

Approximating the Exponential, the Lanczos Method and an $\tilde{O}(m)$ -Time Spectral Algorithm for Balanced Separator

Lorenzo Orecchia
MIT
Cambridge, MA, USA.
orecchia@mit.edu

Sushant Sachdeva*
Princeton University
Princeton, NJ, USA.
sachdeva@cs.princeton.edu

Nisheeth K. Vishnoi
Microsoft Research
Bangalore, India
nisheeth.vishnoi@gmail.com

Abstract

We give a novel spectral approximation algorithm for the balanced separator problem that, given a graph G , a constant balance $b \in (0, 1/2]$, and a parameter γ , either finds an $\Omega(b)$ -balanced cut of conductance $O(\sqrt{\gamma})$ in G , or outputs a certificate that all b -balanced cuts in G have conductance at least γ , and runs in time $\tilde{O}(m)$. This settles the question of designing asymptotically optimal spectral algorithms for balanced separator. Our algorithm relies on a variant of the heat kernel random walk and requires, as a subroutine, an algorithm to compute $\exp(-L)v$ where L is the Laplacian of a graph related to G and v is a vector. Algorithms for computing the matrix-exponential-vector product efficiently comprise our next set of results. Our main result here is a new algorithm which computes a good approximation to $\exp(-A)v$ for a class of symmetric positive semidefinite (PSD) matrices A and a given vector u , in time roughly $\tilde{O}(m_A)$, where m_A is the number of non-zero entries of A . This uses, in a non-trivial way, the breakthrough result of Spielman and Teng on inverting symmetric and diagonally-dominant matrices in $\tilde{O}(m_A)$ time. Finally, we prove that e^{-x} can be uniformly approximated up to a small additive error, in a non-negative interval $[a, b]$ with a polynomial of degree roughly $\sqrt{b-a}$. While this result is of independent interest in approximation theory, we show that, via the Lanczos method from numerical analysis, it yields a simple algorithm to compute $\exp(-A)v$ for symmetric PSD matrices that runs in time roughly $O(t_A \cdot \sqrt{\|A\|})$, where t_A is time required for the computation of the vector Aw for given vector w . As an application, we obtain a simple and practical algorithm, with output conductance $O(\sqrt{\gamma})$, for balanced separator that runs in time $\tilde{O}(m/\sqrt{\gamma})$. This latter algorithm matches the running time, but improves on the approximation guarantee of the Evolving-Sets-based algorithm by Andersen and Peres for balanced separator.

Keywords. Spectral Algorithms, Balanced Graph Partitioning, Matrix Exponential, Lanczos Method, Uniform Approximation.

*This work was done while this author was interning at Microsoft Research India, Bangalore.

Contents

1	Introduction and Our Results	3
1.1	Balanced Separator	3
1.2	The Matrix Exponential, the Lanczos Method and Approximations to e^{-x}	4
2	Organization of the Main Body of the Paper	6
3	Technical Overview of Our Results	7
3.1	Our Spectral Algorithm for Balanced Separator	7
3.2	Our Algorithms to Compute an Approximation to $\exp(-A)v$	10
3.3	Our Uniform Approximation for e^{-x}	14
4	Discussion and Open Problems	15
5	The Algorithm for Balanced Separator	15
5.1	Basic Preliminaries	15
5.2	AHK Random Walks	16
5.3	Algorithm Description	18
5.4	Analysis	20
5.5	Proofs	22
5.6	SDP Interpretation	26
5.7	The FINDCUT Subroutine	26
6	Computing $\exp(-A)v$	30
6.1	Lanczos Method – From Scalars to Matrices	31
6.2	Procedures for Approximating $\exp(-A)v$	34
6.3	Exponentiating PSD Matrices – Proofs of Theorem 1.2 and 3.2	35
6.4	Beyond SDD - Proof of Theorem 3.2	38
6.5	Error Analysis for EXPRATIONAL	41
7	Uniform Approximations to e^{-x}	50
7.1	Preliminaries	50
7.2	Known Approximation Results and Discussion.	51
7.3	Proof of Upper Bound in Theorem 1.5	52
7.4	Proof of Lower Bound in Theorem 1.5	54
7.5	Remaining Proofs	55

1 Introduction and Our Results

1.1 Balanced Separator

The BALANCED SEPARATOR problem (BS) asks the following decision question: given an unweighted graph $G = (V, E)$, $V = [n]$, $|E| = m$, a constant balance parameter $b \in (0, 1/2]$, and a target conductance value $\gamma \in (0, 1)$, does G have a b -balanced cut S such that $\phi(S) \leq \gamma$? Here, the conductance of a cut (S, \bar{S}) is defined to be $\phi(S) \stackrel{\text{def}}{=} |E(S, \bar{S})| / \min\{\text{vol}(S), \text{vol}(\bar{S})\}$, where $\text{vol}(S)$ is the sum of the degrees of the vertices in the set S . Moreover, a cut (S, \bar{S}) is b -balanced if $\min\{\text{vol}(S), \text{vol}(\bar{S})\} \geq b \cdot \text{vol}(V)$. This is a classic NP-hard problem and a central object of study for the development of approximation algorithms, both in theory and in practice. On the theoretical side, BS has far reaching connections to spectral graph theory, the study of random walks and metric embeddings. In practice, algorithms for BS play a crucial role in the design of recursive algorithms [35], clustering [19] and scientific computation [32].

Spectral methods are an important set of techniques in the design of graph-partitioning algorithms and are fundamentally based on the study of the behavior of random walks over the instance graph. Spectral algorithms tend to be conceptually appealing, because of the intuition based on the underlying diffusion process, and easy to implement, as many of the primitives required, such as eigenvector computation, already appear in highly-optimized software packages. The most important spectral algorithm for graph partitioning is the Laplacian Eigenvector (LE) algorithm of Alon and Milman [2], which, given a graph of conductance at most γ , outputs a cut of conductance at most $O(\sqrt{\gamma})$, an approximation guarantee that is asymptotically optimal for spectral algorithms. A consequence of the seminal work of Spielman and Teng [37] is that the LE algorithm can run in time $\tilde{O}(m)$ using the Spielman-Teng solver. Hence, LE is an asymptotically optimal spectral algorithm for the minimum-conductance problem, both for running time (up to polylog factors) and approximation quality. In this paper, we present a simple random-walk-based algorithm that is the first such asymptotically optimal spectral algorithm for BS. Our algorithm can be seen as an analogue to the LE algorithm for the balanced version of the minimum-conductance problem and settles the question of designing spectral algorithms for BS. The following is our main theorem on graph partitioning.

Theorem 1.1 (Spectral Algorithm for Balanced Separator) *Given an unweighted graph $G = (V, E)$, a balance parameter $b \in (0, 1/2]$, $b = \Omega(1)$ and a conductance value $\gamma \in (0, 1)$, we give an algorithm called $\text{BALSEP}(G, b, \gamma)$, that either outputs an $\Omega(b)$ -balanced cut $S \subset V$ such that $\phi(S) \leq O(\sqrt{\gamma})$, or outputs a certificate that no b -balanced cut of conductance γ exists. BALSEP runs in time $O(m \text{ poly}(\log n))$.*

The algorithm for Theorem 1.1 relies on our ability to compute the product of the matrix-exponential of a matrix and an arbitrary vector in time essentially proportional to the sparsity of the matrix. Our contribution to the problem of computing the matrix-exponential-vector product appear in detail in Section 1.2. The algorithm required for Theorem 1.1 runs in time $\tilde{O}(m)$ and, notably, makes use of the Spielman-Teng solver in a non-trivial way. We also prove an alternative novel result on how to perform this matrix-exponential computation, which relies just on matrix-vector products. This result, when combined with our BS algorithm based on random walks, yields a theorem identical to Theorem 1.1 except that the running time now increases to $\tilde{O}(m/\sqrt{\gamma})$, see Theorem 3.1. However, this latter algorithm not only turns out to be almost as simple and practical as the LE algorithm, but it also improves in the approximation factor upon the result of Andersen and Peres [6] who obtain the same running time using Evolving-Sets-based random walk.

1.1.1 Comparison to Previous Work on Balanced Separator

The best known approximation for BS is $O(\sqrt{\log n})$ achieved by the seminal work of Arora, Rao and Vazirani [8] that combines semidefinite programming (SDP) and flow ideas. A rich line of research has centered on reducing the running time of this algorithm using SDP and flow ideas [20, 7, 24]. This effort culminated in Sherman’s work [33], which brings down the required running time to $O(n^\epsilon)$ s - t maximum-flow computations.¹ However, these algorithms are based on advanced theoretical ideas that are not easy to implement or even capture in a principled heuristic. Moreover, they fail to achieve a nearly-linear² running time, which is crucial in many of today’s applications that involve very large graphs. To address these issues, researchers have focused on the design of simple, nearly-linear-time algorithms for BS based on spectral techniques. The simplest spectral algorithm for BS is the Recursive Laplacian Eigenvector (RLE) algorithm (see, for example, [19]). This algorithm iteratively uses LE to remove low-conductance unbalanced cuts from G , until a balanced cut or an induced γ -expander is found. The running time of the RLE algorithm is quadratic in the worst case, as $\Omega(n)$ unbalanced cuts may be found, each requiring a global computation of the eigenvector. Spielman and Teng [36] were the first to design nearly-linear-time algorithms outputting an $\Omega(b)$ -balanced cut of conductance $O(\sqrt{\gamma \text{polylog} n})$, if a b -balanced cut of conductance less than γ exists. Their algorithmic approach is based on local random walks, which are used to remove unbalanced cuts in time proportional to the size of the cut removed, hence avoiding the quadratic dependence of RLE. Using similar ideas, Andersen, Chung and Lang [4], and Andersen and Peres [6] improved the approximation guarantee to $O(\sqrt{\gamma \log n})$ and the running time to $\tilde{O}(m/\sqrt{\gamma})$. More recently, Orecchia and Vishnoi (OV) [25] employed an SDP formulation of the problem, together with the Matrix Multiplicative Weight Update (MMWU) of [7] and a new SDP rounding, to obtain an output conductance of $O(\sqrt{\gamma})$ with running time $\tilde{O}(m/\gamma^2)$, effectively removing unbalanced cuts in $O(\log n/\gamma)$ iterations. In Section 5.6, we give a more detailed comparison with OV and discussion of our novel width-reduction techniques from an optimization point of view. Finally, our algorithm should also be compared to the remarkable results of Madry [22] for BS, which build up on Räcke’s work [28] and on the low-stretch spanning trees from Abraham *et al.* [1], to achieve a trade-off between running time and approximation. For every integer $k \geq 1$, he achieves roughly $O((\log n)^k)$ approximation in time $\tilde{O}(m + 2^k \cdot n^{1+2^{-k}})$. Calculations show that for $\gamma \geq 2^{-(\log \log n)^2}$, our algorithm achieves strictly better running time and approximation than Madry’s for sparse graphs.³ More importantly, we believe that our algorithm is significantly simpler, especially in its second form mentioned above, and likely to find applications in practical settings.

1.2 The Matrix Exponential, the Lanczos Method and Approximations to e^{-x}

We first state a few definitions used in this section. We will work with $n \times n$, symmetric and positive semi-definite (PSD) matrices over \mathbb{R} . For a matrix M , abusing notation, we denote its exponential by $\exp(-M)$, or by e^{-M} , and define it as $\sum_{i \geq 0} \frac{(-1)^i}{i!} M^i$. M is said to be *Symmetric and Diagonally Dominant* (SDD) if, $M_{ij} = M_{ji}$, for all i, j and $M_{ii} \geq \sum_j |M_{ij}|$, for all i . Let m_M denote the number of non-zero entries in M and let t_M denote the time required to multiply the matrix

¹Even though the results of [8] and [33] are stated for the Sparsest Cut problem, the same techniques apply to the conductance problem, e.g. by modifying the underlying flow problems. See for example [5].

²Following the convention of [38], we denote by nearly-linear a running time of $\tilde{O}(m/\text{poly}(\gamma))$.

³In the table on Page 4 of the full version of Madry’s paper, it has been erroneously claimed that using the Spielman-Teng solver, Alon-Milman algorithm runs in time $\tilde{O}(m)$ for BS.

M with a given vector v . In general, t_M depends on how M is given as an input and can be $\Theta(n^2)$. However, it is possible to exploit the special structure of M if given as an input appropriately: It is possible to just multiply the non-zero entries of M , giving $t_M = O(m_M)$. Also, if M is a rank one matrix ww^\top , where w is known, we can multiply with M in $O(n)$ time. We move on to our results.

At the core of our algorithm for BS, and more generally of most MMWU based algorithms, lies an algorithm to quickly compute $\exp(-A)v$ for a PSD matrix A and a unit vector v . It is sufficient to compute an approximation u , to $\exp(-A)v$, in time which is as close as possible to t_A . It can be shown that using about $\|A\|$ terms in the Taylor series expansion of $\exp(-A)$, one can find a vector u that approximates $\exp(-A)v$. Hence, this method runs in time roughly $O(t_A \cdot \|A\|)$. In our application, and certain others [7, 18, 15, 16], this dependence on the norm is prohibitively large. The following remarkable result was cited in Kale [18].

Hypothesis. *Let $A \succeq 0$ and $\varepsilon > 0$. There is an algorithm that requires $O(\log^2 1/\varepsilon)$ iterations to find a vector u such that $\|\exp(-A)v - u\| \leq \|\exp(-A)\| \varepsilon$, for any unit vector v . The time for every iteration is $O(t_A)$.*

This hypothesis would suffice to prove Theorem 1.1. But, to the best of our knowledge, there is no known proof of this result. In fact, the source of this unproved hypothesis can be traced to a paper of Eshof and Hochbruck (EH) [13]. EH suggest that one may use the Lanczos method (described later), and combine it with a rational approximation for e^{-x} due to Saff, Schonhage and Varga [30], to reduce the computation of $\exp(-A)v$ to a number of $(I + \alpha A)^{-1}v$ computations for some $\alpha > 0$. Note that this is insufficient to prove the hypothesis above as there is no known way to compute $(I + \alpha A)^{-1}v$ in time $O(t_A)$. They note this and propose the use of iterative methods to do this computation. They also point out that this will only result in an approximate solution to $(I + \alpha A)^{-1}v$ and make no attempt to analyze the running time or the error of their method when the inverse computation is approximate. We believe that we are quite distant from proving the hypothesis for all PSD matrices and, moreover, that proving such a result may provide valuable insights into a fast (approximate) inversion method for symmetric PSD matrices, an extremely important open problem.

A significant part of this paper is devoted to a proof of the above hypothesis for a class of PSD matrices that turns out to be sufficient for the BS application. For the norm-independent, fast-approximate inverse computation, we appeal to the result of Spielman and Teng [37] (also see improvements by Koutis, Miller and Peng [21]). The theorem we prove is the following.

Theorem 1.2 (SDD Matrix Exponential Computation) *Given an $n \times n$ SDD matrix A , a vector v and a parameter $\delta \leq 1$, there is an algorithm that can compute a vector u such that $\|\exp(-A)v - u\| \leq \delta \|v\|$ in time $\tilde{O}((m_A + n) \log(2 + \|A\|))$. Here the tilde hides $\text{poly}(\log n)$ and $\text{poly}(\log 1/\delta)$ factors.*

First, we note that for our application, the dependence of the running time on the $\log(2 + \|A\|)$ turns out to just contribute an extra $\log n$ factor. Also, for our application $\delta = 1/\text{poly}(n)$. Secondly, for our BS application, the matrix we need to invert is not SDD or sparse. Fortunately, we can combine Spielman-Teng solver with the Sherman-Morrison formula to invert our matrices; see Theorem 3.2. A significant effort goes into analyzing the effect of the error introduced due to approximate matrix inversion. This error can cascade due to the iterative nature of our algorithm that proves this theorem.

Towards proving the hypothesis above, when the only guarantee we know on the matrix is that it is symmetric and PSD, we prove the following theorem, which is the best known algorithm to compute $\exp(-A)v$ for an arbitrary symmetric PSD matrix A , when $\|A\| = \omega(\text{poly}(\log n))$.

Theorem 1.3 (PSD Matrix Exponential Computation) *Given an $n \times n$ symmetric PSD matrix A , a vector v and a parameter $\delta \leq 1$, there is an algorithm that computes a vector u such that $\|\exp(-A)v - u\| \leq \delta \|v\|$ in time $\tilde{O}\left((t_A + n)\sqrt{1 + \|A\|} \log(2 + \|A\|)\right)$. Here the tilde hides $\text{poly}(\log n)$ and $\text{poly}(\log^{1/\delta})$ factors.*

In the symmetric PSD setting we also prove the following theorem which, for our application, gives a result comparable to Theorem 1.3.

Theorem 1.4 (Simple PSD Matrix Exponential Computation) *Given an $n \times n$ symmetric PSD matrix A , a vector v and a parameter $\delta \leq 1$, there is an algorithm that computes a vector u such that $\|\exp(-A)v - u\| \leq \delta \|v\|$, in time $O((t_A + n) \cdot k + k^2)$, where $k \stackrel{\text{def}}{=} \tilde{O}(\sqrt{1 + \|A\|})$. Here the tilde hides $\text{poly}(\log^{1/\delta})$ factors.*

As noted above, t_A can be significantly smaller than m_A . Moreover, it only uses multiplication of a vector with the matrix A as a primitive and does not require matrix inversion. Consequently, it does not need tools like the SDD solver or conjugate gradient, thus obviating the error analysis required for the previous algorithms. Furthermore, this algorithm is very simple and when combined with our random walk-based BALSEP algorithm, results in a very simple and practical $O(\sqrt{\gamma})$ approximation algorithm for BS that runs in time $\tilde{O}(m/\sqrt{\gamma})$.

Theorem 1.4 relies on the Lanczos method which can be used to convert guarantees about polynomial approximation from scalars to matrices. In particular, it uses the following structural result (the upper bound) on the best degree k polynomial δ -uniformly approximating e^{-x} in an interval $[a, b]$. We also prove a lower bound which establishes that the degree cannot be improved beyond lower order terms. This suggests that improving on the $\tilde{O}(m/\sqrt{\gamma})$ running time in Theorem 1.4 requires more advanced techniques. To the best of our knowledge, this theorem is new and is of independent interest in approximation theory.

Theorem 1.5 (Uniform Approximation to e^{-x})

- **Upper Bound.** *For every $0 \leq a < b$, and $0 < \delta \leq 1$, there exists a polynomial p that satisfies, $\sup_{x \in [a, b]} |e^{-x} - p(x)| \leq \delta \cdot e^{-a}$, and has degree $O\left(\sqrt{\max\{\log^2 1/\delta, (b - a) \cdot \log 1/\delta\}} \cdot (\log 1/\delta)^2\right)$.*
- **Lower Bound.** *For every $0 \leq a < b$ such that $a + \log_e 4 \leq b$, and $\delta \in (0, 1/8]$, any polynomial $p(x)$ that approximates e^{-x} uniformly over the interval $[a, b]$ up to an error of $\delta \cdot e^{-a}$, must have degree at least $\frac{1}{2} \cdot \sqrt{b - a}$.*

2 Organization of the Main Body of the Paper

In Section 3 we present a technical overview of our results and in Section 4 we discuss the open problems arising from our work. The main body of the paper follows after it and is divided into three sections, each of which have been written so that they can be read independently. Section 5 contains a complete description and all the proofs related to Theorem 1.1 and Theorem 3.1. Section 6 contains our results on computing the matrix exponential; in particular the proofs of Theorems 1.2, 1.4 and 3.2. Section 7 contains the proof of our structural results on approximating e^{-x} and the proof of Theorem 1.5.

3 Technical Overview of Our Results

3.1 Our Spectral Algorithm for Balanced Separator

In this section, we provide an overview of Theorem 1.1. As pointed out in the introduction, our algorithm, `BALSEP`, when combined with the matrix-exponential-vector algorithm in Theorem 1.4 results in a very simple and practical algorithm for BS. We record the theorem here for completeness and then move on to the overview of `BALSEP` and its proof. The proof of this theorem appears in Section 5.4.

Theorem 3.1 (Simple Spectral Algorithm for Balanced Separator) *Given an unweighted graph $G = (V, E)$, a balance parameter $b \in (0, 1/2]$, $b = \Omega(1)$ and a conductance value $\gamma \in (0, 1)$, we give an algorithm, which runs in time $\tilde{O}(m/\sqrt{\gamma})$, that either outputs an $\Omega(b)$ -balanced cut $S \subset V$ such that $\phi(S) \leq O(\sqrt{\gamma})$ or outputs a certificate that no b -balanced cut of conductance γ exists.*

3.1.1 Comparison with the RLE Algorithm

Before we explain our algorithm, it is useful to review the RLE algorithm. Recall that given G, γ and b , the goal of the BS problem is to either certify that every b -balanced cut in G has conductance at least γ , or produce a $\Omega(b)$ balance cut in G of conductance $O(\sqrt{\gamma})$. RLE does this by applying LE iteratively to remove unbalanced cuts of conductance $O(\sqrt{\gamma})$ from G . The iterations stop and the algorithm outputs a cut, when it either finds a $(b/2)$ -balanced cut of conductance $O(\sqrt{\gamma})$ or the union of all unbalanced cuts found so far is $(b/2)$ -balanced. Otherwise, the algorithm terminates when the residual graph has spectral gap at least 2γ . In the latter case, any b -balanced cut must have at least half of its volume lie within the final residual graph, and hence, has conductance at least γ in the original graph. Unfortunately, this algorithm may require $\Omega(n)$ iterations in the worst case. For instance, this is true if the graph G consists of $\Omega(n)$ components loosely connected to an expander-like core through cuts of low conductance. This example highlights the weakness of the RLE approach: the second eigenvector of the Laplacian may only be correlated with one low-conductance cut and fail to capture at all even cuts of slightly larger conductance. This limitation makes it impossible for RLE to make significant progress at any iteration. We now proceed to show how to fix RLE and present our algorithm at a high level.

3.1.2 High-Level Idea of Our Algorithm

Rather than working with the vertex embedding given by the eigenvector, at iteration t , we will consider the multi-dimensional vector embedding represented by the transition probability matrix $P^{(t)}$ of a certain random walk over the graph. We refer to this kind of walk as an *Accelerated Heat Kernel Walk* (AHK) and we describe it formally in Section 3.1.4. At each iteration $t = 1, 2, \dots$, the current AHK walk is simulated for $\tau = \log n / \gamma$ time to obtain $P^{(t)}$. For any t , this choice of τ ensures that the walk must mix across all cuts of conductance much larger than γ , hence emphasizing cuts of the desired conductance in the embedding $P^{(t)}$. The embedding obtained in this way, can be seen as a weighted combination of multiple eigenvectors, with eigenvectors of low eigenvalue contributing more weight. Hence, the resulting embedding captures not only the cut corresponding to the second eigenvector, but also cuts associated with other eigenvectors of eigenvalue close to γ . This enables our algorithm to potentially find many different low-conductance unbalanced cuts at once. Moreover, the random walk matrix is more stable than the eigenvector under small perturbations of the graph, making it possible to precisely quantify our progress from

one iteration to the next as a function of the mixing of the current random walk. For technical reasons, we are unable to show that we make sufficient progress if we just remove the unbalanced cuts found, as in RLE. Instead, if we find a low-conductance unbalanced cut $S^{(t)}$ at iteration t , we perform a *soft* removal, by modifying the current walk $P^{(t)}$ to accelerate the convergence to stationarity on the set $S^{(t)}$. This ensures that a different cut is found using $P^{(t+1)}$ in the next iteration. In particular, the AHK walks we consider throughout the execution of the algorithm will behave like the standard heat kernel on most of the graph, except on a small unbalanced subset of vertices, where their convergence will be accelerated. We now present our algorithm in more detail. We first recall some definitions.

3.1.3 Definitions

$G = (V, E)$ is the unweighted instance graph, where $V = [n]$ and $|E| = m$. We let $d \in \mathbb{R}^n$ be the degree vector of G , i.e., d_i is the degree of vertex i . For a subset $S \subseteq V$, we define the edge volume as $\text{vol}(S) \stackrel{\text{def}}{=} \sum_{i \in S} d_i$. The total volume of G is $2m$. We denote by K_V the complete graph with weight $d_i d_j / 2m$ between every pair $i, j \in V$. For $i \in V$, S_i is the star graph rooted at i , with edge weight of $d_i d_j / 2m$ between i and j , for all $j \in V$. For an undirected graph $H = (V, E_H)$, let $A(H)$ denote the adjacency matrix of H , and $D(H)$ the diagonal matrix of degrees of H . The (combinatorial) Laplacian of H is defined as $L(H) \stackrel{\text{def}}{=} D(H) - A(H)$. By D and L , we denote $D(G)$ and $L(G)$ respectively for the input graph G . For two matrices A, B of equal dimensions, let $A \bullet B \stackrel{\text{def}}{=} \text{Tr}(A^\top B) = \sum_{ij} A_{ij} \cdot B_{ij}$. $\mathbf{0}$ denotes the all 0s vector.

3.1.4 The AHK Random Walk and its Mixing

We will be interested in continuous-time random walk processes over V that take into account the edge structure of G . The simplest such process is the *heat kernel* process, which has already found many applications in graph partitioning, particularly in the work of Chung [14], and in many of the combinatorial algorithms for solving graph partitioning SDPs [23]. The heat kernel is defined as the continuous-time process having transition rate matrix $-LD^{-1}$ and, hence, at time τ the probability-transition matrix becomes $\exp(-\tau LD^{-1})$. The AHK random walk process has an additional parameter β , which is a non-negative vector in \mathbb{R}^n . The transition rate matrix is then defined to be $-(L + \sum_{i \in V} \beta_i L(S_i))D^{-1}$. The effect of adding the star terms to the transition rate matrix is that of accelerating the convergence of the process to stationarity at vertices i with large value of β_i , since a large fraction of the probability mass that leaves these vertices is distributed uniformly over the edges. We denote by $P_\tau(\beta)$ the probability-transition $\exp(-\tau(L + \sum_{i \in V} \beta_i L(S_i))D^{-1})$. As is useful in the study of spectral properties of non-regular graphs, we will study $D^{-1}P_\tau(\beta)$. This matrix describes the probability distribution over the edges of G and has the advantage of being symmetric and PSD: $D^{-1}P_\tau(\beta) = D^{-1/2} \exp(-\tau D^{-1/2}(L + \sum_{i \in V} \beta_i L(S_i))D^{-1/2})D^{-1/2}$. In particular, $D^{-1}P_\tau(\beta)$ can be seen as the Gram matrix of the embedding given by the columns of its square root, which in this case is just $D^{-1/2}P_{\tau/2}(\beta)$. This property will enable us to use geometric SDP-rounding techniques to analyze AHK walks. Throughout the algorithm, we keep track of the mixing of $P^{(t)}$ by considering the *total deviation* of $P^{(t)}$ from stationarity, i.e., the sum over all vertices $i \in V$ of the ℓ_2^2 -distance from the stationary distribution of $P^{(t)}e_i$. Here, e_i denotes the vector in \mathbb{R}^n which is 1 at the i^{th} coordinate and 0 elsewhere. We denote the contribution of vertex i to this distance by $\Psi(P_\tau(\beta), i)$. Similarly, the total deviation from stationarity over a subset $S \subseteq V$, is given by, $\Psi(P_\tau(\beta), S) \stackrel{\text{def}}{=} \sum_{i \in S} \Psi(P_\tau(\beta), i)$. $\Psi(P_\tau(\beta), V)$ will play the role of potential function in our

algorithm. Moreover, it follows from the definition of Ψ that $\Psi(P_\tau(\beta), V) = L(K_V) \bullet D^{-1}P_\tau(\beta)$. Finally, to connect to the high-level idea described earlier, for each iteration t , we use $P^{(t)} \stackrel{\text{def}}{=} P_\tau(\beta^{(t)})$ for $\tau = \log n / \gamma$ with $\beta^{(t)} \approx 1/\tau \sum_{j=1}^{t-1} \sum_{i \in S^{(j)}} e_i$ and starting with $\beta^{(1)} = \mathbf{0}$. We now provide a slightly more detailed description of our algorithm from Theorem 1.1 (called BALSEP) and its analysis. For reference, BALSEP appears in Figure 1.

3.1.5 Our Algorithm and its Analysis

The algorithm proceeds as follows: At iteration t , it checks if the total deviation of $P^{(t)}$, i.e., $\Psi(P^{(t)}, V)$, is sufficiently small (i.e., $P^{(t)}$ is mixing). In this case, we can guarantee that no balanced cut of conductance less than γ exists in G . In more formal language, it appears below.

(A) (see Lemma 5.6) Let $S = \cup_{i=1}^t S^{(i)}$. For any $t \geq 1$, if $\Psi(P^{(t)}, V) \leq 1/\text{poly}(n)$, and $\text{vol}(S) \leq b/100 \cdot 2m$, then $L + 1/\tau \sum_{i \in S} L(S_i) \succeq \Omega(\gamma) \cdot L(K_V)$. Moreover, no b -balanced cut of conductance less than γ exists in G .

