

On Partitioning Graphs via Single Commodity Flows

Lorenzo Orecchia^{*}
University of California
Berkeley, USA
orecchia@cs.berkeley.edu

Umesh V. Vazirani[‡]
University of California
Berkeley, USA
vazirani@cs.berkeley.edu

Leonard J. Schulman[†]
California Institute of
Technology
Pasadena, USA
schulman@caltech.edu

Nisheeth K. Vishnoi[§]
IBM India Research Lab
New Delhi, India
nisheeth.vishnoi@gmail.com

ABSTRACT

In this paper we obtain improved upper and lower bounds for the best approximation factor for SPARSEST CUT achievable in the cut-matching game framework proposed in Khandekar *et al.* [9]. We show that this simple framework can be used to design combinatorial algorithms that achieve $O(\log n)$ approximation factor and whose running time is dominated by a poly-logarithmic number of single-commodity max-flow computations. This matches the performance of the algorithm of Arora and Kale [2].

Moreover, we also show that it is impossible to get an approximation factor of better than $\Omega(\sqrt{\log n})$ in the cut-matching game framework. These results suggest that the simple and concrete abstraction of the cut-matching game may be powerful enough to capture the essential features of the complexity of SPARSEST CUT.

Categories and Subject Descriptors

F.22 [Analysis of Algorithms and Problem Complexity]: Non-numerical Algorithms and Problems

General Terms

Algorithms, Theory

^{*}Supported by NSF Awards CCF-0635401, CCF-0515304, CCF-0635357 and CCR-0121555.

[†]Supported by NSA Award H98230-06-1-0074, NSF Award CCF-0515342.

[‡]Supported by NSF Award CCF-0635401.

[§]Part of this work was done when author was visiting UC Berkeley and is supported by NSF Grant CCF-0635401.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

STOC'08, May 17–20, 2008, Victoria, British Columbia, Canada.

Copyright 2008 ACM 978-1-60558-047-0/08/05 ...\$5.00.

Keywords

Edge-separator, graph partitioning, single-commodity max-flow, sparsest cut, spectral method, matrix exponential

1. INTRODUCTION

The problem. The SPARSEST CUT problem is the following: Given a graph $G = (V, E)$, find a partition (S, \bar{S}) of the vertices which has the minimum edge expansion. The edge expansion of a cut (S, \bar{S}) is defined to be $\frac{|E(S, \bar{S})|}{\min\{|S|, |\bar{S}|\}}$.

Interest in this problem derives both from its numerous practical applications such as image segmentation, VLSI layout and clustering (see the survey of Shmoys [13]), and from its theoretical connections to spectral methods, linear/semi-definite programming, measure concentration and metric embeddings. In this paper, we provide fast algorithms for this problem using the framework proposed in Khandekar *et al.* [9] (KRV). This framework reduces the NP-hard SPARSEST CUT problem to the computation of a poly-logarithmic number of single-commodity max-flows while, at the same time, keeping the approximation factor poly-logarithmic. We also prove lower bounds on the best approximation factor that can be obtained in this framework.

The cut-matching game. At the heart of the results of KRV lies the following two person game, denoted by $\mathcal{G}(n)$, which starts with an empty graph on n vertices (assume n is even): in each round, the cut player chooses a bisection (S, \bar{S}) of the vertices and in response the matching player chooses a perfect matching that pairs each vertex in S with a unique vertex in \bar{S} . The game ends when the (multi)-graph consisting of the multi-set union of the perfect matchings has edge expansion at least 1, i.e., every subset T of at most $n/2$ vertices has at least $|T|$ edges leaving it. The goal of the cut player is to minimize the number of rounds of play, while the matching player tries to draw the game out for as many rounds as possible. KRV show that any strategy for the cut player that guarantees termination in t rounds yields a $O(t)$ -approximation algorithm for SPARSEST CUT. Moreover the resulting algorithm runs in $O(t(T_c + T_f))$ where T_c is the time to implement the cut player and T_f is the running time of a single commodity max-flow computation on a n -vertex graph. They also give a quasi-linear time implementation of a cut player strategy that achieves $t = O(\log^2 n)$, thereby obtaining an $O(\log^2 n)$

approximation algorithm for SPARSEST CUT whose running time is dominated by $O(\log^2 n)$ single commodity max-flow computations. Subsequently, Arora and Kale [2] gave an algorithm that achieves an approximation ratio of $O(\log n)$ while still running in time dominated by poly-logarithmic single commodity max-flow computations. Their algorithm worked in a more general framework (than the KRV framework) for designing primal-dual algorithms for SDPs.

This work. In this paper we investigate what is the best approximation factor for SPARSEST CUT achievable in this framework of KRV. We start by observing that rather than the number of rounds, a better measure of performance in the cut-matching game is the ratio of the number of rounds to the expansion of the union of matchings produced. This quantity, which we refer to as expansion ratio, upper bounds the approximation ratio achieved by the algorithm. Thus, rather than reducing the number of rounds, we can increase the expansion of the union of matchings. We give two simple modifications of the cut-finding procedure of KRV that result in an expansion of $\Omega(\log n)$, while maintaining number of rounds equal to $O(\log^2 n)$. These result in simple and easy to analyze combinatorial algorithms for SPARSEST CUT that match the Leighton-Rao approximation ratio of $O(\log n)$ [11], while using single-commodity max-flows rather than multi-commodity flows. In theoretical terms, single-commodity max-flows can be computed in time roughly $\tilde{O}(m + n^{1.5})$ by a result of Goldberg and Rao [7]. The central motivation in KRV was to design fast algorithms with provable approximation guarantees for SPARSEST CUT that could compete with practical heuristics such as METIS in running time. They also noted certain similarities of the resulting algorithm with METIS. The algorithms in this paper add to this mix, in terms of the potential for both practical improvements over METIS, as well as a rigorous justification of METIS-like heuristics.

The cut-finding procedures we propose are quite similar to those of KRV. Given a set of matchings, they start with a suitably chosen initial distribution on the vertices, and compute the probability distribution that results from running a random walk on the graph consisting of the union of the matchings. They sort the vertices according to their final probabilities and output the median cut. The difference lies in the kind of random walk. In KRV, the walk mixed along each matching in turn (which we refer to as a round robin walk), whereas our walks more closely resemble the natural random walk on the graph of matchings and allow for more effective bounds on the spectral gap. The analysis of the resulting cut-finding procedures is similar to that of KRV in that it makes use of a potential function related to the mixing of the walk on the current union of matchings: However, bounding the change in potential function is now much trickier. The advantage is that we can exploit the direct connection of the mixing of the walk to the spectral gap of the graph.

Turning to lower bounds, we prove an $\Omega(\sqrt{\log n})$ bound on the rounds to expansion ratio for the cut-matching game. It follows that this framework cannot be used to design an algorithm for SPARSEST CUT with approximation ratio better than $\Omega(\sqrt{\log n})$. By contrast, the best lower bound known for the approach in Arora *et al.* [3] is $\Omega(\log \log n)$ as proved in Devanur *et al.* [6] via an involved and technical argument. This suggests that the cut-matching game provides an attractive, concrete framework in which to study the com-

plexity of finding sparse cuts. An intriguing question here is whether this lower bound is tight and whether there is an efficient cut player strategy that achieves a $O(\sqrt{\log n})$ bound. Towards this goal, it is encouraging to note that the algorithm of Arora *et al.* [3], which achieves a $O(\sqrt{\log n})$ approximation, also outputs, as a certificate of expansion, a union of large matchings.

The proof of the lower bound relies on the following combinatorial statement that is related to the iso-perimetric inequality for the hypercube: Given a bisection of the vertices of the d -dimensional hypercube, there is always a pairing of vertices between the two sides such that the average Hamming distance between paired vertices is at most $O(\sqrt{d})$. As far as we know, this natural property has not been studied, and is worth exploring further.

Finally, we mention that a combinatorial study of the cut-matching game was done in [10], where it was shown that there is a cut-finding procedure, albeit not known to be efficient, which would terminate in $O(\log n)$ rounds. In the same report, the connection between the approach of Arora *et al.* [3] and KRV was made explicit and it was shown that any integrality gap, $\text{gap}(n)$ example for the approach of Arora *et al.* [3] can be used to prove a $\text{gap}(n)$ lower bound on the number of rounds required by any cut-finding procedure.

Organization. The two cut-finding procedures appear in Section 3 and the lower bound appears in Section 4. Proofs of the more technical lemmata are deferred to the Appendix. In the remaining of this paper, we denote by $\phi(G)$ the edge expansion of a graph G , i.e., the minimum edge expansion over all cuts of G . We start by reviewing the framework of KRV.

