

## QUANTUM SPEEDUP FOR GRAPH SPARSIFICATION, CUT APPROXIMATION, AND LAPLACIAN SOLVING\*

SIMON APERS<sup>†</sup> AND RONALD DE WOLF<sup>‡</sup>

**Abstract.** Graph sparsification underlies a large number of algorithms, ranging from approximation algorithms for cut problems to solvers for linear systems in the graph Laplacian. In its strongest form, “spectral sparsification” reduces the number of edges to near-linear in the number of nodes, while approximately preserving the cut and spectral structure of the graph. In this work we demonstrate a polynomial quantum speedup for spectral sparsification and many of its applications. In particular, we give a quantum algorithm that, given a weighted graph with  $n$  nodes and  $m$  edges, outputs a classical description of an  $\epsilon$ -spectral sparsifier in sublinear time  $\tilde{O}(\sqrt{mn}/\epsilon)$ . This contrasts with the optimal classical complexity  $\tilde{O}(m)$ . We also prove that our quantum algorithm is optimal up to polylog-factors. The algorithm builds on a string of existing results on sparsification, graph spanners, quantum algorithms for shortest paths, and efficient constructions for  $k$ -wise independent random strings. Our algorithm implies a quantum speedup for solving Laplacian systems and for approximating a range of cut problems such as min cut and sparsest cut.

**Key words.** quantum computing, quantum algorithms, graph theory

**MSC codes.** 81P68, 68Q12, 05C85

**DOI.** 10.1137/21M1391018

**1. Introduction and summary.** The complexity of many graph problems naturally scales with the number of edges in the graph. As depicted in Figure 1, graph sparsification aims to reduce this number of edges, while preserving certain quantities of interest. When considering, for instance, the approximation of cut problems such as MIN CUT or SPARSEST CUT, the aim is to sparsify the graph while approximately preserving its cut values. This was first shown to be possible in the pioneering work of Karger [62] and later Benczúr and Karger [18]. They introduced the concept of *cut sparsifiers*, which are reweighted subgraphs that  $\epsilon$ -approximate all cuts in the graph. We can then solve cut problems in the hopefully sparser subgraph, yielding an approximate solution to the original problem. Quite surprisingly, they showed that for any undirected graph with  $n$  nodes and  $m$  edges, there always exists a cut sparsifier with as few as  $\tilde{O}(n/\epsilon^2)$  edges, and moreover this sparsifier can be constructed in time  $\tilde{O}(m)$ . This result lies at the basis of  $\tilde{O}(m)$ -time approximation algorithms for, among others, MIN CUT [62], MIN *st*-CUT [65, 84, 89], SPARSEST CUT, and BALANCED SEPARATOR [11, 89]. We refer the interested reader to [86, 90] for surveys on the many applications of cut approximation.

In their breakthrough work on Laplacian solvers, Spielman and Teng [96] strengthened the notion of cut sparsifiers to so-called *spectral sparsifiers*. Rather than

---

\*Received by the editors January 11, 2021; accepted for publication (in revised form) September 6, 2022; published electronically December 20, 2022. An extended abstract of this work appeared in Proceedings of the 61st Annual Symposium on Foundations of Computer Science (FOCS 2020), IEEE Computer Society, Duke, NC, 2020.

<https://doi.org/10.1137/21M1391018>

**Funding:** Most of the work was done while the first author was part of the CWI-Inria International Lab. The second author was partially supported by the Dutch Research Council (NWO) through Gravitation-grant Quantum Software Consortium 024.003.037 and through QuantERA project QuantAlgo 680-91-034.

<sup>†</sup>Université de Paris, CNRS, IRIF, F-75013, Paris, France (smgapers@gmail.com).

<sup>‡</sup>QuSoft, CWI, and University of Amsterdam, and The Netherlands (Ronald.de.Wolf@cwi.nl).

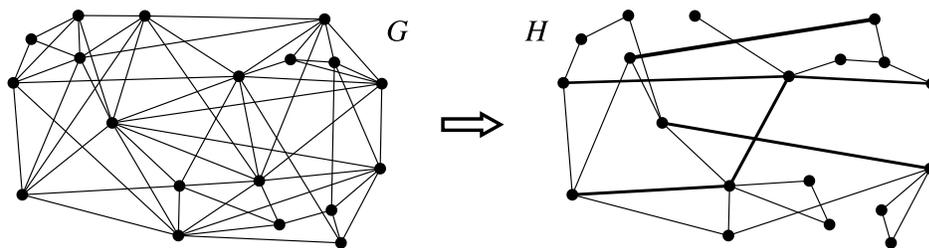


FIG. 1. A sparsifier  $H$  of a graph  $G$  is a sparse, reweighted subgraph that preserves certain quantities such as the cut values (cut sparsifier) or quadratic forms in the graph Laplacian (spectral sparsifier).

preserving the cut structure, these reweighted subgraphs preserve the spectral structure or *quadratic form* of the Laplacian associated to the graph. More specifically,  $H$  is an  $\epsilon$ -spectral sparsifier of  $G$  if

$$(1 - \epsilon)L_G \preceq L_H \preceq (1 + \epsilon)L_G,$$

with  $L_H$  and  $L_G$  the Laplacian matrices associated to  $H$ , respectively,  $G$ . Since the value of any cut can be expressed as a quadratic form in the Laplacian, any spectral sparsifier is necessarily a cut sparsifier. More importantly it implies that Laplacian systems, which are linear systems in the graph Laplacian, can be approximately solved using the Laplacian of the sparsified graph. Similar to the case for cut sparsifiers, Spielman and Teng showed the existence and  $\tilde{O}(m)$ -time construction of  $\epsilon$ -spectral sparsifiers with  $\tilde{O}(n/\epsilon^2)$  edges. This formed a critical cornerstone of their  $\tilde{O}(m)$ -time solver for Laplacian systems (and faster solvers in later papers; see, e.g., [57]), and the string of results and algorithms that followed it, commonly referred to as the “Laplacian paradigm” [98]. Some examples among these are faster algorithms for learning [108, 109], computer vision and image processing [72], spectral clustering [82, 102], computing random walk properties [33], and most recently the breakthrough almost-linear time algorithm for maximum flow and other flow problems [27]. The sparsification results of Spielman and Teng were later refined most notably by Spielman and Srivastava [93] and Batson, Spielman, and Srivastava [14]. In [14], the existence of spectral sparsifiers with only  $O(n/\epsilon^2)$  edges was proved, which later inspired the resolution of the famous Kadison–Singer problem by Marcus, Spielman, and Srivastava [78].

**1.1. Main result and applications.** In this work we give a *quantum* algorithm for spectral sparsification, leading to the theorem below.

**THEOREM 1.1** (quantum algorithm for sparsification). *Fix  $n$ ,  $m$  and let  $\epsilon \geq \sqrt{n/m}$ . There exists a quantum algorithm that, given adjacency-list access to a weighted and undirected  $n$ -node graph  $G$  with  $m$  edges, outputs with high probability the explicit description of an  $\epsilon$ -spectral sparsifier of  $G$  with  $\tilde{O}(n/\epsilon^2)$  edges in time  $\tilde{O}(\sqrt{mn}/\epsilon)$ .*

The algorithm outputs an explicit classical description in the form of the list of  $\tilde{O}(n/\epsilon^2)$  edges of the sparsifier together with their new weights. Note the assumption  $\epsilon \geq \sqrt{n/m}$ . This is because sparsification is only useful when the number of edges of the sparsifier (roughly  $n/\epsilon^2$ ) is at most the number of edges  $m$  of the original graph  $G$ . Note also that  $\tilde{O}(\sqrt{mn}/\epsilon) \in \tilde{O}(m)$  whenever  $\epsilon \geq \sqrt{n/m}$ , and hence our quantum algorithm provides a speedup over classical algorithms, whose  $\tilde{O}(m)$  runtime can be

shown to be optimal.<sup>1</sup> For dense graphs, where  $m \in \Omega(n^2)$ , this improves the time complexity from  $\tilde{O}(n^2)$  classically to  $\tilde{O}(n^{3/2})$  quantumly.