This result has a simple explanation in terms of the AHK random walk $P^{(t)}$. Notice that $P^{(t)}$ is accelerated only on a small unbalanced set S . Hence, if a balanced cut of conductance less than γ existed, its convergence could not be greatly helped by the acceleration over S . Thus, if $P^{(t)}$ is still mixing very well, no such balanced cut can exist. On the other hand, if $P^{(t)}$ has high total deviation (i.e., the walk has not yet mixed), then, intuitively, some cut of low conductance exists in G . Formally, we show that, the embedding $\{v_i^{(t)}\}_{i \in V}$ has low quadratic form with respect to the Laplacian of G .

(B) (see Lemma 5.7) If $\Psi(P^{(t)}, V) \geq \frac{1}{\text{poly}(n)}$, then $L \bullet D^{-1}P^{(t)} \leq O(\gamma)L(K_V) \bullet D^{-1}P^{(t)}$.

From an SDP-rounding perspective, this means that the embedding $P^{(t)}$ can be used to recover a cut $S^{(t)}$ of conductance $O(\sqrt{\gamma})$, using the SDP-rounding techniques from OV. If $S^{(t)}$ or $\cup_{i=1}^t S^{(i)}$ is $\Omega(b)$ -balanced, then we output that cut and terminate. Otherwise, $S^{(t)}$ is unbalanced. In this case, we accelerate the convergence from $S^{(t)}$ in the current AHK walk by increasing $(\beta^{(t)})_i$ for every $i \in S^{(t)}$ to give $\beta^{(t+1)}$, and using $\beta^{(t+1)}$ to produce $P^{(t+1)}$ and move on to the next iteration.

The analysis of our algorithm bounds the number of iterations by using the total deviation of $P^{(t)}$ from stationarity as a potential function. Using the techniques of OV, it is possible to show that, whenever an unbalanced cut $S^{(t)}$ is found, most of the deviation of $P^{(t)}$ can be attributed to $S^{(t)}$. In words, we can think of $S^{(t)}$ as the main reason why $P^{(t)}$ is not mixing. Formally,

(C) (see Corollary 5.9) At iteration t , if $\Psi(P^{(t)}, V) \geq 1/\text{poly}(n)$ and $S^{(t)}$ is not $b/100$ -balanced, then w.h.p. $\Psi(P^{(t)}, S^{(t)}) \geq 1/2 \cdot \Psi(P^{(t)}, V)$.

Moreover, we can show that accelerating the convergence of the walk from $S^{(t)}$ has the effect of removing from $P^{(t+1)}$ a large fraction of the deviation due to $S^{(t)}$. The proof is a simple application of the Golden-Thompson inequality [9] and mirrors the main step in the MMWU analysis. Hence, we can show the total deviation of $P^{(t+1)}$ is just a constant fraction of that of $P^{(t)}$.

(D) (see Theorem 5.10) If $\Psi(P^{(t)}, V) \geq 1/\text{poly}(n)$ and $S^{(t)}$ is not $b/100$ balanced, then w.h.p.

$$\Psi(P^{(t+1)}, V) \leq \Psi(P^{(t)}, V) - 1/3 \cdot \Psi(P^{(t)}, S^{(t)}) \leq 5/6 \cdot \Psi(P^{(t)}, V).$$

This potential reduction at every iteration allows us to argue that after $T = O(\log n)$ iterations, $P^{(T+1)}$ must have a small deviation from stationarity and yields a certificate that no balanced cut of conductance less than γ exists in G . Finally, to ensure that each iteration requires only $\tilde{O}(m)$ time,

we use the Johnson-Lindenstrauss Lemma to compute a $O(\log n)$ -dimensional approximation to the embedding $P^{(t)}$. To compute this approximation, we rely on the results on approximating the matrix exponential discussed in Section 3.2.

3.1.6 Exponential Embeddings of Graph and Proof Ideas

Now, we illustrate the usefulness of the exponential embedding obtained from the AHK random walk, which is key to the proofs for (A) and (B) above. We suppress some details pertinent to our algorithm. Consider the AHK walk in the first iteration, with $\beta^{(1)} = 0$. Letting $C \stackrel{\text{def}}{=} D^{-1/2}LD^{-1/2}$, $P \stackrel{\text{def}}{=} P^{(1)} = \exp(-\tau C)$. For the proof of (A), it follows that if $\Psi(P, V) \leq 1/\text{poly}(n)$, then by definition $L(K_V) \bullet D^{-1}P = \text{Tr}(\exp(-\tau C)) - 1 \leq 1/\text{poly}(n)$. Hence, $\lambda_2(C) \gtrsim \log n/\tau = \gamma$, by the choice of τ . This lower bound on the second eigenvalue certifies that G has no cut of conductance at most γ . For our algorithm, at iteration t , when $\beta^{(t)} \neq 0$, we will ensure that the Laplacians of the stars have small enough weight ($\approx 1/\tau$) and small support ($\text{vol}(\cup_{i=1}^t S^{(i)}) \leq b/100 \cdot 2m$) for the argument above to still yield a lower bound of γ on the conductance of any b -balanced cut.

For (B), observe that $L \bullet D^{-1}P = C \bullet \exp(-\tau C) = \sum_i \lambda_i e^{-\tau \lambda_i}$, where λ_i , for $i = 1, \dots, n$, are the eigenvalues of C . For eigenvalues larger than 2γ , $e^{-\tau \lambda_i}$ is bounded by $1/\text{poly}(n)$ for $\tau = \log n/\gamma$. Since eigenvalues of a normalized Laplacian are $O(1)$, the contribution to the sum above by eigenvalues $\lambda_i > 2\gamma$ is at most $1/\text{poly}(n)$ overall. This can be shown to be a small fraction of the total sum. Hence, the quantity $L \bullet D^{-1}P$ is mostly determined by the eigenvalues of value less than 2γ and we have $L \bullet D^{-1}P \leq O(\gamma) \cdot L(K_V) \bullet D^{-1}P$. The same analysis goes through when $t > 1$.

3.2 Our Algorithms to Compute an Approximation to $\exp(-A)v$

In this section, we give an overview of the algorithms in Theorem 1.2 and Theorem 1.4 and their proofs. The algorithm for Theorem 1.3 is very similar to the one for Theorem 1.2 and we give the details in Section 6.3.2. A few quick definitions: A matrix M is called *Upper Hessenberg* if, $(M)_{ij} = 0$ for $i > j + 1$. M is called *tridiagonal* if $M_{ij} = 0$ for $i > j + 1$ and for $j > i + 1$. Let $\lambda_1(M)$ and $\lambda_n(M)$ denote the largest and smallest eigenvalues of M respectively.

As we mention in the introduction, the matrices that we need to exponentiate for the BS algorithm are no longer sparse or SDD. Thus, Theorem 1.2 is insufficient for our application. Fortunately, the following theorem suffices and its proof is not very different from that of Theorem 1.2, which is explained below. Its proof appears in Section 6.4.

Theorem 3.2 (Matrix Exponential Computation Beyond SDD) *Given a vector v , a parameter $\delta \leq 1$ and an $n \times n$ symmetric matrix $A = \Pi H M H \Pi$ where M is SDD, H is a diagonal matrix with strictly positive entries and Π is a rank $(n - 1)$ projection matrix, $\Pi \stackrel{\text{def}}{=} I - w w^\top$ (w is explicitly known and $\|w\| = 1$), there is an algorithm that computes a vector u such that $\|\exp(-A)v - u\| \leq \delta \|v\|$ in time $\tilde{O}((m_M + n) \log(2 + \|H M H\|))$. The tilde hides $\text{poly}(\log n)$ and $\text{poly}(\log 1/\delta)$ factors.*

Recall from Section 3.1.4 that our algorithm for BS requires us to compute $\exp(-A)v$ for a matrix A of the form $D^{-1/2}(L + \sum_i \beta_i L(S_i))D^{-1/2}$, where $\beta_i \geq 0$. We first note that if we let $\Pi \stackrel{\text{def}}{=} I - 1/2m \cdot (D^{1/2}1)(D^{1/2}1)^\top$, the projection onto the space orthogonal to $1/\sqrt{2m} \cdot D^{1/2}1$, then, for each i , $D^{-1/2}L(S_i)D^{-1/2} = \Pi(d_i/2m \cdot I + e_i e_i^\top)\Pi$. Since $D^{1/2}1$ is an eigenvector of $D^{-1/2}LD^{-1/2}$, we have, $\Pi D^{-1/2}LD^{-1/2}\Pi = D^{-1/2}LD^{-1/2}$. Thus,

$$A = \Pi D^{-1/2}LD^{-1/2}\Pi + \sum_i \beta_i \Pi(d_i/2m \cdot I + e_i e_i^\top)\Pi = \Pi D^{-1/2}(L + \sum_i \beta_i d_i/2m \cdot D + \sum_i \beta_i d_i \cdot e_i e_i^\top)D^{-1/2}\Pi.$$

This is of the form $\Pi H M H \Pi$, where $H \stackrel{\text{def}}{=} D^{-1/2}$ is diagonal and M is SDD. It is worth noting that since A itself may be neither sparse nor SDD, we cannot apply the Spielman-Teng SDD solver to approximate $(I + \alpha A)^{-1}$. The proof of the above theorem uses the Sherman-Morrison formula to extend the SDD solver to fit our requirement. Moreover, to obtain a version of Theorem 1.4 for such matrices, we do not have to do anything additional since multiplication by H and Π take $O(n)$ steps and hence, t_A is still $O(m_M + n)$. The details appear in Section 6.4. Finally, note that in our application, $\|H M H\|$ is $\text{poly}(n)$.

We now give an overview of the proofs of Theorem 1.2 and Theorem 1.4. First, we explain a general method known as the Lanczos method, which is pervasive in numerical linear algebra. We then show how suitable adaptations of this can be combined with (old and new) structural results in approximation theory to obtain our results.

3.2.1 Lanczos Method

Given an $n \times n$ symmetric PSD matrix B and a function $f : \mathbb{R} \mapsto \mathbb{R}$, we can define $f(B)$ as follows: Let u_1, \dots, u_n be eigenvectors of B with eigenvalues $\lambda_1, \dots, \lambda_n$. Define $f(B) \stackrel{\text{def}}{=} \sum_i f(\lambda_i) u_i u_i^\top$. We will reduce both our algorithms to computing $f(B)v$ for a given vector v , albeit with different f 's and B 's. We point out the f 's and B 's required for Theorems 1.2 and 1.4 in Sections 3.2.2 and 3.2.3 respectively.

Since exact computation of $f(B)$ usually requires diagonalization of B , which could take as much as $O(n^3)$ time (see [26]), we seek an approximation to $f(B)v$. The Lanczos method allows us to do exactly that: It looks for an approximation to $f(B)v$ of the form $p(B)v$, where p is a polynomial of small degree, say k . Before we describe how, we note that it computes this approximation in roughly $O((t_B + n)k)$ time plus the time it takes to compute $f(\cdot)$ on a $(k + 1) \times (k + 1)$ tridiagonal matrix, which can often be upper bounded by $O(k^2)$ (see [26]). Hence, the time is reduced to $O((t_B + n)k + k^2)$. What one has lost in this process is accuracy: The candidate vector u output by the Lanczos method, is now only an approximation to $f(B)v$. The quality of approximation, or $\|f(B)v - u\|$, can be upper bounded by the *uniform error* of the best degree k polynomial approximating f in the interval $[\lambda_n(B), \lambda_1(B)]$. Roughly, $\|f(B)v - u\| \approx (\min_{p_k \in \Sigma_k} \sup_{x \in [\lambda_n(B), \lambda_1(B)]} |f(x) - p_k(x)|)$. Here Σ_k is the collection of all real polynomials of degree at most k . Surprisingly, one does not need to know the best polynomial and proving *existence* of good polynomials is sufficient. By increasing k , one can reduce this error and, indeed, if one lets $k = n$, there is no error. Thus, the task is reduced to proving existence of low degree polynomials that approximate f within the error tolerable for the applications.

Computing the Best Polynomial Approximation.

Now, we describe in detail, the Lanczos method and how it achieves the error guarantee claimed above. Notice that for any polynomial p of degree at most k , the vector $p(B)v$ lies in $\mathcal{K} \stackrel{\text{def}}{=} \text{Span}\{v, Bv, \dots, B^k v\}$ – called the *Krylov subspace*. The Lanczos method iteratively creates an orthonormal basis $\{v_i\}_{i=0}^k$ for \mathcal{K} , such that $\forall i \leq k$, $\text{Span}\{v_0, \dots, v_i\} = \text{Span}\{v, \dots, B^i v\}$. Let V_k be the $n \times (k + 1)$ matrix with $\{v_i\}_{i=0}^k$ as its columns. Thus, $V_k V_k^\top$ denotes the projection onto the Krylov subspace. We let T_k be the $(k + 1) \times (k + 1)$ matrix expressing B as an operator restricted to \mathcal{K} in the basis $\{v_i\}_{i=0}^k$, i.e., $T_k \stackrel{\text{def}}{=} V_k^\top B V_k$. Note that this is not just a change of basis, since vectors in \mathcal{K} can be mapped by B to vectors outside \mathcal{K} . Now, since $v, Bv \in \mathcal{K}$, we must have $Bv = (V_k V_k^\top) B (V_k V_k^\top) v = V_k (V_k^\top B V_k) V_k^\top v = V_k T_k V_k^\top v$. Iterating this argument, we get that for

all $i \leq k$, $B^i v = V_k T_k^i V_k^\top v$, and hence, by linearity, $p(B)v = V_k p(T_k) V_k^\top v$, for any polynomial p of degree at most k .

Now, a natural approximation for $f(B)v$ is $V_k f(T_k) V_k^\top v$. Writing $r_k(x) \stackrel{\text{def}}{=} f(x) - p_k(x)$, where p_k is any degree k approximation to $f(x)$, the error in the approximation is $f(B)v - V_k f(T_k) V_k^\top v = r_k(B)v - V_k r_k(T_k) V_k^\top v$, for any choice of p_k . Hence, the norm of the error vector is at most $(\|r_k(B)\| + \|r_k(T_k)\|) \|v\|$, which is bounded by the value of r_k on the eigenvalues of B (eigenvalues of T_k are a subset of eigenvalues of B). More precisely, the norm of the error is bounded by $2 \|v\| \cdot \max_{\lambda \in \text{Spectrum}(B)} |f(\lambda) - p_k(\lambda)|$. Minimizing over p_k gives the error bound claimed above. Note that we do not explicitly need the approximating polynomial. It suffices to prove that there exists a degree k polynomial that uniformly approximates f well on an interval containing the spectrum of B and T_k .

If we construct the basis iteratively as above, $Bv_j \in \text{Span}\{v_0, \dots, v_{j+1}\}$ by construction, and if $i > j + 1$, v_i is orthogonal to this subspace and hence $v_i^\top (Bv_j) = 0$. Thus, T_k is Upper Hessenberg. Moreover, if B is symmetric, $v_j^\top (Bv_i) = v_i^\top (Bv_j)$, and hence T_k is symmetric and tridiagonal. This means that while constructing the basis, at step $i + 1$, it needs to orthonormalize Bv_i only w.r.t. v_{i-1} and v_i . Thus the total time required is $O((t_B + n)k)$, plus the time required for the computation of $f(T_k)$, which can typically be bounded by $O(k^2)$ for a tridiagonal matrix (using [26]). This completes an overview of the Lanczos method. The LANCZOS procedure described in Figure 4 in the main body, implements the Lanczos method. We now move on to describing how we apply it to obtain our two algorithms.

3.2.2 Approximating $\exp(-A)v$ Using a Rational Approximation to e^{-x}

Our Algorithm.

The starting point of the algorithm that underlies Theorem 1.2 is a rather surprising result by Saff, Schönhage and Varga (SSV) [30], which says that for any integer k , there exists a degree k polynomial p_k^* such that, $p_k^*((1 + x/k)^{-1})$ approximates e^{-x} up to an error of $O(k \cdot 2^{-k})$ over the interval $[0, \infty)$ (Theorem 6.8, Corollary 6.9). Then, to compute $\exp(-A)v$, one could apply the Lanczos method with $B \stackrel{\text{def}}{=} (I + A/k)^{-1}$ and $f(x) \stackrel{\text{def}}{=} e^{k(1-1/x)}$. Essentially, this was the method suggested by Eshof and Hochbruck [13]. The strong approximation guarantee of the SSV result along with the guarantee of the Lanczos method from the previous section, would imply that the order of the Krylov subspace for B required would be roughly $\log^{1/\delta}$, and hence, independent of $\|A\|$. The running time is then dominated by the computation $Bv = (I + A/k)^{-1}v$.

EH note that the computation of exact matrix inverse is a costly operation ($O(n^3)$ time in general) and all known faster methods for inverse computation incur some error. They suggest using the Lanczos method with faster iterative methods, e.g. Conjugate Gradient, for computing the inverse (or rather the product of the inverse with a given vector) as a *heuristic*. They make no attempt to give a theoretical justification of why approximate computation suffices. Also note that, even if the computation was error-free, a method such as Conjugate Gradient will have running time which varies with $\sqrt{\lambda_1(A)/\lambda_n(A)}$ in general. Thus, the EH method falls substantially short of resolving the hypothesis mentioned in the introduction.

To be able to prove Theorem 1.2 using the SSV guarantee, we have to adapt the Lanczos method in several ways, and hence, deviate from the method suggested by EH: 1) EH construct T_k as a tridiagonal matrix as Lanczos method suggests, but since the computation is no longer exact, the basis $\{v_i\}_{i=0}^k$ is no longer guaranteed to be orthonormal. As a result, the proofs of the Lanczos method break down. Our algorithm, instead, builds an orthonormal basis, which means

that T_k becomes an Upper Hessenberg matrix instead of tridiagonal and we need to compute k^2 dot products in order to compute T_k . 2) With T_k being asymmetric, several nice spectral properties are lost, *e.g.* real eigenvalues and an orthogonal set of eigenvectors. We overcome this fact by symmetrizing T_k to construct $\hat{T}_k = \frac{T_k + T_k^\top}{2}$ and computing our approximation with \hat{T}_k . This permits us to bound the quality of a polynomial approximation applied to \hat{T}_k by the behavior of the polynomial on the eigenvalues of \hat{T}_k . 3) Our analysis is based on the SSV approximation result, which is better than the variant proved and used by EH. Moreover, for their *shifting* technique, which is the source of the $\|\exp(-A)\|$ factor in the hypothesis, the given proof in EH is incorrect and it is not clear if the given bound could be achieved even under exact computation⁴. 4) Most importantly, since A is SDD, we are able to employ the Spielman-Teng solver (Theorem 6.10) to approximate $(I + A/k)^{-1}v$. This procedure, called EXPRATIONAL, has been described in Figure 5 in the main body.

Error Analysis.

To complete the proof of Theorem 1.2, we need to analyze the role of the error that creeps in due to approximate matrix inversion. The problem is that this error, generated in each iteration of the Krylov basis computation, propagates to the later steps. Thus, small errors in the inverse computation may lead to the basis V_k computed by our algorithm to be quite far from the k -th order Krylov basis for B, v . We first show that, assuming the error in computing the inverse is small, \hat{T}_k can be used to approximate degree k polynomials of $B = (I + A/k)^{-1}$ when restricted to the Krylov subspace, *i.e.* $\|p(B)v - V_k p(\hat{T}_k) V_k^\top v\| \lesssim \|p\|_1$. Here, if $p \stackrel{\text{def}}{=} \sum_{i=0}^k a_i \cdot x^i$, $\|p\|_1 = \sum_{i \geq 0} |a_i|$. This is the most technical part of the error analysis and unfortunately, the only way we know of proving the error bound above is by *tour de force*. A part of this proof is to show that the spectrum of \hat{T}_k cannot shift far from the spectrum of B .

To bound the error in the candidate vector output by the algorithm, *i.e.* $\|f(B)v - V_k f(\hat{T}_k) V_k^\top v\|$, we start by expressing e^{-x} as the sum of a degree k -polynomial p_k in $(1 + x/k)^{-1}$ and a remainder function r_k . We use the analysis from the previous paragraph to upper bound the error in the polynomial part by $\approx \|p\|_1$. We bound the contribution of the remainder term to the error by bounding $\|r_k(B)\|$ and $\|r_k(\hat{T}_k)\|$. This step uses the fact that eigenvalues of $r_k(\hat{T}_k)$ are $\{r_k(\lambda_i)\}_i$, where $\{\lambda_i\}_i$ are eigenvalues of \hat{T}_k . This is the reason our algorithm symmetrizes T_k to \hat{T}_k . To complete the error analysis, we use the polynomials p_k^* from SSV and bound $\|p_k^*\|_1$. Even though we do not know p_k^* explicitly, we can bound its coefficients indirectly by writing it as an interpolation polynomial. All these issues make the error analysis highly technical. However, since the error analysis is crucial for our algorithms, a more illuminating proof is highly desirable.

3.2.3 Approximation Using Our Polynomial Approximation to e^{-x}

More straightforwardly, combining the Lanczos method with the setting $B \stackrel{\text{def}}{=} A$ and $f(x) \stackrel{\text{def}}{=} e^{-x}$ along with the polynomial approximation to e^{-x} that we prove in Theorem 1.5, we get that setting $k \approx \sqrt{\lambda_1(A) - \lambda_n(A)} \cdot \text{poly}(\log 1/\delta)$ suffices to obtain a vector u that satisfies $\|\exp(-A)v - u\| \leq \delta \|v\| \|\exp(-A)\|$. This gives us our second method for approximating $\exp(-A)v$. Note that this

⁴EH show the existence of degree k polynomials in $(1 + vx)^{-1}$ for any constant $v \in (0, 1)$, that approximate e^{-x} up to an error of $\exp(1/2v - \Theta(\sqrt{k(v^{-1} - 1))))$. In order to deduce the claimed hypothesis, it needs to be used for $v \approx 1/\lambda_n(A)$, in which case, there is a factor of $e^{\lambda_n(A)}$ in the error, which could be huge.

algorithm avoids any inverse computation and, as a result, the procedure and the proofs are simpler and the algorithm practical.

3.3 Our Uniform Approximation for e^{-x}

In this section, we give a brief overview of the proof of Theorem 1.5. The details appear in Section 7 and can be read independently of the rest of the paper.

A straightforward approach to approximate e^{-x} over $[a, b]$ is by truncating its series expansion around $\frac{a+b}{2}$. With a degree of the order of $(b-a) + \log 1/\delta$, these polynomials achieve an error of $\delta \cdot e^{-(b+a)/2}$, for any constant $\delta > 0$. This approach is equivalent to approximating e^λ over $[-1, 1]$, for $\lambda \stackrel{\text{def}}{=} (b-a)/2$, by polynomials of degree $O(\lambda + \log 1/\delta)$. On the flip side, it is known that if λ is constant, the above result is optimal (see e.g. [31]). Instead of polynomials, one could consider approximations by rational functions, as in [12, 39]. However, the author in [31] shows that, if both λ and the degree of the denominator of the rational function are constant, the required degree of the numerator is only an additive constant better than that for the polynomials. It might seem that the question of approximating the exponential has been settled and one cannot do much better. However, the result by SSV mentioned before, seems surprising in this light. The lower bound does not apply to their result, since the denominator of their rational function is unbounded. In a similar vein, we ask the following question: If we are looking for weaker error bounds, e.g. $\delta \cdot e^{-a}$ instead of $\delta \cdot e^{-(b+a)/2}$ (recall $b > a$), can we improve on the degree bound of $O((b-a) + \log 1/\delta)$? Theorem 1.5 answers this question in the affirmative and gives a new upper bound and an almost matching lower bound. We give an overview of the proofs of both these results next.

Upper Bound.

We wish to show that there exists a polynomial of degree of the order of $\sqrt{b-a} \cdot \text{poly}(\log 1/\delta)$ that approximates e^{-x} on the interval $[a, b]$, up to an error of $\delta \cdot e^{-a}$ for any $\delta > 0$. Our approach is to approximate $(1 + x/k)^{-1}$ on the interval $[a, b]$, by a polynomial q of degree l , and then compose the polynomial p_k^* from the SSV result with q , to obtain $p_k^*(q(x))$ which is a polynomial of degree $k \cdot l$ approximating e^{-x} over $[a, b]$. Thus, we are looking for polynomials q that minimize $|q(x) - 1/x|$ over $[1 + a/k, 1 + b/k]$. Slightly modifying the optimization, we consider polynomials q that minimize $|x \cdot q(x) - 1|$ over $[1 + a/k, 1 + b/k]$. In Section 7, we show that the solution to this modified optimization can be derived from the well-known Chebyshev polynomials. For the right choice of k and l , the composition of the two polynomials approximates e^{-x} to within an error of $\delta \cdot e^{-a}$ over $[a, b]$, and has degree $\sqrt{b-a} \cdot \text{poly}(\log 1/\delta)$. To bound the error in the composition step, we need to bound the sum of absolute values of coefficients of p_k^* , which we achieve by rewriting p_k^* as an interpolation polynomial. The details appear in Section 7.

Lower Bound.

As already noted, since we consider a weaker error bound $\delta \cdot e^{-a}$ and $\lambda \stackrel{\text{def}}{=} (b-a)/2$ isn't a constant for our requirements, the lower bounds mentioned above no longer hold. Nevertheless, we prove that the square-root dependence on $b-a$ of the required degree is optimal. The proof is simple and we give the details here: Using a theorem of Markov from approximation theory (see [10]), we show that, any polynomial approximating e^{-x} over the interval $[a, b]$ up to an error of $\delta \cdot e^{-a}$, for some constant δ small enough, must have degree of the order of $\sqrt{b-a}$. Markov's theorem says that the absolute value of the derivative of a univariate polynomial p of degree k , which lives

in a box of height h over an interval of width w , is upper bounded by d^2h/w . Let p_k be a polynomial of degree k that $\delta \cdot e^{-a}$ -approximates e^{-x} in the interval $[a, b]$. If b is large enough and, δ a small enough constant, then one can get a lower bound of $\Omega(e^{-a})$ on the derivative of p_k using the Mean Value Theorem. Also, one can obtain an upper bound of $O(e^{-a})$ on the height of the box in which p_k lives. Both these bounds use the fact that p_k approximates e^{-x} and is $\delta \cdot e^{-a}$ close to it. Since the width of the box is $b - a$, these two facts, along with Markov's theorem, immediately imply a lower bound of $\Omega(\sqrt{b - a})$ on k . This shows that our upper bound is tight up to a factor of $\text{poly}(\log 1/\delta)$.

4 Discussion and Open Problems

In this paper, using techniques from disparate areas such as random walks, SDPs, numerical linear algebra and approximation theory, we have settled the question of designing an asymptotically optimal $\tilde{O}(m)$ spectral algorithm for BS (Theorem 1.1) and alongwith provided a simple and practical algorithm (Theorem 3.1). However, there are several outstanding problems that emerge from our work.

The main remaining open question regarding the design of spectral algorithms for BS is whether it is possible to obtain stronger certificates that no sparse balanced cuts exist, in nearly-linear time. This question is of practical importance in the construction of decompositions of the graph into induced graphs that are near-expanders, in nearly-linear time [38]. OV show that their certificate, which is of the same form as that of BALSEP, is stronger than the certificate of Spielman and Teng [38]. In particular, our certificate can be used to produce decompositions into components that are guaranteed to be subsets of induced expanders in G . However, this form of certificate is still much weaker than that given by RLE, which actually outputs an induced expander of large volume.

With regards to approximating the Matrix exponential, a computation which plays an important role in SDP-based algorithms, random walks, numerical linear algebra and quantum computing, settling the hypothesis remains the main open question. Further, as noted earlier, the error analysis plays a crucial role in making Theorem 1.2 and, hence, Theorem 1.1 work, but its proof is rather long and difficult. A more illuminating proof of this would be highly desirable.

Another question is to close the gap between the upper and lower bounds on polynomial approximations to e^{-x} over an interval $[a, b]$ in Theorem 1.5.

5 The Algorithm for Balanced Separator

In this section we provide our spectral algorithm BALSEP and prove Theorem 1.1. We also mention how Theorem 3.1 follows easily from the proof of Theorem 1.1 and Theorem 1.4. We first present the preliminaries for this section.

5.1 Basic Preliminaries

Instance Graph and Edge Volume.

We denote by $G = (V, E)$ the unweighted instance graph, where $V = [n]$ and $|E| = m$. We assume G is connected. We let $d \in \mathbb{R}^n$ be the degree vector of G , i.e. d_i is the degree of vertex i . For a subset $S \subseteq V$, we define the edge volume as $\text{vol}(S) \stackrel{\text{def}}{=} \sum_{i \in S} d_i$. The total volume of G is $2m$. The

conductance of a cut (S, \bar{S}) is defined to be $\phi(S) \stackrel{\text{def}}{=} |E(S, \bar{S})| / \min\{\text{vol}(S), \text{vol}(\bar{S})\}$, where $\text{vol}(S)$ is the sum of the degrees of the vertices in the set S . Moreover, a cut (S, \bar{S}) is b -balanced if $\min\{\text{vol}(S), \text{vol}(\bar{S})\} \geq b \cdot \text{vol}(V)$.

