the cut is at least $(1 - O(\frac{\delta''''}{\varepsilon})) |E|$. Given such guarantee, standard techniques show that such cut may be found in polynomial time by a randomized algorithm, thereby finishing the proof of Theorem 3.

► **Lemma 16.** $\mathbb{E}_H [\text{Cut}(S) \mid H \text{ does not hit } \mathcal{H}] \geq (1 - O(\frac{\delta''''}{\varepsilon})) |E|$.

Proof. The expected size of the cut is

$$\sum_{v,u} \mathbb{E}_H \left[\mathbb{1}_{H \text{ separates } z''_{u,1} \text{ and } z''_{v,1}} \mid H \text{ does not hit } \mathcal{H} \right],$$

so the deficit in the cut size (i.e. $|E|$ minus its size) has expectation

$$\sum_{v,u} \mathbb{E}_H \left[\mathbb{1}_{H \text{ separates } z''_{u,1} \text{ and } z''_{v,-1}} \mid H \text{ does not hit } \mathcal{H} \right].$$

For H to separate $z''_{u,1}$ and $z''_{v,-1}$, by the discussion preceding the lemma, we must have $\|z''_{u,1} - z''_{v,-1}\|_2 \geq \Omega(\varepsilon)$, so we get that the last sum is at most

$$2 \sum_{v,u} \mathbb{E}_H \left[\mathbb{1}_{H \text{ separates } z''_{u,1} \text{ and } z''_{v,-1}} \right] \mathbb{1}_{\|z''_{u,1} - z''_{v,-1}\|_2 \geq \Omega(\varepsilon)}.$$

Here, we used the fact that the probability that H does not hit \mathcal{H} is at least $1/2$. Letting $\theta(z_{u,1}, z_{v,-1})$ be the angle between the two vectors, we have

$$\mathbb{E}_H \left[\mathbb{1}_{H \text{ separates } z''_{u,1} \text{ and } z''_{v,-1}} \right] = \frac{1}{\pi} \theta(z''_{u,1}, z''_{v,-1}) = O(\|z''_{u,1} - z''_{v,-1}\|_2),$$

hence the expected deficit in the size of the cut defined by S is at most

$$O\left(\sum_{v,u} \|z''_{u,1} - z''_{v,-1}\|_2 \mathbb{1}_{\|z_{u,1} - z_{v,-1}\|_2 \geq \Omega(\varepsilon)}\right) = \frac{1}{\varepsilon} O\left(\sum_{v,u} \|z''_{u,1} - z''_{v,-1}\|_2^2\right) = O\left(\frac{\delta''''}{\varepsilon} |E|\right). \blacktriangleleft$$

3 Spectral Partitioning

In this section, we use our techniques to establish stronger algorithmic spectral partitioning results for the second largest eigenvalue of a graph G , proving Theorem 4. We also state an analogous result for the smallest eigenvalue of a small set expander G (but omit the mostly-identical proof).

3.1 Cheeger's Inequality on SSE's

As explained in the introduction, to avoid the square root loss in Cheeger's inequality, we must make sure that we perform the rounding phase of it on vectors in which the ratio between any two coordinates is either 1, or is bounded away from 1. Below, we formalize this notion via a notion we refer to as ε -quantization.

► **Definition 17.** Let $\varepsilon > 0$. We say a vector $x \in \mathbb{R}^n$ is ε -quantized if for each $i, j \in [n]$:

1. $x(i) = 0$ or $x(j) = 0$;
2. otherwise, $|x(i)| = |x(j)|$, $|x(i)| \geq (1 + \varepsilon) |x(j)|$, or $|x(i)| \leq (1 - \varepsilon) |x(j)|$.

26:14 Rounding via Low Dimensional Embeddings

Next, we prove that for an SSE graph G and a vector x such that $\mathbb{E}_{(u,v) \sim E} [(x(u) - x(v))^2]$ is small, one can transform x to a vector y that is quantized while not increasing the square distances by much.

► **Lemma 18.** *Suppose that $G = (V, E)$ is a d -regular, (ε, γ) small set expander. If $x \in \mathbb{R}^n$ is a vector such that $\mathbb{E}_{(u,v) \sim E} [(x(u) - x(v))^2] \leq \delta$, then there is an ε -quantized vector $y \in \mathbb{R}^n$ such that:*

1. $\mathbb{E}_{(u,v) \sim E} [(y(u) - y(v))^2] \leq O\left(\frac{\delta}{\gamma^3 \varepsilon^2}\right)$.
2. For each $v \in V$, $y(v) = 0$ if and only if $x(v) = 0$.
3. For each $v \in V$, $(1 - \varepsilon)x(v) \leq y(v) \leq (1 + \varepsilon)x(v)$ if $x(v) \geq 0$, and $(1 + \varepsilon)x(v) \leq y(v) \leq (1 - \varepsilon)x(v)$ otherwise.

Proof. Let $V_+ = \{v \mid x(v) > 0\}$, $V_- = \{v \mid x(v) < 0\}$, and denote $M = \lceil \max_u |x(u)| \rceil$. We partition V_+ and V_- as

$$V_+ = \bigcup_{j=0}^{\infty} V_{+,j}, \quad V_- = \bigcup_{j=0}^{\infty} V_{-,j},$$

where

$$V_{+,j} = \{v \in V_+ \mid M(1 + \varepsilon)^{-j-1} \leq x(v) < M(1 + \varepsilon)^{-j}\},$$

$$V_{-,j} = \{v \in V_- \mid -M(1 + \varepsilon)^{-j} \leq x(v) < -M(1 + \varepsilon)^{-j-1}\}.$$

The construction of the vector y . To construct the vector y , for each vertex $V \in V$ we find the part $V_{+,j}$ (if $x(v)$ is positive) or $V_{-,j}$ (if $x(v)$ is negative) that v belongs to, and set the value of $y(v)$ to be one of the endpoints of the interval defining that part. Towards this end, for each j we will choose points $p_{j,+}$ and $p_{j,-}$ (in a way that we describe shortly) in the intervals $[M(1 + \varepsilon)^{-j-1}, M(1 + \varepsilon)^{-j}]$ and $[-M(1 + \varepsilon)^{-j}, -M(1 + \varepsilon)^{-j-1}]$ respectively, and use them to define the vector y that satisfies the assertion of the lemma. More specifically: we set $y(u) = 0$ if $x(u) = 0$; otherwise – without loss of generality say that $x(u) > 0$ – we take j such that $u \in V_{+,j}$, and define $y(u) = M(1 + \varepsilon)^{-j}$ if $x(u) > p_{j,+}$ and else $y(u) = M(1 + \varepsilon)^{-j-1}$.

Choosing the cut-off points $p_{j,+}$ and $p_{j,-}$. Next, we describe the choice of the points $p_{j,+}$ and $p_{j,-}$. As the argument is analogous in both cases, we focus on $p_{j,+}$. We define the mass of $V_{+,j}$ as the cardinality of it. We say $V_{+,j}$ is heavy if its mass is at least εn , and otherwise we say it is light. The choice of $p_{j,+}$ is done differently in light intervals and in heavy intervals, and below we elaborate on these two cases. To do so, we consider the cost associated with interval j ,

$$\text{cost}_j = \frac{1}{n} \sum_{u \in V_{+,j}, v \text{ neighbour}} (x(u) - x(v))^2,$$

and show that the cost of $V_{+,j}$ only increases by factor $O(1/\gamma^3)$ as a result of the operation on $V_{+,j}$.

Light intervals. In this case, we take $p_{j,+}$ to be the midpoint of its respective interval, that is

$$p_{j,+} = \frac{1}{2} (M(1 + \varepsilon)^{-j-1} + M(1 + \varepsilon)^{-j}).$$

We denote by L_j the length of the interval, that is

$$L_j = M(1 + \varepsilon)^{-j} - M(1 + \varepsilon)^{-j-1} = M\varepsilon(1 + \varepsilon)^{-j-1}.$$

Let cost_j be the contribution of points from $V_{+,j}$ to the average square distances i.e.

$$\text{cost}_j = \frac{1}{n} \sum_{u \in V_{+,j}, v \text{ neighbour}} (x(u) - x(v))^2.$$

We show that doing the rounding at $p_{j,+}$ as described above only increases cost_j by factor $O\left(\frac{1}{\gamma^3}\right)$.

Let $\ell_1 = \frac{\gamma}{K}$, where K is an absolute constant to be determined later. Consider the interval $B_1 = [p_{j,+} - \ell_1 L_j, p_{j,+} + \ell_1 L_j]$. We will show that $\text{cost}_j > \frac{\gamma \ell_1^2 L_j^2}{100} d |B_1|$, and then noting that rounding at $p_{j,+}$ increases cost_j additively by at most $d |B_1| L_j^2$ it follows that the cost of $V_{+,j}$ increases by at most factor $O(1/\gamma^3)$ as a result of the rounding.