Our algorithm assumes coherent access to the input graph in the form of quantum queries to the adjacency lists. This assumption is very standard and allows us to talk about the query complexity of our algorithm. In order to also be able to talk about “time” complexity, we also assume a QRAM (coherent RAM) memory of  $\tilde{O}(\sqrt{mn}/\epsilon)$  classical bits to which we can do classical writes, and whose bits we can query in superposition. Our algorithm uses just  $O(\log n)$  “actual” qubits. The “time” (complexity) in the above theorem then measures the number of elementary gates, input queries, and QRAM writes and queries. See section 2 for more details about our computational model.

The algorithm builds on a range of quantum and classical results, the most important of which are classical sparsification algorithms by Spielman and Srivastava [93] and Koutis and Xu [73], a spanner algorithm by Thorup and Zwick [99], a quantum algorithm for single-source shortest-path trees by Dürr et al. [41] and an efficient  $k$ -independent hash function by Christiani, Pagh, and Thorup [30].

We prove a matching lower bound, showing that the runtime of our quantum algorithm is optimal up to polylog-factors. In fact, we show that even outputting a weaker *cut* sparsifier requires the same number of queries.

**THEOREM 1.2** (quantum lower bound for sparsification). *Fix  $n, m$  and let  $\epsilon \geq \sqrt{n/m}$ . Any quantum algorithm that, given adjacency-list access to a weighted and undirected  $n$ -node graph  $G$  with  $m$  edges, explicitly constructs with high probability an  $\epsilon$ -cut sparsifier of  $G$  has query complexity  $\tilde{\Omega}(\sqrt{mn}/\epsilon)$ .*

Our algorithm provides a direct speedup for many of the aforementioned applications. In Table 1 we illustrate this speedup for a number of cut approximation problems. All bounds follow by combining our sparsification algorithm with the best classical algorithms, applied to the sparsifier. As far as we know, this is the first quantum speedup for these cut approximation problems.

We can also use a classical Laplacian solver on the sparsifier to find a speedup for Laplacian solving, i.e., solving the linear system  $Lx = b$ , where  $L$  is the Laplacian of the original graph.

**THEOREM 1.3** (quantum Laplacian solver). *Fix  $n, m$  and let  $\epsilon \geq \sqrt{n/m}$ . There exists a quantum algorithm that, given adjacency-list access to a weighted and undi-*

TABLE 1

*Classical and quantum time complexity of cut approximation problems. All quantum bounds follow from combining our quantum sparsification algorithm with the corresponding classical algorithm. Parameter  $\delta$  is an arbitrarily small but positive constant.*

	Classical	Quantum (this work)
.878-MAX CUT	$\tilde{O}(m)$ [10]	$\tilde{O}(\sqrt{mn})$
$\epsilon$ -MIN CUT	$\tilde{O}(m)$ [63]	$\tilde{O}(\sqrt{mn}/\epsilon)$
$\epsilon$ -MIN <i>st</i> -CUT	$\tilde{O}(m + n/\epsilon^5)$ [84]	$\tilde{O}(\sqrt{mn}/\epsilon + n/\epsilon^5)$
$O(\sqrt{\log n})$ -SPARSEST CUT/BAL.SEP.	$\tilde{O}(m + n^{1+\delta})$ [88]	$\tilde{O}(\sqrt{mn} + n^{1+\delta})$

<sup>1</sup>Because there is an  $\Omega(m)$  query lower bound for deciding whether a graph is connected or not (see, for instance, [42, Theorem 4.9,  $k = 1$ ] for a stronger statement), we have the same linear lower bound for finding a cut sparsifier for a given graph, as well as for applications like approximating MIN CUT.

TABLE 2

Classical and quantum time complexity of Laplacian solving and some of its applications. All classical bounds without reference follow from [97]. The quantum bounds essentially follow from combining our quantum sparsification algorithm with the corresponding classical algorithm.

	Classical	Quantum (this work)
$\epsilon$ -Laplacian/SDD Solving	$\tilde{O}(m)$ [97]	$\tilde{O}(\sqrt{mn}/\epsilon)$
$\epsilon$ -Effective Resistance (single)	$\tilde{O}(m)$	$\tilde{O}(\sqrt{mn}/\epsilon)$
$\epsilon$ -Effective Resistances (all)	$\tilde{O}(m + n/\epsilon^4)$ [93]	$\tilde{O}(\sqrt{mn}/\epsilon + n/\epsilon^4)$
$O(1)$ -Cover Time	$\tilde{O}(m)$ [39]	$\tilde{O}(\sqrt{mn})$
$k$ bottom eigenvalues	$\tilde{O}(m + kn/\epsilon^2)$	$\tilde{O}(\sqrt{mn}/\epsilon + kn/\epsilon^2)$
Spectral ( $k$ -means) Clustering	$\tilde{O}(m + npoly(k))$	$\tilde{O}(\sqrt{mn} + npoly(k))$

rected  $n$ -node graph  $G$  with  $m$  edges and Laplacian  $L$ , outputs with high probability an approximate solution  $\tilde{x} \in \mathbb{R}^n$  to the linear system  $Lx = b$  such that  $\|\tilde{x} - x\|_L \leq \epsilon \|x\|_L$  in time  $\tilde{O}(\sqrt{mn}/\epsilon)$ .

In contrast to the well-known HHL algorithm [52], our algorithm outputs an explicit classical description of  $\tilde{x}$  (i.e., a vector of  $n$  real entries), not an  $n$ -dimensional quantum state. Here  $\|v\|_L$  denotes the  $L$ -induced norm  $\|v\|_L = \sqrt{v^\dagger L v} = \|L^{1/2} v\|$ , with  $v^\dagger$  the complex transpose of vector  $v$ . This is the typical norm considered for Laplacian solving. This speeds up the dependency on  $m$  with respect to classical solvers [97], whose runtime is  $\tilde{O}(m \log(1/\epsilon))$ . Similar to the classical case, we show how this also yields a quantum speedup for solving the more general class of *symmetric, weakly diagonally dominant* (SDD) linear systems.

We also find quantum speedups for approximating effective resistances and random walk commute times, creating an approximate “resistance oracle” which allows us to query for the effective resistance of *any* node pair in time  $\tilde{O}(1)$ , for approximating the random walk cover time, and for approximating the bottom eigenvalues of the Laplacian. Finally we discuss how a spectral sparsifier allows us to implement spectral  $k$ -means clustering more efficiently, so that our quantum sparsification algorithm also leads to a speedup for this task. We summarize our speedups in Table 2 and discuss prior work on quantum algorithms for some of these problems in section 1.4. More details are provided in (missing dot) section 7

**1.2. Quantum algorithm.** Our quantum sparsification algorithm starts from the iterative sparsification algorithm by Koutis and Xu [73]. Their algorithm provides a simple combinatorial counterpart to the usual algebraic treatment of spectral sparsification. It crucially relies on the growth of so-called *spanners* of the graph, which are sparse subgraphs that approximately preserve all pairwise distances between nodes. After growing a small number of disjoint spanners in the graph, and keeping these edges, they downsample the remaining edge set by keeping every edge independently with some fixed constant probability, and discarding the rest. This results in a sparsifier with approximately half the number of edges of the original graph. Repeating this procedure a logarithmic number of times results in an  $\epsilon$ -spectral sparsifier with  $\tilde{O}(n/\epsilon^2)$  edges.