Special Graphs

We denote by K_V the complete graph with weight $d_i d_j / 2m$ between every pair $i, j \in V$. For $i \in V$, S_i is the star graph rooted at i , with edge weight of $d_i d_j / 2m$ between i and j , for all $j \in V$.

Graph matrices.

For an undirected graph $H = (V, E_H)$, let $A(H)$ denote the adjacency matrix of H , and $D(H)$ the diagonal matrix of degrees of H . The (combinatorial) Laplacian of H is defined as $L(H) \stackrel{\text{def}}{=} D(H) - A(H)$. Note that for all $x \in \mathbb{R}^V$, $x^\top L(H)x = \sum_{\{i,j\} \in E_H} (x_i - x_j)^2$. By D and L , we denote $D(G)$ and $L(G)$ respectively for the input graph G . Finally, the natural random walk over G has transition matrix $W \stackrel{\text{def}}{=} AD^{-1}$.

Vector and Matrix Notation.

We are working within the vector space \mathbb{R}^n . We will denote by I the identity matrix over this space. For a symmetric matrix A , we will use $A \succeq 0$ to indicate that A is positive semi-definite. The expression $A \succeq B$ is equivalent to $A - B \succeq 0$. For two matrices A, B of equal dimensions, let $A \bullet B \stackrel{\text{def}}{=} \text{Tr}(A^\top B) = \sum_{ij} A_{ij} \cdot B_{ij}$. We denote by $\{e_i\}_{i=1}^n$ the standard basis for \mathbb{R}^n . $\mathbf{0}$ and $\mathbf{1}$ will denote the all 0s and all 1s vectors respectively.

Fact 5.1 $L(K_V) = D - 1/2m \cdot D\mathbf{1}\mathbf{1}^\top D = D^{1/2}(I - 1/2m \cdot D^{1/2}\mathbf{1}\mathbf{1}D^{1/2})D^{1/2}$.

Embedding Notation.

We will deal with vector embeddings of G , where each vertex $i \in V$ is mapped to a vector $v_i \in \mathbb{R}^d$, for some $d \leq n$. For such an embedding $\{v_i\}_{i \in V}$, we denote by v_{avg} the mean vector, i.e. $v_{\text{avg}} \stackrel{\text{def}}{=} \sum_{i \in V} d_i / 2m \cdot v_i$. Given a vector embedding $\{v_i \in \mathbb{R}^d\}_{i \in V}$, recall that X is the Gram matrix of the embedding if $X_{ij} = v_i^\top v_j$. A Gram matrix X is always PSD, i.e., $X \succeq 0$. For any $X \in \mathbb{R}^{n \times n}$, $X \succeq 0$, we call $\{v_i\}_{i \in V}$ the *embedding corresponding to X* if X is the Gram matrix of $\{v_i\}_{i \in V}$. For $i \in V$, we denote by R_i the matrix such that $R_i \bullet X = \|v_i - v_{\text{avg}}\|^2$.

Fact 5.2 $\sum_{i \in V} d_i R_i \bullet X = \sum_{i \in V} d_i \|v_i - v_{\text{avg}}\|^2 = 1/2m \cdot \sum_{i < j} d_j d_i \|v_i - v_j\|^2 = L(K_V) \bullet X$.

5.2 AHK Random Walks

The random-walk processes used by our algorithm are continuous-time Markov processes [27] over V . In these processes, state transitions do not take place at specified discrete intervals, but follow exponential distributions described by a *transition rate matrix* $Q \in \mathbb{R}^{n \times n}$, where Q_{ij} specifies the rate of transition from vertex j to i . More formally, letting $p(\tau) \in \mathbb{R}^n$ be the probability distribution of the process at time $t \geq 0$, we have that $\partial p(\tau) / \partial \tau = Qp(\tau)$. Given a transition rate matrix ⁵ Q , the differential equation for $p(\tau)$ implies that $p(\tau) = e^{\tau Q} p(0)$. In this paper, we will be

⁵A matrix Q is a valid transition rate matrix if its diagonal entries are non-positive and its off-diagonal entries are non-negative. Moreover, it must be that $\mathbf{1}Q = 0$, to ensure that probability mass is conserved.

interested in a class of continuous-time Markov processes over V that take into account the edge structure of G . The simplest such process is the *heat kernel* process, which is defined as having transition rate matrix $Q = -(I - W) = -LD^{-1}$. The heat kernel can also be interpreted as the probability transition matrix of the following discrete-time random walk: sample a number of steps i from a Poisson distribution with mean τ and perform i steps of the natural random walk over G :

$$p(\tau) = e^{-\tau LD^{-1}} p(0) = e^{-\tau(I-W)} p(0) = e^{-\tau} \sum_{i=0}^{\infty} \frac{\tau^i}{i!} W^i p(0).$$

For the construction of our algorithm, we generalize the concept of heat kernel to a larger class of continuous-time Markov processes, which we name *Accelerated Heat Kernel* (AHK) processes. A process $\mathcal{H}(\beta)$ in this class is defined by a non-negative vector $\beta \in \mathbb{R}^n$ and the transition rate matrix of $\mathcal{H}(\beta)$ is $Q(\beta) \stackrel{\text{def}}{=} -(L + \sum_{i \in V} \beta_i L(S_i)) D^{-1}$. As this is the negative of a sum of Laplacian matrices, it is easy to verify that it is a valid transition rate matrix. The effect of adding the star terms to the transition rate matrix is that of accelerating the convergence of the process to stationary at vertices i with large value of β_i , as a large fraction of the probability mass that leaves these vertices is distributed uniformly over the edges. We denote by $P_\tau(\beta)$ the probability-transition matrix of $\mathcal{H}(\beta)$ between time 0 and τ , i.e. $P_\tau(0) = e^{\tau Q(\beta)}$.

Embedding View.

A useful matrix to study $\mathcal{H}(\beta)$ will be $D^{-1} P_{2\tau}(\beta)$. This matrix describes the probability distribution over the edges of G and has the advantage of being symmetric and positive semidefinite:

$$D^{-1} P_{2\tau}(\beta) = D^{-1/2} e^{-(2\tau) D^{-1/2} (L + \sum_{i \in V} \beta_i L(S_i)) D^{-1/2}} D^{-1/2},$$

Moreover, we have the following fact:

Fact 5.3 $D^{-1/2} P_\tau(\beta)$ is a square root of $D^{-1} P_{2\tau}(\beta)$.

Proof:

$$\left(D^{-1/2} P_\tau(\beta) \right)^\top D^{-1/2} P_\tau(\beta) = e^{\tau(Q(\beta))^\top} D^{-1} e^{\tau Q(\beta)} = D^{-1} e^{\tau Q(\beta)} e^{\tau Q(\beta)} = D^{-1} e^{2\tau Q(\beta)}.$$

■

Hence, $D^{-1} P_{2\tau}(\beta)$ is the Gram matrix of the embedding given by the columns of its square root $D^{-1/2} P_\tau(\beta)$. This property will enable us to use geometric SDP techniques to analyze $\mathcal{H}(\beta)$.

Mixing.

Spectral methods for finding low-conductance cuts are based on the idea that random walk processes mix slowly across sparse cuts, so that it is possible to detect such cuts by considering the starting vertices for which the probability distribution of the process strongly deviates from stationary. We measure this deviation for vertex i at time t by the ℓ_2^2 -norm of the distance between $P_\tau(\beta) e_i$ and the uniform distribution over the edges of G . We denote it by $\Psi(P_\tau(\beta), i)$:

$$\Psi(P_\tau(\beta), i) \stackrel{\text{def}}{=} d_i \sum_{j \in V} d_j \left(\frac{e_j^\top P_\tau(\beta) e_i}{d_j} - \frac{1}{2m} \right)^2$$

A fundamental quantity for our algorithm will be the total deviation from stationarity over a subset $S \subseteq V$. We will denote $\Psi(P_t(\beta), S) \stackrel{\text{def}}{=} \sum_{i \in S} \Psi(P_t(\beta), i)$. In particular, $\Psi(P_\tau(\beta), V)$ will play the role of potential function in our algorithm. The following facts express these mixing quantities in the geometric language of the embedding corresponding to $D^{-1}P_{2\tau}(\beta)$.

Fact 5.4 $\Psi(P_\tau(\beta), i) = d_i R_i \bullet D^{-1}P_{2\tau}(\beta)$.

Proof: By Fact 5.3 and the definition of R_i :

$$\begin{aligned} d_i R_i \bullet D^{-1}P_{2\tau}(\beta) &= d_i \left\| D^{-1/2}P_\tau(\beta)e_i - \sum_{j \in V} \frac{d_j}{2m} D^{-1/2}P_\tau(\beta)e_j \right\|^2 = d_i \left\| D^{-1/2}P_\tau(\beta)e_i - \frac{D^{1/2}\mathbf{1}}{2m} \right\|^2 \\ &= d_i \left\| D^{1/2} \left(D^{-1}P_\tau(\beta)e_i - \frac{\mathbf{1}}{2m} \right) \right\|^2 = \Psi(P_\tau(\beta), i). \end{aligned}$$

■

The following is a consequence of Fact 5.2:

Fact 5.5 $\Psi(P_\tau(\beta), V) = \sum_{i \in V} d_i R_i \bullet D^{-1}P_{2\tau}(\beta) = L(K_V) \bullet D^{-1}P_{2\tau}(\beta)$.

5.3 Algorithm Description

Preliminaries

All the random walks in our algorithm will be run for time $\tau \stackrel{\text{def}}{=} O(\log n)/\gamma$. We will consider embeddings given by the columns of $D^{-1/2}P_\tau(\beta)$ for some choice of β . Because we want our algorithm to run in time $\tilde{O}(m)$ and we are only interested in Euclidean distances between vectors in the embedding, we will use the Johnson-Lindenstrauss Lemma (see Lemma 5.18 in Section 5.5) to obtain an $O(\log n)$ -dimensional embedding approximately preserving distances between columns of $D^{-1/2}P_\tau(\beta)$ up to a factor of $(1 + \varepsilon)$, where ε is a constant such that $1 + \varepsilon/1 - \varepsilon \leq 4/3$.

Our algorithm BALSEP will call two subroutines FINDCUT and EXPV. FINDCUT is an SDP-rounding algorithm that uses random projections and radial sweeps to find a low-conductance cut, that is either c -balanced, for some constant $c = \Omega(b) \leq b/100$ defined in OV, or obeys a strong guarantee stated in Theorem 5.8. Such algorithm is implicit in [25] and is described precisely in Section 5.7. EXPV is a generic algorithm that approximately computes products of the form $P_\tau(\beta)u$ for unit vectors u . EXPV can be chosen to be either the algorithm implied by Theorem 3.2, which makes use of the Spielman-Teng solver, or that in Theorem 1.4, which just applies the Lanczos method.

We are now ready to describe BALSEP, which will output a c -balanced cut of conductance $O(\sqrt{\gamma})$ or the string NO, if it finds a certificate that no b -balanced cut of conductance less than γ exists. BALSEP can also fail and output the string FAIL. We will show that this only happens with small probability. The algorithm BALSEP is defined in Figure 1. The constants in this presentation are not optimized and are likely to be higher than what is necessary in practice. They can also be modified to obtain different trade-offs between the approximation guarantee and the output balance.

At iteration $t = 1$, we have $\beta^{(1)} = \mathbf{0}$, so that $P^{(1)}$ is just the probability transition matrix of the heat kernel on G for time τ . In general at iteration t , BALSEP runs EXPV to compute $O(\log n)$

Input: An unweighted connected instance graph $G = (V, E)$, a constant balance value $b \in (0, 1/2]$, a conductance value $\gamma \in [1/n^2, 1)$.

Let $S = \emptyset, \bar{S} = V$. Set $\tau = \log n / 12\gamma$ and $\beta^{(1)} = \mathbf{0}$.

At iteration $t = 1, \dots, T = 12 \log n$:

1. Denote $P^{(t)} \stackrel{\text{def}}{=} P_\tau(\beta^{(t)})$. Pick $k = O(\log n / \varepsilon^2)$ random unit vectors $\{u_1^{(t)}, u_2^{(t)}, \dots, u_k^{(t)} \in \mathbb{R}^n\}$ and use the subroutine EXPV to compute the embedding $\{v_i^{(t)} \in \mathbb{R}^k\}_{i \in V}$ defined as

$$\left(v_i^{(t)}\right)_j = \sqrt{\frac{n}{k}} u_j^\top D^{-1/2} P^{(t)} e_i.$$

Let $X^{(t)}$ be the Gram matrix corresponding to this embedding.

2. If $L(K_V) \bullet X^{(t)} = \sum_{i \in V} d_i \|v_i^{(t)} - v_{\text{avg}}^{(t)}\|^2 \leq \frac{1+\varepsilon}{n}$, output NO and terminate.
3. Otherwise, run $\text{FINDCUT}(G, b, \gamma, \{v_i^{(t)}\}_{i \in V})$. FINDCUT outputs a cut $S^{(t)}$ with $\phi(S^{(t)}) \leq O(\sqrt{\gamma})$ or fails, in which case we also output FAIL and terminate.
4. If $S^{(t)}$ is c -balanced, output $S^{(t)}$ and terminate. If not, update $S \stackrel{\text{def}}{=} S \cup S^{(t)}$. If S is c -balanced, output S and terminate.
5. Otherwise, update $\beta^{(t+1)} = \beta^{(t)} + \frac{72\gamma}{T} \sum_{i \in S^{(t)}} e_i$ and proceed to the next iteration.

Output NO and terminate.

Figure 1: The BALSEP Algorithm

random projections of $P^{(t)}$ and constructs an approximation $\{v_i^{(t)}\}_{i \in V}$ to the embedding given by the columns of $D^{-1/2} P^{(t)}$. This approximate embedding has Gram matrix $X^{(t)}$.

In Step 2, BALSEP computes $L(K_V) \bullet X^{(t)}$, which is an estimate of the total deviation $\Psi(P^{(t)}, V)$ by Fact 5.5. If this deviation is small, the AHK walk $P^{(t)}$ has mixed sufficiently over G to yield a certificate that G cannot have any b -balanced cut of conductance less than γ . This is shown in Lemma 5.6. If the AHK walk $P^{(t)}$ has not mixed sufficiently, we can use FINDCUT to find a cut $S^{(t)}$ of low conductance $O(\sqrt{\gamma})$, which is an obstacle for mixing. If $S^{(t)}$ is c -balanced, we output it and terminate. Similarly, if $S \cup S^{(t)}$ is c -balanced, as $\phi(S \cup S^{(t)}) \leq O(\sqrt{\gamma})$, we can also output $S \cup S^{(t)}$ and exit. Otherwise, $S^{(t)}$ is unbalanced and is potentially preventing BALSEP from detecting balanced cuts in G . We then proceed to modify the AHK walk, by increasing the values of $\beta^{(t+1)}$ for the vertices in $S^{(t)}$. This change ensures that $P^{(t+1)}$ mixes faster from the vertices in $S^{(t)}$ and in particular mixes across $S^{(t)}$. In particular, this means that, at any given iteration t , the support of $\beta^{(t)}$ is $\cup_{r=1}^{t-1} S^{(r)}$, which is an unbalanced set.

The BALSEP algorithm exactly parallels the RLE algorithm, introducing only two fundamental changes. First, we use the embedding given by the AHK random walk $P^{(t)}$ in place of the eigenvector to find cuts in G or in a residual graph. Secondly, rather than fully removing unbalanced low-conductance cuts from the graph, we modify $\beta^{(t)}$ at every iteration t , so $P^{(t+1)}$ at the next iteration mixes across the unbalanced cuts found so far.

5.4 Analysis

The analysis of BALSEP is at heart a modification of the MMWU argument in OV, stated in a random-walk language. This modification allows us to deal with the different embedding used by BALSEP at every iteration with respect to OV.

In this analysis, the quantity $\Psi(P^{(t)}, V)$ plays the role of potential function. Recall that, from a random-walk point of view, $\Psi(P^{(t)}, V)$ is the total deviation from stationarity of $P_\tau(\beta^{(t)})$ over all vertices as starting points. We start by showing that if the potential function is small enough, we obtain a certificate that no b -balanced cut of conductance at most γ exists. In the second step, we show that, if an unbalanced cut $S^{(t)}$ of low conductance is found, the potential decreases by a constant fraction. Unless explicitly stated otherwise, all proofs are found in Section 5.5.

Potential Guarantee.

We argue that, if $\Psi(P^{(t)}, V)$ is sufficiently small, it must be the case that G has no b -balanced cut of conductance less than γ . A similar result is implicit in OV. This theorem has a simple explanation in terms of the AHK random walk $P^{(t)}$. Notice that $P^{(t)}$ is accelerated only on a small unbalanced set S . Hence, if a balanced cut of conductance less than γ existed, its convergence could not be greatly helped by the acceleration over S . Then, if $P^{(t)}$ is still mixing very well, no such balanced cut can exist.

Lemma 5.6 *Let $S = \cup_{i=1}^t S^{(i)}$. If $\Psi(P^{(t)}, V) \leq \frac{4}{3n}$, and $\text{vol}(S) \leq c \cdot 2m \leq b/100 \cdot 2m$, then*

$$L + \sum_{i \in V} \beta_i^{(t)} L(S_i) \geq 3\gamma \cdot L(K_V).$$

Moreover, this implies that no b -balanced cut of conductance less than γ exists in G .

The Deviation of an Unbalanced Cut.

In the next step, we show that, if the walk has not mixed sufficiently, w.h.p. the embedding $\{v_i^{(t)}\}_{i \in V}$, computed by BALSEP, has low quadratic form with respect to the Laplacian of G . From a SDP-rounding perspective, this means that the embedding can be used to recover cuts of value close to γ . This part of the analysis departs from that of OV, as we use our modified definition of the embedding.

Lemma 5.7 *If $\Psi(P^{(t)}, V) \geq \frac{1}{n}$, then w.h.p. $L \bullet X^{(t)} \leq O(\gamma) \cdot L(K_V) \bullet X^{(t)}$.*

This guarantee on the embedding allows us to apply SDP-rounding techniques in the subroutine FINDCUT. The following result is implicit in [25]. Its proof appears in Section 5.7 for completeness.

Theorem 5.8 *Consider an embedding $\{v_i \in \mathbb{R}^d\}_{i \in V}$ with Gram matrix X such that $L \bullet X^{(t)} \leq \alpha L(K_V) \bullet X^{(t)}$, for $\alpha > 0$. On input $(G, b, \alpha, \{v_i\}_{i \in V})$, FINDCUT runs in time $\tilde{O}(md)$ and w.h.p. outputs a cut C with $\phi(C) \leq O(\sqrt{\alpha})$. Moreover, there is a constant $c = \Omega(b) \leq b/100$ such that either C is c -balanced or*

$$\sum_{i \in C} d_i R_i \bullet X \geq 2/3 \cdot L(K_V) \bullet X.$$

The following corollary is a simple consequence of Lemma 5.7 and Theorem 5.8:

Corollary 5.9 *At iteration t of BALSEP, if $\Psi(P^{(t)}, V) \geq \frac{1}{n}$ and $S^{(t)}$ is not c -balanced, then w.h.p. $\Psi(P^{(t)}, S) \geq 1/2 \cdot \Psi(P^{(t)}, V)$.*

In words, at the iteration t of BALSEP, the cut $S^{(t)}$ must either be c -balanced or be an unbalanced cut that contributes a large constant fraction of the total deviation of $P^{(t)}$ from the stationary distribution. In this sense, $S^{(t)}$ is the main reason for the failure of $P^{(t)}$ to achieve better mixing. To eliminate this obstacle and drive the potential further down, $P^{(t)}$ is updated to $P^{(t+1)}$ by accelerating the convergence to stationary from all vertices in $S^{(t)}$. Formally, this is achieved by adding weighted stars rooted at all vertices over $S^{(t)}$ to the transition-rate matrix of the AHK random walk $P^{(t)}$.

Potential Reduction.

The next theorem crucially exploits the stability of the process $\mathcal{H}(\beta^{(t)})$ and Corollary 5.9 to show that the potential decreases by a constant fraction at every iteration in which an unbalanced cut is found. More precisely, the theorem shows that accelerating the convergence from $S^{(t)}$ at iteration t of BALSEP has the effect of eliminating at least a constant fraction of the total deviation due to $S^{(t)}$. The proof is a simple application of the Golden-Thompson inequality [9] and mirrors the main step in the MMWU analysis.

Theorem 5.10 *At iteration t of BALSEP, if $\Psi(P^{(t)}, V) \geq \frac{1}{n}$ and $S^{(t)}$ is not c -balanced, then w.h.p.*

$$\Psi(P^{(t+1)}, V) \leq \Psi(P^{(t)}, V) - 1/3 \cdot \Psi(P^{(t)}, S^{(t)}) \leq 5/6 \cdot \Psi(P^{(t)}, V).$$

We are now ready to prove Theorem 1.1 and Theorem 3.1 by applying Lemma 5.10 to show that after $O(\log n)$ iterations, the potential must be sufficiently low to yield the required certificate according to Lemma 5.6.

Proof: [Proof of Theorem 1.1] If BALSEP outputs a cut S in Step 4, by construction, we have that $\phi(S) \leq O(\sqrt{\gamma})$ and S is $\Omega(b)$ -balanced. Alternatively, at iteration t , if $L(K_V) \bullet X^{(t)} \leq 1+\epsilon/n$, we have by Lemma 5.18 that

$$\Psi(P^{(t)}, V) = L(K_V) \bullet D^{-1}P_{2\tau}(\beta^{(t)}) \leq \frac{1}{1-\epsilon} L(K_V) \bullet X^{(t)} \leq \frac{1+\epsilon}{1-\epsilon} \cdot \frac{1}{1n} \leq \frac{4}{3n}.$$

Therefore, by Lemma 5.6, we have a certificate that no b -balanced cut of conductance less than γ exists in G . Otherwise, we must have $L(K_V) \bullet X^{(t)} \geq 1+\epsilon/n$, which, by Lemma 5.18, implies that

$$\Psi(P^{(t)}, V) \geq 1/n.$$

Then, by Lemma 5.7 and Theorem 5.8, we have w.h.p. that FINDCUT does not fail and outputs a cut $S^{(t)}$ with $\phi(S^{(t)}) \leq O(\sqrt{\gamma})$. As BALSEP has not terminated in Step 4, it must be the case that $S^{(t)}$ is not c -balanced and, by Theorem 5.10, we obtain that w.h.p. $\Psi(P^{(t+1)}, V) \leq 5/6 \cdot \Psi(P^{(t)}, V)$. Now,

$$\Psi(P^{(1)}, V) = L(K_V) \bullet D^{-1}P_{2\tau}(\mathbf{0}) \leq I \bullet P_{2\tau}(\mathbf{0}) \leq n.$$

Hence, after $2 \log n / \log(6/5) \leq 12 \log n = T$ iterations, w.h.p. we have that $\Psi(P^{(T)}, V) \leq 1/n$ and, by Lemma 5.6, no b -balanced cut of conductance less than γ exists.

We now consider the running time required by the algorithm at every iteration. In Step 1, we compute $k = O(\log n)$, products of the form $D^{-1/2}P^{(t)}u$, where u is a unit vector, using the

EXPV algorithm based on the Spielman-Teng solver, given in Theorem 3.2. This application of Theorem 3.2 is explained in Section 3.2. By the definition of $\beta^{(t)}$, at iteration t we have:

$$\begin{aligned} \|HMH\| &= \left\| \tau D^{-1/2} \left(L + \sum_{i \in V} d_i/2m \cdot \beta_i D + \sum_{i \in V} d_i/2m \cdot \beta_i e_i e_i^\top \right) D^{-1/2} \right\| \leq \left\| \tau D^{-1/2} (L + 2 \cdot 72 \cdot \gamma D) D^{-1/2} \right\| \\ &\leq O(\tau) = \text{poly}(n). \end{aligned}$$

Moreover, it is easy to see that our argument is robust up to an error $\delta = 1/\text{poly}(n)$ in this computation and the sparsity of M is $O(m)$ so that the running time of a single matrix-exponential-vector product is $\tilde{O}(m)$. Given the embedding produced by Step 1, $L(K_V) \bullet X^{(t)}$ can be computed in time $\tilde{O}(nk) = \tilde{O}(n)$ by computing the distances $\|v_i^{(t)} - v_{\text{avg}}^{(t)}\|^2$ for all $i \in V$. By Theorem 5.8, Step 3 runs in time $\tilde{O}(mk) = \tilde{O}(m)$. Finally, both Steps 4 and 5 can be performed in time $\tilde{O}(m)$. As there are at most $O(\log n)$ iterations, the theorem follows. ■

Theorem 3.1 is proved similarly. It suffices to show that a single matrix-exponential-vector product requires time $\tilde{O}(m/\sqrt{\gamma})$.

Proof: [Proof of Theorem 3.1] Using the algorithm of Theorem 1.4, we obtain that $\|A\| \leq O(\tau)$, so that $k = \tilde{O}(\sqrt{\tau}) = \tilde{O}(1/\sqrt{\gamma}) \leq \tilde{O}(n)$. Hence, the running time of a single computation for this method is $\tilde{O}(m\sqrt{\tau}) = \tilde{O}(m/\sqrt{\gamma})$. ■

5.5 Proofs

In this section we provide the proofs from the Section 5. We start with some preliminaries.

5.5.1 Preliminaries

Vector and Matrix Notation.

For a symmetric matrix A , denote by $\lambda_i(A)$, the i^{th} smallest eigenvalue of A . For a vector $x \in \mathbb{R}^n$, let $\text{supp}(x)$ be the set of vertices where x is not zero.

Fact 5.11 $L \preceq 2 \cdot D$ and $L(S_i) \preceq 2 \cdot D$.

Fact 5.12 For all $i \in V$, $L(S_i) = d_i/2m \cdot L(K_V) + d_i R_i$. In particular, $L(S_i) \succeq d_i R_i$.

Notation for BALSEP.

At iteration t , we denote

$$C^{(t)} \stackrel{\text{def}}{=} D^{-1/2} Q(\beta^{(t)}) D^{1/2} = D^{-1/2} \left(L + \sum_{i \in V} \beta_i^{(t)} L(S_i) \right) D^{-1/2}.$$

The following are useful facts to record about $C^{(t)}$:

Fact 5.13 The vector $D^{1/2}$ is the eigenvector of $C^{(t)}$ with smallest eigenvalue 0.

Fact 5.14 $C^{(t)} \preceq O(1) \cdot I$.

5.5.2 Useful Lemmata

Lemma 5.15 $\Psi(P^{(t)}, V) = L(K_V) \bullet D^{-1} P_{2\tau}(\beta^{(t)}) = \text{Tr}(e^{-2\tau C^{(t)}}) - 1$.

Proof: By definition, we have

$$\Psi(P^{(t)}, V) = L(K_V) \bullet D^{-1} P_{2\tau}(\beta^{(t)}) = L(K_V) \bullet D^{-1} e^{-2\tau Q(\beta^{(t)})} = L(K_V) \bullet D^{-1/2} e^{-2\tau C^{(t)}} D^{-1/2}.$$

Using Fact 5.1 and the cyclic property of the trace function, we obtain

$$L(K_V) \bullet D^{-1/2} e^{-2\tau C^{(t)}} D^{-1/2} = (I - 1/2m D^{1/2} \mathbf{1} D^{1/2}) \bullet e^{-2\tau C^{(t)}}.$$

Finally, by Fact 5.13, we must have that the right-hand side equals $\text{Tr}(e^{-2\tau C^{(t)}}) - 1$, as required. \blacksquare

The following lemma is a simple consequence of the convexity of e^{-x} . It is proved in [23].

Lemma 5.16 For a symmetric matrix $A \in \mathbb{R}^{n \times n}$ such that $\rho I \succeq A \succeq 0$ and $\tau > 0$, we have

$$e^{-\tau A} \preceq \left(I - \frac{(1 - e^{-\tau \rho})}{\rho} A \right).$$

The following are standard lemmata.