We now prove that $\text{cost}_j > \frac{\gamma \ell_1^2 L_j^2}{100} d |B_1|$, and to do so we assume towards contradiction the contrary inequality holds. Let $\ell_2 = 2\ell_1$, and set $B_2 = [p_{j,+} - \ell_2 L_j, p_{j,+} + \ell_2 L_j]$. We claim that $|B_2| \geq (1 + \frac{\gamma}{2}) |B_1|$. Indeed, otherwise we consider outgoing edges from B_1 , and note that as $|B_1| \leq |V_{j,+}| \leq \varepsilon n$, the small set expansion property implies that at least $d\gamma |B_1|$ of them escape outside B_1 , where d is the degree of G' . As $|B_2| \leq (1 + \frac{\gamma}{2}) |B_1|$, at most $\frac{d\gamma}{2} |B_1|$ of them can go to $B_2 \setminus B_1$, and hence at least $\frac{d\gamma}{2} |B_1|$ of them go outside $B_2 \setminus B_1$. Hence, we conclude that

$$\text{cost}_j \geq \frac{d\gamma}{2} |B_1| (\ell_2 - \ell_1)^2 L_j^2 = \frac{\gamma \ell_1^2}{2} L_j^2 d |B_1|,$$

and contradiction. We thus have that $|B_2| \geq (1 + \frac{\gamma}{2}) |B_1|$.

Next, we set $\ell_3 = \ell_2 + \sqrt{\frac{|B_1|}{|B_2|}} \ell_1$ and consider $B_3 = [p_{j,+} - \ell_3 L_j, p_{j,+} + \ell_3 L_j]$. Similarly to before, if $|B_3| \leq (1 + \frac{\gamma}{2}) |B_2|$ we get a contradiction to the upper bound on cost_j . Continuing this argument, we define iteratively $\ell_{r+1} = \ell_r + \sqrt{\frac{|B_1|}{|B_r|}} \ell_1$ and then $B_{r+1} = [p_{j,+} - \ell_{r+1} L_j, p_{j,+} + \ell_{r+1} L_j]$, and prove that $|B_{r+1}| \geq (1 + \frac{\gamma}{2}) |B_r|$ if $\ell_{r+1} \leq 1/10$. Note that

$$\ell_{j+1} \leq \ell_1 \sum_{k=0}^{\infty} \left(1 + \frac{\gamma}{2}\right)^{-k/2} = O(\ell_1/\gamma) = O(1/K) \leq 1/10$$

provided K is sufficiently large, so we way we keep on increasing $|B_r|$ which is clearly impossible.

This contradiction implies that $\text{cost}_j \geq \frac{\gamma \ell_1^2 L_j^2}{100} d |B_1|$, hence the cost of $V_{+,j}$ increases by at most factor $1/(\gamma \ell_1^2) = O(1/\gamma^3)$. In total, we get that throughout the process the cost of each light interval increases by factor at most $O(1/\gamma^3)$.

Additionally, the cost of each edge (between intervals) may increase by a factor of 4. Therefore, after rounding all light intervals, the cost of each light interval increases by factor at most $O(1/\gamma^3)$, and of any other interval by factor $O(1)$.

Heavy intervals. If $V_{+,j}$ is heavy, say it contains between $r\varepsilon n$ and $(r+1)\varepsilon n$ points for $r \in \mathbb{N}$, we partition $J = [M(1 + \varepsilon)^{-j-1}, M(1 + \varepsilon)^{-j}]$ into 3 equal thirds, consider the middle third, and take equally spaced points v_1, \dots, v_{r+2} in it such that v_1 and v_{r+2} are the endpoints of the middle third. In a formula,

26:16 Rounding via Low Dimensional Embeddings

$$v_i = M(1 + \varepsilon)^{-j-1} + L_j \frac{r + 1 + i}{3r + 5},$$

where again L_j is the length of J , $L_j = M\varepsilon(1 + \varepsilon)^{-j-1}$.

We could repeat the previous argument with each one of the points v_i , but this time since $V_{j,+}$ is heavy, the argument could terminate because small-set expansion no longer holds, in which case we have enlarged the interval around v_i to contain at least εn vertices. We change the definition of ℓ_1 to be $\frac{\gamma}{rK}$, so that if the argument fails for v_i it means we have found an interval around v_i of length at most $L_j/(100r)$ that contains at least εn vertices. Note that these intervals must be disjoint (as the distance between two v_i 's is greater than $L_j/(50r)$), and thus in total these interval cover no more than $|V_{+,j}| \leq (r + 1)\varepsilon n$ vertices, meaning the argument could fail on at most $(r + 1)$ many of the v_i 's. Thus, there is a v_i in which the argument succeeds, so we choose it to be $p_{j,+}$, so that running the previous argument using $p_{j,+}$ gives that cost_j increases by at most factor $O(1/(\gamma\ell_1^2)) = O(\gamma^{-3}\varepsilon^{-2})$.

Besides that, the cost of edges that have one endpoint in $V_{+,j}$ and one outside $V_{+,j}$ may increase by factor $O(1)$. Thus, in total the effect of rounding the heavy intervals is that it may increase the cost of each heavy interval by at most $O(\gamma^{-3}\varepsilon^{-2})$, and each light interval by factor at most $O(1)$. ◀

With Lemma 18 in hand we can prove Theorem 4, restated below.

► **Theorem 19.** *There exists an absolute constant $C > 0$, such that the following holds for all $\varepsilon, \gamma, \delta > 0$. Suppose that G is a d -regular graph whose second normalized eigenvalue is at least $1 - \delta$, and suppose G is an (ε, γ) small set expander. Then there exists a set S of vertices of fractional size at most $n/2$, such that*

$$\Phi(S) \leq C \cdot \frac{\delta}{\gamma^3 \varepsilon^3}.$$

Furthermore, such S can be found efficiently.

Proof. Let x be an eigenvector of G with eigenvalue $\lambda_2 \geq 1 - \delta$ and suppose that $\mathbb{E}[x(v)^2] = 1$; hence $\mathbb{E}_{(u,v) \in E} [(x(u) - x(v))^2] \leq 2\delta$. We assume without loss of generality that $x_1 \leq x_2 \leq \dots \leq x_n$.

Let $z = x + c\vec{1}$ an appropriate constant c so that $z_{n/2} = 0$. We note that this constant is at most 3 in absolute value, otherwise the absolute value of $x_{n/2}$ is greater than 3, say $x_{n/2} > 3$, and then

$$\mathbb{E}_v [x(v)^2] \geq \frac{1}{2} 3^2 > 1.$$

We also note that $\mathbb{E}_v [z(v)^2] = \mathbb{E}_v [x(v)^2] + c^2$ (we used the fact that x is perpendicular to $\vec{1}$), which is at constant between 1 and 10, so we may divide z by an appropriate constant and get its 2-norm back to 1. We note that performing all of these operations changes δ by at most a constant factor.

Using Lemma 18 on z we find an ε -quantized y satisfying the properties of the lemma. We divide the entries of y by a suitable factor A so that $y_1^2 + y_n^2 = 1$, and note that the operation of Lemma 18 preserves orders and 0-entries so that $y_1 \leq \dots \leq y_n$ and $y_{n/2} = 0$.

We now perform the standard Cheeger's inequality rounding scheme. Namely, choose t in $[y_1, y_n]$ according to the density function $2|t|$. We note that

1. If a, b have the same signs, the probability that $t \in [a, b]$ is $|b^2 - a^2|$;
2. If a, b have different signs, the probability that $t \in [a, b]$ is $a^2 + b^2$.
3. Therefore, in either case the probability that $t \in [a, b]$ is at most $|a - b|(|a| + |b|)$.

Let $S = \{v \mid y(v) \leq t\}$. We compute the expectation of the number of edges that escape S , as well as the expectation of $\min(|S|, |V \setminus S|)$.

By the above observation, the probability that a given edge (u, v) crosses the cut of S is at most $|y(u) - y(v|(|y(u)| + |y(v)|)$. If $y(u) = y(v)$ then this is 0 and in particular equal to $|y(u) - y(v)|^2$; if $y(u), y(v)$ are negated in signs this is same as $|y(u) - y(v)|^2$; finally, if either $y(v) = 0$ or $y(u) = 0$ this is again the same as $|y(u) - y(v)|^2$. Otherwise, as y is ε -quantized we have that either $|y(u)| \geq (1 + \varepsilon)|y(v)|$ or $|y(u)| \leq (1 - \varepsilon)|y(v)|$, and then

$$|y(u) - y(v)| |y(v)| = \left| \frac{y(u)}{y(v)} - 1 \right| |y(v)|^2 \leq \varepsilon^{-1} \left| \frac{y(u)}{y(v)} - 1 \right|^2 |y(v)|^2 = \varepsilon^{-1} |y(u) - y(v)|^2.$$

Similarly, flipping the roles of u and v we have that either $|y(v)| \geq (1 + \varepsilon)|y(u)|$ or $|y(v)| \leq (1 - \varepsilon)|y(u)|$, and then

$$|y(u) - y(v)| |y(u)| = \left| \frac{y(v)}{y(u)} - 1 \right| |y(u)|^2 \leq \varepsilon^{-1} \left| \frac{y(v)}{y(u)} - 1 \right|^2 |y(u)|^2 = \varepsilon^{-1} |y(u) - y(v)|^2.$$