The gist of our quantum speedup comes from a faster quantum algorithm for constructing spanners. This algorithm follows essentially by pairing a classical spanner algorithm by Thorup and Zwick [99] with the shortest-paths quantum algorithm by Dürr et al. [41]. More specifically we prove the theorem below, where we call a graph  $H$  a spanner of  $G$  if it is a subgraph with  $O(n \log n)$  edges, and the distance between any pair of nodes in  $H$  is at most  $\log n$  times their original distance in  $G$ .

Our algorithm speeds up the classical  $\tilde{O}(m)$ -time algorithm by Thorup and Zwick, whose runtime is optimal.

**THEOREM 1.4.** *Fix  $n$  and  $m$ . There exists a quantum algorithm that, given adjacency-list access to a weighted and undirected  $n$ -node graph  $G$  with  $m$  edges, outputs with high probability a spanner of  $G$  in time  $\tilde{O}(\sqrt{mn})$ .*

We can now try to plug this faster spanner construction in the Koutis–Xu sparsification algorithm. The problem, however, is that we cannot output the “intermediate” sparsifiers, since after a constant number of iterations these still have  $\Omega(m)$  edges, whereas we aim for a runtime that scales with  $\sqrt{mn}$ . We overcome this issue using two observations, which will allow us to describe the intermediate graphs only implicitly.

First, we show that *if* we were given query access to a uniformly random string of  $\tilde{O}(m)$  bits, then we could implicitly mark the discarded edges and grow spanners in the remaining, unmarked graph without significantly affecting the runtime. Second, we get rid of this long random string by using the fact that any  $(k/2)$ -query quantum algorithm cannot distinguish a uniformly random string from a  $k$ -wise independent string, which only behaves uniformly randomly for subsets of at most  $k$  elements. This is a known result and can be proven, for instance, using the polynomial method [16]. Hence it suffices that we have access to a  $k$ -wise independent random string, allowing us to use the rich literature on  *$k$ -independent hash functions* that aim to simulate access to such random strings. Specifically we require the recent result by Christiani, Pagh, and Thorup [30], which shows that in  $\tilde{O}(k)$  time we can construct a data structure that can simulate queries to a  $k$ -wise independent string, requiring only  $\tilde{O}(1)$  time per query. Prior to their work, all algorithms required preprocessing time  $\tilde{O}(k^{1+\delta})$  for  $\delta > 0$ . Using their construction we can efficiently simulate the random string, which leads to the following claim.

**CLAIM 1.5.** *Consider any quantum algorithm with runtime  $q$  that uses a uniformly random string. Then we can construct a quantum algorithm without random string that has the same output distribution and has a runtime  $\tilde{O}(q)$ .*

Combining these observations remedies the issue of having to store the intermediate graphs and leads to a speedup of the Koutis–Xu algorithm runtime to  $\tilde{O}(\sqrt{mn}/\epsilon^2)$  quantum time.

We then further improve the runtime down to  $\tilde{O}(\sqrt{mn}/\epsilon)$  by combining this quantum sparsification algorithm with the sparsification toolbox of Spielman and Srivastava [93]. In that work, they show that a graph can be sparsified very elegantly by sampling edges with weights roughly proportional to their effective resistances. Complementing this, they propose a near-linear time constructible “resistance oracle,” which allows us to query for effective resistances in logarithmic time. We use our quantum sparsification algorithm to construct an initial, rough sparsifier with a constant error, in time  $\tilde{O}(\sqrt{mn})$ . We then construct an approximate resistance oracle for this sparsifier, which effectively yields an approximate resistance oracle for the original graph. Surprisingly, such rough approximation suffices for constructing an  $\epsilon$ -spectral sparsifier using the Spielman–Srivastava sampling scheme. This finally allows us to sample the  $\tilde{O}(n/\epsilon^2)$  edges of the sparsifier in time  $\tilde{O}(\sqrt{mn}/\epsilon)$ , using Grover’s algorithm. This idea of using a “poor” spectral approximation to compute sampling probabilities to obtain a better spectral approximation is also used in [34, 76].

**1.3. Matching lower bound.** We prove that the  $\tilde{O}(\sqrt{mn}/\epsilon)$ -runtime of our quantum algorithm is optimal, up to polylog-factors, even when we wish to construct

a weaker *cut* sparsifier. The intuition behind this is that an  $\epsilon$ -cut sparsifier of a general graph must contain  $\Omega(n/\epsilon^2)$  edges (and this is tight [14]). If we can appropriately “hide” these edges among the  $m$  edges of a graph, then a quantum search algorithm requires  $\Theta(\sqrt{mn}/\epsilon^2) = \Theta(\sqrt{mn}/\epsilon)$  queries to retrieve them.

Turning this intuition into a concrete lower bound, however, turns out to be rather complicated. We start with a random graph construction by Andoni et al. [6]. This construction describes graphs on  $n$  nodes and  $\tilde{O}(n/\epsilon^2)$  edges, so that any  $\epsilon$ -cut sparsifier must contain a constant fraction of the edges. As such, the constructed graphs are in fact already sparsifiers. We then carefully “hide” these sparsifiers in a larger, denser graph, in such a way that a sparsifier of this graph must retrieve all of the original, hidden sparsifiers. To prove a quantum lower bound for this search problem, we describe it as the composition of the problem of finding a constant fraction of the nonzero bits in a Boolean matrix with the OR-function. Finally, we combine lower bounds for the individual problems using a composition theorem for adversary bounds, applicable to the composition of a relational problem with a function. This composition theorem was very recently proven by Belovs and Lee [17], prompted by our question to them.

**1.4. Prior work.** We are not aware of any prior work on quantum speedups for graph sparsification. In a very different line of work, though, spectral sparsification has been studied in a quantum context with the goal of sparsifying Hamiltonian matrices, which are used to describe many-body systems. Aharonov and Zhou [1] asked whether the *interaction graph* of a many-body system can be sparsified while preserving its spectrum, showing that this is not possible in general. More recently, Herbert and Subramanian [53] considered the weaker notion of sparsifying the Hamiltonian matrix and suggested that sparsification could indeed help in Hamiltonian simulation. They do not consider quantum algorithms for effectively constructing such a sparsifier.

Research on quantum algorithms for cut approximation is also limited. There is recent work by Hamoudi et al. [51] on quantum approximate minimization of submodular functions, which can be used for cut approximation. However, their work was more recently superseded by better classical algorithms [12]. Other recent work by Brandão, Kueng, and Stilck França [20] used quantum SDP solvers to approximate quadratic binary optimization problems, of which MAX CUT is the most notable instance. They do not succeed in finding a speedup for MAX CUT, though, mainly because their algorithm does not benefit from the special structure of this instance.

Concerning our speedup for Laplacian solving, we mention a range of papers on quantum speedups for general linear system solving. Most famous is the work by Harrow, Hassidim, and Lloyd [52], which was later refined in work by Ambainis [5] and Childs, Kothari, and Somma [29]. They describe a quantum algorithm for solving general linear systems  $Ax = b$  in time  $\tilde{O}(d_M \kappa \log(1/\epsilon))$ , with  $d_M$  the row sparsity and  $\kappa$  the condition number of  $A$ . These algorithms are particularly relevant for sparse and well-conditioned systems (in general, however,  $\kappa$  can be as large as  $O(n^3 w_{\max}/w_{\min})$  for graph Laplacians [97, Lemma 6.1]). Crucially, they only output a quantum state that encodes the solution, rather than an explicit description as we do.