Lemma 5.17 (Golden-Thompson inequality [9]) Let $X, Y \in \mathbb{R}^{n \times n}$ be symmetric matrices. Then,

$$\text{Tr} \left(e^{X+Y} \right) \leq \text{Tr} \left(e^X e^Y \right).$$

Lemma 5.18 (Johnson-Lindenstrauss) Given an embedding $\{v_i \in \mathbb{R}^d\}_{i \in V}$, $V = [n]$, let u_1, u_2, \dots, u_k , be vectors sampled independently uniformly from the $n - 1$ -dimensional sphere of radius $\sqrt{n/k}$. Let U be the $k \times t$ matrix having the vector u_i as i^{th} row and let $\tilde{v}_i \stackrel{\text{def}}{=} U v_i$. Then, for $k_\epsilon \stackrel{\text{def}}{=} O(\log n / \epsilon^2)$, for all $i, j \in V$

$$(1 - \epsilon) \cdot \|v_i - v_j\|^2 \leq \|\tilde{v}_i - \tilde{v}_j\|^2 \leq (1 + \epsilon) \cdot \|v_i - v_j\|^2.$$

5.5.3 Proof of Lemma 5.6

Proof:

Let $S = \cup_{i=1}^t S^{(i)}$ and set $\beta \stackrel{\text{def}}{=} \beta^{(t)}$. By Lemma 5.15, we have $\text{Tr}(e^{-2\tau C^{(t)}}) - 1 \leq 4/3n$. Hence, $\lambda_{n-1}(e^{-2\tau C^{(t)}}) \leq 4/3n$, which implies that, by taking logs,

$$\lambda_2(C^{(t)}) \geq \frac{\log n}{4\tau} \geq 3\gamma.$$

This can be rewritten in matrix terms, by Fact 5.1 and Fact 5.13, and because $\text{supp}(\beta) = S$ by the construction of BALSEP:

$$L + \sum_{i \in S} \beta_i^{(t)} L(S_i) \succeq 3\gamma \cdot L(K_V). \quad (1)$$

which proves the first part of the Lemma.

For the second part, we start by noticing that, for $i \in S$, $\beta_i^{(t)} \leq 72\gamma \cdot t/T \leq 72\gamma$. Now for any b -balanced cut U , with $\text{vol}(U) \leq \text{vol}(\bar{U})$, consider the vector x_U defined as

$$(x_U)_i \stackrel{\text{def}}{=} \begin{cases} \sqrt{\frac{1}{2m} \cdot \frac{\text{vol}(\bar{U})}{\text{vol}(U)}} & \text{for } i \in U \\ -\sqrt{\frac{1}{2m} \cdot \frac{\text{vol}(U)}{\text{vol}(\bar{U})}} & \text{for } i \in \bar{U} \end{cases}$$

Applying the guarantee of Equation 1, we obtain

$$x_U^\top L x_U + 72\gamma \cdot \sum_{i \in S} x_U^\top L(S_i) x_U \geq 3\gamma \cdot x_U^\top L(K_V) x_U.$$

Notice that

$$\begin{aligned} x_U^\top L x_U &= \sum_{\{i,j\} \in E} ((x_U)_i - (x_U)_j)^2 = \frac{2m}{\text{vol}(\bar{U})} \cdot \frac{|E(U, \bar{U})|}{\text{vol}(U)} \leq 2 \cdot \phi(U), \\ x_U^\top L(S_i) x_U &= \sum_{j \in V} \frac{d_j}{2m} ((x_U)_i - (x_U)_j)^2 \leq \frac{2m}{\text{vol}(U)} \cdot \frac{d_i}{2m} = \frac{d_i}{\text{vol}(U)}, \\ x_U^\top L(K_V) x_U &= \sum_{i < j \in V} \frac{d_j d_i}{2m} ((x_U)_i - (x_U)_j)^2 = 1. \end{aligned}$$

Hence, our guarantee becomes

$$2\phi(U) + 72\gamma \cdot \frac{\text{vol}(S)}{\text{vol}(U)} \geq 3\gamma.$$

As $\text{vol}(S) \leq b/100 \cdot 2m \leq \text{vol}(U)/100$, we have $\phi(U) \geq \gamma$. ■

5.5.4 Proof of Lemma 5.7

Proof: Consider $L \bullet D^{-1} P_{2\tau}(\beta^{(t)}) = L \bullet D^{-1/2} e^{-2\tau C^{(t)}} D^{-1/2}$. Using the cyclic property of trace and the definition of $C^{(t)}$, we have that

$$L \bullet D^{-1} P_{2\tau}(\beta^{(t)}) \leq C^{(t)} \bullet e^{-2\tau C^{(t)}}.$$

We now consider the spectrum of $C^{(t)}$. By Fact 5.13, the smallest eigenvalue is 0. Let the remaining eigenvalues be $\lambda_2 \leq \lambda_3 \leq \dots \leq \lambda_n$. Then, $C^{(t)} \bullet e^{-2\tau C^{(t)}} = \sum_{i=2}^n \lambda_i e^{-2\tau \lambda_i}$. We will analyze these eigenvalues in two groups. For the first group, we consider eigenvalues smaller than 24γ and use Lemma 5.15, together with the fact that $\gamma \geq 1/n^2$:

$$\sum_{i: \lambda_i \leq 24\gamma} \lambda_i e^{-2\tau \lambda_i} \geq \frac{1}{n^2} \cdot (\text{Tr}(e^{-2\tau C^{(t)}}) - 1) = \frac{1}{n^3}.$$

For the remaining eigenvalues, we have, by Lemma 5.14::

$$\sum_{i: \lambda_i \geq 24\gamma} \lambda_i e^{-2\tau \lambda_i} \leq O(1) \cdot n \cdot e^{-2\tau 24\gamma} \leq \frac{O(1)}{n^3}.$$

Combining these two parts, we have:

$$L \bullet D^{-1} P_{2\tau}(\beta^{(t)}) \leq C^{(t)} \bullet e^{-2\tau C^{(t)}} \leq O(1) \cdot \sum_{i:\lambda_i \leq 24\gamma} \lambda_i e^{-2\tau \lambda_i} \leq O(\gamma) \cdot L(K_V) \bullet D^{-1} P_{2\tau}(\beta^{(t)}).$$

Now, we apply the Johnson-Lindenstrauss Lemma (Lemma 5.18) to both sides of this inequality to obtain:

$$L \bullet X^{(t)} \leq O(\gamma) \cdot L(K_V) \bullet X^{(t)}.$$

5.5.5 Proof of Corollary 5.9

Proof: By Lemma 5.7 and Theorem 5.8, we have that $S^{(t)}$ w.h.p. is either c -balanced or $\sum_{i \in S^{(t)}} d_i R_i \bullet X^{(t)} \geq 2/3 \cdot L(K_V) \bullet X^{(t)}$. By Lemma 5.18 and as $1+\varepsilon/1-\varepsilon \leq 4/3$, we have w.h.p.:

$$\begin{aligned} \Psi(P^{(t)}, S^{(t)}) &= \sum_{i \in S^{(t)}} d_i R_i \bullet D^{-1} P_{2\tau}(\beta^{(t)}) \geq \frac{1}{1+\varepsilon} \cdot \left(\sum_{i \in S^{(t)}} d_i R_i \bullet X^{(t)} \right) \geq \frac{2}{3} \cdot L(K_V) \bullet X^{(t)} \geq \\ &\frac{2}{3} \cdot \frac{1-\varepsilon}{1+\varepsilon} \cdot L(K_V) \bullet D^{-1} P_{2\tau}(\beta^{(t)}) = \frac{2}{3} \cdot \frac{1-\varepsilon}{1+\varepsilon} \cdot \Psi(P^{(t)}, V) \geq \frac{1}{2} \cdot \Psi(P^{(t)}, V). \end{aligned}$$

5.5.6 Proof of Theorem 5.10

Proof: By Lemma 5.15 and the Golden-Thompson inequality in Lemma 5.17:

$$\Psi(P^{(t+1)}, V) = \text{Tr}(e^{-2\tau C^{(t+1)}}) - 1 \leq \text{Tr} \left(e^{-2\tau C^{(t)}} e^{-2\tau D^{-1/2} \left(\frac{\gamma}{T} \sum_{i \in S^{(t)}} L(S_i) \right) D^{-1/2}} \right) - 1.$$

We now apply Lemma 5.16 to the second term under trace. To do this we notice that $\sum_{i \in S^{(t)}} L(S_i) \preceq 2L(K_V) \preceq 2D$, so that

$$D^{-1/2} \left(\frac{72\gamma}{T} \sum_{i \in S^{(t)}} L(S_i) \right) D^{-1/2} \preceq \frac{144\gamma}{T} I.$$

Hence, we obtain

$$\Psi(P^{(t+1)}, V) \leq \text{Tr} \left(e^{-2\tau C^{(t)}} \left(I - (1 - e^{-288 \cdot \tau\gamma/T}) \cdot \frac{1}{2} D^{-1/2} \left(\sum_{i \in S^{(t)}} L(S_i) \right) D^{-1/2} \right) \right) - 1.$$

Applying the cyclic property of trace, we get

$$\Psi(P^{(t+1)}, V) \leq \Psi(P^{(t)}, V) - \frac{(1 - e^{-288 \cdot \tau\gamma/T})}{2} \sum_{i \in S^{(t)}} L(S_i) \bullet D^{-1} P_{2\tau}(\beta^{(t)}).$$

Next, we use Fact 5.12 to replace $L(S_i)$ by R_i and notice that $288 \cdot \tau\gamma/T = 2$:

$$\Psi(P^{(t+1)}, V) \leq \Psi(P^{(t)}, V) - \frac{(1 - e^{-2})}{2} \sum_{i \in S^{(t)}} d_i R_i \bullet D^{-1} P_{2\tau}(\beta^{(t)}).$$

Then, we apply the definition of $\Psi(P^{(t)}, S)$:

$$\Psi(P^{(t+1)}, V) \leq \Psi(P^{(t)}, V) - 1/3 \cdot \Psi(P^{(t)}, S).$$

Finally, by Corollary 5.9, we know that w.h.p. $\Psi(P^{(t)}, S^{(t)}) \geq 1/2 \cdot \Psi(P^{(t)}, V)$ and the required result follows. \blacksquare

5.6 SDP Interpretation

OV designed an algorithm that outputs either a $\Omega(b)$ -balanced cut of conductance $O(\sqrt{\gamma})$ or a certificate that no b -balanced cut of conductance γ exists in time $\tilde{O}(m/\gamma^2)$. This algorithm uses the MMWU of Arora and Kale [7] to approximately solve an SDP formulation of the BS problem. The main technical contribution of their work is the routine `FINDCUT` (implicit in their `ORACLE`), which takes the role of an approximate separation oracle for their SDP. In an iteration of their algorithm, OV use the MMWU update to produce a candidate SDP-solution $Y^{(t)}$. In one scenario, $Y^{(t)}$ does not have sufficiently low Laplacian objective value:

$$L \bullet Y^{(t)} \geq \Omega(\gamma)L(K_V) \bullet Y^{(t)}. \quad (2)$$

In this case, the MMWU uses Equation 2 to produce a candidate solution $Y^{(t+1)}$ with lower objective value. Otherwise, `FINDCUT` is run on the embedding corresponding to $Y^{(t)}$. By Theorem 5.8, this yields either a cut of the required balance or a dual certificate that $Y^{(t)}$ is infeasible. This certificate has the form

$$\gamma \cdot \sum_{i \in S^{(t)}} d_i R_i \bullet Y^{(t)} \geq \Omega(\gamma)L(K_V) \bullet Y^{(t)} \quad (3)$$

and is used by the update to construct the next candidate $Y^{(t+1)}$. The number of iterations necessary is determined by the width of the two possible updates described above. A simple calculation shows that the width of the update for Equation 2 is $\Theta(1)$, while for Equation 3, it is only $O(\gamma)$. Hence, the overall width is $\Theta(1)$, implying that $O(\log n/\gamma)$ iterations are necessary for the algorithm of OV to produce a dual certificate that the SDP is infeasible and therefore no b -balanced cut of conductance γ exists.

Our modification of the update is based on changing the starting candidate solutions from $Y^{(1)} \propto D^{-1}$ to $X^{(1)} \propto D^{-1/2}e^{-2\tau D^{-1/2}LD^{-1/2}}D^{-1/2}$. In Lemma 5.6 and Lemma 5.7, we show that this modification implies that all $X^{(t)}$ must now have $L \bullet X^{(t)} \leq O(\gamma) \cdot L(K_V) \bullet X^{(t)}$ or else we find a dual certificate that the SDP is infeasible. This additional guarantee effectively allows us to bypass the update of Equation 2 and only work with updates of the form given in Equation 3. As a result, our width is now $O(\gamma)$ and we only require $O(\log n)$ iterations.

Another way to interpret our result is that all possible $\tau \approx \log n/\gamma$ updates of the form of Equation 2 in the algorithm of OV are regrouped into a single step, which is performed at the beginning of the algorithm.

5.7 The `FINDCUT` Subroutine

Most of the material in this Section appears in [25] or in [23]. We reproduce it here in the language of this paper for completeness. The constants in these proofs are not optimized.

5.7.1 Preliminaries

Fact 5.19 For a subset $S \subseteq V$,

$$\sum_{i \in S} d_i R_i \succeq \frac{\text{vol}(S)}{2m} (L(K_V) - L(K_S)).$$

Proof: By Fact 5.12,

$$\sum_{i \in \bar{S}} d_i R_i = \sum_{i \in \bar{S}} L(S_i) - \frac{\text{vol}(\bar{S})}{2m} L(K_V).$$

Moreover, by the definitions it is clear that

$$\sum_{i \in \bar{S}} L(S_i) + \frac{\text{vol}(S)}{2m} L(K_S) \succeq L(K_V).$$

Combining these two equations, we obtain the required statement. \blacksquare

The following is a variant of the sweep cut argument of Cheeger's inequality [11], tailored to ensure that a constant fraction of the variance of the embedding is contained inside the output cut.

Lemma 5.20 *Let $x \in \mathbb{R}^n, x \geq 0$, such that $x^\top Lx \leq \lambda$ and $\text{vol}(\text{supp}(x)) \leq 2m/2$. Relabel the vertices so that $x_1 \geq x_2 \geq \dots \geq x_{z-1} > 0$ and $x_z = \dots = x_n = 0$. For $i \in [z-1]$, denote by $S_i \subseteq V$, the sweep cut $\{1, 2, \dots, i\}$. Further, assume that $\sum_{i=1}^n d_i x_i^2 \leq 1$, and, for some fixed $k \in [z-1]$, $\sum_{i=k}^n d_i x_i^2 \geq \sigma$. Then, there is a sweep cut S_h of x such that $z-1 \geq h \geq k$ and $\phi(S_h) \leq 1/\sigma \cdot \sqrt{2\lambda}$.*

We will also need the following simple fact.

Fact 5.21 *Given $v, u, t \in \mathbb{R}^h$, $(\|v-t\| - \|u-t\|)^2 \leq \|v-u\|^2$.*

5.7.2 Roundable Embeddings and Projections

The following definition of *roundable embedding* captures the case in which a vector embedding of the vertices V highlights a balanced cut of conductance close to α in G . Intuitively, in a roundable embedding, a constant fraction of the total variance is spread over a large set R of vertices.

Definition 5.22 (Roundable Embedding) *Given an embedding $\{v_i\}_{i \in V}$ with Gram matrix X , denote by Ψ the total variance of the embedding: $\Psi \stackrel{\text{def}}{=} L(K_V) \bullet X$. Also, let $R = \{i \in V : \|v_i - v_{\text{avg}}\|^2 \leq 32 \cdot (1-b)/b \cdot \frac{\Psi}{2m}\}$. For $\alpha > 0$, we say that $\{v_i\}_{i \in V}$ is roundable for (G, b, α) if:*

- $L \bullet X \leq \alpha \Psi$,
- $L(K_R) \bullet X \geq \frac{\Psi}{128}$.

A roundable embedding can be converted into a balanced cut of conductance $O(\sqrt{\alpha})$ by using a standard projection rounding, which is a simple extension of an argument already appearing in [8] and [7]. The rounding procedure PROJROUND is described in Figure 2 for completeness. It is analyzed in [25] and [23], where the following theorem is proved.

Theorem 5.23 (Rounding Roundable Embeddings) [25, 23] *If $\{v_i \in \mathbb{R}^h\}_{i \in V}$ is roundable for (G, b, α) , then PROJROUND($\{v_i\}_{i \in V}, b$) produces a $\Omega(b)$ -balanced cut of conductance $O(\sqrt{\alpha})$ with high probability in time $\tilde{O}(nh + m)$.*

1. **Input:** An embedding $\{v_i \in \mathbb{R}^h\}_{i \in V}$, $b \in (0, 1/2]$.
2. Let $c = \Omega(b) \leq b/100$ be a constant, fixed in the proof of Theorem 5.23 in [25].
3. For $t = 1, 2, \dots, O(\log n)$:
 - a. Pick a unit vector u uniformly at random from S^{h-1} and let $x \in \mathbb{R}^n$ with $x_i \stackrel{\text{def}}{=} \sqrt{h} \cdot u^\top v_i$.
 - b. Sort the vector x . Assume w.l.o.g. that $x_1 \geq x_2 \geq \dots \geq x_n$. Define $S_i \stackrel{\text{def}}{=} \{j \in [n] : x_j \geq x_i\}$.
 - c. Let $S^{(t)} \stackrel{\text{def}}{=} (S_i, \bar{S}_i)$ which minimizes $\phi(S_i)$ among sweep-cuts for which $\text{vol}(S_i) \in [c \cdot 2m, (1-c) \cdot 2m]$.
4. **Output:** The cut $S^{(t)}$ of least conductance over all choices of t .

Figure 2: PROJROUND

5.7.3 Description of FINDCUT

In this subsection we describe the subroutine FINDCUT and prove Theorem 5.8.

Theorem 5.24 (Theorem 5.8 Restated) Consider an embedding $\{v_i \in \mathbb{R}^d\}_{i \in V}$ with Gram matrix X such that $L \bullet X^{(t)} \leq \alpha L(K_V) \bullet X^{(t)}$, for $\alpha > 0$. On input $(G, b, \alpha, \{v_i\}_{i \in V})$, FINDCUT runs in time $\tilde{O}(md)$ and w.h.p. outputs a cut C with $\phi(C) \leq O(\sqrt{\alpha})$. Moreover, there is a constant $c = \Omega(b) \leq b/100$ such that either C is c -balanced or

$$\sum_{i \in C} d_i R_i \bullet X \geq 2/3 \cdot L(K_V) \bullet X.$$

Proof: By Markov's inequality, $\text{vol}(\bar{R}) \leq b/(32 \cdot (1-b)) \cdot 2m \leq b/16 \cdot 2m \leq 1/32 \cdot 2m$. By assumption, CASE 1 cannot take place. If CASE 2 holds, then the embedding is roundable: by Theorem 5.23, PROJROUND outputs an $\Omega(b)$ -balanced cut C with conductance $O(\sqrt{\alpha})$. If this is not the case, we are in CASE 3.

We then have $L(K_R) \leq \Psi/128$ and, by Fact 5.19:

$$\begin{aligned} \sum_{i \in \bar{R}} d_i R_i \bullet X &= \sum_{i \in \bar{R}} d_i r_i^2 \geq \frac{\text{vol}(\bar{R})}{2m} \cdot \left(1 - \frac{1}{128}\right) \cdot \Psi \geq \left(1 - \frac{1}{32}\right) \cdot \left(1 - \frac{1}{128}\right) \cdot \Psi \\ &\geq \left(1 - \frac{5}{128}\right) \cdot \Psi. \end{aligned}$$

It must be the case that $\bar{R} = S_g$ for some $g \in [n]$, with $g \leq z$ as $\text{vol}(S_g) \leq \text{vol}(S_z)$. Let $k \leq z$ be the vertex in \bar{R} such that $\sum_{j=1}^k d_j r_j^2 \geq 3/4 \cdot (1 - 5/128)$ and $\sum_{j=k}^g d_j r_j^2 \geq 1/4 \cdot (1 - 5/128)$. By the definition of z , we have $k \leq g < z$ and $r_z^2 \leq 4/b \cdot \Psi/2m \leq 8 \cdot (1-b)/b \cdot \Psi/2m$. Hence, we have $r_z \leq 1/2 \cdot r_i$, for all $i \geq g$. Define the vector x as $x_i \stackrel{\text{def}}{=} (r_i - r_z)$ for $i \in S_z$ and $r_i \stackrel{\text{def}}{=} 0$ for $i \notin S_z$. Notice

1. **Input:** Instance graph G , balance b , conductance value α and embedding $\{v_i\}_{i \in V}$, with Gram matrix X .
2. Let $r_i = \|v_i - v_{\text{avg}}\|$ for all $i \in V$. Denote $\Psi \stackrel{\text{def}}{=} L(K_V) \bullet X$ and define the set $R \stackrel{\text{def}}{=} \{i \in V : r_i^2 \leq 32 \cdot (1-b)/b \cdot \Psi/2m\}$.
3. CASE 1: If $L \bullet X > \alpha\Psi$, output FAIL and terminate.
4. CASE 2: If $L(K_R) \bullet X \geq \Psi/128$, the embedding $\{v_i\}_{i \in V}$ is roundable for (G, b, α) . Run PROJROUND, output the resulting cut and terminate.
5. CASE 3: Relabel the vertices of V such that $r_1 \geq r_2 \geq \dots \geq r_n$ and let $S_i = \{1, \dots, i\}$ be the i^{th} sweep cut of r . Let z the smallest index such that $\text{vol}(S_z) \geq b/4 \cdot 2m$. Output the most balanced sweep cut C among $\{S_1, \dots, S_{z-1}\}$, such that $\phi(C) \leq 40 \cdot \sqrt{\gamma}$.

Figure 3: FINDCUT

that:

$$\begin{aligned} x^\top Lx &= \sum_{\{i,j\} \in E} (x_i - x_j)^2 \leq \sum_{\{i,j\} \in E} (r_i - r_j)^2 \\ &\stackrel{\text{Fact 5.21}}{\leq} \sum_{\{i,j\} \in E} \|v_i - v_j\|^2 \leq \alpha\Psi. \end{aligned}$$

Also, $x \geq 0$ and $\text{vol}(\text{supp}(x)) \leq b/4 \cdot 2m \leq 2m/2$, by the definition of z . Moreover,

$$\sum_{i=1}^n d_i x_i^2 = \sum_{i=1}^z d_i (r_i - r_z)^2 \leq \sum_{i=1}^z d_i r_i^2 \leq \Psi,$$

and

$$\begin{aligned} \sum_{i=k}^n d_i x_i^2 &= \sum_{i=k}^z d_i (r_i - r_z)^2 \\ &\geq \sum_{i=k}^g d_i (r_i - 1/2 \cdot r_i)^2 \\ &= 1/4 \cdot \sum_{i=k}^g d_i r_i^2 \\ &\geq 1/16 \cdot (1 - 5/128) \cdot \Psi \geq 1/20 \cdot \Psi. \end{aligned}$$

Hence we can now apply Lemma 5.20 to the vector $1/\Psi \cdot x$. This shows that there exists a sweep cut S_h with $z > h \geq k$, such that $\phi(S_h) \leq 40 \cdot \sqrt{\gamma}$. It also shows that C , as defined in Figure 3, must exist. Moreover, it must be the case that $S_k \subseteq S_h \subseteq C$. As $h \geq k$, we have

$$\sum_{i \in C} d_i R_i \bullet X = \sum_{i \in C} d_i r_i^2 \geq \sum_{i=1}^k d_i r_i^2 \geq \frac{3}{4} \cdot \left(1 - \frac{5}{128}\right) \cdot \Psi \geq \frac{2}{3} \cdot \Psi = \frac{2}{3} \cdot L(K_V) \bullet X.$$

Finally, using the fact that $\{v_i\}_{i \in V}$ is embedded in d dimensions, we can compute $L \bullet \tilde{X}$ in time $O(dm)$. Moreover, $L(K_V) \bullet X$ can be computed in time $O(nd)$ by using the decomposition $L(K_V) \bullet X = \sum_{i \in V} d_i \|v_i - v_{\text{avg}}\|^2$. By the same argument, we can compute $L(K_R) \bullet X$ in time $O(nd)$. The sweep cut over r takes time $\tilde{O}(m)$. And, by Theorem 5.23, PROJROUND runs in time $\tilde{O}(md)$. Hence, the total running time is $\tilde{O}(md)$. \blacksquare

6 Computing $\exp(-A)v$

In this section, we describe procedures for approximating $\exp(-A)v$ up to an ℓ_2 error of $\delta \|v\|$, given a symmetric PSD matrix A and a vector v (w.l.o.g., $\|v\| = 1$). In particular, we give the required procedures and proofs for Theorems 1.2, Theorem 1.4 and Theorem 3.2. For this section, we will assume the upper bound from Theorem 7.1 (which is a more precise version of Theorem 1.5), regarding polynomials approximating e^{-x} . Discussion about this theorem and the proofs are included in Section 7. We restate the basic definitions used in this section for completeness.

Definitions.

We will always work with square $n \times n$ matrices over \mathbb{R} . For a matrix M , abusing notation, we will denote its exponential by $\exp(-M)$ which is defined as $\sum_{i \geq 0} \frac{(-1)^i}{i!} M^i$. $\|M\| \stackrel{\text{def}}{=} \sup_{\|x\|=1} \|Mx\|$ denotes the spectral norm of M . M is said to be *Symmetric and Diagonally Dominant* (SDD) if, $M_{ij} = M_{ji}$, for all i, j and $M_{ii} \geq \sum_j |M_{ij}|$, for all i . M is called *Upper Hessenberg* if, $(M)_{ij} = 0$ for $i > j + 1$. M is called *tridiagonal* if $M_{ij} = 0$ for $i > j + 1$ and for $j > i + 1$. Let $\Lambda(M)$ denote the spectrum of a matrix M and let $\lambda_1(M)$ and $\lambda_n(M)$ denote the largest and the smallest eigenvalues of M respectively. For a matrix M , let m_M denote the number of non-zero entries in M . Further, let t_M denote the time required to multiply the matrix M with a given vector w . In general t_M depends on how M is given as an input and can be $\Theta(n^2)$. However, it is possible to exploit the special structure of M if given as an input appropriately: It is possible to just multiply the non-zero entries of M , giving $t_M = O(m_M)$. Also, if M is a rank one matrix ww^\top , where w is known, we can multiply with M in $O(n)$ time. For any positive integer k , let Σ_k denote the set of all polynomials with degree at most k . Given a degree k polynomial $p \stackrel{\text{def}}{=} \sum_{i=0}^k a_i \cdot x^i$, the ℓ_1 norm of p , denoted as $\|p\|_1$ is defined as $\|p\|_1 = \sum_{i \geq 0} |a_i|$.

Algorithms for Theorem 1.2, 1.3 and 3.2.

Theorem 1.2, 1.3 and 3.2 are based on a common algorithm we describe, called EXPRATIONAL (see Figure 5), which requires a procedure Invert_A with the following guarantee: given a vector y , a positive integer k and $\varepsilon_1 > 0$, $\text{Invert}_A(y, k, \varepsilon_1)$ returns a vector u_1 such that, $\|(I + A/k)^{-1}y - u_1\| \leq \varepsilon_1 \|y\|$. The algorithms for the two theorems differ only in their implementation of Invert_A . We prove the following theorem about EXPRATIONAL.

Theorem 6.1 (Running Time of EXPRATIONAL given Invert_A) *Given a symmetric p.s.d. matrix $A \succeq 0$, a vector v with $\|v\| = 1$, an error parameter $0 < \delta \leq 1$ and oracle access to Invert_A , for parameters $k \stackrel{\text{def}}{=} O(\log 1/\delta)$ and $\varepsilon_1 \stackrel{\text{def}}{=} \exp(-\Theta(k \log k + \log(1 + \|A\|)))$, EXPRATIONAL computes a vector u such that $\|\exp(-A)v - u\| \leq \delta$, in time $O(T_{A,k,\varepsilon_1}^{\text{inv}} \cdot k + n \cdot k^2 + k^3)$, where $T_{A,k,\varepsilon_1}^{\text{inv}}$ is the time required by $\text{Invert}_A(\cdot, k, \varepsilon_1)$.*

The proof of this theorem appears in Section 6.5. Theorem 1.2 will follow from the above theorem by using the Spielman-Teng SDD solver to implement the Invert_A procedure (See Section 6.3.1). For Theorem 3.2, we combine the SDD solver with the Sherman-Morrison formula (for matrix inverse with rank 1 updates) to implement the Invert_A procedure (See Section 6.4).

Algorithm for Theorem 1.4.

The procedure and proof for Theorem 1.4 is based on the well-known Lanczos method. We give a description of the Lanczos method (e.g. see [29]) in Figure 4 and give a proof of a well known theorem about the method that permits us to extend polynomial approximations for a function f over reals to approximating f over matrices (Theorem 6.7). Combining our result on polynomials approximating e^{-x} from the upper bound in Theorem 7.1 with the theorem about the Lanczos method, we give a proof of the following theorem that immediately implies Theorem 1.4.

Theorem 6.2 (Running Time Using LANCZOS) *Given a symmetric p.s.d. matrix A , a vector v with $\|v\| = 1$ and a parameter $0 < \delta \leq 1$, for*

$$k \stackrel{\text{def}}{=} O \left(\sqrt{\max\{\log^2 1/\delta, (\lambda_1(A) - \lambda_n(A)) \cdot \log 1/\delta\}} \cdot (\log 1/\delta) \cdot \log \log 1/\delta \right),$$

and $f(x) = e^{-x}$, the procedure LANCZOS computes a vector u such that $\|\exp(-A)v - u\| \leq \|\exp(-A)\| \delta$. The time taken by LANCZOS is $O((n + t_A)k + k^2)$.