Combining everything, we get that the probability that a given edge (u, v) crosses the cut of S is at most $2\varepsilon^{-1} |y(u) - y(v)|^2$, so the expected number of edges between S and $V \setminus S$ is at most

$$\sum_{(u,v) \in E} 2\varepsilon^{-1} |y(u) - y(v)|^2 \leq O\left(\frac{\delta}{\gamma^3 \varepsilon^3} \frac{1}{A^2} |E|\right).$$

As for $\mathbb{E}[\min(|S|, |V \setminus S|)]$, if $t \leq 0$ the set S is smaller and otherwise $V \setminus S$ is smaller. We show that for each v , the probability that v is in the smaller among S and $V \setminus S$ is at least $y(v)^2$, from which it follows that

$$\mathbb{E}[\min(|S|, |V \setminus S|)] \geq \sum_{v \in V} y(v)^2 \geq \frac{1}{A^2} (1 - \varepsilon)n.$$

This would imply Theorem 4, as then we get that

$$\frac{\mathbb{E}[|\text{Edges}(S, V \setminus S)|]}{d \mathbb{E}[\min(|S|, |V \setminus S|)]} = O\left(\frac{\delta}{\gamma^3 \varepsilon^3}\right),$$

hence there is a choice of t for which

$$\frac{|\text{Edges}(S, V \setminus S)|}{d \min(|S|, |V \setminus S|)} = O\left(\frac{\delta}{\gamma^3 \varepsilon^3}\right).$$

Furthermore, standard techniques show that such t can be found in polynomial time, giving us a set S as desired.

To finish the proof, we argue that for any vertex v , the probability that v is in the smaller among S and $V \setminus S$ is at least $y(v)^2$. We consider the case that $y(v) \leq 0$ and the case that $y(v) > 0$ separately. In the first case, we have

$$\begin{aligned} \Pr[v \text{ in smaller among } S, V \setminus S] &= \Pr[t \leq 0] \Pr[v \in S \mid t \leq 0] \\ &= \Pr[t \leq 0] \Pr[y(v) \leq t \mid t \leq 0], \end{aligned}$$

26:18 Rounding via Low Dimensional Embeddings

which is equal to $\Pr[t \in [y(v), 0]] = y(v)^2$. If $y(v) > 0$,

$$\begin{aligned} \Pr[v \text{ in smaller among } S, V \setminus S] &= \Pr[t > 0] \Pr[v \in V \setminus S \mid t > 0] \\ &= \Pr[t > 0] \Pr[y(v) > t \mid t \leq 0], \end{aligned}$$

which is equal to $\Pr[t \in [0, y(v)]] = y(v)^2$. ◀

3.2 Dense cuts on SSE's

Using the same proof strategy as above, one may establish the following improved spectral partitioning result based on the smallest eigenvalue of G when the graph is a small-set expander [22].

► **Theorem 20.** *There exists an absolute constant $C > 0$, such that the following holds for all $\varepsilon, \gamma, \delta > 0$. Suppose that G is a d -regular graph with $\lambda_n(G) \leq -1 + \delta$, and suppose G is an (ε, γ) small set expander. Then there is $y \in \{-1, 0, 1\}^n$ such that*

$$\frac{\sum_{(u,v) \in E} |y(u) + y(v)|}{d \sum_{u \in V} |y(u)|} \leq C \cdot \frac{\delta}{\gamma^3 \varepsilon^3}.$$

Furthermore, such S can be found efficiently.

As the proof is very close in spirit, we omit the details. We remark that Trevisan [22] proves a version of Theorem 20 for general graphs and uses it to get a spectral approximation algorithm for Max-cut with better approximation ratio than $1/2$. In that context, a solution as in Theorem 20 is to be interpreted as a partial cut, and after finding y Trevisan considers the set of vertices which receive value 0 in y and applies the algorithm on them recursively. One may expect such approach to also work in our case to get an algorithm for Max-cut, however the recursive nature of Trevisan's approach may not preserve the small-set expansion of the induced graph. We were therefore led to seeking a more direct approach for applying the soft-rounding idea to the Max-cut problem, and indeed our proof of Theorem 3 was found in this way.

4 Strong Parallel Repetition for Unique-Games on SSE's

In this section, we prove a more general version of Theorem 6 for the class of projection games. Our proof goes through the information theoretic approach to parallel repetition [20, 10, 19, 4], and we will closely follow the argument in [4]. Our argument differs only towards the end of the argument, but we sketch for completeness. We will prove that for sufficiently large absolute constant $K > 0$, setting $T = \lceil \frac{K}{\varepsilon^3 \gamma^3 \delta} \rceil$, for $t \geq T$ it holds that $\text{val}(\Psi^t) \leq (1 - \Omega(\delta \varepsilon^3 \gamma^3))^t$. The result for any other t then follows automatically: for any $t \leq T$, letting $\ell = \lfloor T/t \rfloor$ we have that $T \geq \ell t$ so

$$1 - \Omega(1) \geq \text{val}(\Psi^T) \geq \text{val}(\Psi^{\ell t}) \geq \text{val}(\Psi^t)^\ell,$$

so

$$\text{val}(\Psi^t) \leq (1 - \Omega(1))^{1/\ell} = e^{-\Omega(t/T)} = (e^{-\Omega(1/T)})^t = (1 - \Omega(1/T))^t,$$

as desired. From now on, we assume that $t \geq T$.

Notations

We will denote random variables by boldface capital letters, and instantiations of them by small letters. We use standard information theoretic tools presented in the appendix for completeness (see also [7]). The KL-divergence from Q to P is denoted by $D_{\text{KL}}(P \parallel Q)$, however for space considerations, we will sometimes denote it by

$$D_{\text{KL}} \left(\begin{matrix} P \\ = \\ Q \end{matrix} \right).$$

4.1 Parallel Repetition: the Information Theoretic Set-up

We will refer to the players in the game as Alice and Bob. The challenges of Alice are in the t -fold repeated game $\Psi^{\otimes n}$ are denoted by $\mathbf{X}_1, \dots, \mathbf{X}_t$ and her answers are denoted by $(\mathbf{A}_1, \dots, \mathbf{A}_t) = f(\mathbf{X}_1, \dots, \mathbf{X}_t)$. Similarly, the challenges of Bob are denoted by $\mathbf{Y}_1, \dots, \mathbf{Y}_t$, and his answers are denoted by $(\mathbf{B}_1, \dots, \mathbf{B}_t) = g(\mathbf{Y}_1, \dots, \mathbf{Y}_t)$. Finally, we denote by W the event that Alice and Bob win on all of the t challenges. We will prove the counter positive statement, namely that if $\Pr[W] > (1 - c\delta\varepsilon^3\gamma^3)^t$ for sufficiently small absolute constant c , then $\text{val}(\Psi) > 1 - \delta$ (we recall that $t \geq T$).

4.1.1 The information theoretic approach

Let $\mathbf{s}_g, \mathbf{s}_h$ be uniformly chosen integers in $\{3/4t + 1, \dots, t\}$, and let $\sigma \in S_t$ be a uniformly chosen permutation. Denote

$$\mathbf{H} = \sigma([\mathbf{s}_h]) = \{\sigma(i) \mid i \in [\mathbf{s}_h]\}, \quad \mathbf{G} = \sigma(\{t - \mathbf{s}_g + 1, \dots, t\}) = \{\sigma(i) \mid i \in \{t - \mathbf{s}_g + 1, \dots, t\}\}.$$

Let $\mathbf{I} \in \mathbf{G} \cap \mathbf{H}$ be chosen uniformly, and let $\ell \in [t/4]$ be chosen uniformly. Let $\mathbf{S} \subseteq (\mathbf{G} \cap \mathbf{H}) \setminus \{\mathbf{I}\}$ be chosen uniformly of size ℓ . Denote

$$\mathbf{L}_{\mathbf{S}, \mathbf{G}, \mathbf{H}, \mathbf{I}} = (\mathbf{X}_{\mathbf{G} \setminus \{\mathbf{I}\}}, \mathbf{Y}_{\mathbf{H} \setminus \{\mathbf{I}\}}, \mathbf{B}_{\mathbf{S}}).$$

Below we state Lemmas 5.2, 5.3, 5.4 and 5.5 from [4]. We omit the proofs as they are identical to the proofs therein.

$$\triangleright \text{Claim 21. } \mathbb{E}_{\mathbf{S}, \mathbf{G}, \mathbf{H}, \mathbf{I}} [I(\mathbf{A}_{\mathbf{I}}; \mathbf{Y}_{\mathbf{I}} | \mathbf{X}_{\mathbf{I}}, \mathbf{L}_{\mathbf{S}, \mathbf{G}, \mathbf{H}, \mathbf{I}}, W)] \leq \frac{4}{t} \log(1/\Pr[W]).$$

$$\triangleright \text{Claim 22. } \mathbb{E}_{\mathbf{S}, \mathbf{G}, \mathbf{H}, \mathbf{I}} [I(\mathbf{L}_{\mathbf{S}, \mathbf{G}, \mathbf{H}, \mathbf{I}}; \mathbf{Y}_{\mathbf{I}} | \mathbf{X}_{\mathbf{I}}, W)] \leq \frac{8}{t} \log(1/\Pr[W]).$$

$$\triangleright \text{Claim 23. } \mathbb{E}_{\mathbf{S}, \mathbf{G}, \mathbf{H}, \mathbf{I}} [I(\mathbf{L}_{\mathbf{S}, \mathbf{G}, \mathbf{H}, \mathbf{I}}; \mathbf{X}_{\mathbf{I}} | \mathbf{Y}_{\mathbf{I}}, W)] \leq \frac{8}{t} \log(1/\Pr[W]).$$

$$\triangleright \text{Claim 24. } \mathbb{E}_{\mathbf{S}, \mathbf{G}, \mathbf{H}, \mathbf{I}} [I(\mathbf{B}_{\mathbf{I}}; 1_W | \mathbf{X}_{\mathbf{I}}, \mathbf{Y}_{\mathbf{I}}, \mathbf{L}_{\mathbf{S}, \mathbf{G}, \mathbf{H}, \mathbf{I}}, \mathbf{A}_{\mathbf{I}})] \leq \frac{4}{t} H(1_W).$$