2. REVIEW OF KRV FRAMEWORK

Certifying expansion. One way to certify that a graph G has no sparse cut is to appeal to the expander flow formalism of Arora *et al.* [3]. This consists of constructing a graph H of known expansion on n vertices and embedding it as a flow in G such that the flow routed through each edge in G is at most unit. Then the expansion of H is a lower bound on the expansion of G . Indeed, the seminal result of Leighton and Rao [11] may be viewed as a multi-commodity flow based algorithm to embed the scaled complete graph $\frac{\alpha}{n \log n} K_n$ in any n -vertex graph G or produce a cut in G of expansion at most α . The expander flow formulation of the $O(\sqrt{\log n})$ factor approximation in Arora *et al.* [3] established that for every graph G , either one can find, using an SDP, a cut of expansion at most α , or there is an $\Omega(\alpha/\sqrt{\log n})$ expander $H(G)$ that one can embed as a flow in it. An efficient implementation of this algorithm was given by Arora, Hazan and Kale [1], showing how to combinatorially solve this SDP in time comparable to that of computing multi-commodity flows.

Converting a cut player strategy into a SPARSEST CUT algorithm. KRV show that a cut player strategy producing an expander in $t(n)$ rounds can be converted into a $O(t(n))$ approximation algorithm using single-commodity max-flows. Here is how: Given a bisection (S_t, \bar{S}_t) , the matching player performs a single-commodity max-flow computation in an attempt to route a perfect matching between S_t and \bar{S}_t with minimum congestion c_t and outputs the matching and the min-cut obtained. It is not hard to show that the expansion of the min-cut is at most $\frac{1}{c_t}$. After $t(n)$

rounds, the algorithm outputs the worst such cut, which has expansion $\frac{1}{c_{\max}}$. At this point, the union of the matchings returned by the matching player has constant expansion and can be embedded in G with congestion equal to $\sum c_t \leq c_{\max} t(n)$. This implies that G must have expansion at least $\Omega\left(\frac{1}{c_{\max} t(n)}\right)$. Hence, the cut that was output is within $O(t(n))$ of the optimal. Moreover, the running time of the algorithm is $O(t(n)(T_c + T_f))$, where T_c is the time required by the cut player to compute his response in each round, and T_f is the time for a single commodity max-flow computation.

The cut-finding procedure \mathcal{C}_{KRV} . Let us recall the cut player strategy \mathcal{C}_{KRV} of KRV. Given a set of perfect matchings (M_1, \dots, M_t) on n vertices, such that their union is not an expander yet, the goal of \mathcal{C}_{KRV} is to find a non-expanding cut. The procedure starts with a random bisection on n vertices, which can be thought of as assigning to a randomly chosen set of half of the vertices a “charge” of 1 and to the other half a charge of -1 . It then iterates through each matching in turn to “mix” the charge across the matchings. At time j , \mathcal{C}_{KRV} averages the charge across the edges of M_j . The total charge is always zero. More precisely, after $j-1$ steps, let \mathbf{u}_{j-1} be the vector whose i -th coordinate indicates the charge on vertex i . Then $\mathbf{u}_j := \left(\frac{I}{2} + \frac{M_j}{2}\right) \mathbf{u}_{j-1}$. After all the t matchings have been processed, \mathcal{C}_{KRV} sorts the vertices by the charge remaining on them, indicated by the coordinates of vector \mathbf{u}_t , and then outputs a bisection whose two parts are the first half and the second half of vertices in this order respectively.

Intuitively, the reason this procedure finds a bisection containing a non-expanding cut is the following: Suppose (S, \bar{S}) is a cut across which there are very few edges, say none. Then the averaging procedure using matching edges never transfers any charge across this cut. Also, the initial random bisection will create a $\Theta(\sqrt{n})$ charge differential between S and \bar{S} , i.e., for all j , $|\sum_{i \in S} \mathbf{u}_j(i) - \sum_{i \in \bar{S}} \mathbf{u}_j(i)| = \Theta(\sqrt{n})$. Hence, in the bisection output after mixing and then sorting based on \mathbf{u}_t , the two sides of the bisection will have non-trivial correlations with S and \bar{S} respectively. Thus, the matching added in the next iteration will add some edges across the sparse cut (S, \bar{S}) . It is remarkable that strategy \mathcal{C}_{KRV} , after $O(\log^2 n)$ of such simplistic steps, is able to ensure that there is no non-expanding cut with constant probability.

3. UPPER BOUNDS

3.1 Preliminaries

Graphs, Random Walks and Eigenvalues. For a symmetric matrix $M \in \mathbb{R}^{n \times n}$, let $L(M) \in \mathbb{R}^{n \times n}$ be its combinatorial Laplacian which is defined as $L(M)_{ii} := \sum_{j \neq i} M_{ij}$, and $L(M)_{ij} := -M_{ij}$ if $i \neq j$. The eigenvalues of a symmetric matrix $M \in \mathbb{R}^{n \times n}$, are denoted by $\lambda_1(M) \leq \dots \leq \lambda_n(M)$. In the following, we identify a graph G with its adjacency matrix and let $L(G)$ denote the combinatorial Laplacian of the graph G . Note that $L(G)$ is symmetric and positive semi-definite. For all graphs G , we then have $\lambda_1(L(G)) = 0$, with corresponding eigenvector $\mathbf{1}$. The following is a fundamental relation between $\lambda_2(L(G))$ and the expansion of a d -regular graph [5].

THEOREM 3.1 (Cheeger’s Bound). *For an undirected d -regular graph G :*

$$\phi(G) \geq \frac{1}{2} \lambda_2(L(G))$$

All graphs considered in this paper are undirected, they may have parallel edges but have no loops, and are assumed to have an even number of vertices. Some of the graphs will have weighted edges. We extend the definition of expansion to these edge-weighted graphs by replacing the cardinality of the cut with the sum of the weights of the edges in the cut. It is also possible to show that Cheeger’s Bound holds in the same form for such weighted graphs, if we let the degree of a vertex be the sum of the weights of the edges incident to it.

In the remainder of this section, we describe and analyze our two new cut player strategies. Before doing so, let us understand why the procedure \mathcal{C}_{KRV} described in the previous section cannot guarantee edge expansion more than $O(1)$ at termination. Suppose the union of matchings (M_1, \dots, M_t) , have no edges crossing some bisection (S, \bar{S}) . Now suppose that the round-robin walk on matchings M_1, \dots, M_t mixes the charges on each side of (S, \bar{S}) perfectly (or very well). The next cut selected by \mathcal{C}_{KRV} is necessarily (S, \bar{S}) . Moreover, presented with any perfect matching crossing \mathcal{C}_{KRV} , when \mathcal{C}_{KRV} mixes across that perfect matching, the charge at each vertex becomes zero, and \mathcal{C}_{KRV} will terminate. So at termination, the edge expansion across bisection (S, \bar{S}) is 1.

3.2 The Cut-Finding Procedure \mathcal{C}_{nat}

After j rounds of the cut-matching game, we have seen matchings M_1, \dots, M_j and the cut player must pick a bisection of the graph $G_j := \sum_{i=1}^j M_i$.

Ideally we would like to analyze the following cut player strategy:

- Pick the initial vector of charges \mathbf{u} as in the previously described cut-finding strategy \mathcal{C}_{KRV} .
- Repeat for t iterations: At the j -th iteration, given the graph G_j , apply t steps of the lazy natural random walk on G_j to \mathbf{u} , i.e., compute $\mathbf{y} := \left(\frac{jI + G_j}{2j}\right)^t \mathbf{u}$.
- Finishing up: As in strategy \mathcal{C}_{KRV} , at every iteration, sort the coordinates of \mathbf{y} and output the median cut.

Can we analyze this new strategy along the lines of the analysis of strategy \mathcal{C}_{KRV} in KRV? The broad outline of the analysis there was to introduce a potential function measuring how “well mixed” the walk is after j iterations, and show that the new matching M_{j+1} significantly reduces this potential function. The round robin nature of the random walk greatly facilitates this analysis, since the process through $j+1$ iterations is exactly the process through j iterations followed by averaging across matching M_{j+1} . Therefore computing the potential reduction in the $j+1$ -th iteration simply boils down to analyzing how much the potential drops due to averaging across matching M_{j+1} .

Things are much more complicated for the lazy natural walk in the procedure sketched above. This is because the process through $j+1$ iterations no longer has such a simple decomposition: all steps of the random walk include the possibility of using edges from the matching M_{j+1} . The

actual analysis of the random walk is based on the matrix inequality $\|(ABA)^k\|_F \leq \|A^k B^k A^k\|_F$, which facilitates the desired decomposition. To make all this work out we have to modify the actual random walk as follows: fix a large d and let $N_t := \frac{d}{d+1}I + \frac{1}{d+1}M_t$. A single step of the modified random walk after t iterations is given by $R_t := N_t N_{t-1} \cdots N_2 N_1 N_2 \cdots N_{t-1} N_t$. The t -th iteration consists of applying this random walk $d+1$ times. Intuitively each step of the random walk R_t corresponds to a step of the natural random walk. It might be possible to show that the analysis of the cut player strategy with this modified random walk directly implies an analysis with the natural random walk. We leave this as an open question.