Quantum speedups for the problems of estimating effective resistances and spectral gaps have also been studied in other work. Very recently and independently from our work, Piddock [87] constructed a quantum walk algorithm that  $\epsilon$ -approximates the effective resistance  $R_{s,t}$  using  $\tilde{O}(\sqrt{mR_{s,t}}/\epsilon^2) \in \tilde{O}(\sqrt{mn}/\epsilon^2)$  quantum walk steps. He also argued how this could possibly be further improved to  $\tilde{O}(\sqrt{mR_{s,t}}/\epsilon) \in \tilde{O}(\sqrt{mn}/\epsilon)$ . While the quantum walk model is different from our model, this ten-

tative bound would agree with our runtime. In addition, however, we can effectively approximate *all* effective resistances simultaneously in the graph in the same complexity. The problem was also studied in slightly different settings in [25, 56, 104]. A quantum walk algorithm for estimating the second bottom eigenvalue  $\lambda_2$  of the Laplacian in the adjacency-matrix model was studied by Jarret et al. [58]. They give a multiplicative  $\epsilon$ -approximation of  $\lambda_2$  in time  $\tilde{O}(n/(\sqrt{\lambda_2}\epsilon))$ , which is  $\tilde{O}(n^2/\epsilon)$  in the worst case. We improve the worst-case complexity to  $\tilde{O}(\sqrt{mn}/\epsilon) \in \tilde{O}(n^{3/2}/\epsilon)$ .

We also mention some past and concurrent work on quantum speedups for clustering. One paper by Daskin [38] describes a quantum algorithm for spectral clustering, but no direct speedup is found with respect to classical algorithms. Concurrent to our work is a paper by Kerenidis and Landman [67] which describes a quantum algorithm for outputting the centroids of a  $k$ -means spectral clustering. In contrast to our work, they start from quantum access to a data set, which they then use to query an associated Laplacian. They find a quantum speedup under certain assumptions (e.g., that the input data is appropriately clustered), and as such is incomparable to our quantum algorithm. Less directly related, there exists a number of papers [3, 68, 77, 105] on quantum speedups for  $k$ -means clustering and the construction of a neighborhood graph. These tasks are complementary to our work on finding a spectral embedding, given a similarity graph of the data. It does seem interesting to try and use these algorithms to further speed up our spectral clustering algorithm.

Finally, we mention some classical work on sublinear algorithms for Laplacian solving and spectral sparsification. First, the work by Andoni, Krauthgamer, and Pogrow [7] describes a sublinear algorithm for Laplacian solving, with the aim of approximating a single coordinate of the output. Their algorithm is inspired by quantum algorithms for linear system solving, and similarly only finds a speedup for sparse and well-conditioned systems. The second work is by Lee [74], who proposes a classical algorithm for spectral sparsification of unweighted graphs which is sublinear in  $m$ . He succeeds in bypassing the  $\Omega(m)$  lower bound on classical sparsification by only achieving a weaker, additive error in the approximation. As such this work is incomparable to ours.

**1.5. Follow-up work.** After the initial appearance of our results, a number of works have appeared that build on our results.

Apers and Lee [9] used our quantum algorithm for graph sparsification to design a quantum algorithm for finding an *exact* minimum cut in a graph. They obtain quantum speedups when the ratio of the maximum weight over the minimum weight is bounded, and show that the quantum algorithm is optimal under such a condition. Using similar techniques and under similar constraints, Apers, Auza, and Lee [8] proved a quantum speedup for finding the exact minimum  $s$ - $t$  cut of a graph.

In a rather different setting, van Apeldoorn et al. [101] recently proved a quantum speedup for approximate matrix scaling. Their algorithm builds on quantum approximate counting and returns an  $\epsilon$ -approximate matrix scaling of an  $n \times n$  matrix with  $m$  nonzero entries in time  $\tilde{O}(\sqrt{mn}/\epsilon^4)$ . For entrywise positive matrices, the upper bound was later improved to  $\tilde{O}(n^{3/2}/\epsilon^3)$  by Gribling and Nieuwboer [49]. Their algorithm combines a second-order method for matrix scaling [35] with our quantum algorithm for approximately solving Laplacian systems.

In another work, Cade, Labib, and Niesen [23] described a quantum algorithm for *motif clustering*, which is a graph clustering method that is based on higher-order patterns or motifs. For dense instances, their algorithm builds on our quantum sparsification algorithm.

Finally, a recent work by Chen and de Wolf [28] describes quantum algorithms and quantum lower bounds for linear regression with norm constraints. The lower bounds rely on a strategy similar to ours in that they use the composition property of the adversary bound, which was recently proved by Belovs and Lee [17].

**1.6. Open questions.** Our work raises a number of interesting questions and future directions, some of which we summarize below.

- We prove a matching  $\tilde{\Omega}(\sqrt{mn}/\epsilon)$  lower bound on the quantum query complexity of spectral sparsification. Can we extend this to a tight lower bound for any of the resulting applications, e.g., for  $\epsilon$ -approximating the min cut or effective resistance? Since these problems can be reduced to constructing an  $\epsilon$ -spectral sparsifier, this would yield a stronger lower bound.
- The runtime of our quantum algorithm is tight, up to polylogarithmic factors, for sparsification of weighted graphs. Can we potentially improve the runtime for *unweighted* graphs?
- Graph sparsification is a key technique in efficient algorithms for calculating minimum cuts [62, 63], and the use of fast Laplacian solvers played a key in the recent resolution of the long-standing question of computing maximum flows in graphs in almost-linear time [27, 31, 65, 84, 89]. In follow-up work to this one, Apers and Lee [9] used our quantum algorithm for graph sparsification to obtain a quantum speedup for finding a minimum cut in a graph. We leave it as an open question whether a quantum speedup for finding a maximum flow can be obtained as well.
- Spectral sparsification of graphs and Laplacians has been extended in different directions such as sparsification of hypergraphs [13, 92], sparsification of sums of positive semidefinite matrices [91, 92], and sparsification in a streaming setting [60, 66]. It is also closely related to concepts such as spectral sketching [6] and linear data regression using leverage scores [40]. It seems likely that we can also find quantum speedups for these related problems. Similarly we might hope to solve more “quantum” tasks, such as sparsifying density operators or POVMs.

**1.7. Road map.** To round up the introduction, we give a road map of the main parts of our paper. Section 2 introduces some necessary preliminaries on graphs, the computational model, and the quantum algorithmic routines that we use. It also describes the classical sparsification algorithm that our quantum algorithm is based on. In section 3 we describe a first “rough” quantum sparsification algorithm, and we show how to simulate its access to a random string using  $k$ -wise independent hash functions. We improve the error dependency of the algorithm in section 4. In section 5 we describe a quantum algorithm for constructing graph spanners, which our sparsification algorithms are based on. In section 6 we prove a lower bound that matches the performance of our quantum algorithm, and finally in section 7 we elaborate on the applications of our quantum algorithm for sparsification.

**2. Preliminaries.** Throughout the paper we say that something holds “with high probability” if it holds with probability at least  $1 - O(1/n)$ .

**2.1. Computational model and quantum algorithms.** We assume as our computational model a quantum-accessible classical control system that

1. can run quantum subroutines on at most  $O(\log N)$  qubits, where  $N$  is the size of the problem instance;

2. can make quantum queries to the input; and
3. has access to a quantum-read/classical-write RAM (QRAM)<sup>2</sup> of  $\tilde{O}(\sqrt{mn}/\epsilon)$  classical bits, where a single QRAM operation corresponds to either classically writing a bit to the QRAM or making a quantum query (a read operation) to bits stored in QRAM, possibly in superposition.