Remark 6.3 *Note the k^3 term in the running time for Theorem 6.1 and the k^2 term in the running time for Theorem 6.2. This is the time required for computing the eigendecomposition of a $(k + 1) \times (k + 1)$ symmetric matrix. While this process requires $O(k^3)$ time in general, as in Theorem 6.1; in case of Theorem 6.2, the matrix is tridiagonal and hence the time required is $O(k^2)$ (see [26]).*

Organization.

We first describe the Lanczos method and prove some of its properties in Section 6.1. Then, we give descriptions of the LANCZOS and the EXPRATIONAL procedures in Section 6.2. Assuming Theorem 6.1, we give proofs of Theorem 1.2 and Theorem 3.2 in Section 6.3.1 and Section 6.4 respectively by implementing the respective Invert_A procedures. Finally, we give the error analysis for EXPRATIONAL and a proof for Theorem 6.1 in Section 6.5.

6.1 Lanczos Method – From Scalars to Matrices

Suppose one is given a symmetric PSD matrix B , and a function $f : \mathbb{R} \mapsto \mathbb{R}$. Then one can define $f(B)$ as follows: Let u_1, \dots, u_n be eigenvectors of B with eigenvalues $\lambda_1, \dots, \lambda_n$. Define $f(B) \stackrel{\text{def}}{=} \sum_i f(\lambda_i) u_i u_i^\top$. Given a vector v , we wish to compute $f(B)v$. Since exact computation of $f(B)$ usually requires diagonalization of B , which is costly, we seek an approximation to $f(B)v$.

For a given positive integer k , the Lanczos method looks for an approximation to $f(B)v$ of the form $p(B)v$, where p is a polynomial of degree k . Note that for any polynomial p of degree at most k , the vector $p(B)v$ is a linear combination of the vectors $\{v, Bv, \dots, B^k v\}$. The span of these vectors is referred to as the *Krylov Subspace* and is defined below.

Definition 6.4 (Krylov Subspace) *Given a matrix B and a vector v , the Krylov subspace of order k , denoted by $\mathcal{K}(B, v, k)$, is defined as the subspace that is spanned by the vectors $\{v, Bv, \dots, B^k v\}$.*

Note that any vector in $\mathcal{K}(B, v, k)$ has to be of the form $p(B)v$, where p is some degree k polynomial. The Lanczos method starts by generating an orthonormal basis for $\mathcal{K}(B, v, k)$. Let v_0, \dots, v_k be any orthonormal basis for $\mathcal{K}(B, v, k)$, and let V_k be the $n \times (k+1)$ matrix with $\{v_i\}_{i=0}^k$ as its columns. Thus, $V_k^\top V_k = I_k$ and $V_k V_k^\top$ denotes the projection onto the subspace. Also, let T_k be the operator B in the basis $\{v_i\}_{i=0}^k$, restricted to this subspace, i.e., $T_k \stackrel{\text{def}}{=} V_k^\top B V_k$. Since, all the vectors $v, Bv, \dots, B^k v$ are in the subspace, any of these vectors (or a linear combination of them) can be obtained by applying T_k to v (after a change of basis), instead of B . The following lemma proves this formally.

Lemma 6.5 (Exact Computation with Polynomials. See e.g. [29]) *Let V_k be the orthonormal basis, and T_k be the operator B restricted to $\mathcal{K}(B, v, k)$ where $\|v\| = 1$, i.e., $T_k = V_k^\top B V_k$. Let p be a polynomial of degree at most k . Then,*

$$p(B)v = V_k p(T_k) V_k^\top v.$$

Proof: Recall that $V_k V_k^\top$ is the orthogonal projection onto the subspace $\mathcal{K}(B, v, k)$. By linearity, it suffices to prove this when p is x^t for $t \leq k$. This is true for $t = 0$ since $V_k V_k^\top v = v$. For any $j \leq k$, $B^j v$ lies in $\mathcal{K}(B, v, k)$, thus, $\forall j \leq k$, $V_k V_k^\top B^j v = B^j v$. Hence,

$$\begin{aligned} B^t v &= (V_k V_k^\top) B (V_k V_k^\top) B \cdots B (V_k V_k^\top) v \\ &= V_k (V_k^\top B V_k) (V_k^\top B V_k) \cdots (V_k^\top B V_k) V_k^\top v = V_k T_k^t V_k^\top v \end{aligned}$$

■

The following lemma shows that $V_k f(T_k) V_k^\top v$ approximates $f(B)v$ as well as the *best* degree k polynomial that uniformly approximates f . The proof is based on the observation that if we express f as a sum of *any* degree k polynomial and an *error* function, the above lemma shows that the polynomial part is exactly computed in this approximation.

Lemma 6.6 (Approximation by Best Polynomial (Lemma 4.1, [29])) *Let V_k be the orthonormal basis, and T_k be the operator B restricted to $\mathcal{K}(B, v, k)$ where $\|v\| = 1$, i.e., $T_k = V_k^\top B V_k$. Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be any function such that $f(B)$ and $f(T_k)$ are well-defined. Then,*

$$\left\| f(B)v - V_k f(T_k) V_k^\top v \right\| \leq \min_{p_k \in \Sigma_k} \left(\max_{\lambda \in \Lambda(B)} |f(\lambda) - p_k(\lambda)| + \max_{\lambda \in \Lambda(T_k)} |f(\lambda) - p_k(\lambda)| \right).$$

Proof: Let p_k be any degree k polynomial. Let $r_k \stackrel{\text{def}}{=} f - p_k$. Then,

$$\begin{aligned} \left\| f(B)v - V_k f(T_k) V_k^\top v \right\| &\leq \left\| p_k(B)v - V_k p_k(T_k) V_k^\top v \right\| + \left\| r_k(B)v - V_k r_k(T_k) V_k^\top v \right\| \\ &\leq 0 + \|r_k(B)\| + \left\| V_k r_k(T_k) V_k^\top \right\| && \text{(Using Lemma 6.5)} \\ &= \max_{\lambda \in \Lambda(B)} |r_k(\lambda)| + \max_{\lambda \in \Lambda(T_k)} |r_k(\lambda)|. \end{aligned}$$

Minimizing over p_k gives us our lemma. ■

Observe that in order to compute this approximation, we do not need to know the polynomial explicitly. It suffices to prove that there exists a degree k polynomial that uniformly approximates f well on an interval containing the spectrum of B and T_k (For exact computation, $\Lambda(T_k) \subseteq \Lambda(B)$.) Moreover, if $k \ll n$, the computation has been reduced to a much smaller matrix. We now show that an orthonormal basis for the Krylov Subspace, V_k , can be computed quickly and then describe the LANCZOS procedure.

Input: A symmetric matrix $B \succeq 0$, a vector v such that $\|v\| = 1$, a positive integer k , and a function $f : \mathbb{R} \rightarrow \mathbb{R}$.

Output: A vector u that is an approximation to $f(B)v$.

1. Initialize $v_0 \stackrel{\text{def}}{=} v$.
 2. For $i = 0$ to $k - 1$, (Construct an orthonormal basis to Krylov subspace of order k)
 - a. If $i = 0$, compute $w_0 \stackrel{\text{def}}{=} Bv_0$. Else, compute $w_i = Bv_i - \beta_i v_{i-1}$. (Orthogonalize w.r.t. v_{i-1})
 - b. Define $\alpha_i \stackrel{\text{def}}{=} v_i^\top w_i$ and $w'_i \stackrel{\text{def}}{=} w_i - \alpha_i v_i$ *. (Orthogonalize w.r.t. v_i)
 - c. Define $\beta_{i+1} \stackrel{\text{def}}{=} \|w'_i\|$ and $v_{i+1} \stackrel{\text{def}}{=} w'_i / \beta_{i+1}$. (Scaling it to norm 1)
 3. Let V_k be the $n \times (k + 1)$ matrix whose columns are v_0, \dots, v_k respectively.
 4. Let T_k be the $(k + 1) \times (k + 1)$ matrix such that for all i , $(T_k)_{ii} = v_i^\top Bv_i = \alpha_i$, $(T_k)_{i,i+1} = (T_k)_{i+1,i} = v_{i+1}^\top Bv_i = \beta_{i+1}$ and all other entries are 0. (Compute $T_k \stackrel{\text{def}}{=} V_k^\top B V_k$)
 5. Compute $\mathcal{B} \stackrel{\text{def}}{=} f(T_k)$ exactly via eigendecomposition. Output the vector $V_k \mathcal{B} V_k^\top v$.
- * If $w'_i = 0$, compute the approximation with the matrices T_{i-1} and V_{i-1} , instead of T_k and V_k . The error bounds still hold.

Figure 4: The LANCZOS algorithm for approximating $f(B)v$

6.1.1 Efficiently Computing a Basis for the Krylov Subspace

In this section, we show that if we construct the basis $\{v_i\}_{i=0}^k$ in a particular way, the matrix T_k has extra structure. In particular, if B is symmetric, we show that T_k must be *tridiagonal*. This will help us speed up the construction of the basis.

Suppose we compute the orthonormal basis $\{v_i\}_{i=0}^k$ iteratively, starting from $v_0 = v$: For $i = 0, \dots, k$, we compute Bv_i and remove the components along the vectors $\{v_0, \dots, v_i\}$ to obtain a new vector that is orthogonal to the previous vectors. This vector, scaled to norm 1, is defined to be v_{i+1} . These vectors, by construction, satisfy that for all $i \leq k$, $\text{Span}\{v_0, \dots, v_i\} = \text{Span}\{v, Bv, \dots, B^k v\}$. Note that $(T_k)_{ij} = v_i^\top Bv_j$.

If we construct the basis iteratively as above, $Bv_j \in \text{Span}\{v_0, \dots, v_{j+1}\}$ by construction, and if $i > j + 1$, v_i is orthogonal to this subspace and hence $v_i^\top (Bv_j) = 0$. Thus, T_k is *Upper Hessenberg*, i.e., $(T_k)_{ij} = 0$ for $i > j + 1$.

Moreover, if B is symmetric, $v_j^\top (Bv_i) = v_i^\top (Bv_j)$, and hence T_k is symmetric and tridiagonal. This means that at most three coefficients are non-zero in each row. Thus, while constructing the basis, at step $i + 1$, it needs to orthonormalize Bv_i only w.r.t. v_{i-1} and v_i . This fact is used for efficient computation of T_k . The algorithm LANCZOS appears in Figure 4 and the following meta-theorem summarizes the main result regarding this method.

Theorem 6.7 (LANCZOS Theorem) *Given a symmetric p.s.d. matrix B , a vector v with $\|v\| = 1$, a function f and a positive integer parameter k as inputs, the procedure LANCZOS computes a vector u such*

that,

$$\|f(B)v - u\| \leq 2 \cdot \min_{p_k \in \Sigma_k} \max_{\lambda \in \Lambda(B)} |f(\lambda) - p_k(\lambda)|.$$

Here Σ_k denotes the set of all degree k polynomials and $\Lambda(B)$ denotes the spectrum of B . The time taken by LANCZOS is $O((n + t_B)k + k^2)$.

Proof: The algorithm LANCZOS implements the Lanczos method we've discussed here. The guarantee on u follows from Lemma 6.6 and the fact that $\Lambda(T_k) \subseteq \Lambda(B)$. We use the fact that $(T_k)_{ij} = v_i^\top B v_j$ and that T_k must be *tridiagonal* to reduce our work to just computing $O(k)$ entries in T_k . The total running time is dominated by k multiplications of B with a vector, $O(k)$ dot-products and the eigendecomposition of the tridiagonal matrix T_k to compute $f(T_k)$ (which can be done in $O(k^2)$ time [26]), giving a total running time of $O((n + t_B)k + k^2)$. ■

6.2 Procedures for Approximating $\exp(-A)v$.

Having introduced the Lanczos method, we describe the algorithms we use for approximating the matrix exponential.

6.2.1 Using LANCZOS for Approximating $\exp(-A)v$ – Proof of Theorem 1.4

Theorem 1.4 follows from Theorem 6.2, which is proved by combining Theorem 6.7 about the approximation guarantee of the LANCZOS algorithm and Theorem 7.1 (a more precise version of Theorem 1.5) about polynomials approximating e^{-x} . We now give a proof of Theorem 6.2.

Proof: We are given a matrix A , a unit vector v and an error parameter δ . Let $p_{\lambda_n(A), \lambda_1(A), \delta/2}(x)$ be the polynomial given by Theorem 7.1 and let k be its degree. We know from the theorem that $p_{\lambda_n(A), \lambda_1(A), \delta/2}(x)$ satisfies $\sup_{x \in [\lambda_n(A), \lambda_1(A)]} |e^{-x} - p_{\lambda_n(A), \lambda_1(A), \delta/2}(x)| \leq \delta/2 \cdot e^{-\lambda_n(A)}$, and that its degree is,

$$k \stackrel{\text{def}}{=} O\left(\sqrt{\max\{\log^2 1/\delta, (\lambda_1(A) - \lambda_n(A)) \cdot \log 1/\delta\}} \cdot (\log 1/\delta) \cdot \log \log 1/\delta\right).$$

Now, we run the LANCZOS procedure with the matrix A , the vector v , function $f(x) = e^{-x}$ and parameter k as inputs, and output the vector u returned by the procedure. In order to prove the error guarantee, we use Theorem 6.7 and bound the error using the polynomial $p_{\lambda_n(A), \lambda_1(A), \delta/2}$. Let $r(x) \stackrel{\text{def}}{=} \exp(-x) - p_{\lambda_n(A), \lambda_1(A), \delta/2}(x)$. We get,

$$\|\exp(-A)v - u\| \stackrel{\text{Thm. 6.7}}{\leq} 2 \max_{\lambda \in \Lambda(A)} |r_k(\lambda)| \stackrel{\Lambda(A) \subseteq [\lambda_n(A), \lambda_1(A)]}{\leq} 2 \max_{\lambda \in [\lambda_n(A), \lambda_1(A)]} |r_k(\lambda)| \stackrel{\text{Thm. 7.1}}{\leq} \delta \cdot e^{-\lambda_n(A)}$$

By Theorem 6.7, the total running time is $O((n + t_A)k + k^2)$. ■

6.2.2 The EXPRATIONAL Algorithm.

Now we move on to applying the Lanczos method in a way that was suggested as a heuristic by Eshof and Hochbruck [13]. The starting point here, is the following result by Saff, Schonhage and Varga [30], that shows that simple rational functions provide uniform approximations to e^{-x} over $[0, \infty)$ where the error term decays exponentially with the degree. Asymptotically, this result is best possible, see [17].

Theorem 6.8 (Rational Approximation [30]) *There exists constants $c_1 \geq 1$ and k_0 such that, for any integer $k \geq k_0$, there exists a polynomial $P_k(x)$ of degree $k - 1$ such that,*

$$\sup_{x \in [0, \infty)} \left| \exp(-x) - \frac{P_k(x)}{(1 + x/k)^k} \right| \leq c_1 k \cdot 2^{-k}.$$

Note that the rational function given by the above lemma can be written as a polynomial in $(1 + x/k)^{-1}$. The following corollary makes this formal.

Corollary 6.9 (Polynomial in $(1 + x/k)^{-1}$) *There exists constants $c_1 \geq 1$ and k_0 such that, for any integer $k \geq k_0$, there exists a polynomial $p_k^*(x)$ of degree k such that $p_k^*(0) = 0$, and,*

$$\sup_{t \in (0, 1]} \left| e^{-k/t+k} - p_k^*(t) \right| = \sup_{x \in [0, \infty)} \left| e^{-x} - p_k^* \left((1 + x/k)^{-1} \right) \right| \leq c_1 k \cdot 2^{-k}. \quad (4)$$

Proof: Define p_k^* as $p_k^*(t) \stackrel{\text{def}}{=} t^k \cdot P_k(k/t - k)$, where P_k is the polynomial from Theorem 6.8. Note that since P_k is a polynomial of degree $k - 1$, p_k^* is a polynomial of degree k with the constant term being zero, i.e., $p_k^*(0) = 0$. Also, for any $k \geq k_0$,

$$\sup_{t \in (0, 1]} \left| e^{-k/t+k} - p_k^*(t) \right| = \sup_{x \in [0, \infty)} \left| e^{-x} - p_k^* \left((1 + x/k)^{-1} \right) \right| = \sup_{x \in [0, \infty)} \left| e^{-x} - \frac{P_k(x)}{(1 + x/k)^k} \right| \leq c_1 k \cdot 2^{-k}.$$

■

The corollary above inspires the application of the Lanczos method to obtain the EXPRATIONAL algorithm that appears in Figure 5. We would like to work with the function $f(x) = e^{k(1-1/x)}$ and the matrix $B \stackrel{\text{def}}{=} (I + A/k)^{-1}$ for some positive integer k and use the Lanczos method to compute approximation to $\exp(-A)v$ in the Krylov subspace $\mathcal{K}(B, v, k)$, for small k . This is equivalent to looking for uniform approximations to $\exp(-y)$ that are degree k polynomials in $(1 + y/k)^{-1}$.

Unfortunately, we can't afford to exactly compute the vector $(I + A/k)^{-1}y$ for a given vector y . Instead, we will resort to a fast but error-prone solver, e.g. the Conjugate Gradient method and the Spielman-Teng SDD solver (Theorem 6.10). Since the computation is now approximate, the results for Lanczos method no longer apply. Dealing with the error poses a significant challenge as the Lanczos method is iterative and the error can propagate quite rapidly. A significant new and technical part of the paper is devoted to carrying out the error analysis in this setting. The details appear in Section 6.5.

Moreover, due to inexact computation, we can no longer assume B is symmetric. Hence, we perform complete orthonormalization while computing the basis $\{v_i\}_{i=0}^k$. We also define the symmetric matrix $\widehat{T}_k \stackrel{\text{def}}{=} 1/2 \cdot (T_k^\top + T_k)$ and compute our approximation using this matrix. The complete procedure EXPRATIONAL, with the exception of specifying the choice of parameters, is described in Figure 5. We give a proof of Theorem 6.1 in Section 6.5.

6.3 Exponentiating PSD Matrices – Proofs of Theorem 1.2 and 3.2

In this section, we give a proof of Theorem 1.2 and Theorem 1.3, assuming Theorem 6.1. Our algorithms for these theorems are based on the combining the EXPRATIONAL algorithm with appropriate Invert_A procedures.

Input: A Matrix $A \succeq 0$, a vector v such that $\|v\| = 1$, and an approximation parameter ε .

Output: A vector u such that $\|\exp(-A)v - u\| \leq \varepsilon$.

Parameters: Let $k \stackrel{\text{def}}{=} O(\log 1/\varepsilon)$ and $\varepsilon_1 \stackrel{\text{def}}{=} \exp(-\Theta(k \log k + \log(1 + \|A\|)))$.

1. Initialize $v_0 \stackrel{\text{def}}{=} v$.
2. For $i = 0$ to $k - 1$, (Construct an orthonormal basis to Krylov subspace of order k)
 - a. Call the procedure $\text{Invert}_A(v_i, k, \varepsilon_1)$. The procedure returns a vector w_i , such that, (Approximate $(I + A/k)^{-1}v_i$)
 $\|(I + A/k)^{-1}v_i - w_i\| \leq \varepsilon_1 \|v_i\|$.
 - b. For $j = 0, \dots, i$,
 - i. Let $\alpha_{j,i} \stackrel{\text{def}}{=} v_j^\top w_i$. (Compute projection onto w_i)
 - c. Define $w'_i \stackrel{\text{def}}{=} w_i - \sum_{j=0}^i \alpha_{j,i} v_j$. (Orthogonalize w.r.t. v_j for $j \leq i$)
 - d. Let $\alpha_{i+1,i} \stackrel{\text{def}}{=} \|w'_i\|$ * and $v_{i+1} \stackrel{\text{def}}{=} w'_i / \alpha_{i+1,i}$. (Scaling it to norm 1)
 - e. For $j = i + 2, \dots, k$,
 - i. Let $\alpha_{j,i} \stackrel{\text{def}}{=} 0$.
3. Let V_k be the $n \times (k + 1)$ matrix whose columns are v_0, \dots, v_k respectively.
4. Let T_k be the $(k + 1) \times (k + 1)$ matrix $(\alpha_{i,j})_{i,j \in \{0, \dots, k\}}$ and $\widehat{T}_k \stackrel{\text{def}}{=} 1/2(T_k^\top + T_k)$. (Symmetrize T_k)
5. Compute $\mathcal{B} \stackrel{\text{def}}{=} \exp\left(k \cdot (I - \widehat{T}_k^{-1})\right)$ exactly and output the vector $V_k \mathcal{B} e_1$.

* If $w'_i = 0$, compute the approximation the matrices T_{i-1} and V_{i-1} , instead of T_k and V_k . The error bounds still hold.

Figure 5: The EXPRATIONAL algorithm for approximating $\exp(-A)v$

6.3.1 SDD Matrices – Proof of Theorems 1.2

For Theorem 1.2 about exponentiating SDD matrices, we implement the Invert_A procedure using the Spielman-Teng SDD solver [37]. Here, we state an improvement on the Spielman-Teng result by Koutis, Miller and Peng [21].

Theorem 6.10 (SDD Solver [21]) *Given a system of linear equations $Mx = b$, where the matrix M is SDD, and an error parameter $\varepsilon > 0$, it is possible to obtain a vector u that is an approximate solution to the system, in the sense that*

$$\|u - M^{-1}b\|_M \leq \varepsilon \|M^{-1}b\|_M.$$

The time required for this computation is $\tilde{O}(m_M \log n \log^{1/\varepsilon})$, where M is an $n \times n$ matrix. (The tilde hides $\log n$ factors.)

We restate Theorem 1.2 for completeness.

Theorem 6.11 (Theorem 1.2 Restated) *Given an $n \times n$ symmetric matrix A which is SDD, a vector v and a parameter $\delta \leq 1$, there is an algorithm that can compute a vector u such that $\|\exp(-A)v - u\| \leq \delta \|v\|$ in time $\tilde{O}((m_A + n) \log(2 + \|A\|))$. The tilde hides $\text{poly}(\log n)$ and $\text{poly}(\log^{1/\delta})$ factors.*

Proof: We use the EXPRATIONAL procedure to approximate the exponential. We only need to describe how to implement the Invert_A procedure for an SDD matrix A . Recall that the procedure Invert_A , given a vector y , a positive integer k and real parameter $\varepsilon_1 > 0$, is supposed to return a vector u_1 such that $\|(I + A/k)^{-1}y - u_1\| \leq \varepsilon_1 \|y\|$, in time $T_{A,k,\varepsilon_1}^{\text{inv}}$. Also, observe that this is equivalent to approximately solving the linear system $(I + A/k)z = y$ for the vector z .

If the matrix A is SDD, $(I + A/k)$ is also SDD, and hence, we can use the Spielman-Teng SDD solver to implement Invert_A . We use Theorem 6.10 with inputs $(I + A/k)$, the vector y and error parameter ε_1 . It returns a vector u_1 such that,

$$\|(I + A/k)^{-1}y - u_1\|_{(I+A/k)} \leq \varepsilon_1 \|(I + A/k)^{-1}y\|_{(I+A/k)}.$$

This implies that,

$$\begin{aligned} \|(I + A/k)^{-1}y - u_1\|^2 &= ((I + A/k)^{-1}y - u_1)^\top ((I + A/k)^{-1}y - u_1) \\ &\leq ((I + A/k)^{-1}y - u_1)^\top (I + A/k) ((I + A/k)^{-1}y - u_1) \\ &\leq \left\| (I + A/k)^{-1}y - u_1 \right\|_{(I+A/k)}^2 \leq \varepsilon_1^2 \cdot \|(I + A/k)^{-1}y\|_{(I+A/k)}^2 \\ &= \varepsilon_1^2 \cdot y^\top (I + A/k)^{-1}y \leq \varepsilon_1^2 \cdot y^\top y, \end{aligned}$$

which gives us $\|(I + A/k)^{-1}y - u_1\| \leq \varepsilon_1 \|y\|$, as required for Invert_A . Thus, Theorem 6.1 implies that the procedure EXPRATIONAL computes a vector u approximating $e^{-A}v$, as desired.

The time required for the computation of u_1 is $T_{A,k,\varepsilon_1}^{\text{inv}} = \tilde{O}((m_A + n) \log n \log^{1/\varepsilon_1})$, and hence from Theorem 6.1, the total running time is $\tilde{O}((m_A + n) \log n (\log^{1/\delta} + \log(1 + \|A\|)) \log^{1/\delta} + (\log^{1/\delta})^3)$, where the tilde hides polynomial factors in $\log \log n$ and $\log \log^{1/\delta}$. \blacksquare

6.3.2 General PSD Matrices – Proof of Theorem 1.3

For Theorem 1.3 about exponentiating general PSD matrices, we implement the Invert_A procedure using the Conjugate Gradient method. We use the following theorem.

Theorem 6.12 (Conjugate Gradient Method. See [34]) *Given a system of linear equations $Mx = b$ and an error parameter $\varepsilon > 0$, it is possible to obtain a vector u that is an approximate solution to the system, in the sense that*

$$\|u - M^{-1}b\|_M \leq \varepsilon \|M^{-1}b\|_M.$$

The time required for this computation is $O\left(t_M \sqrt{\kappa(M)} \log 1/\varepsilon\right)$, – where $\kappa(M)$ denotes the condition number of M .

We restate Theorem 1.3 for completeness.

Theorem 6.13 (Theorem 1.3 Restated) *Given an $n \times n$ symmetric PSD matrix A , a vector v and a parameter $\delta \leq 1$, there is an algorithm that can compute a vector u such that $\|\exp(-A)v - u\| \leq \delta \|v\|$ in time $\tilde{O}\left((t_A + n)\sqrt{1 + \|A\|} \log(2 + \|A\|)\right)$. Here the tilde hides $\text{poly}(\log n)$ and $\text{poly}(\log 1/\delta)$ factors.*

Proof: We use the EXPRATIONAL procedure to approximate the exponential. We run the Conjugate Gradient method with the on input $(I + A/k)$, the vector y and error parameter ε_1 . The method returns a vector u_1 with the same guarantee as the SDD solver. As in Theorem 1.2, this implies $\|(I + A/k)^{-1}y - u_1\| \leq \varepsilon_1 \|y\|$, as required for Invert_A . Thus, Theorem 6.1 implies that the procedure EXPRATIONAL computes a vector u approximating $e^{-A}v$, as desired.

We can compute u_1 in time $T_{A,k,\varepsilon_1}^{\text{inv}} = O\left(t_A \sqrt{\frac{1+1/k \cdot \lambda_1(A)}{1+1/k \cdot \lambda_n(A)}} \log 1/\varepsilon_1\right) = O\left(t_A \sqrt{1 + \|A\|} \log 1/\varepsilon_1\right)$, and hence from Theorem 6.1, the total running time is

$$\tilde{O}\left(t_A \sqrt{1 + \|A\|} (\log 1/\delta + \log(1 + \|A\|)) \log 1/\delta + (\log 1/\delta)^2\right),$$

where the tilde hides polynomial factors in $\log \log n$ and $\log \log 1/\delta$. ■

6.4 Beyond SDD - Proof of Theorem 3.2

In this section, we give a proof of Theorem 3.2, which we restate below.

Theorem 6.14 (Theorem 3.2 Restated) *Given an $n \times n$ symmetric matrix $A = \Pi H M H \Pi$ where M is SDD, H is a diagonal matrix with strictly positive entries and Π is a rank $(n - 1)$ projection matrix $= 1 - ww^\top$ (w is explicitly known and $\|w\| = 1$), a vector v and a parameter $\delta \leq 1$, there is an algorithm that can compute a vector u such that $\|\exp(-A)v - u\| \leq \delta \|v\|$ in time $\tilde{O}((m_M + n) \log(2 + \|HMH\|))$. The tilde hides $\text{poly}(\log n)$ and $\text{poly}(\log 1/\delta)$ factors.*

Proof: In order to prove this, we will use the EXPRATIONAL procedure. For $A = \Pi H M H \Pi$, Lemma 6.15 given below implements the required Invert_A procedure. A proof of this lemma is given later in this section.