The random walk \mathcal{C}_{nat} . Let $d+1$ be such that it is equal to 2^k for some integer $k > 0$ to be fixed later. $d+1$ will be the duration of the game. Fix a matching player \mathcal{M} . Let M_t be the perfect matching on $[n]$ output by \mathcal{M} at round t when presented by the cut (S_t, \bar{S}_t) by \mathcal{C}_{nat} . The following probability transition matrices (in $\mathbb{R}^{n \times n}$), defined recursively, are used by the cut-finding procedure: $R_0 := I$. For $1 \leq t \leq d+1$,

$$R_t := N_t R_{t-1} N_t,$$

where $N_t := \frac{d}{d+1}I + \frac{1}{d+1}M_t$. Note that, for all t , R_t is a symmetric doubly stochastic matrix with largest eigenvalue 1 (corresponding to the all 1s eigenvector $\mathbf{1}$). Define:

$$W_t := R_t^{d+1}$$

The cut-finding procedure, \mathcal{C}_{nat} , for step $t+1$ is as follows. Let (M_1, \dots, M_t) be fixed for some $t \geq 0$.

1. Pick a random unit vector $\mathbf{u}_{t+1} \in \mathbb{S}^{n-1} \perp \mathbf{1}$ independent of all previous choices.
2. Compute $\mathbf{y}_{t+1} := W_t \mathbf{u}_{t+1}$.
3. Sort the entries of $\mathbf{y}_{t+1} = (y_1, \dots, y_n)$ as $y_{i_1} \leq \dots \leq y_{i_{n/2}} \leq y_{i_{n/2+1}} \leq \dots \leq y_{i_n}$.
4. Let $S_{t+1} := \{i_1, \dots, i_{n/2}\}$ and $\bar{S}_{t+1} := [n] \setminus S_{t+1}$, and output (S_{t+1}, \bar{S}_{t+1}) .

Analysis

THEOREM 3.2. \mathcal{C}_{nat} achieves an expansion of $\Omega(\log n)$ in $O(\log^2 n)$ rounds against any matching player \mathcal{M} with constant probability.

As in KRV, we measure the progress of \mathcal{C}_{nat} by a simple potential function which is the sum of squares of all the eigenvalues (or Frobenius norm), except the largest one, of the walk matrix used to mix the charge. Formally, if for a symmetric matrix X , we let $\text{Tr}_1[X]$ denote the sum of all the eigenvalues of X except the largest one, then the potential function for the analysis is $\text{Tr}_1[W_t^2] = \sum_{i=1}^{n-1} \lambda_i^2(W_t) = \|W_t - \frac{1}{n} \mathbf{1} \mathbf{1}^T\|_F^2$. Now we analyze the reduction in the potential by comparing the potentials at time t and $t+1$ respectively. Let M_{t+1} be the perfect matching across (S_{t+1}, \bar{S}_{t+1}) obtained at the $(t+1)$ -th step. The following lemma proves that once we fix (M_1, \dots, M_t) , for every perfect matching M_{t+1} across the bisection (S_{t+1}, \bar{S}_{t+1}) , the potential reduces by a multiplicative factor of $1 - \frac{c}{\log n}$, for some fixed constant $c > 0$, in expectation over the choice of \mathbf{u}_{t+1} .

LEMMA 3.3 (POTENTIAL REDUCTION).

Let (M_1, \dots, M_t) be fixed. Then there is a constant $c > 0$ such that, for any perfect matching M_{t+1} across the bisection (S_{t+1}, \bar{S}_{t+1}) (which is selected using \mathbf{u}_{t+1}),

$$\mathbf{E}_{\mathbf{u}_{t+1}} [\text{Tr}_1[W_{t+1}^2]] \leq \left(1 - \frac{c}{\log n}\right) \text{Tr}_1[W_t^2].$$

PROOF. Recall $W_t = R_t^{d+1}$ and $R_{t+1} := N_{t+1} R_t N_{t+1}$. Hence, by the matrix inequality in Theorem A.2 in the Appendix, it follows that

$$\begin{aligned} \text{Tr}_1[W_{t+1}^2] &= \text{Tr}_1 \left[(N_{t+1} R_t N_{t+1})^{2d+2} \right] \\ &\leq \text{Tr}_1 \left[N_{t+1}^{2d+2} R_t^{2d+2} N_{t+1}^{2d+2} \right] = \text{Tr}_1 \left[N_{t+1}^{4d+4} W_t^2 \right]. \end{aligned}$$

Here the last equality follows from Fact A.1 in the Appendix and from the definition of W_t .

Now, N_{t+1}^{4d+4} is just a power of a lazy walk over a matching and it must be equal to a similar lazy walk $(1-\lambda)I + \lambda M_{t+1}$ with constant probability λ of making a transition. A simple calculation, detailed in Fact A.3 in the Appendix, shows that $\lambda \geq \frac{1-e^{-8}}{2}$.

Using this, one can write:

$$\begin{aligned} \text{Tr} \left[N_{t+1}^{4d+4} W_t^2 \right] &= \text{Tr} \left[(I - \lambda(I - M_{t+1})) W_t^2 \right] \\ &= \text{Tr} [W_t^2] - \lambda \text{Tr} [L(M_{t+1}) W_t^2]. \end{aligned}$$

The decrease in potential $\text{Tr}_1 [W_t^2] - \text{Tr}_1 [W_{t+1}^2]$ is

$$\lambda \text{Tr} [L(M_{t+1}) W_t^2] = \lambda \sum_{\{i,j\} \in M_{t+1}} \|\mathbf{w}_i - \mathbf{w}_j\|^2,$$

where \mathbf{w}_i is the i -th row vector of W_t .

Now we need to show that the above potential reduction is large in expectation when M_{t+1} is a matching across the bisection (S_{t+1}, \bar{S}_{t+1}) given by the cut-finding procedure \mathcal{C}_{nat} . Notice that \mathcal{C}_{nat} picks this bisection by projecting the vectors $\{\mathbf{w}_i\}_{i \in V}$ onto a random line \mathbf{u} , in the hope that vectors which are far apart when projected onto \mathbf{u} may also be far apart before projection and hence contribute a large quantity to the potential reduction when matched. This intuition is formalized by the following lemma, which is based on the Gaussian nature of projections and appears in the same form in KRV. The proof follows from the Projection Lemma (Lemma B.2 in the Appendix) applied to $\{v_i := \mathbf{w}_i - \frac{1}{n} \mathbf{1}\}$.

LEMMA 3.4. There is a constant $c' > 0$ such that

$$\mathbf{E}_{\mathbf{u}_{t+1}} \left[\sum_{\{i,j\} \in M_{t+1}} \|\mathbf{w}_i - \mathbf{w}_j\|^2 \right] \geq \frac{c'}{\log n} \cdot \text{Tr}_1 [W_t^2]$$

Notice that $\lambda \geq \frac{1-e^{-8}}{2}$. Hence, it follows from the lemma that

$$\mathbf{E}_{\mathbf{u}_{t+1}} [\text{Tr}_1[W_{t+1}^2]] \leq \left(1 - \frac{(1-e^{-8})c'}{2 \log n}\right) \text{Tr}_1[W_t^2].$$

This completes the proof of Lemma 3.3. \square

This implies that starting with a potential of $n-1$, the potential drops at a rate of roughly $(1-c'/\log n)$ in expectation. Hence, combining the recursive application of the previous lemma with Markov's Inequality, we obtain the following corollary.

COROLLARY 3.5 (TOTAL POTENTIAL).

With constant probability over the choices of $(\mathbf{u}_1, \dots, \mathbf{u}_{d+1})$,

$$\mathrm{Tr}_1[W_{d+1}^2] \leq \left(1 - \frac{c}{\log n}\right)^{d+1} O(n).$$

Hence, for $d = O(\log^2 n)$, the potential drops below $\frac{1}{n}$, and we have that $\lambda_{n-1}(W_{d+1}^2) \leq \frac{1}{n}$.

Now, differently from KRV, we embed R_{d+1} (and not the walk W_{d+1} itself) in G_{d+1} to show that $\phi(G_{d+1}) = \Omega(\log n)$. In particular, we use the bound on $\lambda_{n-1}(W_{d+1}^2)$ to show a lower bound on the spectral gap of the 1-regular weighted graph R_{d+1} . A simple application of Cheeger's Bound then yields the final result.