In this model, an algorithm has *time complexity*  $T$  if it uses at most  $T$  elementary classical and quantum gates, quantum queries to the input, and QRAM operations. The *query complexity* of an algorithm only measures the number of queries to the input. If we only care about query complexity, the assumption of having QRAM may be dropped at the expense of a polynomial increase in the number of gates.

An important quantum subroutine in our work is Grover’s algorithm [50] for searching sets of marked elements, which is summarized in the claim below.

CLAIM 2.1 (repeated Grover search). *Let  $f : [N] \rightarrow \{0, 1\}$  be a function that marks a set of elements  $S = \{i \in [N] \mid f(i) = 1\}$ . Then there is a quantum algorithm that finds  $S$  with probability at least  $2/3$  in  $\tilde{O}(\sqrt{N}|S|)$  elementary operations and queries to  $f$ , and uses  $O(\log N)$  qubits and a QRAM of  $\tilde{O}(|S|)$  bits.*

We also use a quantum algorithm for finding shortest-path trees by Dürr et al. [41], but the quantum routines in this algorithm can be reduced to Grover search.

**2.2. Graphs, queries, and spanners.** We consider weighted and undirected graphs  $G = (V, E, w)$  with  $|V| = n$  nodes and  $|E| = m$  edges, and edge weights  $w : E \rightarrow \mathbb{R}_{\geq 0}$ . We are given *adjacency-list access* to  $G$ , as is considered in, e.g., [41, 47]. This allows us to query for the degree of a node, its  $k$ th neighbor (according to some unknown but fixed ordering), or the weight of an edge. This model is more restrictive than both the “general graph model,” which in addition allows for adjacency matrix queries [46], and the “sparse-access model,” in which the neighbors are ordered lexicographically, as is commonly assumed in quantum algorithms for linear system solving and Hamiltonian simulation [29, 52].

We define the *distance*  $\delta_G(u, v)$  between nodes  $u$  and  $v$  with respect to  $G$  as

$$\delta_G(u, v) = \min_{u-v \text{ path } P} \sum_{e \in P} \frac{1}{w_e}.$$

This definition is in accordance with the interpretation of  $G$  as an electrical network, in which an edge  $e$  corresponds to a link of conductance  $w_e$  (and hence resistance or “cost”  $1/w_e$ ), as is common in the literature on spectral sparsification. A span-

<sup>2</sup> Another name for this type of memory is “coherent RAM.” We feel a QRAM containing a classical  $k$ -bit string  $z$  and allowing efficient queries of the form  $|i, b\rangle \mapsto |i, b \oplus z_i\rangle$  (where  $i \in [k]$  and  $b \in \{0, 1\}$ ) is a reasonable generalization of classical RAM; if one believes in classical RAM and in quantum superposition, then QRAM is quite natural. Like classical RAM, the physical hardware of such a QRAM necessarily requires the size to be at least proportional to  $k$  because it contains  $k$  bits of information, but answering a query would have a cost proportional to  $\log k$ , even when querying multiple stored bits in superposition. This could be realized, for instance, by laying out the  $k$  bits as the leaves of a binary tree of depth  $\log k$ ; the  $\log k$  bits of the binary representation of an address  $i \in [k]$  would chart a path from the root to the addressed bit  $z_i$ , allowing for efficient lookup of the addressed bit. Note that running this on a superposition of different addresses  $i$  would involve going down different paths in superposition, but would still only use a superposition of  $O(\log k)$  qubits. Assuming such a QRAM is very common in quantum algorithms for graph problems. It should be noted, though, that the notion of QRAM is a bit controversial, because (1) the term is sometimes used in the literature for a different and stronger kind of memory (allowing for efficient conversion of a classically stored unit vector of  $k$  numbers into the corresponding state of  $\log k$  qubits), and (2) implementing it on noisy hardware might require  $O(k)$  work to error-correct a quantum query, rather than  $O(\log k)$  work.

ner of  $G$  is a sparse subgraph  $H$  that approximately preserves all pairwise distances. Specifically, we will call  $H$  a  $t$ -spanner of  $G$  if for any pair  $u, v \in V$  it holds that

$$\delta_G(u, v) \leq \delta_H(u, v) \leq t\delta_G(u, v).$$

Note that the first inequality is trivially satisfied since  $H$  is a subgraph. It is well known that every weighted graph has a  $(2k - 1)$ -spanner with  $O(n^{1+1/k})$  edges [4]. Throughout the paper we will use the shorthand *spanner* to denote a  $t$ -spanner with  $t = 2 \log n$  and  $O(n)$  edges. An  $r$ -packing of spanners of  $G$  is an ordered set  $H = (H_1, H_2, \dots, H_r)$  of  $r$  edge-disjoint spanners such that  $H_j$  is a spanner for  $G - \cup_{i < j} H_i$ , which is the remaining graph after removal of the edges of all previous spanners. Note that such an  $r$ -packing always exists for every  $r$ , though once  $G - \cup_{i < j} H_i$  has no edges left anymore, the subsequent spanners  $H_j, H_{j+1}, \dots, H_r$  in the packing will all be empty.

The *Laplacian*  $L$  of a weighted graph  $G$  is given by  $L = D - A$ , with  $A$  the weighted adjacency matrix  $(A_{ij}) = w_{ij}$  and  $D$  the diagonal weighted degree matrix  $(D_{ii}) = \sum_j w_{ij}$ . Alternatively, we can rewrite the Laplacian as

$$L = \sum_{e \in E} w_e \chi_e \chi_e^T,$$

where we let  $\chi_e = \chi_u - \chi_v$  denote a vector associated to the edge  $e = (u, v)$ , with  $\chi_u, \chi_v$  indicator vectors of the nodes  $u, v$  (we fix an arbitrary orientation of the edges). If  $G$  is connected, then  $L_G$  has a trivial kernel consisting only of the all-ones vector. Moreover,  $L_G$  is a real, symmetric, diagonally dominant matrix with nonnegative diagonal entries and is hence positive semidefinite.

**2.3. Spectral sparsification using spanner packings.** A *cut sparsifier*  $H$  of a graph  $G$  is a sparse, reweighted subgraph that preserves the value of all cuts. Specifically,  $H$  is called an  $\epsilon$ -cut sparsifier if for any  $S \subseteq V$  it holds that

$$(2.1) \quad (1 - \epsilon)\text{val}_G(S) \leq \text{val}_H(S) \leq (1 + \epsilon)\text{val}_G(S),$$

where  $\text{val}_G(S) = \sum_{i \in S, j \notin S} w_{(i,j)}$  denotes the total weight of the edges leaving  $S$ .