The next claim is Lemma 4.6 from [4].

$$\triangleright \text{Claim 25. } \mathbb{E}_{\mathbf{I}} [D_{\text{KL}}(\mathbf{X}_{\mathbf{I}}, \mathbf{Y}_{\mathbf{I}} \mid W \parallel \mathbf{X}_{\mathbf{I}}, \mathbf{Y}_{\mathbf{I}})] \leq \frac{1}{t} \log\left(\frac{1}{\Pr[W]}\right).$$

The next claim is Lemma 4.9 from [4].

\triangleright Claim 26. Suppose $G, H, S_a, S_b \subseteq [t]$ and $i \in [t]$ such that $G \cup H = [t] \setminus \{i\}$. Then for any $\bar{x}, \bar{y}, \bar{a}, \bar{b}, x, y$ in the support of $\mathbf{X}_G, \mathbf{Y}_H, \mathbf{A}_{S_a}, \mathbf{B}_{S_b}, \mathbf{X}_i$ and \mathbf{Y}_i , for all a, b it holds that

$$\begin{aligned} & \Pr[\mathbf{A}_i = a, \mathbf{B}_i = b \mid (\mathbf{X}_G, \mathbf{Y}_H, \mathbf{A}_{S_a}, \mathbf{B}_{S_b}, \mathbf{X}_i, \mathbf{Y}_i) = (\bar{x}, \bar{y}, \bar{a}, \bar{b}, x, y)] \\ &= \Pr[\mathbf{A}_i = a \mid (\mathbf{X}_G, \mathbf{Y}_H, \mathbf{A}_{S_a}, \mathbf{X}_i) = (\bar{x}, \bar{y}, \bar{a}, x)] \\ & \cdot \Pr[\mathbf{B}_i = b \mid (\mathbf{X}_G, \mathbf{Y}_H, \mathbf{B}_{S_b}, \mathbf{Y}_i) = (\bar{x}, \bar{y}, \bar{b}, x)]. \end{aligned}$$

26:20 Rounding via Low Dimensional Embeddings

Next, we state and prove an analog to Lemma 5.6 from [4]. The Lemma therein applies to the case that the probability of W is small, and below we make a small adjustment to it (using the fact that the number of repetitions is assumed to be large in our case) to remove that assumption.

► **Lemma 27.** *Suppose $\zeta = \delta\varepsilon^3\gamma^3$ and that $\Pr[W] \geq 2^{-\zeta t/K}$. There exist fixings $\mathbf{G} = G$, $\mathbf{H} = H$, $\mathbf{S} = S$ and $\mathbf{I} = I$ such that the following properties hold:*

1.

$$\mathbb{E}_{(x,y) \sim (\mathbf{X}_I, \mathbf{Y}_I) | W} [\text{D}_{KL}(\mathbf{L}_{S,G,H,I} | \mathbf{X}_I = x, \mathbf{Y}_I = y, W \parallel \mathbf{L}_{S,G,H,I} | \mathbf{X}_I = x, W)] \leq O\left(\frac{\zeta}{K}\right).$$

2.

$$\mathbb{E}_{(x,y) \sim (\mathbf{X}_I, \mathbf{Y}_I) | W} [\text{D}_{KL}(\mathbf{L}_{S,G,H,I} | \mathbf{X}_I = x, \mathbf{Y}_I = y, W \parallel \mathbf{L}_{S,G,H,I} | \mathbf{Y}_I = y, W)] \leq O\left(\frac{\zeta}{K}\right).$$

3.

$$\text{D}_{KL}(\mathbf{X}_I, \mathbf{Y}_I | W \parallel \mathbf{X}_I, \mathbf{Y}_I) \leq O\left(\frac{\zeta}{K}\right).$$

4. *Sampling $(x, y) \sim (\mathbf{X}_I, \mathbf{Y}_I) | W$ and $L \sim \mathbf{L}_{S,G,H,I} | \mathbf{X}_I = x, \mathbf{Y}_I = y, W$, we have*

$$\mathbb{E} \left[\text{D}_{KL}(\mathbf{A}_I | \mathbf{X}_I = x, \mathbf{Y}_I = y, \mathbf{L}_{S,G,H,I} = L, W \parallel \mathbf{A}_I | \mathbf{X}_I = x, \mathbf{L}_{S,G,H,I} = L, W) \right] \leq O\left(\frac{\zeta}{K}\right).$$

5. *Sampling $(x, y) \sim (\mathbf{X}_I, \mathbf{Y}_I) | W$ and $L \sim \mathbf{L}_{S,G,H,I} | \mathbf{X}_I = x, \mathbf{Y}_I = y, W$, we have*

$$\mathbb{E} \left[\text{D}_{KL} \left(\frac{\mathbf{A}_I, \mathbf{B}_I | \mathbf{X}_I = x, \mathbf{Y}_I = y, \mathbf{L}_{S,G,H,I} = L, W}{\mathbf{A}_I | \mathbf{X}_I = x, \mathbf{L}_{S,G,H,I} = L, W \otimes \mathbf{B}_I | \mathbf{Y}_I = y, \mathbf{L}_{S,G,H,I} = L, W} \right) \right] \leq O\left(\frac{\zeta}{K}\right).$$

Proof. We calculate the expected value of each one of these over the choice of $\mathbf{S}, \mathbf{G}, \mathbf{H}, \mathbf{I}$, show that each one of these expectations is at most $O\left(\frac{1}{t} \log\left(\frac{1}{\Pr[W]}\right)\right)$, and then the result follows from Markov's inequality together with the union bound.

The expectation of the first item is, by Fact 36, equal to the mutual information in Claim 22, so the bound follows from there. The expectation of the second item is, by Fact 36, equal to the mutual information in Claim 23, so the bound follows from there. The expectation of the third item is computed in Claim 25. The expectation of the fourth item is equal to the mutual information in Claim 21, and the bound follows from there.

For the fifth expectation, using the chain rule we can write it as

$$\underbrace{\mathbb{E}_{\substack{\mathbf{G}, \mathbf{H}, \mathbf{S}, \mathbf{I} \\ (x,y) \sim (\mathbf{X}_I, \mathbf{Y}_I) | W \\ L \sim \mathbf{L}_{S,G,H,I} | x,y,W}} [\text{D}_{KL}(\mathbf{A}_I | x, y, \mathbf{L}_{S,G,H,I} = L, W \parallel \mathbf{A}_I | x, \mathbf{L}_{S,G,H,I} = L, W)]}_{(I)} + \underbrace{\mathbb{E}_{\substack{a \sim \mathbf{A}_I \\ \mathbf{X}_I = x, \mathbf{Y}_I = y \\ \mathbf{L}_{S,G,H,I} = L, W}} [\text{D}_{KL}(\mathbf{B}_I | x, y, \mathbf{L}_{S,G,H,I} = L, \mathbf{A}_I = a, W \parallel \mathbf{B}_I | y, \mathbf{L}_{S,G,H,I} = L)]}_{(II)}$$

The expectation of (I) is the expectation of the fourth item in the lemma, hence at most $O\left(\frac{1}{t} \log\left(\frac{1}{\Pr[W]}\right)\right)$. For (II), first looking at Claim 26, summing over a and setting $S_a = \{\mathbf{I}\}$, $S_b = \mathbf{S}$ we get that the distributions $\mathbf{B}_I | \mathbf{Y}_I = y, \mathbf{L}_{S,G,H,I} = L$ and $\mathbf{B}_I | \mathbf{X}_I = x, \mathbf{Y}_I = y, \mathbf{L}_{S,G,H,I} = L, \mathbf{A}_I = a$ are identical, so our goal is to upper bound the expectation of

$$D_{\text{KL}}(\mathbf{B}_I | x, y, \mathbf{L}_{S,G,H,I} = L, W, \mathbf{A}_I = a, W \parallel \mathbf{B}_I | x, y, \mathbf{L}_{S,G,H,I} = L, \mathbf{A}_I = a).$$

Letting $w \sim 1_W$ and sampling the rest of the random variables conditioned on it, we see that the expectation of the above expression is at most $\frac{1}{\Pr[W]}$ times the expectation of

$$D_{\text{KL}}\left(\frac{\mathbf{B}_I | X_I = x, Y_I = y, \mathbf{L}_{S,G,H,I} = L, 1_W = w, \mathbf{A}_I = a}{\mathbf{B}_I | X_I = x, Y_I = y, \mathbf{L}_{S,G,H,I} = L, \mathbf{A}_I = a}\right)$$

since the probability that $w = 1$ is $\Pr[W]$. By Fact 36, the expectation of the last KL-divergence is the mutual information in Claim 24, hence at most $\frac{4}{t} H[1_W]$.