LEMMA 3.6 (ESTIMATING EXPANSION).

Consider the graph G_{d+1} on the vertex set $[n]$ formed by the union of matchings (M_1, \dots, M_{d+1}) . Then, for $d = O(\log^2 n)$, $\phi(G_{d+1}) \geq \Omega(\log n)$ with constant probability.

PROOF. By the argument above: $\lambda_{n-1}(W_{d+1}^2) = (1 - \lambda_2(L(R_{d+1})))^{2d+2} \leq \frac{1}{n}$. Hence, $\lambda_2(L(R_{d+1})) \geq \Omega\left(\frac{1}{\log n}\right)$.

By Cheeger's Bound, this implies that $\phi(R_{d+1}) = \Omega\left(\frac{1}{\log n}\right)$. Finally, as R_{d+1} embeds in G_{d+1} with congestion $\frac{2}{d+1}$, we have $\phi(G_{d+1}) \geq \frac{d+1}{2} \phi(R_{d+1}) = \Omega(\log n)$. \square

Hence, with constant probability, after $d = O(\log^2 n)$ rounds one can ensure that the graph obtained by the union of matchings (M_1, \dots, M_{d+1}) has expansion at least $\Omega(\log n)$. This completes the proof of Theorem 3.2.

Running time

Note that we do not need to compute W_t explicitly, as we only use W_t to compute $W_t \mathbf{u}_t$. Hence, at iteration t we only need to perform $O(2t \log^2 n)$ matrix-vector multiplications. As each matrix is a step of a lazy random walk along a matching, each of these operations takes time $O(n)$. As t varies from 0 to $O(\log^2 n)$, the total running time throughout the game of the cut-finding procedure \mathcal{C}_{nat} is $O(n \log^6 n) = \tilde{O}(n)$.

3.3 The Cut-Finding Procedure \mathcal{C}_{exp}

Preliminaries: Matrix Exponential

DEFINITION 3.7 (MATRIX EXPONENTIAL).

For a matrix $X \in \mathbb{R}^{n \times n}$, the matrix exponential e^X of X is defined as the following Taylor Series:

$$e^X := \sum_{j=0}^{\infty} \frac{X^j}{j!}$$

If X is symmetric, then X has an ortho-normal basis of eigenvectors and, hence, X and e^X are simultaneously diagonalizable. Thus, for all i : $\lambda_i(e^X) = e^{\lambda_i(X)}$.

A basic identity $e^X e^Y = e^{X+Y}$, that is true for scalar exponentiation, holds for matrix exponentiation only under the condition that X and Y commute. The following theorem shows that something weaker is true in general, and it turns out to be sufficient for our purpose.

THEOREM 3.8 (GOLDEN-THOMPSON INEQUALITY). Let $X, Y \in \mathbb{R}^{n \times n}$ be symmetric matrices. Then, $\mathrm{Tr}[e^{X+Y}] \leq \mathrm{Tr}[e^X e^Y]$.

It should be noted that this is only true for *two* matrices.

Finding the Cut

Fix a matching player \mathcal{M} . Let M_t be the perfect matching on $[n]$ output by \mathcal{M} at round t when presented by the cut (S_t, \bar{S}_t) by \mathcal{C}_{exp} . In the same way \mathcal{C}_{nat} considered the transition matrix of a walk across the matchings to produce the next bisection to present, \mathcal{C}_{exp} makes use of the matrix $W_t := e^{-\frac{1}{2}L(G_t)} = e^{-\frac{1}{2}\sum_{i=1}^t L(M_i)}$. We note that, for all t , W_t is a doubly stochastic positive semi-definite matrix.

The bisection finding procedure for \mathcal{C}_{exp} is exactly the same as that described in Subsection 3.2 for \mathcal{C}_{nat} if we take $W_t = e^{-\frac{1}{2}L(G_t)}$.

Overview. It is worthwhile pointing out that \mathcal{C}_{KRV} can also be seen as using matrix exponentials. In particular, it is not difficult to show that $e^{-\eta(I-M)}$ is the matrix representing a kind of slowed down lazy random walk along matching M . As $\left(\frac{I+M}{2}\right)^i = \frac{(I+M)^i}{2^i}$ for $i \geq 1$, the Taylor Series of the exponential yields: $e^{-\eta(I-M)} = \frac{1+e^{-\eta}}{2}I + \frac{1-e^{-\eta}}{2}M$. Hence, the \mathcal{C}_{KRV} mixing procedure is close to the walk

$$e^{-\eta(I-M_t)} e^{-\eta(I-M_{t-1})} \dots e^{-\eta(I-M_1)}$$

for small η . Our second cut-finding procedure mixes the charge using the matrix $e^{-\eta(tI - (M_1 + \dots + M_t))}$ on the union of matchings to find a non-expanding cut. If it were true that $e^{-\eta(tI - (M_1 + \dots + M_t))} \approx e^{-\eta(I-M_t)} e^{-\eta(I-M_{t-1})} \dots e^{-\eta(I-M_1)}$, then it would establish that the two walks are almost the same. In fact, even under the weaker condition that

$$\mathrm{Tr} \left[e^{-\eta(tI - (M_1 + \dots + M_t))} \right] \leq \mathrm{Tr} \left[e^{-\eta(I-M_t)} \dots e^{-\eta(I-M_1)} \right]$$

we could extend the analysis of \mathcal{C}_{exp} to \mathcal{C}_{KRV} and improve \mathcal{C}_{KRV} 's approximation guarantee. Unfortunately, we cannot show this to be true.

We also note that the random walk in the cut-finding procedure of KRV is best thought of as a "discrete gradient descent" method while the random walk based on matrix exponential updates as a "continuous gradient descent" method with the same "loss" functions, i.e., those defined by the matchings.

Analysis

THEOREM 3.9. \mathcal{C}_{exp} achieves an expansion of $\Omega(\log n)$ in $O(\log^2 n)$ rounds against any matching player \mathcal{M} with constant probability.

The structure of the proof is the same as for \mathcal{C}_{nat} . Using the exponential walk means that we need to apply the Golden-Thompson Inequality to analyze the potential reduction. Moreover, our certificate of expansion is also different. The potential function is again $\mathrm{Tr}_1[W_t^2]$.

LEMMA 3.10 (POTENTIAL REDUCTION).

Let (M_1, \dots, M_t) be fixed. Then there is a constant $c > 0$ such that, for any perfect matching M_{t+1} across the bisection (S_{t+1}, \bar{S}_{t+1}) (which is selected using \mathbf{u}_{t+1}),

$$\mathbf{E}_{\mathbf{u}_{t+1}} [\mathrm{Tr}_1[W_{t+1}^2]] \leq \left(1 - \frac{c}{\log n}\right) \mathrm{Tr}_1[W_t^2].$$

PROOF. We have:

$$\begin{aligned} \mathrm{Tr}[W_{t+1}^2] &= \mathrm{Tr} \left[e^{-\sum_{i=1}^{t+1} L(M_i)} \right] \leq \mathrm{Tr} \left[W_t^2 e^{-L(M_{t+1})} \right] \\ &\leq \mathrm{Tr} \left[W_t^2 \right] - \left(\frac{1 - e^{-2}}{2} \right) \mathrm{Tr} \left[W_t^2 L(M_{t+1}) \right]. \end{aligned}$$

The first inequality follows from the Golden-Thompson Inequality (3.8) and the second inequality follows from Fact A.4 in the Appendix. Hence, the decrease in potential at round $t + 1$ is:

$$\mathrm{Tr}_1 [W_t^2] - \mathrm{Tr}_1 [W_{t+1}^2] \leq \left(\frac{1 - e^{-2}}{2} \right) \mathrm{Tr} [W_t^2 L(M_{t+1})].$$

Let \mathbf{w}_i denote the i -th row of W_t . Then:

$$\mathrm{Tr}_1 [W_t^2] - \mathrm{Tr}_1 [W_{t+1}^2] \leq \left(\frac{1 - e^{-2}}{2} \right) \sum_{\{i,j\} \in M_{t+1}} \|\mathbf{w}_i - \mathbf{w}_j\|^2.$$

Now we need to show that the above potential reduction is large. As in the analysis of $\mathcal{C}_{\mathrm{nat}}$, this is a simple application of the Lemma B.2 with $\{\mathbf{v}_i := \mathbf{w}_i - \frac{1}{n}\mathbf{1}\}$. We get:

$$\begin{aligned} \mathbf{E}_{\mathbf{u}_{t+1}} [\mathrm{Tr}_1 [W_{t+1}^2]] &\leq \mathrm{Tr}_1 [W_t^2] \left(1 - \left(\frac{c'(1 - e^{-2})}{2 \log n} \right) \right) \\ &= \mathrm{Tr}_1 [W_t^2] \left(1 - \Omega \left(\frac{1}{\log n} \right) \right) \end{aligned}$$

□

By recursively applying the lemma above and by Markov's Inequality, we obtain the following.