A *spectral sparsifier*  $H$  of a graph  $G$  is a sparse, reweighted subgraph that preserves the quadratic form  $x^T L_G x$  associated to the Laplacian  $L_G$  of  $G$  for any vector  $x \in \mathbb{C}^n$ . Specifically,  $H$  is called an  $\epsilon$ -spectral sparsifier if for any  $x \in \mathbb{C}^n$  it holds that

$$(2.2) \quad (1 - \epsilon)x^T L_G x \leq x^T L_H x \leq (1 + \epsilon)x^T L_G x.$$

Alternatively, we can rewrite this as  $(1 - \epsilon)L_G \preceq L_H \preceq (1 + \epsilon)L_G$ , where  $A \preceq B$  denotes that  $B - A$  is positive semidefinite. This condition implies, for instance, that all eigenvalues of  $H$   $\epsilon$ -approximate the eigenvalues of  $G$  [15], and all cuts in  $H$   $\epsilon$ -approximate those in  $G$ . To see the latter, consider a subset  $S \subseteq V$  and let  $\chi_S$  denote the indicator on  $S$ ; then

$$\chi_S^T L_G \chi_S = \sum_{(u,v) \in E} w_e (\chi_S(u) - \chi_S(v))^2 = \text{val}_G(S).$$

This shows that the cut value can be described by a quadratic form in the Laplacian, and hence (2.2) implies that  $(1 - \epsilon)\text{val}_G(S) \leq \text{val}_H(S) \leq (1 + \epsilon)\text{val}_G(S)$  for all  $S \subseteq V$ . Any  $\epsilon$ -spectral sparsifier is therefore also an  $\epsilon$ -cut sparsifier.

Spectral sparsifiers can be constructed by using spanners to identify the “important” edges in the graph. This was first noticed by Kapralov and Panigrahy [61],

and further refined by Koutis and Xu [73]. We will build on the latter work, which describes a very elegant approach for constructing spectral sparsifiers from spanner packings. Their algorithm iteratively invokes the routine described below, which creates a spectral sparsifier with approximately half the number of edges of the original graph.

---

**Algorithm 2.1**  $H = \text{Half Sparsify}(G, \epsilon)$ .

---

- 1: construct an  $O(\log^2(n)/\epsilon^2)$ -packing of spanners of  $G$
  - 2: let  $P$  be their union and set  $H = P$
  - 3: **for** each edge  $e \notin P$  **do**
  - 4:     with probability  $1/4$ , add  $e$  to  $H$  with weight  $4w_e$
  - 5: **end for**
  - 6: return  $H$
- 

**THEOREM 2.2** ([73, Theorem 3.2]). *The graph  $H = \text{Half Sparsify}(G, \epsilon)$  is, with probability at least  $1 - 1/n^2$ , an  $\epsilon$ -spectral sparsifier of  $G$  with at most  $m/2 + \tilde{O}(n/\epsilon^2)$  edges.*

Now consider a fixed  $\epsilon > 0$ . If we iterate  $T \in O(\log(m/n))$  times the routine **Half Sparsify** $(G, \epsilon')$ , with  $\epsilon' \in O(\epsilon/T)$ , then we retrieve with high probability an  $\epsilon$ -spectral sparsifier with  $\tilde{O}(n/\epsilon^2)$  edges. By [6] this is optimal up to log-factors. Classically the complexity is dominated by the construction of  $\tilde{O}(1/\epsilon^2)$  spanners, each of which requires time  $\tilde{O}(m)$  [99], giving a total time complexity  $\tilde{O}(m/\epsilon^2)$ .

**3. Quantum sparsification algorithm.** In this section we describe our quantum algorithm for constructing spectral sparsifiers. The algorithm is based on the scheme by Koutis and Xu. We use as a black box a quantum algorithm for constructing a spanner in time  $\tilde{O}(\sqrt{mn})$ , whose description we postpone to section 5.

---

**Algorithm 3.1**  $H = \text{Quantum Sparsify}(G, \epsilon)$ .

---

- 1: let  $\{w'_e = w_e\}$  and  $\ell = \lceil \log(m/n) \rceil$
  - 2: **for**  $i = 1, 2, \dots, \ell$  **do**
  - 3:     create an  $O(\log^2(n)/\epsilon^2)$ -packing of spanners of  $G' = (V, E, w')$ , let  $P_i$  denote its union
  - 4:     **for** each edge  $e \notin P_i$  **do**
  - 5:         **if**  $r_i(e) = 1$  **then** set  $w'_e = 4w'_e$  **else** set  $w'_e = 0$
  - 6:     **end for**
  - 7: **end for**
  - 8: use repeated Grover search to find  $H = \{e \in E \mid w'_e > 0\}$
- 

As we already mentioned in the introduction, we cannot simply plug this quantum spanner algorithm into the Koutis and Xu algorithm. Indeed, after a single iteration of their algorithm this would require outputting a graph with up to  $m/2$  edges, which is much too costly since we aim at a runtime that scales as  $\sqrt{mn}$ . We resolve this issue in two stages. First, we assume that we have access to a random string of length  $\tilde{O}(m)$ . We use this string to mark edges that have been discarded at some iteration by 0-bits, which we later use to implicitly set their weight equal to zero. By its construction, the spanner algorithm can then construct a spanner in the remaining graph. At the end we use Grover search to explicitly retrieve the remaining  $\tilde{O}(n/\epsilon^2)$  edges, whose union forms the spectral sparsifier. We then get rid of the random string. To this end

we use efficient  $k$ -independent hash functions that allow us to simulate queries to a  $k$ -wise independent random string. This suffices since by standard results a  $k$ -query quantum algorithm cannot distinguish  $2k$ -wise independent strings from a uniformly random one.

**3.1. Using a random string.** We first assume access to a family of independent, random strings  $r_i \in \{0, 1\}^m$ , with indices  $i \in [\log(m/n)]$ , such that all bits are independent and equal to 1 with probability  $1/4$ . For different indices  $i$ , the strings  $r_i$  will function as consecutive “sieves” of the edge set.

Algorithm 3.1 describes the sparsification algorithm using such random strings. A critical remark is that steps 4 and 5 of the algorithm are only performed implicitly, as mentioned before. Rather than keeping an explicit list of updated edge weights, we maintain an implicit “weight oracle.” Only when an edge weight is queried does this weight oracle calculate its weight by consulting the necessary random strings. We show how to do this efficiently in the proof of Theorem 3.1.

**THEOREM 3.1.** *Given access to independent, uniformly random strings  $r_i \in \{0, 1\}^m$  for  $i \in [\log(m/n)]$ , algorithm **Quantum Sparsify** $(G, \epsilon)$  returns with probability  $1 - O(\log(n)/n^2)$  an  $\epsilon$ -spectral sparsifier of  $G$  with  $\tilde{O}(n/\epsilon^2)$  edges. There is a quantum algorithm that implements it in time  $\tilde{O}(\sqrt{mn}/\epsilon^2)$  and using a QRAM of  $\tilde{O}(\sqrt{mn}/\epsilon^2)$  bits.*

*Proof.* Correctness easily follows from Theorem 2.2: in every iteration we “half-sparsify” the remaining graph (induced by all edges of weight  $w_e > 0$ ). The probability that all  $\log(m/n)$  iterations succeed is  $1 - O(\log(n)/n^2)$ . Below we discuss how steps 4 and 5 can be implemented efficiently, so that the runtime of the for-loop is dominated by the construction of  $\tilde{O}(1/\epsilon^2)$  spanners. By Theorem 5.3 this takes time  $\tilde{O}(\sqrt{mn}/\epsilon^2)$ . By standard results [81], the repeated Grover search routine in the final step takes time  $\tilde{O}(\sqrt{mn}/\epsilon)$ , which is the time needed to find  $n/\epsilon^2$  edges among  $m$  edges.