All in all, we get that the expectation of (II) is upper bounded by

$$\frac{4H[1_W]}{t\Pr[W]}.$$

Denote $q = \Pr[W]$. If $q \leq 1/2$, then we may bound $H[1_W] \leq O(q \log(1/q))$, and then we get the bound $O\left(\frac{\log(1/\Pr[W])}{t}\right) \leq O\left(\frac{\zeta}{K}\right)$ on (II). Otherwise, $q > 1/2$ and we have the bound $H[1_W] \leq 1$, and we get the bound $O(1/t)$. By the assumption in the beginning of this section, $t \geq T \geq \Omega(K/\zeta)$, so $O(1/t) \leq O\left(\frac{\zeta}{K}\right)$. ◀

4.1.2 Departing from [4]

In the next part of the proof, we insert an additional ingredient on top of [4], and before that we quickly explain how the proof there proceeds using Lemma 27.

First, a protocol is designed so that if the players received challenges from the distribution $\mathbf{X}_I, \mathbf{Y}_I | W$, then the players succeed with probability close to 1; given that, the third item in Lemma 27 shows that the players succeed with probability close to 1 given challenges distributed as $\mathbf{X}_I, \mathbf{Y}_I$.

The correctness of the protocol is argued by appealing to Pinsker's inequality on the third item of Lemma 27, to get that Alice and Bob can jointly sample from distributions that are $O(\sqrt{\zeta/K})$ close to $\mathbf{L}_{S,G,H,I} | X_I = x, Y_I = y, W$, i.e. with probability $1 - O(\sqrt{\zeta/K})$ they get a joint sample from that distribution. Assuming the joint sampling was successful, the players can sample the answer A_I, B_I conditioned on the information they have so far. Note that in the distribution $\mathbf{A}_I, \mathbf{B}_I | X_I = x, Y_I = y, \mathbf{L}_{S,G,H,I} = L, W$ the players win the coordinate I with probability 1 (since W was conditioned on), and by the fifth item in Lemma 27 the KL-divergence between that distribution and the joint distribution of the answers of the players in the designed protocol is at most $O(\zeta/K)$, and one conclude that the winning probability of the players in the protocol (conditioned on the joint sampling being successful) is at least $2^{-O(\zeta/K)} \geq 1 - O(\zeta/K)$. Overall, the designed protocol wins with probability $1 - O(\sqrt{\zeta/K})$.

The source of the square loss above is thus due to the application of Pinsker's inequality, and we will circumvent that by appealing to the small set expansion property. More precisely, we will use the first and second item in Lemma 27 in order to come up with distributions $\tilde{\mathbf{L}}_{S,G,H,I} | X_I = x, W, \tilde{\mathbf{L}}_{S,G,H,I} | X_I = y, W$ that are very close to the original distribution

26:22 Rounding via Low Dimensional Embeddings

(the probability of each atom changes by at most factor $(1 \pm \varepsilon)$, but from which the players can jointly sample without losing the square root in Pinsker's inequality (and instead loses some factors depending on the small set expansion parameters of the graph). From that point, the rest of the proof proceeds in the same way.

4.2 Quantizing the Random Variables

Throughout this section, we have a graph $G = (V, E)$, and we associate a distribution \mathcal{D}_v with each vertex v . We will also need to consider pseudo-distributions, which we define next.

► **Definition 28.** A pseudo-distribution \mathcal{D} over a finite domain Ω is a map $\mathcal{D}: \Omega \rightarrow [0, \infty)$.

Sometimes we will want to sample from a pseudo-distribution; by that, we mean that we first normalize \mathcal{D} so that the sum of its values is 1, and then sample from it. One may define Hellinger distance as well as statistical distance for pseudo-distributions as well (though one has to be careful with using properties of them that only hold for distributions).

► **Definition 29.** A collection of pseudo-distributions $\mathcal{D}_{v \in V}$ over a finite domain Ω is called ε -quantized if for any $w \in \Omega$ and $u, v \in V$ we either have that $\mathcal{D}_v(w) = \mathcal{D}_u(w)$, or $\mathcal{D}_v(w) \leq (1 - \varepsilon)\mathcal{D}_u(w)$, or $\mathcal{D}_v(w) \geq (1 + \varepsilon)\mathcal{D}_u(w)$.

The following lemma explains the benefit of quantized pseudo-distributions: it allows us to move from Hellinger distance to statistical distance without square root loss.

► **Lemma 30.** If P, Q are pseudo-distributions such that the collection $\{P, Q\}$ is ε -quantized. Then $\text{SD}(P, Q) = O\left(\frac{1}{\varepsilon} \text{Hellinger}(P, Q)^2\right)$.

Proof. By definitions, the inequality we wish to show is that for large enough absolute constant $C > 0$ it holds that

$$\sum_{w \in \Omega} |P(w) - Q(w)| \leq \frac{C}{\varepsilon} \sum_{w \in \Omega} \left| \sqrt{P(w)} - \sqrt{Q(w)} \right|^2.$$

We show that the inequality in fact holds term by term. Without loss of generality, $P(w) \geq Q(w)$. If $Q(w) = 0$ it is clear. Otherwise, letting $t = P(w)/Q(w)$, the inequality we wish to show is that $t - 1 \leq \frac{C}{\varepsilon}(\sqrt{t} - 1)^2$. If $t = 1$ the inequality is clear, and otherwise it is equivalent to $\sqrt{t} + 1 \leq \frac{C}{\varepsilon}(\sqrt{t} - 1)$. Solving for t , we get that

$$\sqrt{t} \geq \frac{C/\varepsilon + 1}{C/\varepsilon - 1} = 1 + \frac{2}{C/\varepsilon - 1} = 1 + \frac{2\varepsilon}{C - \varepsilon}$$

By the fact that $\{P, Q\}$ is ε -quantized and $t \neq 1$, it follows that $t \geq 1 + \varepsilon$ and hence $\sqrt{t} \geq 1 + \varepsilon/4$, and for large enough C ($C = 10$ will do) we get that this is indeed more than $1 + \frac{2\varepsilon}{C - \varepsilon}$. ◀

The next lemma shows that one may modify distributions associated with vertices of small set expanders by small multiplicative modifications to make them quantized, and only incur a constant multiplicative increase in Hellinger distance.

► **Lemma 31.** There exists an absolute constant $C > 0$, such that the following holds. Suppose that $G = (V, E)$ is a regular graph as well as a (ε, γ) small set expander, and $(\mathcal{D}_v)_{v \in V}$ is a collection of distributions over domain Ω such that

$$\mathbb{E}_{(u,v) \in E} [\text{Hellinger}(\mathcal{D}_u, \mathcal{D}_v)^2] \leq \eta.$$

Then there exists a collection of pseudo-distributions $(\tilde{\mathcal{D}}_v)_{v \in V}$ over Ω satisfying the following properties:

1. $\mathbb{E}_{(u,v) \in E} [\text{Hellinger}(\tilde{\mathcal{D}}_u, \tilde{\mathcal{D}}_v)^2] \leq C \frac{\eta}{\gamma^3 \varepsilon^2}$.
2. For all $w \in \Omega$ and $v \in V$ it holds that $(1 - \varepsilon)\mathcal{D}_v(w) \leq \tilde{\mathcal{D}}_v(w) \leq (1 + \varepsilon)\mathcal{D}_v(w)$.
3. The collection $(\tilde{\mathcal{D}}_v)_{v \in V}$ is $\varepsilon/2$ -quantized.

Proof. For each $w \in \Omega$, define the contribution of w to the Hellinger distance squared as

$$c_w = \mathbb{E}_{(u,v) \in E} \left[\left(\sqrt{\mathcal{D}_u(w)} - \sqrt{\mathcal{D}_v(w)} \right)^2 \right],$$

and define the vector $x_w \in [0, 1]^V$ by $x_w(v) = \sqrt{\mathcal{D}_v(w)}$. Then by Lemma 18 we may find a vector y_w that is ε -quantized and

$$c_w \stackrel{\text{def}}{=} \mathbb{E}_{(u,v) \in E} [(y_w(u) - y_w(v))^2] \leq C \frac{c_w}{\gamma^3 \varepsilon^2}.$$

For each v , define the pseudo-distribution $\tilde{\mathcal{D}}_v(w) = y_w(v)^2$. Then as y_w is ε -quantized, it is easy to see that this collection of distributions is $\varepsilon/2$ quantized. Also, the second item follows immediately from the corresponding properties of y_w . Finally,

$$\mathbb{E}_{(u,v) \in E} [\text{Hellinger}(\tilde{\mathcal{D}}_u, \tilde{\mathcal{D}}_v)^2] = \sum_w \mathbb{E}_{(u,v) \in E} [(y_w(u) - y_w(v))^2] \leq \sum_w C \frac{c_w}{\gamma^3 \varepsilon^2} = \frac{C}{\gamma^3 \varepsilon^2} \sum_w c_w,$$

which is at most $C \frac{\eta}{\gamma^3 \varepsilon^2}$ as the sum of c_w is exactly $\mathbb{E}_{(u,v) \in E} [\text{Hellinger}(\tilde{\mathcal{D}}_u, \tilde{\mathcal{D}}_v)^2]$. \blacktriangleleft

4.3 Using the Random Variables to Derive a Strategy

We now have all of the ingredients to prove Theorem 6. We will show that assuming that $\Pr[W] \geq 2^{-\zeta t/K}$ for $\zeta = \delta \varepsilon^3 \gamma^3$, we can design a protocol for the players to win the game with probability $1 - O(\delta/K)$, so taking K large enough yields a contradiction.