COROLLARY 3.11 (TOTAL POTENTIAL). *With constant probability over the choices of $(\mathbf{u}_1, \dots, \mathbf{u}_d)$,*

$$\mathrm{Tr}_1 [W_d^2] \leq \left(1 - \frac{c}{\log n} \right)^d O(n).$$

LEMMA 3.12 (ESTIMATING EXPANSION). *Consider the graph G_d on the vertex set $[n]$ formed by the union of matchings (M_1, \dots, M_d) . Then, for $d = O(\log^2 n)$,*

$$\phi(G_d) = \Omega(\log n)$$

with constant probability.

PROOF. Let G_d be the union of matchings (M_1, \dots, M_d) obtained through the cut-finding procedure $\mathcal{C}_{\mathrm{exp}}$. From Corollary 3.11 it follows that for $d = O(\log^2 n)$, $\lambda_{n-1}(W_d) = \lambda_{n-1}(e^{-L(G_d)}) \leq \frac{1}{n}$. By Theorem 3.1, $\phi(G_t) \geq \frac{\lambda_2(L(G_t))}{2}$. Hence:

$$\phi(G_d) \geq \frac{\lambda_2(G_d)}{2} \geq \frac{-\log(\lambda_{n-1}(e^{-L(G_d)}))}{2} = \frac{\log n}{2}.$$

□

This completes the proof of Theorem 3.9.

Running time

We approximate the exponential $e^{-\frac{1}{2}L(G_t)}$ by truncating its Taylor series as it is also done in Arora and Kale [2]. All we need to compute is $e^{-\frac{1}{2}L(G_t)} \mathbf{u}$ for some random unit vector \mathbf{u} . We define the approximation \mathbf{v}_k as

$$\mathbf{v}_k := \sum_{j=0}^k \frac{1}{j!} \left(\frac{L(G_t)}{2} \right)^j \mathbf{u}.$$

For \mathbf{v}_k to be a good approximation for the purposes of the algorithm we require that

$$\|\mathbf{v}_k - e^{-\frac{1}{2}L(G_t)} \mathbf{u}\|^2 \leq O \left(\frac{\mathrm{Tr}_1 [W_t^2]}{\log n} \right).$$

Arora and Kale [2] show that this inequality is satisfied for

$$k \geq \max \left\{ \frac{e^2}{2} t, \Omega(\log n) \right\}.$$

Finally, to compute \mathbf{v}_k we just need to perform k matrix-vector multiplications. Each of these takes time $O(tn)$ as G_t is t regular. As $t \leq O(\log^2 n)$ and by the bound on k , we get a running time of $O(n \log^6 n) = \tilde{O}(n)$ for the procedure $\mathcal{C}_{\mathrm{exp}}$.

4. LOWER BOUND

In this section we prove the following lower bound for the cut-matching game.

THEOREM 4.1. *There is a matching player \mathcal{M}^* such that the expansion ratio for any cut player is at least $\Omega(\sqrt{\log n})$.*

This result establishes a limit on the upper bound arguments in the framework of the cut-matching game, i.e., that there do not exist cut players, even computationally unbounded ones, which can achieve a rounds to expansion ratio better than $\Omega(\sqrt{\log n})$ against \mathcal{M}^* . In this section we focus on the description and the analysis of \mathcal{M}^* as claimed by the theorem.

Main Lemma. The main technical result of this section is the following claim about the hypercube. The lower bound is a simple consequence of this result.

LEMMA 4.2 (MAIN LEMMA). *Given a unit embedding of the d -dimensional hypercube $\left(\frac{\pm 1}{\sqrt{d}}, \frac{-1}{\sqrt{d}} \right)^d$, for any bisection of its vertices, there exists a perfect matching of the vertices across the bisection such that the average ℓ_2 -distance of matched vertices is $O\left(\frac{1}{\sqrt{d}}\right)$.*

Notice that here a matching just indicates a pairing of the vertices and has no relation with the edges of the hypercube graph.

The Main Lemma will be shown to be a consequence of the vertex iso-perimetry of the hypercube. Intuitively, if the vertex iso-perimetry is large, no two large sets of vertices can be a large distance apart. The proof shows how to apply the same reasoning to show the existence of a “short” perfect matching across any bisection. Moreover, as the ℓ_2 distance is just a scaling of the Hamming distance on the hypercube, a corollary of the lemma is the following combinatorial statement.

COROLLARY 4.3. *For any bisection of the vertex set of the d -dimensional hypercube, there exists a perfect matching of the vertices across the bisection such that the average Hamming distance between matched vertices is $O(\sqrt{d})$.*

This simple combinatorial fact, which, to the best of our knowledge, was not known before, opens a number of questions about the existence of similar, possibly stronger statements about the hypercube or other polytopes. We discuss this in Section 5.

4.1 Proof Idea

A matching player \mathcal{M} could keep the game $\mathcal{G}(n)$ going for $\Omega(\sqrt{\log n})$ rounds if, at each of the first $\Omega(\sqrt{\log n})$ rounds, \mathcal{M} is able to exhibit a cut with expansion less than $\frac{1}{4}$ in the graph formed by the union of the matchings thus far.

A simple way for \mathcal{M} to do this would be to pick a fixed cut (D, \overline{D}) at the beginning of the game and keep this cut as sparse as possible round after round. However, if the cut player guesses one bisection containing or equal to D , any perfect matching that \mathcal{M} adds across this bisection will make the cut (D, \overline{D}) have constant expansion immediately.

To overcome this problem, the matching player \mathcal{M}^* first identifies the vertex set $[n]$ with the vertex set of a hypercube with d coordinates, $\{-1, 1\}^d$. (Assume $n = 2^d$.) Then, rather than trying to keep one bisection sparse, it tries to keep $d = \log n$ “orthogonal” bisections sparse on an average. The natural choice for such orthogonal bisections for the hypercube vertex set are those induced by the coordinate cuts. Formally, denote this set of bisections by $\mathcal{D} := \{(D_1, \overline{D}_1), \dots, (D_d, \overline{D}_d)\}$. Here, $D_i := \{(x_1, \dots, x_d) \in \{-1, 1\}^d \mid x_i = 1\}$, and $\overline{D}_i := \{-1, 1\}^d \setminus D_i$. The orthogonality makes it possible to add edges across one (D_i, \overline{D}_i) without increasing the expansion of other bisections in \mathcal{D} by too much. The Main Lemma formalizes this intuition by showing that at any iteration t , and for any choice of a bisection (S_t, \overline{S}_t) (by any cut player), there exists a matching M_t across (S_t, \overline{S}_t) which increases the average expansion over \mathcal{D} by at most $O\left(\frac{1}{\sqrt{\log n}}\right)$. This implies that $\Omega(\sqrt{\log n})$ rounds are necessary before the union of matchings has a constant expansion.

To establish this lemma, we first encode the task of finding a matching across the given bisection with minimum ℓ_2^2 length as a min-cost perfect-matching LP. Then, we show that the dual of this LP can be interpreted as a non-expanding embedding of the hypercube into ℓ_1 . This allows us to use the hypercube vertex iso-perimetry to upper bound its optimal value.

4.2 Preliminaries

Cut vectors

For any cut (S, \overline{S}) of $[n]$, we define the cut vector $\vec{x}_S \in \mathbb{R}^n$ by:

$$(\vec{x}_S)_i = \begin{cases} +1 & \text{if } i \in S \\ -1 & \text{if } i \notin S \end{cases}$$

Hence, for any cut (S, \overline{S}) :

$$\vec{x}_S^\top L(G) \vec{x}_S = 4|E(S, \overline{S})|.$$

Vertex iso-perimetry of the hypercube

For any graph $G = (V, E)$, let $\gamma(G)$ denote the vertex iso-perimetry number of G . $\gamma(G)$ is the minimum ratio among all cuts $S \subseteq V$, with $|S| \leq \frac{|V|}{2}$, of the number of neighbors of S outside of S to that of the size of S . That is

$$\gamma(G) := \min_{S \subseteq V, |S| \leq \frac{|V|}{2}} \frac{|\{i \in V \setminus S : \exists j \in S : \{i, j\} \in E\}|}{|S|}.$$

The following is a standard fact about the vertex iso-perimetry of the hypercube [5].

$$\text{FACT 4.4. } \gamma(H_d) = \Theta\left(\frac{1}{\sqrt{d}}\right).$$

4.3 Proof of Theorem 4.1

Let $n := 2^d$ for a positive integer d . Let H_d denote the d -dimensional hypercube. This is the graph with $V(H_d) := \{-1, 1\}^d$ and $\{i, j\} \in E(H_d)$ if and only if i and j differ

in exactly one coordinate. At the start of the game, \mathcal{M}^* picks an arbitrary bijection $f : V \rightarrow H_d$. Let U_d be the unit embedding of H_d , i.e., $U_d := \frac{H_d}{\sqrt{d}}$ and, for all $v \in V$, denote by \mathbf{u}_v the point $\frac{f(v)}{\sqrt{d}}$ of U_d . Each dimension cut in H_d corresponds to a cut in V through the mapping f . In particular, we denote by D_i the cut $\{v \in V : f(v)_i = +1\}$, and $\overline{D}_i := V \setminus D_i$. This defines a set $\mathcal{D} := \{D_1, \dots, D_d\}$ of bisections of V .