What remains to prove is that there exists an efficient oracle that keeps track of the weight updates in steps 4 and 5. Consider the  $i$ th iteration. Given an edge  $e$ , let  $k$  denote the number of spanners before this iteration in which  $e$  occurs so far. If  $k = 0$ , return  $w'_e = 4^i w_e$  if  $(r_i r_{i-1} \dots r_1)(e) = 1$ , and  $w'_e = 0$  if  $(r_i r_{i-1} \dots r_1)(e) = 0$ . If  $k > 0$ , let  $j < i$  denote the last spanner packing in which it occurs. Now return  $w'_e = 4^{i-k} w_e$  if  $(r_i r_{i-1} \dots r_{j+1})(e) = 1$ , and  $w'_e = 0$  otherwise. This takes  $\tilde{O}(1)$  searches through the set of spanners (which we may assume is sorted) and at most  $O(i) \in \tilde{O}(1)$  evaluations of the random oracle.  $\square$

The space complexity of the algorithm requires  $O(\log n)$  qubits and a QRAM of  $\tilde{O}(n/\epsilon^2)$  bits. The number of qubits follows from the space complexity of the quantum spanner algorithm and the Grover search routine. The classical space complexity is dominated by the output size.

**3.2. Using  $k$ -independent hash functions.** In order to get rid of the random strings  $\{r_i\}$ , we build on the following fact, which is an easy consequence of the polynomial method [16]. It seems that this was first used in the proof of [22, Theorem 19] and is stated explicitly in, for instance, [107, Theorem 3.1].

**FACT 1.** *The output distribution of a quantum algorithm making  $q$  queries to a uniformly random string is identical to the same algorithm making  $q$  queries to a  $2q$ -wise independent string.*

As a consequence, we can replace the uniformly random strings of length  $m$  by a  $k$ -wise independent string with  $k \in \tilde{O}(\sqrt{mn}/\epsilon^2)$ . Surely we also cannot explicitly construct a  $k$ -wise independent string of length  $\tilde{O}(m)$  in time  $\tilde{O}(\sqrt{mn}/\epsilon^2)$ , but we can use hash functions to simulate queries to such a string. A family of hash functions  $F = \{h : [u] \rightarrow [r]\}$  for  $u, r \in \mathbb{N}$  is called *k-independent* if, for any subset  $S \subseteq [u]$  of size  $|S| \leq k$  and a uniformly random function  $h$  in the family, the image of  $h$  on  $S$  behaves uniformly randomly in  $[r]^{|S|}$ . This implies that the image of a random member of  $F$ , which we will refer to as a *k-independent hash function*, describes a  $k$ -wise independent string over  $[r]^u$ . Elegant constructions of such functions have long been known, the most famous example being random degree- $k$  polynomials, as proposed by Carter and Wegman [24]. Crucial to our cause, however, is that we can evaluate the hash function in  $\tilde{O}(1)$  time, potentially allowing  $\tilde{O}(k)$  preprocessing time. Fortunately, such a result was established very recently by Christiani, Pagh, and Thorup [30], who proved the theorem below. We note that this is a purely classical construction.

**THEOREM 3.2** ([30]). *It is possible to construct in time  $\tilde{O}(k)$  a data structure of size  $\tilde{O}(k)$  that allows one to simulate queries to a  $k$ -independent hash function in  $\tilde{O}(1)$  time per query.*

With  $k = 2q$  and  $[r] = \{0, 1\}$ , we can combine this with Fact 1 to give the corollary below.

**COROLLARY 3.3.** *Consider any quantum algorithm with runtime  $q$  that makes queries to a uniformly random string. We can simulate this algorithm with a quantum algorithm with runtime  $\tilde{O}(q)$  without random string, using an additional QRAM of  $\tilde{O}(q)$  bits.*

This shows that we can efficiently simulate the random string in Algorithm 3.1, leading to at most a polylogarithmic overhead in the runtime. The classical space complexity of the algorithm does increase from  $\tilde{O}(n/\epsilon^2)$  to  $\tilde{O}(\sqrt{mn}/\epsilon^2)$ . The following theorem is immediate by combining Theorem 3.1 with Corollary 3.3.

**THEOREM 3.4.** *There exists a quantum algorithm that, given adjacency-list access to a weighted and undirected graph  $G$ , constructs with high probability an  $\epsilon$ -spectral sparsifier of  $G$  with  $\tilde{O}(n/\epsilon^2)$  edges in time  $\tilde{O}(\sqrt{mn}/\epsilon^2)$ . The algorithm uses  $O(\log n)$  qubits and a QRAM of  $\tilde{O}(\sqrt{mn}/\epsilon^2)$  bits.*

**4. Refined quantum sparsification algorithm.** In the last section we proposed a quantum algorithm for constructing an  $\epsilon$ -spectral sparsifier in time  $\tilde{O}(\sqrt{mn}/\epsilon^2)$ . Here we show how to improve the runtime of this algorithm to  $\tilde{O}(\sqrt{mn}/\epsilon)$ , which we will later show is optimal up to polylog-factors. The improvement essentially follows from combining our previous algorithm with the seminal results on spectral sparsification by Spielman and Srivastava [93]. In that work, they first showed that sampling edges with probabilities approximately proportional to their effective resistances results in a spectral sparsifier (the Koutis–Xu algorithm is derived from their result). Then they showed how Laplacian solvers could be used to efficiently estimate these effective resistances. We will use our quantum sparsification algorithm to first construct a “rough”  $\epsilon$ -sparsifier, for some constant  $\epsilon$ , which we only use to approximate the effective resistances in the original graph. Surprisingly such approximation suffices to implement the Spielman–Srivastava sampling scheme on the original graph. We then use a quantum sampling routine to efficiently implement this sampling scheme, finally leading to an  $\epsilon$ -spectral sparsifier for arbitrary  $\epsilon > 0$  in time  $\tilde{O}(\sqrt{mn}/\epsilon)$ . This idea

of using a “poor” spectral sparsifier for computing sampling probabilities to obtain a better spectral sparsifier is also present in, for instance, [34, 76].

**4.1. Spielman–Srivastava toolbox and quantum sampling.** Here we formally introduce the main tools that we use. These are an efficiently constructible “resistance oracle” and a sparsification algorithm based on this oracle from [93], and a quantum sampling routine for implementing this sparsification algorithm.

**4.1.1. Approximate resistance oracle.** The *effective resistance* in a graph  $G$  between a pair of nodes  $s$  and  $t$  is defined as the effective resistance between  $s$  and  $t$  after replacing every edge  $e$  by a resistor of value  $1/w_e$ . It can be expressed algebraically as  $R_{s,t} = (\chi_s - \chi_t)^T L_G^+(\chi_s - \chi_t)$ , where  $L_G^+$  is the pseudoinverse of  $L_G$  (i.e., the inverse of its image). A Laplacian solver hence allows one to efficiently compute  $R_{s,t}$ . Spielman and Srivastava proved that in some sense one can efficiently compute *all* effective resistances in roughly the same time. More specifically, they showed that it is possible to construct in near-linear time a data structure of size  $\tilde{O}(n/\epsilon^2)$  that allows one to efficiently approximate  $R_{s,t}$  for any  $s, t$ .

**THEOREM 4.1 ([93]).** *Consider a weighted and undirected graph  $G$ . There is an  $\tilde{O}(m/\epsilon^2)$ -time algorithm which computes a  $(24 \log(n)/\epsilon^2) \times n$  matrix  $Z$  such that with probability at least  $1 - 1/n$ , for every pair  $s, t \in V$ , it holds that*

$$(1 - \epsilon)R_{s,t} \leq \|Z(\chi_s - \chi_t)\|^2 \leq (1 + \epsilon)R_{s,t}.$$

Hence the matrix  $Z$  represents a data structure which allows us to  $\epsilon$ -approximate  $R_{s,t}$  for any pair  $s, t$  by calculating the 2-norm distance between two columns, each of dimension  $\tilde{O}(1/\epsilon^2)$ .