Apply Lemma 27 to find a fixing of S, H, G, I . In the protocol we design, the input of the players will be a pair of challenges distribution $(x, y) \sim (\mathbf{X}_I, \mathbf{Y}_I)$. We will analyze the protocol under a different distribution of challenges, namely under $(x, y) \sim (\mathbf{X}_I, \mathbf{Y}_I) \mid W$, and prove that the players succeed with probability at least $1 - O(\delta/K)$. From this, Lemma 42 together with the third item in Lemma 27 gives that the players win under the original distribution of challenges with probability at least $1 - O(\delta/K)$, as $\zeta \leq \delta$.

We set up some notations. Given a challenge x to Alice and y to Bob (which are vertices in Ψ), they define the random variable

$$\mathcal{D}_x = \mathbf{L}_{S,G,H,I} \mid \mathbf{X}_I = x, W \quad \mathcal{D}_y = \mathbf{L}_{S,G,H,I} \mid \mathbf{Y}_I = y, W$$

respectively. We also define $\mathcal{D}_{x,y} = \mathbf{L}_{S,G,H,I} \mid \mathbf{X}_I = x, \mathbf{Y}_I = y, W$ for convenience (though we stress no player has access to it). Then the first and second items in Lemma 27,

$$\mathbb{E}_{(x,y) \sim (\mathbf{X}_I, \mathbf{Y}_I) \mid W} [\text{D}_{\text{KL}}(\mathcal{D}_{x,y} \parallel \mathcal{D}_x) + \text{D}_{\text{KL}}(\mathcal{D}_{x,y} \parallel \mathcal{D}_y)] \leq O(\zeta/K),$$

and applying Lemma 41 we get that

$$\mathbb{E}_{(x,y) \sim (\mathbf{X}_I, \mathbf{Y}_I) \mid W} [\text{Hellinger}(\mathcal{D}_x, \mathcal{D}_y)^2] \leq O(\zeta/K).$$

26:24 Rounding via Low Dimensional Embeddings

Thus, by Lemma 44

$$\mathbb{E}_{(x,y) \sim (\mathbf{X}_I, \mathbf{Y}_I)} [\text{Hellinger}(\mathcal{D}_x, \mathcal{D}_y)^2] \leq O(\zeta/K),$$

and from Lemma 31 there are pseudo-distributions $(\tilde{\mathcal{D}}_x)$ as in the lemma such that

$$\mathbb{E}_{(x,y) \sim (\mathbf{X}_I, \mathbf{Y}_I)} [\text{Hellinger}(\tilde{\mathcal{D}}_x, \tilde{\mathcal{D}}_y)^2] \leq O\left(\frac{\zeta}{K\varepsilon^2\gamma^3}\right) = O\left(\frac{\delta\varepsilon}{K}\right).$$

We note that $\text{Hellinger}(\tilde{\mathcal{D}}_x, \tilde{\mathcal{D}}_y)^2 \leq 2$ for all x, y (it can be more than 1 as these are pseudo-distributions, but by the second item in Lemma 31 their sum of values is at most $(1 + \varepsilon)$). Thus we may apply Lemma 44 in the other direction to conclude that

$$\mathbb{E}_{(x,y) \sim (\mathbf{X}_I, \mathbf{Y}_I) \mid W} [\text{Hellinger}(\tilde{\mathcal{D}}_x, \tilde{\mathcal{D}}_y)^2] \leq O\left(\frac{\delta\varepsilon}{K}\right).$$

Thus, from Lemma 30 we have that

$$\mathbb{E}_{(x,y) \sim (\mathbf{X}_I, \mathbf{Y}_I) \mid W} [\text{SD}(\tilde{\mathcal{D}}_x, \tilde{\mathcal{D}}_y)] \leq O\left(\frac{\delta}{K}\right)$$

From Lemma 45, it thus follows that the players can jointly sample $d_x \sim \tilde{\mathcal{D}}_x$, $d_y \sim \tilde{\mathcal{D}}_y$ such that

$$\Pr_{(x,y) \sim (\mathbf{X}_I, \mathbf{Y}_I) \mid W, d_x, d_y} [d_x \neq d_y] = O(\delta/K).$$

We can now state the protocol for the players.

Input: Alice is given x , Bob is given y such that $(x, y) \sim (\mathbf{X}_I, \mathbf{Y}_I) \mid W$.

Protocol:

1. Alice and Bob use correlated sampling to jointly sample $d_x \sim \tilde{\mathcal{D}}_x$, $d_y \sim \tilde{\mathcal{D}}_y$.
2. Alice samples $a \sim \mathbf{A}_I \mid \mathbf{X}_I = x, \mathcal{D}_x = d_x, W$ and sends it to the referee.
3. Bob samples $b \sim \mathbf{B}_I \mid \mathbf{Y}_I = y, \mathcal{D}_y = d_y, W$ and sends it to the referee.

The following claim finishes the proof.

▷ **Claim 32.** The above protocol is a strategy for the players that succeeds with probability at least $1 - O(\delta/K)$.

Proof. Consider the alternative protocol, where instead of the first step, Alice samples d_x and then communicates it to Bob. Clearly, as the probability that $d_x \neq d_y$ is at most $O(\zeta/K)$, this change increases the success probability of the players by at most $O(\zeta/K)$, so it suffices to show that this modified protocol has success probability at least $1 - O(\delta/K)$.

We first analyze the protocol where we modify the distribution of d_x , and then argue the implication to the protocol with the correct distribution over d_x . Suppose Alice sampled $d_x \sim \mathcal{D}_x$; we show that in this case, the success probability is at least $1 - O(\zeta/K)$. To show that, by the first item of Lemma 27 and Lemma 42, it suffices to show that the success probability of the protocol is at least $1 - O(\zeta/K)$ if we sample $d_x \sim \mathcal{D}_{x,y}$. In that case, the answers of the players are distributed according to $(\mathbf{A}_I \mid \mathbf{X}_I = x, \mathcal{D}_{x,y} = d_x, W) \otimes (\mathbf{B}_I \mid \mathbf{Y}_I = y, \mathcal{D}_y = d_y, W)$. Thus, by the fifth item in Lemma 27 and Lemma 42 it suffices to show that the success probability of the players with $(a, b) \sim (\mathbf{A}_I, \mathbf{B}_I \mid \mathbf{X}_I = x, \mathbf{Y}_I = y, \mathcal{D}_{x,y} = d_x, W)$ is at least $1 - O(\zeta/K)$, which is clear since we have conditioned on the event W so all coordinates are won.

Thus, letting $Z(d)$ be a random variable denoting the failure probability of the modified protocol with $d_x = d$, we have that $\mathbb{E}_{x,y} [\mathbb{E}_{d \sim \mathcal{D}_x} [Z(d)]] \leq O(\zeta/K)$, hence

$$\mathbb{E}_{x,y} \left[\mathbb{E}_{d \sim \tilde{\mathcal{D}}_x} [Z(d)] \right] = \mathbb{E}_{x,y} \left[\sum_d \Pr [\tilde{\mathcal{D}}_x = d] Z(d) \right] \leq \mathbb{E}_{x,y} \left[(1 + \varepsilon) \sum_d \Pr [\mathcal{D}_x = d] Z(d) \right],$$

which is $O\left(\frac{\zeta}{K}\right)$. Therefore the success probability of the protocol when $d_x \sim \tilde{\mathcal{D}}_x$ is at least $1 - O\left(\frac{\zeta}{K}\right)$, finishing the proof. \triangleleft