Lemma 6.15 (Invert_A Procedure for Theorem 3.2) *Given a positive integer k , vector y , an error parameter ε_1 , a rank $(n - 1)$ projection matrix $\Pi = I - ww^\top$ (where $\|w\| = 1$ and w is explicitly known), a diagonal matrix H with strictly positive entries, and an invertible SDD matrix M with m_M non-zero entries; we can compute a vector u such that $\|(I + 1/k \cdot \Pi H M H \Pi)^{-1}y - u\| \leq \varepsilon_1 \|(I + 1/k \cdot \Pi H M H \Pi)^{-1}y\|$, in time $\tilde{O}((m_M + n) \log n \log \frac{1 + \|HMH\|}{\varepsilon_1})$. (The tilde hides $\text{poly}(\log \log n)$ factors.)*

Assuming this lemma, we prove our theorem by combining this lemma with Theorem 6.1 about the EXPRATIONAL procedure, we get that we can compute the desired vector u approximating $e^{-A}v$ in total time

$$\tilde{O}\left((m_M + n) \log n (\log^{1/\delta} + \log(1 + \|HMH\|)) \log^{1/\delta} + (\log^{1/\delta})^3\right),$$

where the tilde hides polynomial factors in $\log \log n$ and $\log \log^{1/\delta}$. \blacksquare

In order to prove Lemma 6.15, we need to show how to approximate the inverse of a matrix of the form HMH , where H is diagonal and M is SDD. The following lemma achieves this.

Lemma 6.16 *Given a vector y , an error parameter ε_1 , a diagonal matrix H with strictly positive entries, and an invertible SDD matrix M with m_m non-zero entries; we can compute a vector u such that $\|(HMH)^{-1}y - u\|_{HMH} \leq \varepsilon_1 \|(HMH)^{-1}y\|_{HMH}$, in time $\tilde{O}((m_M + n) \log n \log^{1/\varepsilon_1})$. (The tilde hides factors of $\log \log n$.)*

Proof: Observe that $(HMH)^{-1}y = H^{-1}M^{-1}H^{-1}y$. Use the SDD solver (Theorem 6.10) with inputs M , vector $H^{-1}y$ and parameter ε_1 to obtain a vector u_1 such that,

$$\left\|M^{-1}(H^{-1}y) - u_1\right\|_M \leq \varepsilon_1 \left\|M^{-1}H^{-1}y\right\|_M.$$

Return the vector $u \stackrel{\text{def}}{=} H^{-1}u_1$. We can bound the error in the output vector u as follows,

$$\begin{aligned} \left\|(HMH)^{-1}y - u\right\|_{HMH}^2 &= \left\|H^{-1}M^{-1}H^{-1}y - H^{-1}u_1\right\|_{HMH}^2 \\ &= (H^{-1}M^{-1}H^{-1}y - H^{-1}u_1)^\top (HMH)(H^{-1}M^{-1}H^{-1}y - H^{-1}u_1) \\ &= (M^{-1}H^{-1}y - u_1)^\top M(M^{-1}H^{-1}y - u_1) \\ &= \left\|M^{-1}(H^{-1}y) - u_1\right\|_M^2 \\ &\leq \varepsilon_1^2 \left\|M^{-1}H^{-1}y\right\|_M^2 = \varepsilon_1^2 (M^{-1}H^{-1}y)^\top M(M^{-1}H^{-1}y) \\ &= \varepsilon_1^2 (H^{-1}M^{-1}H^{-1}y)^\top (HMH)(H^{-1}M^{-1}H^{-1}y) \\ &= \varepsilon_1^2 \left\|H^{-1}M^{-1}H^{-1}y\right\|_{HMH}^2 = \varepsilon_1^2 \left\|(HMH)^{-1}y\right\|_{HMH}^2 \end{aligned}$$

Thus, $\|(HMH)^{-1}y - u\|_{HMH} \leq \varepsilon_1 \|(HMH)^{-1}y\|_{HMH}$. Since H is diagonal, multiplication by H^{-1} requires $O(n)$ time. Hence, the total time is dominated by the SDD solver, giving a total running time of $\tilde{O}((m_M + n) \log n \log^{1/\varepsilon_1})$. \blacksquare

Now, we prove Lemma 6.15.

Lemma 6.17 (Lemma 6.15 Restated) *Given a positive integer k , vector y , an error parameter ε_1 , a rank $(n - 1)$ projection matrix $\Pi = I - ww^\top$ (where $\|w\| = 1$ and w is explicitly known), a diagonal matrix H with strictly positive entries, and an invertible SDD matrix M with m_M non-zero entries; we can compute a vector u such that $\left\|(I + 1/k \cdot \Pi H M H \Pi)^{-1}y - u\right\| \leq \varepsilon_1 \left\|(I + 1/k \cdot \Pi H M H \Pi)^{-1}y\right\|$, in time $\tilde{O}((m_M + n) \log n \log \frac{1 + \|HMH\|}{\varepsilon_1})$. (The tilde hides $\text{poly}(\log \log n)$ factors.)*

1. Compute $z \stackrel{\text{def}}{=} y - (w^\top y)w$.
2. Estimate $(I + M_1)^{-1}z$ with error parameter $\frac{\varepsilon_1}{6(1+\|M_1\|)}$. Denote the vector returned by β_1 .
3. Estimate $(I + M_1)^{-1}w$ with error parameter $\frac{\varepsilon_1}{6(1+\|M_1\|)}$. Denote the vector returned by β_2 .

4. Compute

$$u_1 \stackrel{\text{def}}{=} \beta_1 - \frac{w^\top M_1 \beta_1}{1 + w^\top M_1 \beta_2} \beta_2 + (w^\top y)w. \quad (7)$$

Return u_1 .

Figure 6: The Invert_A procedure for Theorem 3.2

Proof: We sketch the proof idea first. Using the fact that w is an eigenvector of our matrix, we will split y into two components – one along w and one orthogonal. Along w , we can easily compute the component of the required vector. Among the orthogonal component, we will write our matrix as the sum of $I + 1/k \cdot HMH$ and a rank one matrix, and use the Sherman-Morrison formula to express its inverse. Note that we can use Lemma 6.15 to compute the inverse of $I + 1/k \cdot HMH$. The procedure is described in Figure 6 and the proof for the error analysis is given below.

Let $M_1 \stackrel{\text{def}}{=} 1/k \cdot HMH$. Then, $I + 1/k \cdot \Pi HMH \Pi = I + \Pi M_1 \Pi$. Without loss of generality, we will assume that $\|y\| = 1$. Note that $I + \Pi M_1 \Pi \succ 0$, and hence is invertible. Let $z \stackrel{\text{def}}{=} y - (w^\top y)w$. Thus, $w^\top z = 0$. Since w is an eigenvector of $(I + \Pi M_1 \Pi)$ with eigenvalue 1, we get,

$$(I + \Pi M_1 \Pi)^{-1}y = (I + \Pi M_1 \Pi)^{-1}z + (w^\top y)w. \quad (5)$$

Let's say $t \stackrel{\text{def}}{=} (I + \Pi M_1 \Pi)^{-1}z$. Then, $t + \Pi M_1 \Pi t = z$. Left-multiplying by w^\top , we get, $w^\top t = w^\top z = 0$. Thus, $\Pi t = t$, and hence $(I + \Pi M_1)t = z$, or equivalently, $t = (I + \Pi M_1)^{-1}z$.

$$\begin{aligned} (I + \Pi M_1 \Pi)^{-1}z &= (I + \Pi M_1)^{-1}z = (I + M_1 - ww^\top M_1)^{-1}z = (I + M_1 - w(M_1 w)^\top)^{-1}z \\ &= \left((I + M_1)^{-1} - \frac{(I + M_1)^{-1}ww^\top M_1(I + M_1)^{-1}}{1 + w^\top M_1(I + M_1)^{-1}w} \right) z \\ &\quad \text{(Sherman-Morrison formula)} \\ &= (I + M_1)^{-1}z - \frac{w^\top M_1(I + M_1)^{-1}z}{1 + w^\top M_1(I + M_1)^{-1}w} (I + M_1)^{-1}w \end{aligned} \quad (6)$$

Since we can write $I + M_1 = I + HMH = H(H^{-2} + M)H$, we can use Lemma 6.16 to estimate $(I + M_1)^{-1}z$ and $(I + M_1)^{-1}w$. Using Equation (6), the procedure for estimating $(I + \Pi M_1 \Pi)^{-1}x$ is described in Figure 6.

We need to upper bound the error in the above estimation procedure. From the assumption, we know that $\beta_1 = (I + M_1)^{-1}z - e_1$, where $\|e_1\|_{(I+M_1)} \leq \frac{\varepsilon_1}{6(1+\|M_1\|)} \|(I + M_1)^{-1}z\|_{(I+M_1)}$, and $\beta_2 = (I + M_1)^{-1}w - e_2$, where $\|e_2\|_{(I+M_1)} \leq \frac{\varepsilon_1}{6(1+\|M_1\|)} \|(I + M_1)^{-1}w\|_{(I+M_1)}$. Combining Equations (5)

and (6) and subtracting Equation (7), we can write the error as,

$$\begin{aligned}
(I + \Pi M_1 \Pi)^{-1} y - u_1 &= (I + M_1)^{-1} z - \beta_1 - \frac{w^\top M_1 (I + M_1)^{-1} z}{1 + w^\top M_1 (I + M_1)^{-1} w} (I + M_1)^{-1} w + \frac{w^\top M_1 \beta_1}{1 + w^\top M_1 \beta_2} \beta_2 \\
&= e_1 - \frac{w^\top M_1 (I + M_1)^{-1} z}{1 + w^\top M_1 (I + M_1)^{-1} w} (I + M_1)^{-1} w + \frac{w^\top M_1 [(I + M_1)^{-1} z - e_1]}{1 + w^\top M_1 [(I + M_1)^{-1} w - e_2]} [(I + M_1)^{-1} w - e_2] \\
&= e_1 + \frac{w^\top M_1 (I + M_1)^{-1} z \cdot w^\top M_1 e_2}{(1 + w^\top M_1 (I + M_1)^{-1} w)(1 + w^\top M_1 [(I + M_1)^{-1} w - e_2])} (I + M_1)^{-1} w \\
&\quad - \frac{w^\top M_1 e_1}{1 + w^\top M_1 [(I + M_1)^{-1} w - e_2]} [(I + M_1)^{-1} w - e_2] - \frac{w^\top M_1 (I + M_1)^{-1} z}{1 + w^\top M_1 [(I + M_1)^{-1} w - e_2]} e_2
\end{aligned}$$

Let us first bound the scalar terms. Note that $\|z\| \leq \|y\| = 1$.

$$|w^\top M_1 (I + M_1)^{-1} z| \leq \|w\| \left\| M_1 (I + M_1)^{-1} z \right\| \leq \left\| M_1 (I + M_1)^{-1} \right\| \leq 1,$$

$$\begin{aligned}
w^\top M_1 e_1 &\leq \|w\| \|M_1\| \|e_1\| \leq \|M_1\| \|e_1\|_{(I+M_1)} \leq \varepsilon_1/6 \cdot \left\| (I + M_1)^{-1} z \right\|_{(I+M_1)} \\
&= \varepsilon_1/6 \cdot \left\| (I + M_1)^{-1/2} z \right\| \leq \varepsilon_1/6.
\end{aligned}$$

Similarly, $w^\top M_1 e_2 \leq \varepsilon_1/6$. Also, $M_1 (I + M_1)^{-1} \succeq 0$ and hence $w^\top M_1 (I + M_1)^{-1} w \geq 0$. Thus,

$$\begin{aligned}
\left\| (I + \Pi M_1 \Pi)^{-1} y - u_1 \right\| &\leq \left\| (I + \Pi M_1 \Pi)^{-1} y - u_1 \right\|_{(I+M_1)} \\
&\leq \|e_1\|_{(I+M_1)} + \frac{1 \cdot \varepsilon_1/6}{1 \cdot (1 - \varepsilon_1/6)} \left\| (I + M_1)^{-1} w \right\|_{(I+M_1)} \\
&\quad + \frac{\varepsilon_1/6}{(1 - \varepsilon_1/6)} \left(\left\| (I + M_1)^{-1} w \right\|_{(I+M_1)} + \|e_2\|_{(I+M_1)} \right) + \frac{1}{1 - \varepsilon_1/6} \|e_2\|_{(I+M_1)} \\
&\leq \frac{\varepsilon_1}{6(1 + \|M_1\|)} \left\| (I + M_1)^{-1} z \right\|_{(I+M_1)} + \frac{4 \cdot \varepsilon_1/6}{1 - \varepsilon_1/6} \left\| (I + M_1)^{-1} w \right\|_{(I+M_1)} \\
&\leq \varepsilon_1/6 \left\| (I + M_1)^{-1/2} z \right\| + 4\varepsilon_1/5 \left\| (I + M_1)^{-1/2} w \right\| \leq \varepsilon_1
\end{aligned}$$

Other than the estimation of $(I + M_1)^{-1} z$ and $(I + M_1)^{-1} w$, we need to compute a constant number of dot products and a constant number of matrix-vector products with the matrix M_1 . Multiplying a vector with $M_1 = 1/k \cdot H M H$ takes time $O(m_M + n)$, giving a total time of $\tilde{O}\left((m_M + n) \log n \log \frac{1 + \|H M H\|}{\varepsilon_1}\right)$ ■

6.5 Error Analysis for EXP-RATIONAL

In this section, we give the proof of Theorem 6.1, except for the proof of a few lemmas, which have been presented in the Section 6.5.1 for better readability.

Proof Overview.

At a very high-level, the proof follows the outline of the proof for Lanczos method. We first show that assuming the error in computing the inverse is small, \hat{T}_k can be used to approximate small

powers of $B = (I + A/k)^{-1}$ when restricted to the Krylov subspace, *i.e.* for all $i \leq k$, $\|B^i v - V_k \widehat{T}_k^i V_k^\top v\| \lesssim \varepsilon_2$, for some small ε_2 . This implies that we can bound the error in approximating $p((I + A/k)^{-1})$ using $p(\widehat{T}_k)$, by $\varepsilon_2 \|p\|_1$, where p is a polynomial of degree at most k . This is the most technical part of the error analysis because we need to capture the propagation of error through the various iterations of the algorithm. We overcome this difficulty by expressing the final error as a sum of k terms, with the i^{th} term expressing how much error is introduced in the final candidate vector because of the error in the inverse computation during the i^{th} iteration. Unfortunately, the only way we know of bounding each of these terms is by *tour de force*. A part of this proof is to show that the spectrum of \widehat{T}_k cannot shift far from the spectrum of B .

To bound the error in the candidate vector output by the algorithm, *i.e.* $\|f(B)v - V_k f(\widehat{T}_k) V_k^\top v\|$, we start by expressing e^{-x} as the sum of a degree k -polynomial p_k in $(1 + x/k)^{-1}$ and a remainder function r_k . We use the analysis from the previous paragraph to upper bound the error in the polynomial part by $\varepsilon_2 \|p\|_1$. We bound the contribution of the remainder term to the error by bounding $\|r_k(B)\|$ and $\|r_k(\widehat{T}_k)\|$. This step uses the fact that eigenvalues of $r_k(\widehat{T}_k)$ are $\{r_k(\lambda_i)\}_i$, where $\{\lambda_i\}_i$ are eigenvalues of \widehat{T}_k . To complete the error analysis, we use the polynomials p_k^* from Corollary 6.9 and bound its ℓ_1 norm. Even though we do not know p_k^* explicitly, we can bound $\|p_k^*\|_1$ indirectly by writing it as an interpolation polynomial and using that the values it assumes in $[0, 1]$ have to be small in magnitude.

Proof: For notational convenience, define $B \stackrel{\text{def}}{=} (I + A/k)^{-1}$. Since the computation of Bv_i is not exact in each iteration, the eigenvalues of \widehat{T}_k need not be eigenvalues of B . Also, Lemma 6.5 no longer holds, *i.e.*, we can't guarantee that $V_k \widehat{T}_k^t e_1$ is identical to $B^t v_0$. However, we can prove the following lemma that proves bounds on the spectrum of \widehat{T}_k and also bounds the norm of the difference between the vectors $V_k \widehat{T}_k^t e_1$ and $B^t v_0$. This is the most important and technically challenging part of the proof.

Lemma 6.18 (Approximate Computation with \widehat{T}_k . Proof in Sec. 6.5.1) *The coefficient matrix \widehat{T}_k generated satisfies the following:*

1. The eigenvalues of \widehat{T}_k lie in $\left[\left(1 + \frac{\lambda_1(A)}{k}\right)^{-1} - \varepsilon_1 \sqrt{k+1}, \left(1 + \frac{\lambda_n(A)}{k}\right)^{-1} + \varepsilon_1 \sqrt{k+1} \right]$.
2. For any $t \leq k$, if $\varepsilon_1 \leq \varepsilon_2 / (8(k+1)^{5/2})$ and $\varepsilon_2 \leq 1$, we have, $\|B^t v_0 - V_k \widehat{T}_k^t e_1\| \leq \varepsilon_2$.

Here is an idea of the proof of the above lemma: Since, during every iteration of the algorithm, the computation of Bv_i is approximate, we will express BV_k in terms of T_k and an error matrix E . This will allow us to express \widehat{T}_k in terms of T_k and a different error matrix. The first part of the lemma will follow immediately from the guarantee of the Invert_A procedure.

For the Second part, we first express $BV_k - V_k \widehat{T}_k$ in terms of the error matrices defined above. Using this, we can write the telescoping sum $B^t V_k - V_k \widehat{T}_k^t = \sum_{j=1}^t B^{t-j} (BV_k - V_k \widehat{T}_k) \widehat{T}_k^{j-1}$. We use triangle inequality and a *tour de force* calculation to bound each term. A complete proof is included in Section 6.5.1.

As a simple corollary, we can bound the error in the computation of the polynomial, in terms of the ℓ_1 norm of the polynomial being computed.

Corollary 6.19 (Approximate Polynomial Computation. Proof in Sec. 6.5.1) *For any polynomial p of degree at most k , if $\varepsilon_1 \leq \varepsilon_2 / (2(k+1)^{3/2})$ and $\varepsilon_2 \leq 1$,*

$$\|p(B)v_0 - V_k p(\widehat{T}_k) e_1\| \leq \varepsilon_2 \|p\|_1.$$

Using this corollary, we can prove an analogue of Lemma 6.6, giving error bounds on the procedure in terms of degree k polynomial approximations. The proof is very similar and is based on writing f as a sum of a degree k polynomial and an *error function*.

Lemma 6.20 (Polynomial Approximation for EXPRATIONAL. Proof in Sec. 6.5.1) *Let V_k be the orthonormal basis and \widehat{T}_k be the matrix of coefficients generated by EXPRATIONAL. Let f be any function such that $f(B)$ and $f(T_k)$ are defined. Define $r_k(x) \stackrel{\text{def}}{=} f(x) - p(x)$. Then,*

$$\left\| f(B)v_0 - V_k f(\widehat{T}_k)e_1 \right\| \leq \min_{p \in \Sigma_k} \left(\varepsilon_2 \|p\|_1 + \max_{\lambda \in \Lambda(B)} |r_k(\lambda)| + \max_{\lambda \in \Lambda(\widehat{T}_k)} |r_k(\lambda)| \right). \quad (8)$$

In order to control the second error term in the above lemma, we need to bound the eigenvalues of \widehat{T}_k , which is provided by Lemma 6.18.

For our application, $f(t) = f_k(t) \stackrel{\text{def}}{=} \exp(k \cdot (1 - 1/t))$ so that $f_k((1 + x/k)^{-1}) = \exp(-x)$. This function is discontinuous at $t = 0$. Under exact computation of the inverse, the eigenvalues of \widehat{T}_k would be the same as the eigenvalues of B and hence would lie in $(0, 1]$. Unfortunately, due to the errors, the eigenvalues of \widehat{T}_k could be outside the interval. Since f is discontinuous at 0, and goes to infinity for small negative values, in order to get a reasonable approximation to f , we will ensure that the eigenvalues of \widehat{T}_k are strictly positive, *i.e.*, $\varepsilon_1 \sqrt{k+1} < (1 + 1/k \cdot \lambda_1(A))^{-1}$.

We will use the polynomials p_k^* from Corollary 6.9 in Lemma 6.20 to bound the final error. We will require the following lemma to bound the ℓ_1 -norm of p_k^* .

Lemma 6.21 (ℓ_1 -norm Bound. Proof in Sec. 6.5.1) *Given a polynomial p of degree k such that $p(0) = 0$ and*

$$\sup_{t \in (0,1]} \left| e^{-k/t+k} - p(t) \right| = \sup_{x \in [0,\infty)} \left| e^{-x} - p\left((1 + x/k)^{-1}\right) \right| \leq 1,$$

we must have $\|p\|_1 \leq (2k)^{k+1}$.

This lemma is proven by expressing p as the interpolation polynomial on the values attained by p at the $k+1$ points $0, 1/k, \dots, k/k$, which allows us to express the coefficients in terms of these values. We can bound these values, and hence, the coefficients, since we know that p isn't too far from the exponential function. A complete proof is included in Section 6.5.1.

Corollary 6.9 shows that $p_k^*(t)$ is a good uniform approximation to $e^{-k/t+k}$ over the interval $(0, 1]$. Since $\Lambda(B) \subseteq (0, 1]$, this will help us bound the second error term in Equation (8). Since \widehat{T}_k can have eigenvalues larger than 1, we need to bound the error in approximating $f_k(t)$ by $p_k^*(t)$ over an interval $(0, \beta]$, where $\beta \geq 1$. The following lemma, gives us the required error bound. This proof for this lemma bounds the error over $[1, \beta]$ by applying triangle inequality and bounding the change in f_k and p over $[1, \beta]$ separately.

Lemma 6.22 (Approximation on Extended Interval. Proof in Sec. 6.5.1) *For any $\beta \geq 1$, any degree k polynomial p satisfies,*

$$\sup_{t \in (0,\beta]} |p(t) - f_k(t)| \leq \|p\|_1 \cdot (\beta^k - 1) + (f_k(\beta) - f_k(1)) + \sup_{t \in (0,1]} |p(t) - f_k(t)|.$$

We bound the final error using the polynomial p_k^* in Equation (8). We will use the above lemma for $\beta \stackrel{\text{def}}{=} 1 + \varepsilon_1 \sqrt{k+1}$ and assume that $\varepsilon_1 \sqrt{k+1} < (1 + 1/k \cdot \lambda_1(A))^{-1}$.

$$\begin{aligned}
\left\| f(B)v_0 - V_k f(\widehat{T}_k) e_1 \right\| &\leq \varepsilon_2 \|p_k^*\|_1 + \max_{\lambda \in \Lambda(B)} |r_k(\lambda)| + \max_{\lambda \in \Lambda(\widehat{T}_k)} |r_k(\lambda)| \\
&\leq \varepsilon_2 \|p_k^*\|_1 + \sup_{\lambda \in (0,1]} |(f_k - p_k^*)(\lambda)| + \sup_{\lambda \in (0,\beta]} |(f_k - p_k^*)(\lambda)|. \\
&\quad (\text{Since } \Lambda(B) \subseteq (0,1] \text{ and } \Lambda(\widehat{T}_k) \subseteq (0,\beta]) \\
&\leq \varepsilon_2 \|p_k^*\|_1 + \sup_{t \in (0,1]} |p_k^*(t) - f_k(t)| + \\
&\quad \|p_k^*\|_1 \cdot (\beta^k - 1) + (f_k(\beta) - f_k(1)) + \sup_{t \in (0,1]} |p_k^*(t) - f_k(t)| \\
&= \|p_k^*\|_1 \cdot (\varepsilon_2 + \beta^k - 1) + (\exp(k(\beta-1)/\beta) - 1) + 2 \sup_{t \in (0,1]} |p_k^*(t) - f_k(t)|.
\end{aligned}$$

Given $\delta < 1$, we plug in the following parameters,

$$k \stackrel{\text{def}}{=} \max\{k_0, \log_2^{8c_1/\delta} + 2 \log_2 \log_2^{8c_1/\delta}\} = O(\log^{1/\delta}),$$

$$\varepsilon_1 \stackrel{\text{def}}{=} \delta/32 \cdot (k+1)^{-5/2} \cdot (1 + 1/k \cdot \lambda_1(A))^{-1} \cdot (2k)^{-k-1}, \quad \beta \stackrel{\text{def}}{=} 1 + \varepsilon_1 \sqrt{k+1}, \quad \varepsilon_2 \stackrel{\text{def}}{=} 8(k+1)^{5/2} \varepsilon_1,$$

where k_0, c_1 are the constants given by Corollary 6.9. Note that these parameters satisfy the condition $\varepsilon_1 \sqrt{k+1} < (1 + 1/k \cdot \lambda_1(A))^{-1}$. Corollary 6.9 implies that $p_k^*(0) = 0$ and

$$\begin{aligned}
\sup_{t \in (0,1]} |p_k^*(t) - f_k(t)| &\leq \frac{\delta}{8} \cdot \frac{\log_2^{8c_1/\delta} + 2 \log_2 \log_2^{8c_1/\delta}}{(\log_2^{8c_1/\delta})^2} \\
&\leq \frac{\delta}{8} \cdot \frac{1}{\log_2^{8c_1/\delta}} \left(1 + 2 \cdot \frac{\log_2 \log_2^{8c_1/\delta}}{\log_2^{8c_1/\delta}} \right) \leq \frac{\delta}{8} \cdot \frac{1}{3} \cdot 3 \leq \frac{\delta}{8}, \tag{9}
\end{aligned}$$

where the last inequality uses $\delta \leq 1 \leq c_1$ and $\log_2 x \leq x, \forall x \geq 0$. Thus, we can use Lemma 6.21 to conclude that $\|p_k^*\|_1 \leq (2k)^{k+1}$.

We can simplify the following expressions,

$$\exp(k(\beta-1)/\beta) - 1 \leq \exp(k\varepsilon_1 \cdot \sqrt{k+1}) - 1 \leq \exp(\varepsilon_2/8) - 1 \leq (1 + \varepsilon_2/4) - 1 = \varepsilon_2/4,$$

$$\beta^k - 1 = (1 + \varepsilon_1 \sqrt{k+1})^k - 1 \leq \exp(k \cdot \varepsilon_1 \sqrt{k+1}) - 1 \leq \varepsilon_2/4.$$

Thus the total error $\|u - \exp(-A)v\| = \left\| f(B)v_0 - V_k f(\widehat{T}_k) e_1 \right\| \leq (2k)^{k+1} \cdot 2\varepsilon_2 + \varepsilon_2 + \delta/4 \leq \delta$.

Running Time.

The running time for the procedure is dominated by k calls to the Invert_A procedure with parameters k and ε_1 , computation of at most k^2 dot-products and the exponentiation of \widehat{T}_k . The exponentiation of \widehat{T}_k can be done in time $O(k^3)$ [26]. Thus the total running time is $O(T_{A,k,\varepsilon_1}^{\text{inv}} \cdot k + n \cdot k^2 + k^3)$. This completes the proof of the Theorem 6.1. \blacksquare

6.5.1 Remaining Proofs

In this section, we give the remaining proofs in Section 6.5.

Lemma 6.23 (Lemma 6.18 Restated) *The coefficient matrix \widehat{T}_k generated satisfies the following:*

1. The eigenvalues of \widehat{T}_k lie in the interval $\left[\left(1 + \frac{\lambda_1(A)}{k}\right)^{-1} - \varepsilon_1 \sqrt{k+1}, \left(1 + \frac{\lambda_n(A)}{k}\right)^{-1} + \varepsilon_1 \sqrt{k+1} \right]$.
2. For any $t \leq k$, if $\varepsilon_1 \leq \varepsilon_2 / (8(k+1)^{5/2})$ and $\varepsilon_2 \leq 1$, we have, $\left\| B^t v_0 - V_k \widehat{T}_k^t e_1 \right\| \leq \varepsilon_2$.

Proof: Given a vector y , a positive integer k and real parameter $\varepsilon_1 > 0$, $\text{Invert}_A(y, k, \varepsilon_1)$ returns a vector u_1 such that $\|By - u_1\| \leq \varepsilon_1 \|y\|$, in time $T_{A,k,\varepsilon_1}^{\text{inv}}$. Thus, for each i , the vector w_i satisfies $\|Bv_i - w_i\| \leq \varepsilon_1 \|v_i\| = \varepsilon_1$. Also define u_i as $u_i \stackrel{\text{def}}{=} Bv_i - w_i$. Thus, we get, $\|u_i\| \leq \varepsilon_1$. Let E be the $n \times (k+1)$ matrix with its columns being u_0, \dots, u_k . We can write the following recurrence,

$$BV_k = V_k T_k + E + \alpha_{k,k+1} v_{k+1} e_{k+1}^\top, \quad (10)$$

where each column of E has ℓ_2 norm at most ε_1 . Note that we continue to do complete orthonormalization, so $V_k^\top V_k = I_k$. Thus, T_k is not tridiagonal, but rather *Upper Hessenberg*, i.e., $(T_k)_{ij} = 0$ whenever $i > j + 1$.