Fix an arbitrary cut player \mathcal{C} which at every round presents a bisection to the matching player \mathcal{M}^* to which \mathcal{M}^* must add a perfect matching. Let $G_t := ([n], E_t)$ denote the graph formed by the union of matchings M_1, \dots, M_t for some integer $t \geq 0$. $G_0 := ([n], \emptyset)$. Define a potential function $\Phi_t := \mathbf{E}_{D_i \leftarrow \mathcal{D}} \left[\frac{E_t(D_i, \overline{D}_i)}{|D_i|} \right]$ to be the expected expansion in G_t of a cut sampled uniformly at random from \mathcal{D} . Note that the matching player \mathcal{M}^* is not random. These expectations are just averages and it is convenient to use this language for the proof.

$$\text{FACT 4.5. } \Phi_t = \frac{1}{2n} \sum_{\{h,k\} \in E_t} \|\mathbf{u}_h - \mathbf{u}_k\|^2$$

PROOF.

$$\begin{aligned} \Phi_t &= \mathbf{E}_{D_i \leftarrow \mathcal{D}} \left[\frac{|E_t(D_i, \overline{D}_i)|}{|D_i|} \right] = \mathbf{E}_{D_i \leftarrow \mathcal{D}} \left[\frac{\vec{x}_{D_i}^\top L(G_t) \vec{x}_{D_i}}{4|D_i|} \right] \\ &= \frac{1}{d} \sum_{i=1}^d \frac{\vec{x}_{D_i}^\top L(G_t) \vec{x}_{D_i}}{2n} \\ &= \frac{1}{d} \sum_{i=1}^d \frac{\sum_{\{h,k\} \in E_t} ((\vec{x}_{D_i})_h - (\vec{x}_{D_i})_k)^2}{2n} \\ &= \frac{1}{2n} \sum_{\{h,k\} \in E_t} \sum_{i=1}^d \left(\frac{1}{\sqrt{d}} (\vec{x}_{D_i})_h - \frac{1}{\sqrt{d}} (\vec{x}_{D_i})_k \right)^2 \\ &= \frac{1}{2n} \sum_{\{h,k\} \in E_t} \|\mathbf{u}_h - \mathbf{u}_k\|^2. \end{aligned}$$

Notice that in the last inequality we used the definition of the cuts D_1, \dots, D_d as the coordinate cuts of H_d . \square

This shows that Φ_t equals the sum of the squared length of the edges of G_t in the hypercube representation U_d of V , scaled by $2n$. Hence, for any $t \geq 1$, we can rewrite the increase in potential at round t as:

$$\Phi_t - \Phi_{t-1} = \sum_{\{i,j\} \in E_t \setminus E_{t-1}} \|\mathbf{u}_i - \mathbf{u}_j\|^2 = \sum_{\{i,j\} \in M_t} \|\mathbf{u}_i - \mathbf{u}_j\|^2$$

At every iteration t , given \mathcal{C} 's choice of (S_t, \overline{S}_t) , \mathcal{M}^* adds the matching M_t across (S_t, \overline{S}_t) which minimizes $\sum_{\{i,j\} \in M_t} \|\mathbf{u}_i - \mathbf{u}_j\|^2$. This only requires a minimum cost matching computation on the complete bipartite graph induced by (S_t, \overline{S}_t) . The Main Lemma (Lemma 4.2) can now be restated as follows.

LEMMA 4.6 (MAIN LEMMA). *For all bisections (S, \overline{S}) , there exists a perfect matching M across (S, \overline{S}) such that $\sum_{\{i,j\} \in M} \|\mathbf{u}_i - \mathbf{u}_j\|^2 = O\left(\frac{n}{\sqrt{d}}\right)$.*

The proof of this Lemma will be given at the end of the Section. Here we see how the Main Lemma implies the theorem.

PROOF OF THEOREM 4.1. By the Main Lemma, the potential increase in one round is: $\Phi_t - \Phi_{t-1} = O\left(\frac{1}{\sqrt{d}}\right)$. Hence $\Phi_t = O\left(\frac{t}{\sqrt{d}}\right)$. This implies that $\mathbf{E}_{D_i \leftarrow \mathcal{D}} \left[\frac{|E_t(D_i, \overline{D_i})|}{|D_i|} \right] = O\left(\frac{t}{\sqrt{d}}\right)$. Hence, there exists a cut D_i with $\frac{|E_t(D_i, \overline{D_i})|}{|D_i|} = O\left(\frac{t}{\sqrt{d}}\right)$. This shows that $\phi(G_t) = O\left(\frac{t}{\sqrt{d}}\right)$. As the game stops after t iterations only if $\phi(G_t)$ is constant, we must have $t = \Omega\left(\sqrt{d}\right) = \Omega\left(\sqrt{\log n}\right)$. Moreover, this also shows that $\frac{t}{\Phi(G_t)} = \Omega\left(\sqrt{d}\right) = \Omega\left(\sqrt{\log n}\right)$ for all t . \square

We now proceed to prove the Main Lemma.

PROOF OF MAIN LEMMA 4.6. Let $c_{ij} := \|\mathbf{u}_i - \mathbf{u}_j\|^2$. Consider the LP relaxation of Figure 1 for computing the minimum cost perfect matching across the cut (S, \overline{S}) .

$$\begin{aligned} & \text{Minimize} && \sum_{i \in S, j \in \overline{S}} c_{ij} x_{ij} \\ \text{Subject to} &&& \\ & \forall i \in S, && \sum_{j \in \overline{S}} x_{ij} = 1 \\ & \forall j \in \overline{S}, && \sum_{i \in S} x_{ij} = 1 \\ & \forall i \in S, j \in \overline{S}, && x_{ij} \geq 0 \end{aligned}$$

Figure 1: LP for Bipartite Min-Cost Matching

By the integrality of the bipartite perfect matching polytope (see [12]), the objective of this program is the minimum of $\sum_{\{i,j\} \in M} \|\mathbf{u}_i - \mathbf{u}_j\|^2$ over all perfect matchings M across (S, \overline{S}) . In Figure 2 we consider a formulation of the dual of this LP.

$$\begin{aligned} & \text{Maximize} && \sum_{i \in S} y_i - \sum_{j \in \overline{S}} y_j \\ \text{Subject to} &&& \\ & \forall i \in S, j \in \overline{S}, && y_i - y_j \leq c_{ij} \\ & \forall i \in V, && y_i \in \mathbb{R} \end{aligned}$$

Figure 2: The dual of the LP for Bipartite Min-Cost Matching

A feasible solution for this LP can be seen an embedding $\{y_i\}_{i \in [n]}$ of $[n]$ on the real line such that no pair i, j with $i \in S$ and $j \in \overline{S}$ and $y_i \geq y_j$ can be further away in ℓ_1 distance than its ℓ_2^2 distance in the hypercube embedding U_d . We now prove the following two properties of solutions to the dual LP:

1. If $\{y_i\}_{i \in [n]}$ is a feasible solution of value Y , then for any $c \in \mathbb{R}$, $\{y'_i = y_i + c\}_{i \in [n]}$ is a feasible solution of value $Y' = Y$.
2. In any optimal dual solution, we must have, for all pairs $i, j \in [n]$, $|y_i - y_j| \leq c_{ij} = \|\mathbf{u}_i - \mathbf{u}_j\|^2$.