References

- 1 Noga Alon and Bo'az Klartag. Economical toric spines via cheeger's inequality. *Journal of Topology and Analysis*, 1(02):101–111, 2009.
- 2 Sanjeev Arora, Subhash Khot, Alexandra Kolla, David Steurer, Madhur Tulsiani, and Nish-eeth K. Vishnoi. Unique games on expanding constraint graphs are easy: extended abstract. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing, Victoria, British Columbia, Canada, May 17-20, 2008*, pages 21–28, 2008. doi:10.1145/1374376.1374380.
- 3 Boaz Barak, Moritz Hardt, Ishay Haviv, Anup Rao, Oded Regev, and David Steurer. Rounding parallel repetitions of unique games. In *49th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2008, October 25-28, 2008, Philadelphia, PA, USA*, pages 374–383, 2008. doi:10.1109/FOCS.2008.55.
- 4 Mark Braverman and Ankit Garg. Small value parallel repetition for general games. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015*, pages 335–340, 2015. doi:10.1145/2746539.2746565.
- 5 Mark Braverman, Konstantin Makarychev, Yury Makarychev, and Assaf Naor. The grothendieck constant is strictly smaller than krivine's bound. In *IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS 2011, Palm Springs, CA, USA, October 22-25, 2011*, pages 453–462, 2011. doi:10.1109/FOCS.2011.77.
- 6 Mark Braverman and Dor Minzer. Optimal tiling of the euclidean space using permutation-symmetric bodies. In *36th Computational Complexity Conference, CCC 2021, July 20-23, 2021, Toronto, Ontario, Canada (Virtual Conference)*, pages 5:1–5:48, 2021. doi:10.4230/LIPIcs.CCC.2021.5.
- 7 Thomas M Cover. *Elements of information theory*. John Wiley & Sons, 1999.
- 8 Uriel Feige, Guy Kindler, and Ryan O'Donnell. Understanding parallel repetition requires understanding foams. In *Twenty-Second Annual IEEE Conference on Computational Complexity (CCC'07)*, pages 179–192. IEEE, 2007.
- 9 Michel X. Goemans and David P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. ACM*, 42(6):1115–1145, 1995. doi:10.1145/227683.227684.
- 10 Thomas Holenstein. Parallel repetition: Simplification and the no-signaling case. *Theory Comput.*, 5(1):141–172, 2009. doi:10.4086/toc.2009.v005a008.
- 11 Shlomo Hoory, Nathan Linial, and Avi Wigderson. Expander graphs and their applications. *Bulletin of the American Mathematical Society*, 43(4):439–561, 2006.
- 12 Subhash Khot. On the power of unique 2-prover 1-round games. In *Proceedings on 34th Annual ACM Symposium on Theory of Computing, May 19-21, 2002, Montréal, Québec, Canada*, pages 767–775, 2002. doi:10.1145/509907.510017.
- 13 Guy Kindler, Anup Rao, Ryan O'Donnell, and Avi Wigderson. Spherical cubes: optimal foams from computational hardness amplification. *Commun. ACM*, 55(10):90–97, 2012. doi:10.1145/2347736.2347757.

- 14 Tsz Chiu Kwok, Lap Chi Lau, Yin Tat Lee, Shayan Oveis Gharan, and Luca Trevisan. Improved cheeger's inequality: Analysis of spectral partitioning algorithms through higher order spectral gap. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 11–20, 2013.
- 15 James R Lee, Shayan Oveis Gharan, and Luca Trevisan. Multiway spectral partitioning and higher-order cheeger inequalities. *Journal of the ACM (JACM)*, 61(6):1–30, 2014.
- 16 Anand Louis, Prasad Raghavendra, Prasad Tetali, and Santosh Vempala. Many sparse cuts via higher eigenvalues. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 1131–1140, 2012.
- 17 Dana Moshkovitz. Strong parallel repetition for unique games on small set expanders. *CoRR*, abs/2103.08743, 2021. [arXiv:2103.08743](https://arxiv.org/abs/2103.08743).
- 18 Prasad Raghavendra and David Steurer. Graph expansion and the unique games conjecture. In *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010*, pages 755–764, 2010. doi:10.1145/1806689.1806792.
- 19 Anup Rao. Parallel repetition in projection games and a concentration bound. *SIAM J. Comput.*, 40(6):1871–1891, 2011. doi:10.1137/080734042.
- 20 Ran Raz. A parallel repetition theorem. *SIAM J. Comput.*, 27(3):763–803, 1998. doi:10.1137/S0097539795280895.
- 21 Shmuel Safra and Oded Schwartz. On parallel-repetition, Unique-Games and Max-Cut, 2007.
- 22 Luca Trevisan. Max cut and the smallest eigenvalue. *SIAM J. Comput.*, 41(6):1769–1786, 2012. doi:10.1137/090773714.
- 23 Madhur Tulsiani, John Wright, and Yuan Zhou. Optimal strong parallel repetition for projection games on low threshold rank graphs. In *International Colloquium on Automata, Languages, and Programming*, pages 1003–1014. Springer, 2014.

A Information Theory

In this section, we present a few basic notions from information theory. First, we define the notions of Shannon entropies and conditional Shannon entropies.

► **Definition 33.** Let X, Y be random variables with a finite support.

1. The Shannon entropy of X is $H[X] = \sum_x \Pr[X = x] \log\left(\frac{1}{\Pr[X=x]}\right)$.
2. The Shannon entropy of X conditioned on Y is $H[X | Y] = \mathbb{E}_{y \sim Y} [H[X | Y = y]]$, where $H[X | Y = y] = \sum_x \Pr[X = x | Y = y] \log\left(\frac{1}{\Pr[X=x | Y=y]}\right)$.

Second, we define mutual information between random variables as well as conditional mutual information.

► **Definition 34.** Let X, Y, Z be random variables with a finite support.

1. The mutual information between X and Y is $I[X; Y] = H[X] - H[X|Y]$.
2. The mutual information between X, Y conditioned on Z is $I[X; Y | Z] = H[X | Z] - H[X | Y, Z]$.

Third, we define the KL-divergence between random variables.

► **Definition 35.** Let X, Y be random variables with a finite support. The KL-divergence from Y to X is $D_{KL}(X || Y) = \sum_{x,y} \Pr[X = x] \log\left(\frac{\Pr[X=x]}{\Pr[Y=y]}\right)$.

We will need the following standard facts from information theory (for proofs, see [7] for example).

► **Fact 36.** Let X, Y, Z be random variables. Then

$$I[X, Y; Z] = \mathbb{E}_{(x,y) \sim (X,Y)} [D_{KL}(Z|_{X=x, Y=y} \| Z)].$$

► **Fact 37.** Let X, Y_1, \dots, Y_n be random variables. Then

$$I[X; Y_1, \dots, Y_n] = \sum_{i=1}^n I[X; Y_i | Y_{<i}].$$

► **Fact 38.** Let X, Y, Z be random variables. Then $I[X; Y | Z] \leq I[X; Y, Z]$.

B Probability Theory

We will use several standard distance measures between probability measures as well as relations between them that we give in this section.

► **Definition 39.** Let X, Y be random variables supported over a finite set Ω . The statistical distance between X and Y is defined as

$$SD(X, Y) = \frac{1}{2} \sum_{w \in \Omega} |\Pr[X = w] - \Pr[Y = w]|$$

► **Definition 40.** Let X, Y be random variables supported over a finite set Ω . The Hellinger distance between X and Y is defined as

$$\text{Hellinger}(X, Y) = \sqrt{\frac{1}{2} \sum_{w \in \Omega} \left| \sqrt{\Pr[X = w]} - \sqrt{\Pr[Y = w]} \right|^2}$$

B.1 A Relation Between KL-divergence and Hellinger Distance

The KL divergence metric does not obey a triangle inequality, and the following lemma replaces it in our intended application. It asserts that if the KL divergence between X and Z is small, and the KL divergence between Y and Z is small, then the Hellinger distance between X and Y is small.

► **Lemma 41.** Let X, Y and Z be random variables supported on a finite set Ω . Then

$$2\text{Hellinger}(X, Y)^2 \leq D_{KL}(Z \| X) + D_{KL}(Z \| Y).$$

Proof. Let us think of X and Y as fixed random variables, and attempt to minimize $D_{KL}(Z \| X) + D_{KL}(Z \| Y)$ over all random variables Z . In other words, we wish find non-negative numbers $(p_Z(w))_{w \in \Omega}$ summing up to 1 minimizing the form

$$\sum_{w \in \Omega} p_Z(w) \log \left(\frac{p_Z(w)}{p_X(w)} \right) + p_Z(w) \log \left(\frac{p_Z(w)}{p_Y(w)} \right) = \sum_{w \in \Omega} p_Z(w) \log \left(\frac{p_Z(w)^2}{p_X(w)p_Y(w)} \right)$$

Using Lagrange multipliers, we get that there exists $\lambda \in \mathbb{R}$ such that the optimum satisfies the equations

$$\log \left(\frac{p_Z(w)^2}{p_X(w)p_Y(w)} \right) - \frac{1}{\ln 2} - \lambda = 0 \quad \forall w \in \Omega.$$

26:28 Rounding via Low Dimensional Embeddings

In other words, $p_Z(W) = c\sqrt{p_X(w)p_Y(w)}$ for some $c > 0$, and we next compute the constant c . By the constraint that the sum is 1, we get

$$1 = \frac{c}{2} \sum_{w \in \Omega} 2\sqrt{p_X(w)p_Y(w)} = \frac{c}{2} \left(2 - \sum_{w \in \Omega} (\sqrt{p_X(w)} - \sqrt{p_Y(w)})^2 \right) = c(1 - \text{Hellinger}(X, Y)^2),$$

so $c = \frac{1}{1 - \text{Hellinger}(X, Y)^2}$. Plugging that into $p_Z(w)$, we get that the minimum of $D_{\text{KL}}(Z \| X) + D_{\text{KL}}(Z \| Y)$ is at least

$$\sum_{w \in \Omega} p_Z(w) \log \left(\frac{p_Z(w)^2}{p_X(w)p_Y(w)} \right) = \sum_{w \in \Omega} c\sqrt{p_X(w)p_Y(w)} \log \left(\frac{c^2 p_X(w)p_Y(w)}{p_X(w)p_Y(w)} \right).$$

Simplifying, we get that this is equal to

$$c \log c^2 \sum_{w \in \Omega} \sqrt{p_X(w)p_Y(w)} = 2 \log c = -2 \log(1 - \text{Hellinger}(X, Y)^2) \geq 2 \text{Hellinger}(X, Y)^2,$$

where the last inequality uses $\log(s) \leq s - 1$ for all $s > 0$. ◀

B.2 Small KL-divergence and high probability events

The next few lemmas are concerned with a pair of distributions that have a small KL-divergence. They show that an event has probability close to 1 with respect to one distribution if and only if it has probability close to 1 with respect to the other distributions; we then generalize this facts to bounded functions.