Multiplying both sides of Equation (10) by V_k^\top , we get $T_k = V_k^\top BV_k - V_k^\top E$. This implies,

$$\widehat{T}_k = V_k^\top BV_k - 1/2 \cdot (V_k^\top E + E^\top V_k) \quad (11)$$

$$= V_k^\top (V_k T_k + E + \alpha_{k,k+1} v_{k+1} e_{k+1}^\top) - 1/2 \cdot (V_k^\top E + E^\top V_k) \quad (\text{Using (10)})$$

$$= T_k + 1/2 \cdot (V_k^\top E - E^\top V_k). \quad (12)$$

Define $E_1 \stackrel{\text{def}}{=} 1/2 \cdot (V_k^\top E + E^\top V_k)$. Thus, using Equation (11), $\widehat{T}_k = V_k^\top BV_k - E_1$. Let us first bound the norm of E_1 .

$$\|E_1\| \leq 1/2 \cdot (\|V_k^\top E\| + \|E^\top V_k\|) \leq 1/2 \cdot (\|E\| + \|E^\top\|) \leq \|E\|_F \leq \varepsilon_1 \sqrt{k+1}.$$

Since $\widehat{T}_k = V_k^\top BV_k - E_1$. We have,

$$\lambda_{\max}(\widehat{T}_k) \leq \lambda_1(B) + \|E_1\| \leq (1 + 1/k \cdot \lambda_n(A))^{-1} + \varepsilon_1 \sqrt{k+1},$$

$$\lambda_{\min}(\widehat{T}_k) \geq \lambda_n(B) - \|E_1\| \geq (1 + 1/k \cdot \lambda_1(A))^{-1} - \varepsilon_1 \sqrt{k+1}.$$

(We use λ_{\max} and λ_{\min} for the largest and smallest eigenvalues of \widehat{T}_k respectively in order to avoid confusion since \widehat{T}_k is a $(k+1) \times (k+1)$ matrix and not an $n \times n$ matrix.)

First, let us compute $BV_k - V_k \widehat{T}_k$.

$$\begin{aligned} BV_k - V_k \widehat{T}_k &\stackrel{(10),(12)}{=} V_k T_k + E + \alpha_{k,k+1} v_{k+1} e_{k+1}^\top - V_k \left(T_k + 1/2 \cdot (V_k^\top E - E^\top V_k) \right) \\ &= \left(I - 1/2 \cdot V_k V_k^\top \right) E + 1/2 \cdot V_k E^\top V_k + \alpha_{k,k+1} v_{k+1} e_{k+1}^\top. \end{aligned} \quad (13)$$

Now,

$$\begin{aligned}
\left\| B^t v_0 - V_k \widehat{T}_k^t e_1 \right\| &= \left\| \sum_{j=1}^t B^{t-j} (B V_k - V_k \widehat{T}_k) \widehat{T}_k^{j-1} e_1 \right\| && \text{(Telescoping sum)} \\
&\stackrel{(13)}{=} \left\| \sum_{j=1}^t B^{t-j} \left((I - 1/2 \cdot V_k V_k^\top) E + 1/2 \cdot V_k E^\top V_k + \alpha_{k,k+1} v_{k+1} e_{k+1}^\top \right) \widehat{T}_k^{j-1} e_1 \right\| \\
&\stackrel{\Delta\text{-ineq.}}{\leq} \sum_{j=1}^t \left\| B^{t-j} \left((I - 1/2 \cdot V_k V_k^\top) E + 1/2 \cdot V_k E^\top V_k \right) \widehat{T}_k^{j-1} e_1 \right\| \\
&\quad + \left\| \sum_{j=1}^t B^{t-j} \left(\alpha_{k,k+1} v_{k+1} e_{k+1}^\top \right) \widehat{T}_k^{j-1} e_1 \right\|. \tag{14}
\end{aligned}$$

We can bound the first term in Equation (14) as follows.

$$\begin{aligned}
\sum_{j=1}^t \left\| B^{t-j} \left((I - 1/2 \cdot V_k V_k^\top) E + 1/2 \cdot V_k E^\top V_k \right) \widehat{T}_k^{j-1} e_1 \right\| &\leq \sum_{j=1}^t \left\| (I - 1/2 \cdot V_k V_k^\top) E + 1/2 \cdot V_k E^\top V_k \right\| \left\| \widehat{T}_k \right\|^{j-1} \\
&\quad \text{(Using } \|B\| \leq 1) \\
&\leq \left(\sum_{j=1}^t (1 + \varepsilon_1 \sqrt{k+1})^{j-1} \right) \left\| (I - 1/2 \cdot V_k V_k^\top) E + 1/2 \cdot V_k E^\top V_k \right\| \\
&\quad \text{(Using } \left\| \widehat{T}_k \right\| \leq 1 + \varepsilon_1 \sqrt{k+1}) \\
&\leq t(1 + \varepsilon_1 \sqrt{k+1})^{t-1} \left(\left\| (I - 1/2 \cdot V_k V_k^\top) \right\| \|E\| + 1/2 \cdot \|V_k\| \left\| E^\top \right\| \|V_k\| \right) \\
&\leq 2t\varepsilon_1 \sqrt{k+1} (1 + \varepsilon_1 \sqrt{k+1})^{t-1}. \tag{15}
\end{aligned}$$

The second term in Equation (14) can be bounded as follows.

$$\begin{aligned}
& \left\| \sum_{j=1}^t B^{t-j} \left(\alpha_{k,k+1} v_{k+1} e_{k+1}^\top \right) \widehat{T}_k^{j-1} e_1 \right\| \leq |\alpha_{k,k+1}| \sum_{j=1}^t \|B\|^{t-j} \left\| v_{k+1} e_{k+1}^\top \widehat{T}_k^{j-1} e_1 \right\| \\
& \leq (1 + \varepsilon_1) \sum_{j=1}^t \|B\|^{t-j} \left\| v_{k+1} e_{k+1}^\top \left(T_k + 1/2 \cdot (V_k^\top E - E^\top V_k) \right)^{j-1} e_1 \right\| \\
& \quad \text{(Using } |\alpha_{k,k+1}| \leq \|w_k\| \leq \|Bv_k\| + \varepsilon_1 \leq 1 + \varepsilon_1 \text{ and (12))} \\
& \leq (1 + \varepsilon_1) \sum_{j=1}^t \|B\|^{t-j} \left\| v_{k+1} e_{k+1}^\top \left(\left(T_k + 1/2 \cdot (V_k^\top E - E^\top V_k) \right)^{j-1} - T_k^{j-1} \right) e_1 \right\| \\
& \quad \text{(Using } e_{k+1}^\top T_k^r e_1 = 0 \text{ for } r < k \text{ as } T_k \text{ is Upper Hessenberg)} \\
& \leq (1 + \varepsilon_1) \sum_{j=1}^t \left\| v_{k+1} e_{k+1}^\top \left(\sum_{i=1}^{j-1} \binom{j-1}{i} \left\| 1/2 \cdot (V_k^\top E - E^\top V_k) \right\|^i \|T_k\|^{j-1-i} \right) \right\| \\
& \quad \text{(Using sub-multiplicity of } \|\cdot\| \text{ and } \|B\| \leq 1) \\
& \leq (1 + \varepsilon_1) \sum_{j=1}^t \left(\sum_{i=1}^{j-1} \binom{j-1}{i} (\varepsilon_1 \sqrt{k+1})^i (1 + \varepsilon_1 \sqrt{k+1})^{j-1-i} \right) \\
& \quad \text{(Using } \left\| v_{k+1} e_{k+1}^\top \right\| \leq 1, \|T_k\| \leq (1 + \varepsilon_1 \sqrt{k+1}) \\
& \quad \quad \text{and } \left\| 1/2 \cdot (V_k^\top E - E^\top V_k) \right\| \leq \varepsilon_1 \sqrt{k+1}) \\
& \leq (1 + \varepsilon_1) \sum_{j=1}^t ((1 + 2\varepsilon_1 \sqrt{k+1})^{j-1} - 1) \\
& \leq t(1 + \varepsilon_1)((1 + 2\varepsilon_1 \sqrt{k+1})^{t-1} - 1). \tag{16}
\end{aligned}$$

Combining Equations (14),(15) and (16), we get,

$$\begin{aligned}
\left\| B^t v_0 - V_k \widehat{T}_k^t e_1 \right\| & \leq 2t\varepsilon_1 \sqrt{k+1} (1 + \varepsilon_1 \sqrt{k+1})^{t-1} + t(1 + \varepsilon_1)((1 + 2\varepsilon_1 \sqrt{k+1})^{t-1} - 1) \\
& \leq 2t\varepsilon_1 \sqrt{k+1} e^{\varepsilon_1(t-1)\sqrt{k+1}} + t e^{\varepsilon_1} (e^{2\varepsilon_1(t-1)\sqrt{k+1}} - 1) \quad \text{(Using } 1 + x \leq e^x) \\
& \leq 2t\varepsilon_1 \sqrt{k+1} e^{\varepsilon_1(t-1)\sqrt{k+1}} + t(e^{2\varepsilon_1 t \sqrt{k+1}} - 1) \\
& \leq \varepsilon_2/4 \cdot e^{\varepsilon_2/8(k+1)} + k(e^{\varepsilon_2/4(k+1)} - 1) \quad \text{(Using } \varepsilon_1 \leq \varepsilon_2/8(k+1)^{5/2} \text{ and } t \leq k) \\
& \leq \varepsilon_2/4 \cdot e^{\varepsilon_2/8(k+1)} + k\varepsilon_2/4(k+1) \cdot (1 + \varepsilon_2/4(k+1)) \quad \text{(Using } e^x \leq 1 + x + x^2 \text{ for } 0 \leq x \leq 1) \\
& \leq \varepsilon_2/2 + \varepsilon_2/2 \leq \varepsilon_2 \quad \text{(Using } \varepsilon_2 \leq 1 \text{ and } k \geq 0).
\end{aligned}$$

This proves the lemma. ■

Corollary 6.24 (Corollary 6.19 Restated) *For any polynomial p of degree at most k , if $\varepsilon_1 \leq \varepsilon_2/(2(k+1)^{3/2})$ and $\varepsilon_2 \leq 1$,*

$$\left\| p(B)v_0 - V_k p(\widehat{T}_k) e_1 \right\| \leq \varepsilon_2 \|p\|_1.$$

Proof: Suppose $p(x)$ is the polynomial $\sum_{t=0}^k a_t \cdot x^t$,

$$\begin{aligned} \left\| p(B)v_0 - V_k p(\widehat{T}_k)e_1 \right\| &= \left\| \sum_{t=0}^k a_t \cdot B^t v_0 - V_k \sum_{t=0}^k a_t \cdot \widehat{T}_k^t e_1 \right\| \\ &\leq \sum_{t=0}^k |a_t| \cdot \left\| B^t v_0 - V_k \widehat{T}_k^t e_1 \right\| \leq \varepsilon_2 \sum_{t=0}^k |a_t| = \varepsilon_2 \|p\|_1, \end{aligned}$$

where the last inequality follows from the previous lemma as $\varepsilon_2, \varepsilon_1$ satisfy the required conditions. \blacksquare

Lemma 6.25 (Lemma 6.20 Restated) Let V_k be the orthonormal basis and \widehat{T}_k be the matrix of coefficients generated by the above procedure. Let f be any function such that $f(B)$ and $f(T_k)$ are defined. Then,

$$\left\| f(B)v_0 - V_k f(\widehat{T}_k)e_1 \right\| \leq \min_{p \in \Sigma_k} \left(\varepsilon_2 \|p\|_1 + \max_{\lambda \in \Lambda(B)} |r_k(\lambda)| + \max_{\lambda \in \Lambda(\widehat{T}_k)} |r_k(\lambda)| \right). \quad (17)$$

Proof: Let p be any degree k polynomial. Let $r_k \stackrel{\text{def}}{=} f - p$. We express f as $p + r_k$ and use the previous lemma to bound the error in approximating $p(B)v_0$ by $V_k f(\widehat{T}_k)e_1$.

$$\begin{aligned} \left\| f(B)v_0 - V_k f(\widehat{T}_k)e_1 \right\| &\leq \left\| p(B)V_k e_1 - V_k p(\widehat{T}_k)e_1 \right\| + \left\| V_k r_k(B)e_1 - V_k r_k(\widehat{T}_k)e_1 \right\| \\ &\leq \varepsilon_2 \|p\|_1 + \|V_k r_k(B)e_1\| + \left\| V_k r_k(\widehat{T}_k)e_1 \right\| \\ &\leq \varepsilon_2 \|p\|_1 + \|r_k(B)\| + \left\| r_k(\widehat{T}_k) \right\| \\ &\leq \varepsilon_2 \|p\|_1 + \max_{\lambda \in \Lambda(B)} |r_k(\lambda)| + \max_{\lambda \in \Lambda(\widehat{T}_k)} |r_k(\lambda)|. \end{aligned}$$

Minimizing over p gives us our lemma. \blacksquare

Lemma 6.26 (Lemma 6.21 Restated) Given a polynomial p of degree k such that $p(0) = 0$ and

$$\sup_{t \in (0,1]} \left| e^{-k/t+k} - p(t) \right| = \sup_{x \in [0,\infty)} \left| e^{-x} - p\left((1+x/k)^{-1}\right) \right| \leq 1,$$

we must have $\|p\|_1 \leq (2k)^{k+1}$.

Proof: We know that $p(0) = 0$. Interpolating at the $k+1$ points $t = 0, 1/k, 2/k, \dots, 1$, we can use Lagrange's interpolation formula to give,

$$p(x) \equiv \sum_{i=1}^k \frac{\prod_{0 \leq j \leq k, j \neq i} (x - j/k)}{\prod_{0 \leq j \leq k, j \neq i} (i/k - j/k)} p(i/k).$$

The above identity is easily verified by evaluating the expression at the interpolation points and noting that it is a degree k polynomial agreeing with p at $k+1$ points. Thus, if we were to write $p(x) = \sum_{l=1}^k a_l \cdot x^l$ (note that $a_0 = 0$), we can express the coefficients a_l as follows.

$$a_l = \sum_{i=1}^k \frac{\sum_{\substack{0 \leq j_1 < \dots < j_{k-l} \leq k \\ j_1, \dots, j_{k-l} \neq i}} (-1)^{k-l-j_1/k - \dots - j_{k-l}/k}}{\prod_{0 \leq j \leq k, j \neq i} (i/k - j/k)} p(i/k).$$

Applying triangle inequality, and noting that $p(t)$ is a 1-uniform approximation to $e^{-k/t+k}$ for $t \in (0, 1]$, we get,

$$|a_l| \leq \sum_{i=1}^k \frac{\binom{k}{k-l}}{(1/k)^k} |p(i/k)| \leq \sum_{i=1}^k \frac{\binom{k}{k-l}}{(1/k)^k} \left(e^{-\frac{k(k-i)}{i}} + \delta/2 \right) \leq 2 \cdot k^{k+1} \binom{k}{k-l}.$$

Thus, we can bound the ℓ_1 norm of p as follows,

$$\|p\|_1 = \sum_{l=1}^k |a_l| \leq \sum_{l=1}^k 2 \cdot k^{k+1} \binom{k}{k-l} \leq (2k)^{k+1}$$

■

Lemma 6.27 (Lemma 6.22 Restated) *For any $\beta \geq 1$, any degree k polynomial p satisfies,*

$$\sup_{t \in (0, \beta]} |p(t) - f_k(t)| \leq \|p\|_1 \cdot (\beta^k - 1) + (f_k(\beta) - f_k(1)) + \sup_{t \in (0, 1]} |p(t) - f_k(t)|$$

Proof: Given a degree k polynomial p that approximates f_k over $(0, 1]$, we wish to bound the approximation error over $(0, \beta]$ for $\beta \geq 1$. We will split the error bound over $(0, 1]$ and $[1, \beta]$. Since we know that $f_k(1) - p(1)$ is small, we will bound the error over $[1, \beta]$ by applying triangle inequality and bounding the change in f_k and p over $[1, \beta]$ separately.

Let $\beta > 0$. First, let us calculate $\sup_{t \in [1, \beta]} |p(t) - f_k(t)|$.

$$\begin{aligned} \sup_{t \in [1, \beta]} |p(t) - f_k(t)| &\leq \sup_{t \in [1, \beta]} (|p(t) - p(1)| + |p(1) - f_k(1)| + |f_k(1) - f_k(t)|) \\ &\leq \sup_{t \in [1, \beta]} (\|p\|_1 \cdot \max_{0 \leq i \leq k} |t^i - 1^i| + |p(1) - f_k(1)| + |f_k(1) - f_k(t)|) \\ &\leq \sup_{t \in [1, \beta]} (\|p\|_1 \cdot \max_{0 \leq i \leq k} |t^i - 1^i|) + |p(1) - f_k(1)| + \sup_{t \in [1, \beta]} |f_k(1) - f_k(t)| \\ &\leq \|p\|_1 \cdot (\beta^k - 1) + |p(1) - f_k(1)| + \sup_{t \in [1, \beta]} |f_k(1) - f_k(t)| \\ &\quad \text{(Since } t \geq 1 \text{ and } t^k \text{ is increasing for } t \geq 0) \\ &\leq \|p\|_1 \cdot (\beta^k - 1) + |p(1) - f_k(1)| + (f_k(\beta) - f_k(1)) \\ &\quad \text{(Since } f_k(t) \text{ is an increasing function for } t \geq 0). \end{aligned}$$

Now, we can bound the error over the whole interval as follows.

$$\begin{aligned} \sup_{t \in (0, \beta]} |p(t) - f_k(t)| &= \max\left\{ \sup_{t \in (0, 1]} |p(t) - f_k(t)|, \sup_{t \in [1, \beta]} |p(t) - f_k(t)| \right\} \\ &\leq \max\left\{ \sup_{t \in (0, 1]} |p(t) - f_k(t)|, \right. \\ &\quad \left. \|p\|_1 \cdot (\beta^k - 1) + |p(1) - f_k(1)| + (f_k(\beta) - f_k(1)) \right\} \\ &= \|p\|_1 \cdot (\beta^k - 1) + (f_k(\beta) - f_k(1)) + \sup_{t \in (0, 1]} |p(t) - f_k(t)|. \end{aligned}$$

■

7 Uniform Approximations to e^{-x}

In this section, we discuss uniform approximations to e^{-x} and prove give a proof of Theorem 1.5 that shows the existence of polynomials that approximate e^{-x} uniformly over the interval $[a, b]$, whose degree grows as $\sqrt{b-a}$ and also gives a lower bound stating that this dependence is necessary. We restate a more precise version of the theorem here for completeness.

Theorem 7.1 (Uniform Approximation to e^{-x})

- **Upper Bound.** For every $0 \leq a < b$, and a given error parameter $0 < \delta \leq 1$, there exists a polynomial $p_{a,b,\delta}$ that satisfies,

$$\sup_{x \in [a,b]} |e^{-x} - p_{a,b,\delta}(x)| \leq \delta \cdot e^{-a},$$

and has degree $O\left(\sqrt{\max\{\log^2 1/\delta, (b-a) \cdot \log 1/\delta\}} \cdot (\log 1/\delta) \cdot \log \log 1/\delta\right)$.

- **Lower Bound.** For every $0 \leq a < b$ such that $a + \log_e 4 \leq b$, and $\delta \in (0, 1/8]$, any polynomial $p(x)$ that approximates e^{-x} uniformly over the interval $[a, b]$ up to an error of $\delta \cdot e^{-a}$, must have degree at least $\frac{1}{2} \cdot \sqrt{b-a}$.

Organization

We first discuss a few preliminaries (Section 7.1) and discuss relevant results that were already known and compare our result to the existing lower bounds (Section 7.2). Finally, we give a proof of the upper bound in Theorem 1.5 in Section 7.3 and of the lower bound in Section 7.4. Readers familiar with standard results in approximation theory can skip directly to the proofs in Section 7.3 and 7.4.

7.1 Preliminaries

Given an interval $[a, b]$, we are looking for low-degree polynomials (or rational functions) that approximate the function e^{-x} in the sup norm over the interval.

Definition 7.2 (δ -Approximation) A function g is called a δ -approximation to a function f over an interval \mathcal{I} , if, $\sup_{x \in \mathcal{I}} |f(x) - g(x)| \leq \delta$.

Such approximations are known as *uniform approximations* in approximation theory and have been studied quite extensively. We will consider both finite and infinite intervals \mathcal{I} .

For any positive integer k , let Σ_k denote the set of all degree k polynomials. We also need to define the ℓ_1 norm of a polynomial.

Definition 7.3 (ℓ_1 Norm of a Polynomial) Given a degree k polynomial $p \stackrel{\text{def}}{=} \sum_{i=0}^k a_i \cdot x^i$, the ℓ_1 norm of p , denoted as $\|p\|_1$ is defined as $\|p\|_1 = \sum_{i \geq 0} |a_i|$.

7.2 Known Approximation Results and Discussion.

Approximating the exponential is a classic question in Approximation Theory, see e.g. [10]. We ask the following question:

Question: Given $\delta \leq 1$ and $a < b$, what is the smallest degree of a polynomial that is an $\delta \cdot e^{-a}$ -approximation to e^{-x} over the interval $[a, b]$?

This question has been studied in the following form: Given λ , what is the best low degree polynomial (or rational function) approximation to $e^{\lambda x}$ over $[-1, 1]$? In a sense, these questions are equivalent, as is shown by the following lemma, proved using a linear shift of variables. A proof is included in Section 7.5.

Lemma 7.4 (Linear Variable Shift for Approximation) For any non-negative integer k and l and real numbers $b > a$,

$$\min_{p_k \in \Sigma_k, q_l \in \Sigma_l} \sup_{t \in [a, b]} \left| e^{-t} - \frac{p_k(t)}{q_l(t)} \right| = e^{-\frac{b+a}{2}} \cdot \min_{p_k \in \Sigma_k, q_l \in \Sigma_l} \sup_{x \in [-1, 1]} \left| e^{\frac{(b-a)}{2}x} - \frac{p_k(x)}{q_l(x)} \right|$$

Using the above lemma, we can translate the known results to our setting. As a starting point, we could approximate e^{-x} by truncating the Taylor series expansion of the exponential. We state the approximation achieved in the following lemma. A proof is included in Section 7.5.

Lemma 7.5 (Taylor Approximation) The degree k polynomial obtained by truncating Taylor's expansion of e^{-t} around the point $\frac{b+a}{2}$ is a uniform approximation to e^{-t} on the interval $[a, b]$ up to an error of

$$e^{-\frac{b+a}{2}} \cdot \sum_{i=k+1}^{\infty} \frac{1}{i!} \left(\frac{b-a}{2} \right)^i,$$

which is smaller than $\delta \cdot e^{-\frac{b+a}{2}}$ for $k \geq \max\left\{\frac{e^2(b-a)}{2}, \log 1/\delta\right\}$

A lower bound is known in the case where the size of the interval is fixed, i.e., $b - a = O(1)$.

Proposition 7.6 (Lower Bound for Polynomials over Fixed Interval, [3, 31]) For any $a, b \in \mathbb{R}$ such that $b - a$ is fixed, as k goes to infinity, the best approximation achieved by a degree k polynomial has error

$$(1 + o(1)) \frac{1}{(k+1)!} \left(\frac{b-a}{2} \right)^{k+1} \cdot e^{-\frac{b+a}{2}}.$$

In essence, this theorem states that if the size of the interval is fixed, the polynomials obtained by truncating the Taylor series expansion achieve asymptotically the least error possible and hence, the best asymptotic degree for achieving a $\delta \cdot e^{-\frac{b+a}{2}}$ -approximation. In addition, Saff [31] also shows that if, instead of polynomials, we allow rational functions where the degree of the denominator is a constant, the degree required for achieving a $\delta \cdot e^{-\frac{b+a}{2}}$ -approximation changes at most by a constant.

These results indicate that tight bounds on the answer to our question should be already known. In fact, at first thought, the optimality of the Taylor series polynomials seems to be in contradiction with our results. However, note the two important differences:

1. The error in our theorem is $e^{-a} \cdot \delta$, whereas, the Taylor series approximation involves error $e^{-\frac{b+a}{2}} \cdot \delta$, which is smaller, and hence requires larger degree.

2. Moreover, the lower bound applies only when the length of the interval $(b - a)$ is constant, in which case, our theorem says that the required degree is $\text{poly}(\log 1/\delta)$, which is $\Omega(\log 1/\delta)$, in accordance with the lower bound.

If the length of the interval $[a, b]$ grows unbounded (as is the case for our applications to the Balanced Separator problem in the previous sections), the main advantage of using polynomials from Theorem 7.1 is the improvement in the degree from linear in $(b - a)$ to $\sqrt{b - a}$.

7.3 Proof of Upper Bound in Theorem 1.5

In this section, we use Theorem 6.8 by Saff, Schonhage and Varga [30], rather, more specifically, Corollary 6.9 to give a proof of the upper bound result in Theorem 1.5. We restate Corollary 6.9 for completeness.

Corollary 7.7 (Corollary 6.9 Restated, [30]) *There exists constants $c_1 \geq 1$ and k_0 such that, for any integer $k \geq k_0$, there exists a polynomial $p_k^*(x)$ of degree k such that $p_k^*(0) = 0$, and,*

$$\sup_{t \in (0,1)} \left| e^{-k/t+k} - p_k^*(t) \right| = \sup_{x \in [0,\infty)} \left| e^{-x} - p_k^* \left((1 + x/k)^{-1} \right) \right| \leq c_1 k \cdot 2^{-k}. \quad (18)$$

Our approach is to compose the polynomial p_k^* given by Corollary 7.7 with polynomials approximating $(1 + x/k)^{-1}$, to construct polynomials approximating e^{-x} . We first show the existence of polynomials approximating x^{-1} , and from these polynomials, we will derive approximations to $(1 + x/k)^{-1}$.

Our goal is to find a polynomial q of degree k , that minimizes $\sup_{x \in [a,b]} |q(x) - 1/x|$. We slightly modify this optimization to minimizing $\sup_{x \in [a,b]} |x \cdot q(x) - 1|$. Note that $x \cdot q(x) - 1$ is a polynomial of degree $k + 1$ which evaluates to -1 at $x = 0$, and conversely every polynomial that evaluates to -1 at 0 can be written as $x \cdot q(x) - 1$ for some q . So, this is equivalent to minimizing $\sup_{x \in [a,b]} |q_1(x)|$, for a degree $k + 1$ polynomial q_1 such that $q_1(0) = -1$. By scaling and multiplying by -1 , this is equivalent to finding a polynomial q_2 , that maximizes $q_2(0)$, subject to $\sup_{x \in [a,b]} |q_2(x)| \leq 1$. If we shift and scale the interval $[a, b]$ to $[-1, 1]$, the optimal solution to this problem is known to be given by the well known Chebyshev polynomials. We put all these ideas together to prove the following lemma. A complete proof is included in Section 7.5.

Lemma 7.8 (Approximating x^{-1}) *For every $\varepsilon > 0$, $b > a > 0$, there exists a polynomial $q_{a,b,\varepsilon}(x)$ of degree $\left\lceil \sqrt{\frac{b}{a}} \log \frac{2}{\varepsilon} \right\rceil$ such that $\sup_{x \in [a,b]} |x \cdot q_{a,b,\varepsilon}(x) - 1| \leq \varepsilon$.*

As a simple corollary, we can approximate $(1 + x/k)^{-1}$, or rather generally, $(1 + \nu x)^{-1}$ for some $\nu > 0$, by polynomials. A proof is included in Section 7.5.

Corollary 7.9 (Approximating $(1 + \nu x)^{-1}$) *For every $\nu > 0, \varepsilon > 0$ and $b > a \geq 0$, there exists a polynomial $q_{\nu,a,b,\varepsilon}^*(x)$ of degree $\left\lceil \sqrt{\frac{1+\nu b}{1+\nu a}} \log \frac{2}{\varepsilon} \right\rceil$ such that $\sup_{x \in [a,b]} |(1 + \nu x) \cdot q_{\nu,a,b,\varepsilon}^*(x) - 1| \leq \varepsilon$.*

The above corollary implies that the expression $(1 + \nu x) \cdot q^*$ is within $1 \pm \varepsilon$ on $[a, b]$. If ε is small, for a small positive integer t , $[(1 + \nu x) \cdot q^*]^t$ should be at most $1 \pm O(t\varepsilon)$. The following lemma, proved using the binomial theorem proves this formally. A proof is included in Section 7.5.

Lemma 7.10 (Approximating $(1 + vx)^{-t}$) For all real $\varepsilon > 0$, $b > a \geq 0$ and positive integer t ; if $t\varepsilon \leq 1$, then,

$$\sup_{x \in [a, b]} |((1 + vx) \cdot q_{v, a, b, \varepsilon}^*(x))^t - 1| \leq 2t\varepsilon,$$

where $q_{v, a, b, \varepsilon}^*$ is the polynomial given by Corollary 7.9.

Since q^* is an approximation to $(1 + x/k)^{-1}$, in order to bound the error for the composition $p_k^*(q^*)$, we need to bound how the value of the polynomial p_k^* changes on small perturbations in the input. We will use the following crude bound in terms of the ℓ_1 norm of the polynomial.