Proof of Property 1: The shifted solution is feasible as for all $i \in S, j \in \overline{S}$:

$$y'_i - y'_j = y_i + c - y_j - c = y_i - y_j \leq c_{ij}$$

The value of this solution is:

$$\begin{aligned} Y' &= \sum_{i \in S} y'_i - \sum_{j \in \overline{S}} y'_j = \sum_{i \in S} (y_i + c) - \sum_{j \in \overline{S}} (y_j + c) \\ &= \sum_{i \in S} y_i + \frac{cn}{2} - \sum_{j \in \overline{S}} y_j - \frac{cn}{2} = \sum_{i \in S} y_i - \sum_{j \in \overline{S}} y_j = Y \end{aligned}$$

Proof of Property 2: Notice that the costs c_{ij} 's respect the triangle inequality as the ℓ_2^2 -distance on the hypercube is a metric. To prove the statement, we need to handle the three remaining cases:

1. $i \in S, j \in \overline{S}$. Assume $y_i \geq y_j$ without loss of generality. As the solution is optimal, it is not possible to increase y_j to obtain a larger dual objective. This implies that there must exist $k \in \overline{S}$ such that $y_j - y_k = c_{jk}$. But we must have $c_{ik} \geq y_i - y_k = (y_i - y_j) + (y_j - y_k) = y_i - y_j + c_{jk}$. As $c_{ik} \leq c_{ij} + c_{jk}$, we have $y_i - y_j \leq c_{ij}$.
2. $i \in \overline{S}, j \in \overline{S}$. This is handled as in the previous case.
3. $i \in S, j \in \overline{S}$ such that $y_j \geq y_i$. Because the solution is optimal there must exist $j' \in S$ and $i' \in \overline{S}$ such that $y_{j'} - y_j = c_{j'j}$ and $y_i - y_{i'} = c_{ii'}$. But, by the dual constraint, we must have $c_{i'j'} \geq y_{j'} - y_{i'} = (y_{j'} - y_j) + (y_j - y_i) + (y_i - y_{i'}) = c_{j'j} + y_j - y_i + c_{ii'}$. By triangle inequality, $c_{i'j'} \leq c_{j'j} + c_{ji} + c_{ii'}$, so that $y_j - y_i \leq c_{ij}$ as required.

Application of vertex iso-perimetry of H_d : Now we use these properties of an optimal dual solution together with the vertex iso-perimetry of the hypercube to obtain an upper bound on the dual optimal. By Property 1, we can translate any optimal dual solution preserving optimality. Hence, we may consider an optimal solution $\{y_i\}_{i \in [n]}$ such that at most $\frac{n}{2}$ vertices are mapped to positive value and at most $\frac{n}{2}$ are mapped to negative values. Notice that, as $\max_{i,j} \|\mathbf{u}_i - \mathbf{u}_j\|^2 = 4$, we have $y_k \in [-4, 4]$ for all $k \in [n]$. Now define sets $R_1, \dots, R_{4d} \subseteq [n]$ as follows:

$$R_i := \left\{ k \in [n] : y_k \in \left(\frac{i-1}{d}, \frac{i}{d} \right) \right\}.$$

Similarly, for the negative side we can define L_1, \dots, L_{4d} :

$$L_i := \left\{ k \in [n] : y_k \in \left[-\frac{i}{d}, -\frac{i-1}{d} \right) \right\}.$$

We also define $A_i := \bigcup_{k=i}^{4d} R_k$ and $B_i := \bigcup_{k=i}^{4d} L_k$. By our assumption on $\{y_i\}_{i \in [n]}$ we know that, for all i , $|A_i|, |B_i| \leq \frac{n}{2}$.

Consider now any $k \in A_i$ for $i \geq 2$. Consider any $h \notin A_i$ such that $\|\mathbf{u}_h - \mathbf{u}_k\|^2 = \frac{1}{d}$, i.e., \mathbf{u}_h is a neighbor of \mathbf{u}_k in the hypercube graph. Notice that h must lie in L_{i-1} , as, by Property 2, $|y_k - y_h| \leq \frac{1}{d}$ and $h \notin A_i$. Hence, all vertices which are outside of A_i and adjacent to A_i in the hypercube must belong to L_{i-1} . Because $|A_i| \leq \frac{n}{2}$, by the vertex iso-perimetry of the hypercube, there are at least $\gamma(H_d)|A_i|$ such vertices and, for $i \geq 2$:

$$|L_{i-1}| \geq \Omega\left(\frac{1}{\sqrt{d}}\right) |A_i|.$$

This implies that for $i \geq 2$,

$$|A_{i-1}| \geq \left(1 + \Omega\left(\frac{1}{\sqrt{d}}\right)\right) |A_i|.$$

Since $|A_1| \leq \frac{n}{2}$,

$$|A_i| \leq \frac{n}{2} \left(1 + \Omega\left(\frac{1}{\sqrt{d}}\right)\right)^{-(i-1)}.$$

The same reasoning can be applied to B_i to deduce that

$$|B_i| \leq \frac{n}{2} \left(1 + \Omega\left(\frac{1}{\sqrt{d}}\right)\right)^{-(i-1)}.$$

Now notice that the cost of the dual solution $\{y_k\}_{k \in [n]}$ is upper bounded by

$$\begin{aligned} \frac{1}{d} \left(\sum_{i=1}^{4d} i|L_i| + \sum_{i=1}^{4d} i|R_i| \right) &\leq \frac{1}{d} \left(\sum_{i=1}^{4d} |A_i| + \sum_{i=1}^{4d} |B_i| \right) \\ &= \frac{n}{d} \sum_{i=1}^{4d} \left(1 + \Omega\left(\frac{1}{\sqrt{d}}\right)\right)^{-(i-1)} \\ &= \frac{n}{d} O(\sqrt{d}) = O\left(\frac{n}{\sqrt{d}}\right). \end{aligned}$$

But, by strong duality, the primal optimum equals the the dual optimum. Hence, we complete the proof by noticing that

$$\sum_{\{i,j\} \in M} \|\mathbf{u}_i - \mathbf{u}_j\|^2 = O\left(\frac{n}{\sqrt{d}}\right).$$

□

5. OPEN PROBLEMS

The main remaining open question is whether it is possible to construct a cut player running in almost linear time and achieving an expansion ratio of $O(\sqrt{\log n})$. Another direction of study is to show that the natural random walk which inspires \mathcal{C}_{nat} can actually work without the modifications required by our analysis. It may be possible to show that the natural random walk approximates the exponential random walk sufficiently well for the same analysis to apply. Along the same lines, it is, for example, possible to see \mathcal{C}_{nat} as an approximation to \mathcal{C}_{exp} . We will include the details in the full version of the paper.

Finally, the Main Lemma used in the proof of the lower bound raises the following interesting questions, which, to our knowledge, are unresolved: What is the bisection of the hypercube which maximizes the minimum over all pairings of vertices across the bisection of the average Hamming distance between paired vertices? Is it possible to derive similar ℓ_2^2 -distance results for polytopes other than the hypercube?

Acknowledgments

We would like to thank Satish Rao for several illuminating discussions and for sharing his insights into the realm of graph partitioning.

6. REFERENCES

- [1] S. Arora, E. Hazan, and S. Kale. $O(\sqrt{\log n})$ approximation to sparsest cut in $\tilde{O}(n^2)$ time. *Proceedings, IEEE Symposium on Foundations of Computer Science*, 00:238–247, 2004.
- [2] S. Arora and S. Kale. A combinatorial, primal-dual approach to semidefinite programs. In *STOC '07: Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 227–236, New York, NY, USA, 2007. ACM.
- [3] S. Arora, S. Rao, and U. Vazirani. Expander flows, geometric embeddings and graph partitioning. In *STOC '04: Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, pages 222–231, New York, NY, USA, 2004. ACM.
- [4] R. Bhatia. *Matrix Analysis (Graduate Texts in Mathematics)*. Springer, 1996.
- [5] F. R. Chung. *Spectral Graph Theory (CBMS Regional Conference Series in Mathematics, No. 92)*. American Mathematical Society, 1997.
- [6] N. R. Devanur, S. A. Khot, R. Saket, and N. K. Vishnoi. Integrality gaps for sparsest cut and minimum linear arrangement problems. In *STOC '06: Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 537–546, New York, NY, USA, 2006. ACM.
- [7] A. V. Goldberg and S. Rao. Beating the flow decomposition barrier. *Journal of the ACM (JACM)*, 45:783–797, 1998.
- [8] G. Karypis and V. Kumar. A fast and high quality multilevel scheme for partitioning irregular graphs. *SIAM Journal on Scientific Computing*, 20:359–392, 1999.
- [9] R. Khandekar, S. Rao, and U. Vazirani. Graph partitioning using single commodity flows. In *STOC '06: Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 385–390, New York, NY, USA, 2006. ACM.
- [10] R. M. Khandekar, S. Khot, L. Orecchia, and N. K. Vishnoi. On a cut-matching game for the sparsest cut problem. Technical Report EECS-2007-177, EECS Department, University of California, Berkeley, CA, 2007.
- [11] T. Leighton and S. Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *J. ACM*, 46(6):787–832, 1999.
- [12] C. H. Papadimitriou and K. Steiglitz. *Combinatorial Optimization: Algorithms and Complexity*. Prentice-Hall, 1982.
- [13] D. Shmoys. Cut problems and their application to divide and conquer. In D. Hochbaum, editor, *Approximation algorithms for NP-hard problems*, pages 192–235. PWS Publishing Co., Boston, MA, USA, 1996.