► **Lemma 42.** *Suppose P and Q are distributions such that $D_{\text{KL}}(P \| Q) \leq \eta \leq 1/100$, and suppose that E is an event.*

1. *If $P[E] \geq 1 - \eta$, then $Q(E) \geq 1 - 10\eta$.*
2. *If $Q(E) \geq 1 - \eta$, then $P(E) \geq 1 - 10\eta$.*

Proof. By the data processing inequality, $\eta \geq D_{\text{KL}}(P \| Q) \geq D_{\text{KL}}(\text{Ber}(P(E)) \| \text{Ber}(Q(E)))$.

For the first item, if we assume $Q(E) \leq 1 - 10\eta$ then we get that

$$\begin{aligned} \eta &\geq D_{\text{KL}}(\text{Ber}(1 - \eta) \| \text{Ber}(1 - 10\eta)) = (1 - \eta) \log \left(\frac{1 - \eta}{1 - 10\eta} \right) + \eta \log \left(\frac{\eta}{10\eta} \right) \\ &= (1 - \eta) \log \left(1 + \frac{9\eta}{1 - 10\eta} \right) - \eta \log(10) \\ &\geq (1 - \eta) \log(1 + 9\eta) - \eta \log(10). \end{aligned}$$

Using the fact that $\log(1 + s) \geq s/2$ for $s \leq 1$, we get that the last expression is at least $4.5\eta(1 - \eta) - \eta \log(10)$, and as $\log(10) \leq 3.4$ we get that it is at least $1.1\eta - 4.5\eta^2 \geq 1.1\eta - 0.045\eta > \eta$, and contradiction.

For the second item, if we assume that $P(E) \leq 1 - 10\eta$, then we get that

$$\begin{aligned} \eta &\geq D_{\text{KL}}(\text{Ber}(1 - 10\eta) \| \text{Ber}(1 - \eta)) = (1 - 10\eta) \log \left(\frac{1 - 10\eta}{1 - \eta} \right) + 10\eta \log \left(\frac{10\eta}{\eta} \right) \\ &= (1 - 10\eta) \log \left(1 - \frac{9\eta}{1 - \eta} \right) + 10 \log(10) \eta \\ &\geq -(1 - 10\eta) \frac{18\eta}{1 - \eta} + 10 \log(10) \eta, \end{aligned}$$

where we used $\log(1-s) \geq -2s$ for $s \leq 0.1$. As $1/(1-\eta) \leq 1+2\eta$, we get that the last expression is at least

$$(10 \log(10) - 18)\eta - 216\eta^2 \geq 10\eta - \frac{216}{100}\eta^2 > \eta,$$

and contradiction. \blacktriangleleft

► **Lemma 43.** *Suppose X, Y are random variables in $[0, 1]$, and $D_{\text{KL}}(X \parallel Y) \leq \eta$, for $\eta \leq 1/100$. Then*

1. *If $\mathbb{E}[X] \geq 1 - \eta$, then $\mathbb{E}[Y] \geq 1 - 10\eta$.*
2. *If $\mathbb{E}[Y] \geq 1 - \eta$, then $\mathbb{E}[X] \geq 1 - 10\eta$.*

Proof. We prove the first item, and the second item follows similarly.

Consider the following randomized process \mathcal{T} : given a number $x \in [0, 1]$ sample a Bernoulli random variable b such that $\Pr[b = 1] = x$. Then by the data processing inequality we have that

$$\eta \geq D_{\text{KL}}(X \parallel Y) \geq D_{\text{KL}}(\mathcal{X} \parallel \mathcal{Y}).$$

Let E be the event that the output of the process is 1. Then

$$\Pr[\mathcal{X} \in E] = \mathbb{E}[X] \geq 1 - \eta,$$

so we get by Lemma 42 that

$$\mathbb{E}[Y] = \Pr[\mathcal{Y} \in E] \geq 1 - 10\eta. \quad \blacktriangleleft$$

► **Lemma 44.** *Suppose $f: \Omega \rightarrow [0, 2]$, $\eta \leq 1/100$ and P, Q are distributions such that $D_{\text{KL}}(P \parallel Q) \leq \eta$. Then*

1. *If $\mathbb{E}_{z \sim P}[f(z)] \leq \eta$, then $\mathbb{E}_{z \sim Q}[f(z)] \leq 20\eta$.*
2. *If $\mathbb{E}_{z \sim Q}[f(z)] \leq \eta$, then $\mathbb{E}_{z \sim P}[f(z)] \leq 20\eta$.*

Proof. We prove the first item, and the second item follows similarly.

Consider the random variables $X = 1 - \frac{1}{2}f(z)$ where $z \sim P$, and $Y = 1 - \frac{1}{2}f(z)$ where $z \sim Q$. By the Data processing inequality $D_{\text{KL}}(X \parallel Y) \leq D_{\text{KL}}(P \parallel Q) \leq \eta$, and by assumption $\mathbb{E}[X] \geq 1 - \eta$, so by Lemma 43 we get that $0.5\mathbb{E}_{z \sim Q}[f(z)] = 1 - \mathbb{E}[Y] \leq 10\eta$. \blacktriangleleft

B.3 Correlated sampling

Correlated sampling is an important motive in parallel repetition theorems that has been introduced in [10]. Below, we state a version of it for pseudo-distributions, and include a proof for completeness.

► **Lemma 45.** *There exists a randomized procedure \mathcal{R} .*

Let $\eta \leq 1/100$ and $0 < \varepsilon < 1/2$. Suppose P and Q are distributions, and \tilde{P}, \tilde{Q} are pseudo-distributions all over Ω such that

1. $\text{SD}(\tilde{P}, \tilde{Q}) \leq \eta$
 2. *for all $w \in \Omega$, $(1 - \varepsilon)P(w) \leq \tilde{P}(w) \leq (1 + \varepsilon)P(w)$, $(1 - \varepsilon)Q(w) \leq \tilde{Q}(w) \leq (1 + \varepsilon)Q(w)$.*
- Then, Alice generates a sample $p \sim \tilde{P}$, Bob generates a sample $q \sim \tilde{Q}$ such that*

$$\Pr[p \neq q] \leq O(\xi).$$

26:30 Rounding via Low Dimensional Embeddings

Proof. Alice and Bob consider, as public randomness, an infinite string of uniform tuples from $\Omega \times [0, 1 + \varepsilon]$, say $(w_1, r_1), (w_2, r_2)$ and so on. Alice picks the first i such that

$$r_i \leq \tilde{P}(w_i),$$

and outputs w_i . Similarly Bob picks the first j such that

$$r_j \leq \tilde{Q}(w_j),$$

and outputs w_j . It is clear that $w_i \sim \tilde{P}$ and $w_j \sim \tilde{Q}$, and we next bound the probability that $i \neq j$. This happens if for the smallest k such that $r_k \leq \max(\tilde{P}(w_k), \tilde{Q}(w_k))$, it holds that $r_k > \min(\tilde{P}(w_k), \tilde{Q}(w_k))$. The probability for that is

$$\begin{aligned} \Pr_{(r,w)} [r > \min(\tilde{P}(w), \tilde{Q}(w)) \mid r \leq \max(\tilde{P}(w), \tilde{Q}(w))] \\ = \frac{\sum_w \Pr_r [\min(\tilde{P}(w), \tilde{Q}(w)) < r < \max(\tilde{P}(w), \tilde{Q}(w))] }{\sum_w \Pr_r [r \leq \max(\tilde{P}(w), \tilde{Q}(w))]} . \end{aligned}$$

We finish the proof by upper bounding the numerator by η , and lower bounding the denominator by $1/4$. The numerator may be bounded as

$$\sum_w |\max(\tilde{P}(w), \tilde{Q}(w)) - \min(\tilde{P}(w), \tilde{Q}(w))| = \sum_w |\tilde{P}(w) - \tilde{Q}(w)| \leq \eta.$$

As for the denominator, it may be lower bounded by

$$\sum_w \Pr_r [r \leq \tilde{P}(w)] = \sum_w \frac{\tilde{P}(w)}{1 + \varepsilon} \geq \sum_w \frac{(1 - \varepsilon)P(w)}{1 + \varepsilon} = \frac{1 - \varepsilon}{1 + \varepsilon} \geq \frac{1}{4}. \quad \blacktriangleleft$$