Lemma 7.11 (Error in Polynomial) For any polynomial p of degree k , and any $x, y \in \mathbb{R}$, $|p(x) - p(y)| \leq \|p\|_1 \cdot \max_{0 \leq i \leq k} |x^i - y^i|$.

In order to utilize the above lemma, we will need a bound on the ℓ_1 norm of p_k^* , which is provided by Lemma 6.21, that bounds the ℓ_1 norm of any polynomial in $(1 + x/k)^{-1}$ that approximates the exponential function and has no constant term. We restate the lemma here for completeness.

Lemma 7.12 (ℓ_1 -norm Bound. Lemma 6.21 Restated) Given a polynomial p of degree k such that $p(0) = 0$ and

$$\sup_{t \in (0, 1]} \left| e^{-k/t+k} - p(t) \right| = \sup_{x \in [0, \infty)} \left| e^{-x} - p\left((1 + x/k)^{-1}\right) \right| \leq 1,$$

we must have $\|p\|_1 \leq (2k)^{k+1}$.

We can now analyze the error in approximating e^{-x} by the polynomial $p_k^*(q^*)$ and give a proof for Theorem 1.5.

Proof: Given $\delta \leq 1$, let $k = \max\{k_0, \log_2 4c_1/\delta + 2 \log_2 \log_2 4c_1/\delta\} = O(\log 1/\delta)$, where k_0, c_1 are the constants given by Corollary 7.7. Moreover, p_k^* is the degree k polynomial given by Corollary 7.7, which gives, $p_k^*(0) = 0$, and,

$$\begin{aligned} \sup_{x \in [0, \infty)} \left| e^{-x} - p_k^*\left((1 + x/k)^{-1}\right) \right| &\leq \frac{\delta}{4} \cdot \frac{\log_2 4c_1/\delta + 2 \log_2 \log_2 4c_1/\delta}{(\log_2 4c_1/\delta)^2} \\ &\leq \frac{\delta}{4} \cdot \frac{1}{\log_2 4c_1/\delta} \left(1 + 2 \cdot \frac{\log_2 \log_2 4c_1/\delta}{\log_2 4c_1/\delta} \right) \leq \frac{\delta}{4} \cdot \frac{1}{2} \cdot 3 \leq \frac{\delta}{2}, \end{aligned} \quad (19)$$

where the last inequality uses $\delta \leq 1 \leq c_1$ and $\log_2 x \leq x, \forall x \geq 0$. Thus, we can use Lemma 7.12 to conclude that $\|p_k^*\|_1 \leq (2k)^{k+1}$.

Let $\nu \stackrel{\text{def}}{=} 1/k$. Define ε as $\varepsilon \stackrel{\text{def}}{=} \frac{\delta}{2(2k)^{k+2}}$. Let $p_{a, b, \delta}(x) \stackrel{\text{def}}{=} e^{-a} \cdot p_k^*\left(q_{\nu, 0, b-a, \varepsilon}^*(x-a)\right)$, where $q_{\nu, 0, b-a, \varepsilon}^*$ is the polynomial of degree $\left\lceil \sqrt{1 + \nu(b-a)} \log \frac{2}{\varepsilon} \right\rceil$ given by Corollary 7.9. Observe that $p_{a, b, \delta}(x)$ is a polynomial of degree that is the product of the degrees of p_k^* and $q_{\nu, 0, b-a, \varepsilon}^*$, i.e., $k \left\lceil \sqrt{1 + \nu(b-a)} \log \frac{2}{\varepsilon} \right\rceil$. Also note that $k\varepsilon < 1$ and hence we can use Lemma 7.10. We show that $p_{a, b, \delta}$ is a uniform δ -

approximation to e^{-x} on the interval $[a, b]$.

$$\begin{aligned}
\sup_{x \in [a, b]} |e^{-x} - p_{a, b, \delta}(x)| &= e^{-a} \cdot \sup_{x \in [a, b]} \left| e^{-(x-a)} - e^a \cdot p_{a, b, \delta}(x) \right| = e^{-a} \cdot \sup_{y \in [0, b-a]} \left| e^{-y} - e^a \cdot p_{a, b, \delta}(y+a) \right| \\
&\stackrel{\text{by def}}{=} e^{-a} \cdot \sup_{y \in [0, b-a]} \left| e^{-y} - p_k^*(q_{v, 0, b-a, \varepsilon}^*(y)) \right| \\
&\stackrel{\Delta\text{-ineq.}}{\leq} e^{-a} \cdot \sup_{y \in [0, b-a]} \left(\left| e^{-y} - p_k^*((1+vy)^{-1}) \right| + \left| p_k^*((1+vy)^{-1}) - p_k^*(q_{v, 0, b-a, \varepsilon}^*(y)) \right| \right) \\
&\leq e^{-a} \cdot \sup_{y \in [0, b-a]} \left| e^{-y} - p_k^*((1+vy)^{-1}) \right| + e^{-a} \cdot \sup_{y \in [0, b-a]} \left| p_k^*((1+vy)^{-1}) - p_k^*(q_{v, 0, b-a, \varepsilon}^*(y)) \right| \\
&\stackrel{\text{Lem. 7.11}}{\leq} e^{-a} \cdot \sup_{y \in [0, \infty)} \left| e^{-y} - p_k^*((1+vy)^{-1}) \right| \\
&\quad + e^{-a} \cdot \|p_k^*\|_1 \cdot \max_{0 \leq i \leq k} \sup_{y \in [0, b-a]} \left| (1+vy)^{-i} - (q_{v, 0, b-a, \varepsilon}^*(y))^i \right| \\
&\stackrel{\text{Eq. (19)}}{\leq} e^{-a} \cdot \frac{\delta}{2} + e^{-a} \cdot \|p_k^*\|_1 \cdot \max_{0 \leq i \leq k} \sup_{y \in [0, b-a]} (1+vy)^{-i} \left| 1 - ((1+vy) \cdot q_{v, 0, b-a, \varepsilon}^*(y))^i \right| \\
&\stackrel{\text{Lem. 7.10, 7.12}}{\leq} e^{-a} \cdot \frac{\delta}{2} + 2k\varepsilon \cdot e^{-a} \cdot (2k)^{k+1} \leq \delta \cdot e^{-a}
\end{aligned}$$

The degree of the polynomial $p_{a, b, \delta}$ is

$$\begin{aligned}
k \left\lceil \sqrt{1 + v(b-a) \log \frac{2}{\varepsilon}} \right\rceil &= O \left(\sqrt{k^2 + k(b-a)} \cdot (k \log k + \log 1/\delta) \right) \\
&= O \left(\sqrt{\max\{\log^2 1/\delta, \log 1/\delta(b-a)\}} \cdot (\log 1/\delta) \cdot \log \log 1/\delta \right)
\end{aligned}$$

7.4 Proof of Lower Bound in Theorem 1.5

In this section, we will use the following well known theorem of Markov from approximation theory to give a proof of the lower bound result in Theorem 1.5.

Theorem 7.13 (Markov, See [10]) *Let $p : \mathbb{R} \rightarrow \mathbb{R}$ be a univariate polynomial of degree d such that any real number $a_1 \leq x \leq a_2$, satisfies $b_1 \leq p(x) \leq b_2$. Then, for all $a_1 \leq x \leq a_2$, the derivative of p satisfies $|p'(x)| \leq d^2 \cdot \frac{b_2 - b_1}{a_2 - a_1}$.*

The idea is to first use uniform approximation bound to bound the value of the polynomial within the interval of approximation. Next, we use the approximation bound and the *Mean Value theorem* to show that there must exist a point t in the interval where $|p'(t)|$ is large. We plug both these bounds into Markov's theorem to deduce our lower bound.

Proof: Suppose p is a degree k polynomial that is a uniform approximation to e^{-x} over the interval $[a, b]$ up to an error of $\delta \cdot e^{-a}$. For any $x \in [a, b]$, this bounds the values p can take at x . Since p is a uniform approximation to e^{-x} over $[a, b]$ up to an error of $\delta \cdot e^{-a}$, we know that for all $x \in [a, b]$, $e^{-x} - \delta \cdot e^{-a} \leq p(x) \leq e^{-x} + \delta \cdot e^{-a}$. Thus, $\max_{x \in [a, b]} p(x) \leq e^{-a} + \delta \cdot e^{-a}$ and $\min_{x \in [a, b]} p(x) \geq e^{-b} - \delta \cdot e^{-a}$.

Assume that $\delta \leq 1/8$, and $b \geq a + \log_e 4 \geq a + \log_e 2/(1-4\delta)$. Applying the *Mean Value theorem* on the interval $[a, a + \log_e 2/(1-4\delta)]$, we know that there exists $t \in [a, a + \log_e 2/(1-4\delta)]$, such that,

$$\begin{aligned} |p'(t)| &= \left| \frac{p(a + \log_e 2/(1-4\delta)) - p(a)}{\log_e 2/(1-4\delta)} \right| \geq \frac{(e^{-a} - \delta \cdot e^{-a}) - (e^{-a - \log_e 2/(1-4\delta)} + \delta \cdot e^{-a})}{\log_e 2/(1-4\delta)} \\ &\geq e^{-a} \frac{1 - 2\delta - \frac{(1-4\delta)}{2}}{\log_e 2/(1-4\delta)} = e^{-a} \frac{1}{2 \log_e 2/(1-4\delta)} \end{aligned}$$

We plug this in Markov's theorem (Theorem 7.13) stated above to deduce,

$$e^{-a} \frac{1}{2 \log_e 2/(1-4\delta)} \leq k^2 \frac{(e^{-a} + \delta \cdot e^{-a}) - (e^{-b} - \delta \cdot e^{-a})}{b - a} \leq k^2 \cdot e^{-a} \cdot \frac{1 + 2\delta}{b - a}.$$

Rearranging, we get,

$$k \geq \sqrt{\frac{b - a}{2 \cdot (1 + 2\delta) \cdot \log_e 2/(1-4\delta)}} \geq \sqrt{\frac{b - a}{2 \cdot 5/4 \cdot \log_e 4}} \geq \frac{1}{2} \cdot \sqrt{b - a},$$

where the second inequality uses $\delta \leq 1/8$. ■

7.5 Remaining Proofs

Lemma 7.14 (Lemma 7.4 Restated) For any non-negative integer k and l and real numbers $b > a$,

$$\min_{p_k \in \Sigma_k, q_l \in \Sigma_l} \sup_{t \in [a, b]} \left| e^{-t} - \frac{p_k(t)}{q_l(t)} \right| = e^{-\frac{b+a}{2}} \cdot \min_{p_k \in \Sigma_k, q_l \in \Sigma_l} \sup_{x \in [-1, 1]} \left| e^{\frac{(b-a)}{2}x} - \frac{p_k(x)}{q_l(x)} \right|$$

Proof: Using the substitution $t \stackrel{\text{def}}{=} \frac{(b+a)}{2} - \frac{(b-a)}{2}x$,

$$\begin{aligned} \min_{p_k \in \Sigma_k, q_l \in \Sigma_l} \sup_{t \in [a, b]} \left| e^{-t} - \frac{p_k(t)}{q_l(t)} \right| &= \min_{p_k \in \Sigma_k, q_l \in \Sigma_l} \sup_{x \in [-1, 1]} \left| e^{-\frac{(b+a)}{2} + \frac{(b-a)}{2}x} - \frac{p_k\left(\frac{(b+a)}{2} - \frac{(b-a)}{2}x\right)}{q_l\left(\frac{(b+a)}{2} - \frac{(b-a)}{2}x\right)} \right| \\ &= e^{-\frac{b+a}{2}} \cdot \min_{p'_k \in \Sigma_k, q'_l \in \Sigma_l} \sup_{x \in [-1, 1]} \left| e^{\frac{(b-a)}{2}x} - \frac{p'_k(x)}{q'_l(x)} \right| \end{aligned}$$

Lemma 7.15 (Lemma 7.5 Restated) The degree k polynomial obtained by truncating Taylor's expansion of e^{-t} around the point $(b+a)/2$ is a uniform approximation to e^{-t} on the interval $[a, b]$ up to an error of

$$e^{-\frac{b+a}{2}} \cdot \sum_{i=k+1}^{\infty} \frac{1}{i!} \left(\frac{(b-a)}{2} \right)^i,$$

which is smaller than δ for $k \geq \max\left\{\frac{e^2(b-a)}{2}, \log 1/\delta\right\}$

Proof: Let $q_k(t)$ be the degree k Taylor approximation of the function e^{-t} around the point $(b+a)/2$, i.e., $q_k(t) \stackrel{\text{def}}{=} e^{-\frac{(b+a)}{2}} \sum_{i=0}^k \frac{1}{i!} \left(t - \frac{(b+a)}{2}\right)^i$.

$$\sup_{t \in [a, b]} |e^{-t} - q_k(t)| = \sup_{t \in [a, b]} e^{-\frac{b+a}{2}} \cdot \left| \sum_{i=k+1}^{\infty} \frac{1}{i!} \left(t - \frac{(b+a)}{2}\right)^i \right| = e^{-\frac{b+a}{2}} \cdot \sum_{i=k+1}^{\infty} \frac{1}{i!} \left(\frac{(b-a)}{2}\right)^i$$

Using the inequality $i! > \left(\frac{i}{e}\right)^i$, for all i , and assuming $k \geq \frac{e^2(b-a)}{2}$, we get,

$$\sum_{i=k+1}^{\infty} \frac{1}{i!} \left(\frac{(b-a)}{2}\right)^i \leq \sum_{i=k+1}^{\infty} \left(\frac{e(b-a)}{2i}\right)^i \leq \sum_{i=k+1}^{\infty} e^{-i} = \frac{1}{e-1} e^{-k},$$

which is smaller than δ for $k \geq \log^{1/\delta}$. ■

Lemma 7.16 (Lemma 7.8 Restated) For every $\varepsilon > 0$, $b > a > 0$, there exists a polynomial $q_{a,b,\varepsilon}(x)$ of degree $\left\lceil \sqrt{\frac{b}{a}} \log \frac{2}{\varepsilon} \right\rceil$ such that

$$\sup_{x \in [a,b]} |x \cdot q_{a,b,\varepsilon}(x) - 1| \leq \varepsilon.$$

Proof: If $T_{k+1}(x)$ denotes the degree $k+1$ Chebyshev polynomial, consider the function,

$$q_{a,b,\varepsilon}(x) \stackrel{\text{def}}{=} \frac{1}{x} \left(1 - \frac{T_{k+1}\left(\frac{b+a-2x}{b-a}\right)}{T_{k+1}\left(\frac{b+a}{b-a}\right)} \right).$$

First, we need to prove that the above expression is a polynomial. Clearly $1 - \frac{T_{k+1}\left(\frac{b+a-2x}{b-a}\right)}{T_{k+1}\left(\frac{b+a}{b-a}\right)}$ is a polynomial and evaluates to 0 at $x = 0$. Thus, it must have x as a factor. Thus $q_{a,b,\varepsilon}$ is a polynomial of degree k . Let $\kappa = b/a$ and note that $\kappa > 1$. Thus,

$$\begin{aligned} \sup_{x \in [a,b]} |x \cdot q_{a,b,\varepsilon}(x) - 1| &= \sup_{x \in [a,b]} \left| \frac{T_{k+1}\left(\frac{b+a-2x}{b-a}\right)}{T_{k+1}\left(\frac{b+a}{b-a}\right)} \right| \\ &\leq T_{k+1}\left(\frac{b+a}{b-a}\right)^{-1} \quad (\text{Since } |T_{k+1}(y)| \leq 1 \text{ for } |y| \leq 1) \\ &= 2 \left(\left(\frac{\sqrt{\kappa}+1}{\sqrt{\kappa}-1}\right)^{k+1} + \left(\frac{\sqrt{\kappa}-1}{\sqrt{\kappa}+1}\right)^{k+1} \right)^{-1} \quad (\text{By def.}) \\ &\leq 2 \left(\frac{\sqrt{\kappa}+1}{\sqrt{\kappa}-1}\right)^{-k-1} \quad (\text{Each term is positive since } \sqrt{\kappa} > 1) \\ &= 2 \left(\frac{1-1/\sqrt{\kappa}}{1+1/\sqrt{\kappa}}\right)^{k+1} \\ &\leq 2 \cdot (1-1/\sqrt{\kappa})^{k+1} \leq 2 \cdot e^{-(k+1)/\sqrt{\kappa}} \leq \varepsilon, \end{aligned}$$

for $k = \left\lceil \sqrt{\kappa} \log \frac{2}{\varepsilon} \right\rceil$. The first inequality follows from the fact that $|T_{k+1}(x)| \leq 1$ for all $|x| \leq 1$. ■

Corollary 7.17 (Corollary 7.9 Restated) For every $\nu > 0$, $\varepsilon > 0$ and $b > a \geq 0$, there exists a polynomial $q_{\nu,a,b,\varepsilon}^*(x)$ of degree $\left\lceil \sqrt{\frac{1+\nu b}{1+\nu a}} \log \frac{2}{\varepsilon} \right\rceil$ such that

$$\sup_{x \in [a,b]} |(1+\nu x) \cdot q_{\nu,a,b,\varepsilon}^*(x) - 1| \leq \varepsilon.$$

Proof: Consider the polynomial $q_{v,a,b,\varepsilon}^*(x) \stackrel{\text{def}}{=} q_{1+va,1+vb,\varepsilon}(1+vx)$, where $q_{1+va,1+vb,\varepsilon}$ is given by the previous lemma.

$$\begin{aligned} \sup_{x \in [a,b]} |(1+vx) \cdot q_{v,a,b,\varepsilon}^*(x) - 1| &= \sup_{x \in [a,b]} |(1+vx) \cdot q_{1+va,1+vb,\varepsilon}(1+vx) - 1| \\ &\stackrel{\text{def}}{=} \sup_{t \in [1+va,1+vb]} |t \cdot q_{1+va,1+vb,\varepsilon}(t) - 1| \stackrel{\text{Lem. 7.8}}{\leq} \varepsilon. \end{aligned}$$

Since $1+vx$ is a linear transformation, the degree of $q_{v,a,b,\varepsilon}^*$ is the same as that of $q_{1+va,1+vb,\varepsilon}$, which is, $\left\lceil \sqrt{\frac{1+vb}{1+va}} \log \frac{2}{\varepsilon} \right\rceil$. ■

Lemma 7.18 (Lemma 7.10 Restated) For all real $\varepsilon > 0$, $b > a \geq 0$ and positive integer t ; if $t\varepsilon \leq 1$, then,

$$\sup_{x \in [a,b]} |((1+vx) \cdot q_{v,a,b,\varepsilon}^*(x))^t - 1| \leq 2t\varepsilon,$$

where $q_{v,a,b,\varepsilon}^*$ is the polynomial given by Corollary 7.9.

Proof: We write the expression $(1+vx) \cdot q_{v,a,b,\varepsilon}^*(x)$ as 1 plus an error term and then use the Binomial Theorem to expand the t^{th} power.

$$\begin{aligned} \sup_{x \in [a,b]} |((1+vx) \cdot q_{v,a,b,\varepsilon}^*(x))^t - 1| &= \sup_{x \in [a,b]} \left| (1 - [1 - (1+vx) \cdot q_{v,a,b,\varepsilon}^*(x)])^t - 1 \right| \\ &= \sup_{x \in [a,b]} \left| \sum_{i=1}^t \binom{t}{i} (1 - (1+vx) \cdot q_{v,a,b,\varepsilon}^*(x))^i \right| \\ &\leq \sup_{x \in [a,b]} \sum_{i=1}^t \binom{t}{i} |1 - (1+vx) \cdot q_{v,a,b,\varepsilon}^*(x)|^i \\ &\leq \sum_{i=1}^t \binom{t}{i} \sup_{x \in [a,b]} |1 - (1+vx) \cdot q_{v,a,b,\varepsilon}^*(x)|^i \\ &\stackrel{\text{Cor. 7.9}}{\leq} \sum_{i=1}^t \binom{t}{i} \varepsilon^i = (1+\varepsilon)^t - 1 \\ &\leq \exp(t\varepsilon) - 1 \leq 1 + t\varepsilon + (t\varepsilon)^2 - 1 \leq 2t\varepsilon, \end{aligned}$$

where the second last inequality uses $e^x \leq 1 + x + x^2$ for $x \in [0, 1]$.⁶ ■

Lemma 7.19 (Lemma 7.11 Restated) For any polynomial p of degree k , and any $x, y \in \mathbb{R}$

$$|p(x) - p(y)| \leq \|p\|_1 \cdot \max_{0 \leq i \leq k} |x^i - y^i|$$

Proof: Suppose $p(t)$ is the polynomial $\sum_{i=0}^k a_i \cdot t^i$, where $a_i \in \mathbb{R}$. Then,

$$\begin{aligned} |p(x) - p(y)| &= \left| \sum_{i=0}^k a_i \cdot x^i - \sum_{i=0}^k a_i \cdot y^i \right| \leq \sum_{i=0}^k |a_i| |x^i - y^i| \\ &\leq \left(\sum_{i=0}^k |a_i| \right) \max_{0 \leq i \leq k} |x^i - y^i| = \|p\|_1 \cdot \max_{0 \leq i \leq k} |x^i - y^i| \end{aligned}$$

⁶For $x \in [0, 1]$, $e^x = \sum_{i \geq 0} \frac{x^i}{i!} = 1 + x + x^2 \left(\frac{1}{2!} + \frac{x}{3!} + \dots \right) \leq 1 + x + x^2 \left(\frac{1}{2} + \frac{x}{2^2} + \frac{x}{2^3} + \dots \right) \leq 1 + x + x^2$

References

- [1] Ittai Abraham, Yair Bartal, and Ofer Neiman. Advances in metric embedding theory. *Advances in Mathematics*, 228(6):3026 – 3126, 2011.
- [2] Noga Alon and V. D. Milman. λ_1 , isoperimetric inequalities for graphs, and superconcentrators. *J. Comb. Theory, Ser. B*, 38(1):73–88, 1985.
- [3] S. Ja. Al’per. Asymptotic values of best approximation of analytic functions in a complex domain. *Uspehi Mat. Nauk*, 14(1 (85)):131–134, 1959.
- [4] Reid Andersen, Fan R. K. Chung, and Kevin J. Lang. Local graph partitioning using pagerank vectors. In *FOCS’06: Proc. 47th Ann. IEEE Symp. Foundations of Computer Science*, pages 475–486, 2006.
- [5] Reid Andersen and Kevin J. Lang. An algorithm for improving graph partitions. In *Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms, SODA ’08*, pages 651–660, 2008.
- [6] Reid Andersen and Yuval Peres. Finding sparse cuts locally using evolving sets. In *STOC ’09: Proc. 41st Ann. ACM Symp. Theory of Computing*, pages 235–244, 2009.
- [7] Sanjeev Arora and Satyen Kale. A combinatorial, primal-dual approach to semidefinite programs. In *STOC ’07: Proc. 39th Ann. ACM Symp. Theory of Computing*, pages 227–236, 2007.
- [8] Sanjeev Arora, Satish Rao, and Umesh Vazirani. Expander flows, geometric embeddings and graph partitioning. In *STOC ’04: Proc. 36th Ann. ACM Symp. Theory of Computing*, pages 222–231, New York, NY, USA, 2004. ACM.
- [9] Rajendra Bhatia. *Matrix Analysis (Graduate Texts in Mathematics)*. Springer, 1996.
- [10] E. W. Cheney. *Introduction to approximation theory / E.W. Cheney*. McGraw-Hill, New York :, 1966.
- [11] Fan R.K. Chung. *Spectral Graph Theory (CBMS Regional Conference Series in Mathematics, No. 92)*. American Mathematical Society, 1997.
- [12] Byron L. Ehle. A-stable methods and padé approximations to the exponential. *Siam J. on Mathematical Analysis*, 4(4):671–680, 1973.
- [13] Jasper vanden Eshof and Marlis Hochbruck. Preconditioning lanczos approximations to the matrix exponential. *SIAM J. Sci. Comput.*, 27:1438–1457, November 2005.
- [14] Fan Chung Graham. A local graph partitioning algorithm using heat kernel pagerank. In *WAW’09*, pages 62–75, 2009.
- [15] G. Iyengar, D. J. Phillips, and C. Stein. Approximating semidefinite packing programs. *SIAM Journal on Optimization*, 21(1):231–268, 2011.
- [16] G. Iyengar, David J. Phillips, and Clifford Stein. Approximation algorithms for semidefinite packing problems with applications to maxcut and graph coloring. In *IPCO’05: Proc. 11th Conf. Integer Programming and Combinatorial Optimization*, pages 152–166, 2005.

- [17] Donald J and Newman. Rational approximation to e^{-x} . *Journal of Approximation Theory*, 10(4):301 – 303, 1974.
- [18] Satyen Kale. Efficient algorithms using the multiplicative weights update method. Technical report, Princeton University, Department of Computer Science, 2007.
- [19] R. Kannan, S. Vempala, and A. Vetta. On clusterings-good, bad and spectral. In *FOCS'00: Proc. 41th Ann. IEEE Symp. Foundations of Computer Science*, page 367, Washington, DC, USA, 2000. IEEE Computer Society.
- [20] Rohit Khandekar, Satish Rao, and Umesh Vazirani. Graph partitioning using single commodity flows. In *STOC '06: Proc. 38th Ann. ACM Symp. Theory of Computing*, pages 385–390, New York, NY, USA, 2006. ACM.
- [21] Ioannis Koutis, Gary L. Miller, and Richard Peng. Approaching optimality for solving sdd systems. *CoRR*, abs/1003.2958, 2010.
- [22] Aleksander Madry. Fast approximation algorithms for cut-based problems in undirected graphs. In *FOCS*, pages 245–254, 2010.
- [23] Lorenzo Orecchia. *Fast Approximation Algorithms for Graph Partitioning using Spectral and Semidefinite-Programming Techniques*. PhD thesis, EECS Department, University of California, Berkeley, May 2011.
- [24] Lorenzo Orecchia, Leonard J. Schulman, Umesh V. Vazirani, and Nisheeth K. Vishnoi. On partitioning graphs via single commodity flows. In *STOC '08: Proc. 40th Ann. ACM Symp. Theory of Computing*, pages 461–470, 2008.
- [25] Lorenzo Orecchia and Nisheeth K. Vishnoi. Towards an sdp-based approach to spectral methods: A nearly-linear-time algorithm for graph partitioning and decomposition. In *SODA*, pages 532–545, 2011.
- [26] Victor Y. Pan and Zhao Q. Chen. The complexity of the matrix eigenproblem. In *STOC'99*, pages 507–516, 1999.
- [27] E. Parzen. *Stochastic processes*. Classics in applied mathematics. Society for Industrial and Applied Mathematics, 1999.
- [28] Harald Räcke. Optimal hierarchical decompositions for congestion minimization in networks. In *Proceedings of the 40th annual ACM symposium on Theory of computing*, STOC '08, pages 255–264, New York, NY, USA, 2008. ACM.
- [29] Y. Saad. Analysis of some krylov subspace approximations to the matrix exponential operator. *SIAM J. Numer. Anal.*, 29:209–228, February 1992.
- [30] E. B. Saff, A. Schonhage, and R. S. Varga. Geometric convergence to e^{-z} by rational functions with real poles. *Numerische Mathematik*, 25:307–322, 1975.
- [31] E.B Saff. On the degree of best rational approximation to the exponential function. *Journal of Approximation Theory*, 9(2):97 – 101, 1973.

- [32] Kirk Schloegel, George Karypis, Vipin Kumar, J. Dongarra, I. Foster, G. Fox, K. Kennedy, A. White, and Morgan Kaufmann. Graph partitioning for high performance scientific simulations, 2000.
- [33] Jonah Sherman. Breaking the multicommodity flow barrier for $O(\sqrt{\log n})$ -approximations to sparsest cut. In *FOCS'09: Proc. 50th Ann. IEEE Symp. Foundations of Computer Science*, 2009.
- [34] Jonathan Shewchuk. An introduction to the conjugate gradient method without the agonizing pain. 1994. <http://www.cs.cmu.edu/~quake-papers/painless-conjugate-gradient.pdf>.
- [35] David B. Shmoys. *Cut problems and their application to divide-and-conquer*, pages 192–235. 1997.
- [36] Daniel A. Spielman and Shang-Hua Teng. Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In *STOC '04: Proc. 36th Ann. ACM Symp. Theory of Computing*, pages 81–90, New York, NY, USA, 2004. ACM.
- [37] Daniel A. Spielman and Shang-Hua Teng. Nearly-linear time algorithms for preconditioning and solving symmetric, diagonally dominant linear systems. *CoRR*, abs/cs/0607105, 2006.
- [38] Daniel A. Spielman and Shang-Hua Teng. A local clustering algorithm for massive graphs and its application to nearly-linear time graph partitioning. *CoRR*, abs/0809.3232, 2008.
- [39] C. Underhill and A. Wragg. Convergence properties of padé approximants to $\exp(z)$ and their derivatives. *IMA Journal of Applied Mathematics*, 11(3):361–367, 1973.