APPENDIX

A. MATRIX INEQUALITIES

The following facts follow from the definition of the trace function. They are proved in [4].

FACT A.1 (CYCLIC SHIFT INVARIANCE OF TRACE).
Let $X, Y \in \mathbb{R}^{n \times n}$. Then, $\text{Tr}[XY] = \text{Tr}[YX]$.

THEOREM A.2 (SYMMETRIC REARRANGEMENT).
Let $X, Y \in \mathbb{R}^{n \times n}$ be symmetric matrices. Then for any positive integer k ,

$$\text{Tr} \left[(XYX)^{2^k} \right] \leq \text{Tr} \left[X^{2^k} Y^{2^k} X^{2^k} \right].$$

The following claim uses the fact that, for a matching M , $M^2 = I$ to simplify the expression for a power of a lazy random walk across a single matching.

FACT A.3. *Let $\lambda := \frac{1}{2} - \frac{1}{2} \left(1 - \frac{2}{d+1}\right)^{4d+4}$. Then, $N_{t+1}^{4d+4} = I - \lambda(I - M_{t+1})$.*

PROOF. For this proof, let $N := N_{t+1}$, $M := M_{t+1}$. Note that since $M^2 = I$, $(I - M)^2 = 2(I - M)$. Hence, $(I - M)^j = 2^{j-1}(I - M)$ for $j \geq 1$. We use this observation below to complete the proof of the fact.

$$\begin{aligned} N^{4d+4} &= \left(I - \frac{1}{d+1}(I - M)\right)^{4d+4} \\ &= I + \sum_{j=1}^{4d+4} (-1)^j \frac{1}{(d+1)^j} \binom{4d+4}{j} (I - M)^j \\ &= I + \frac{I - M}{2} \sum_{j=1}^{4d+4} (-1)^j \frac{2^j}{(d+1)^j} \binom{4d+4}{j} \\ &= I - \frac{I - M}{2} \left(1 - \left(1 - \frac{2}{d+1}\right)^{4d+4}\right) \end{aligned}$$

□

The following fact is a matrix version of the simple scalar inequality $e^{-\varepsilon x} \leq (1 - (1 - e^{-\varepsilon})x)$ for all $x \in [0, 1]$ and $\varepsilon \in (0, 1)$.

FACT A.4. *Let $X \in \mathbb{R}^{n \times n}$ be symmetric matrix such that $0 \leq X \leq I$ and $\varepsilon \in (0, 1)$. Then,*

$$e^{-\varepsilon X} \preceq (I - (1 - e^{-\varepsilon})X).$$

B. PROJECTION LEMMA

The results in this section essentially appear in KRV. We include them here for completeness.

FACT B.1 (GAUSSIAN BEHAVIOR OF PROJECTIONS).

If \mathbf{v} is a vector of length l in \mathbb{R}^m and \mathbf{u} is a random vector in \mathbb{S}^{m-1} . Then

1. $\mathbf{E}_{\mathbf{u}} [\langle \mathbf{v}, \mathbf{u} \rangle^2] = \frac{l^2}{m}$,
2. For $x \leq m/16$, $\Pr_{\mathbf{u}} [\langle \mathbf{v}, \mathbf{u} \rangle^2 \geq xl^2/m] \leq e^{-x/4}$.

LEMMA B.2. *Let $\{\mathbf{v}_i\}_{i=1}^n$ be vectors in \mathbb{R}^{n-1} such that*

1. $\sum_i \mathbf{v}_i = 0$.
2. For all i , $\|\mathbf{v}_i\|^2 \leq B$ for some constant B .
3. $\Phi := \sum \|\mathbf{v}_i\|^2 \geq \frac{1}{\text{poly}(n)}$.

Let \mathbf{r} be a random unit vector in \mathbb{R}^{n-1} and for all i set $u_i := \langle \mathbf{v}_i, \mathbf{r} \rangle$. Let S be the partition of $[n]$ such $|S| = n/2$ and for all $i \in S$ and $j \in \bar{S}$ $u_i \geq u_j$. Consider any matching M of the indices $[n]$ across (S, \bar{S}) . Then,

$$\mathbf{E}_{\mathbf{r}} \left[\sum_{\{i,j\} \in M} \|\mathbf{v}_i - \mathbf{v}_j\|^2 \right] = \Omega \left(\frac{1}{\log n} \right) \Phi.$$

PROOF. Define the event

$$\mathcal{E}_{ij} := \left\{ (u_i - u_j)^2 \leq \frac{c \log n}{n-1} \|\mathbf{v}_i - \mathbf{v}_j\|^2 \right\}$$

for some constant $c > 0$. Let $\mathcal{E} := \bigcap_{i,j} \mathcal{E}_{ij}$. By the Fact B.1 we have that $\Pr[\bar{\mathcal{E}}_{ij}] \leq n^{-c/4}$. Hence, by a union bound, $\Pr[\bar{\mathcal{E}}] \leq n^{-c/4+2}$. Then:

$$\begin{aligned} \mathbf{E}_{\mathbf{r}} \left[\sum_{\{i,j\} \in M} \|\mathbf{v}_i - \mathbf{v}_j\|^2 \right] &\geq \\ \frac{n-1}{c \log n} \mathbf{E}_{\mathbf{r}} \left[\sum_{\{i,j\} \in M} (u_i - u_j)^2 \mid \mathcal{E} \right] &\cdot \Pr[\mathcal{E}]. \end{aligned}$$

Let a be the real number such that $u_i \geq a \geq u_j$ for all $i \in S, j \in \bar{S}$. Hence, for $\mathbf{u} := (u_1, \dots, u_n)$, $\sum_{i=1}^n u_i = \sum_{i=1}^n \langle \mathbf{v}_i, \mathbf{r} \rangle = \langle (\sum_{i=1}^n \mathbf{v}_i), \mathbf{r} \rangle = 0$. Hence,

$$\begin{aligned} \sum_{\{i,j\} \in M} (u_i - u_j)^2 &\geq \sum_{i=1}^n (u_i - a)^2 \\ &= \|\mathbf{u}\|^2 - 2a \left(\sum_{i=1}^n u_i \right) + na^2 \\ &= \|\mathbf{u}\|^2 + na^2 \geq \|\mathbf{u}\|^2. \end{aligned}$$

Hence, we have

$$\mathbf{E}_{\mathbf{r}} \left[\sum_{\{i,j\} \in M} \|\mathbf{v}_i - \mathbf{v}_j\|^2 \right] \geq \frac{n-1}{c \log n} \mathbf{E}_{\mathbf{r}} [\|\mathbf{u}\|^2 \mid \mathcal{E}] \cdot \Pr[\mathcal{E}].$$

To obtain a lower bound on the r.h.s., notice that

$$\mathbf{E}_{\mathbf{r}} [\|\mathbf{u}\|^2] = \sum_{i=1}^n \mathbf{E}_{\mathbf{r}} [u_i^2] = \sum_{i=1}^n \frac{\|\mathbf{v}_i\|^2}{n-1} = \frac{\Phi}{n-1}.$$

Since $|u_i| \leq B$ for all i , $\|\mathbf{u}\|^2 \leq Bn$. Hence,

$$\begin{aligned} \mathbf{E}_{\mathbf{r}} [\|\mathbf{u}\|^2 \mid \mathcal{E}] \cdot \Pr[\mathcal{E}] &\geq \mathbf{E}_{\mathbf{r}} [\|\mathbf{u}\|^2] - \mathbf{E}_{\mathbf{r}} [\|\mathbf{u}\|^2 \mid \bar{\mathcal{E}}] \cdot \Pr[\bar{\mathcal{E}}] \\ &\geq \frac{\Phi}{n-1} - Bn^{-c/4+3}. \end{aligned}$$

Since $\Phi \geq \frac{1}{\text{poly}(n)}$, by picking c to be a large enough constant, one obtains

$$\mathbf{E}_{\mathbf{r}} [\|\mathbf{u}\|^2 \mid \mathcal{E}] \cdot \Pr[\mathcal{E}] \geq \frac{\Phi}{2(n-1)}.$$

Hence, we can complete the proof of the lemma by concluding that

$$\mathbf{E}_{\mathbf{r}} \left[\sum_{\{i,j\} \in M} \|\mathbf{v}_i - \mathbf{v}_j\|^2 \right] \geq \frac{\Phi}{2c \log n}.$$

□