**4.1.2. Spectral sparsification with edge scores.** In the same paper, Spielman and Srivastava proved the following theorem, which shows that a spectral sparsifier can be constructed by independently keeping edges with weights roughly proportional to their effective resistances.

**THEOREM 4.2 ([93]).** *Let  $2R_e \geq \tilde{R}_e \geq R_e/2$  for each edge  $e \in E$ , and let  $p_e = \min(1, Cw_e\tilde{R}_e \log(n)/\epsilon^2)$  for some universal constant  $C$ . Then keeping every edge  $e$  independently with probability  $p_e$ , and rescaling its weight with  $1/p_e$ , yields with probability at least  $1 - 1/n$  an  $\epsilon$ -spectral sparsifier of  $G$  with  $O(n \log(n)/\epsilon^2)$  edges.*

Note that  $\sum_e p_e \gg 1$  is the expected number of edges of the sparsifier. Since  $\sum_e w_e R_e = n - 1$  [19, Theorem 25], this yields the claimed number of edges. Also note that by slightly tweaking our estimate  $\tilde{R}_e$ , we can ensure that the sampling probabilities  $p_e$  have an  $O(\log n)$ -bit description. We will use this later on.

We note that, in fact, Spielman and Srivastava describe a slightly different scheme. They propose to draw  $\tilde{O}(n/\epsilon^2)$  independent and identically distributed edge samples from the edge set, with probability proportional to their effective resistance. It is well known that both schemes give the same performance bound; see, e.g., [43, Remark 1].

**4.1.3. Quantum sampling.** Assuming access to an approximate resistance oracle that gives approximations  $\tilde{R}_e$  to  $R_e$ , we wish to implement the Spielman–Srivastava sparsification scheme. While classically this requires time  $\tilde{O}(m + \sum_e p_e)$ , we can use quantum algorithms to do so more efficiently.

**CLAIM 4.3.** *Assume we have query access to a list of probabilities  $\{p_e\}_{e \in E}$ , each of which is described with  $\tilde{O}(1)$  bits of precision. Then there is a quantum algorithm that*

---

**Algorithm 4.1**  $H = \text{Quantum Sparsify}(G, \epsilon)$ .

---

- 1: use quantum sparsification (Theorem 3.4) to construct a  $(1/100)$ -spectral sparsifier  $H$  of  $G$
  - 2: create a  $(1/100)$ -approximate resistance oracle of  $H$  using Theorem 4.2, yielding estimates  $\{\tilde{R}_e\}$
  - 3: use quantum sampling (Claim 4.3) to sample a subset of the edges, keeping every edge with probability  $p_e = \min(1, Cw_e\tilde{R}_e \log(n)/\epsilon^2)$
- 

samples a subset  $S \subseteq E$ , such that  $S$  contains every  $e$  independently with probability  $p_e$ , in expected time  $\tilde{O}(\sqrt{m(\sum_e p_e)})$  and using a QRAM of  $\tilde{O}(\sqrt{m(\sum_e p_e)})$  bits.

*Proof.* Let  $\ell \in \tilde{O}(1)$  denote the number of bits to describe each of the probabilities  $p_e$ . We can assume access to a uniformly random  $\ell m \in \tilde{O}(m)$ -bit string  $r$ , since by Corollary 3.3 this implies that there also exists a quantum algorithm with no random string. For all  $e$ , we can derive from this random string a random  $\ell$ -bit number  $q_e \in [0, 1]$  such that  $q_e \leq p_e$  with probability exactly  $p_e$ . We can implement the mapping  $|e\rangle |0\rangle \mapsto |e\rangle |q_e\rangle$  in time  $\tilde{O}(1)$ .

We combine this with a query to the list of probabilities to implement the mapping  $|e\rangle |0\rangle |0\rangle \mapsto |e\rangle |p_e\rangle |q_e\rangle$ . Now we use repeated Grover search to find all edges  $e$  such that  $q_e \leq p_e$ —for every edge  $e$ , this set contains  $e$  with probability exactly  $p_e$ . By Claim 2.1 this routine takes time  $\tilde{O}(\sqrt{m(\sum_e p_e)})$  in expectation, which proves the lemma.  $\square$

**4.2. Refined quantum sparsification.** Now we will combine the Spielman–Srivastava toolbox, the quantum sampling routine, and our quantum sparsification algorithm from the last section to improve the runtime of the latter from  $\tilde{O}(\sqrt{mn}/\epsilon^2)$  to  $\tilde{O}(\sqrt{mn}/\epsilon)$ .

**THEOREM 4.4** (quantum spectral sparsification). **Quantum Sparsify** $(G, \epsilon)$  returns with high probability an  $\epsilon$ -spectral sparsifier  $H$  with  $\tilde{O}(n/\epsilon^2)$  edges and has runtime  $\tilde{O}(\sqrt{mn}/\epsilon)$ . The algorithm uses  $O(\log n)$  qubits and a QRAM of  $\tilde{O}(\sqrt{mn}/\epsilon^2)$  bits.

*Proof.* First we prove correctness. Since  $H$  is a spectral sparsifier of  $G$ , and effective resistances correspond to quadratic forms in the inverse of the Laplacian, we know that the effective resistances of  $H$  approximate those of  $G$ :  $(1 - 1/100)R_{s,t}^G \leq R_{s,t}^H \leq (1 + 1/100)R_{s,t}^G$  for all  $s, t \in V$ . By Theorem 4.1 we know that the approximate resistance oracle yields estimates  $\{\tilde{R}_{s,t}^H\}$  such that  $(1 - 1/100)R_{s,t}^H \leq \tilde{R}_{s,t}^H \leq (1 + 1/100)R_{s,t}^H$ . Combining these inequalities shows that

$$(1 - 1/100)^2 R_{s,t}^G \leq \tilde{R}_{s,t}^H \leq (1 + 1/100)^2 R_{s,t}^G.$$

If we now keep each edge with probability  $p_e = \min(1, Cw_e\tilde{R}_e^H \log(n)/\epsilon^2)$ , then by Theorem 4.2 with probability  $1 - 1/n$  we find an  $\epsilon$ -spectral sparsifier with  $O(n \log(n)/\epsilon^2)$  edges. Combining this success probability with those of the quantum sparsification algorithm and the construction of the resistance oracle, we find a total success probability of at least  $(1 - 1/n)^3 = 1 - O(1/n)$ .

The bound on the runtime follows from summing the  $\tilde{O}(\sqrt{mn})$  runtime of the quantum sparsification algorithm, the  $\tilde{O}(n)$  runtime for creating the resistance oracle of the sparsifier with  $\tilde{O}(n)$  edges, and the  $\tilde{O}(\sqrt{m(\sum_e p_e)})$  expected runtime of

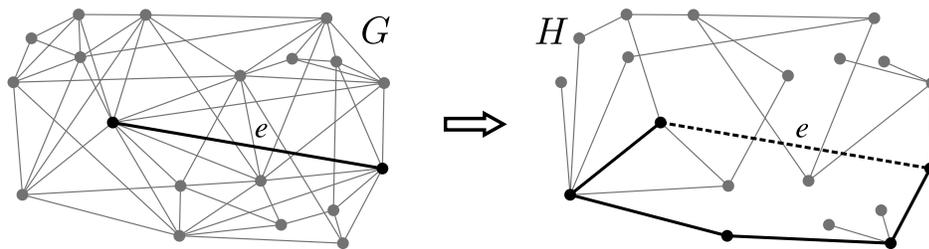


FIG. 2. A  $t$ -spanner  $H$  of a graph  $G$  is a sparse subgraph that preserves all shortest-path distances between pairs of vertices up to a stretch factor  $t$ . Equivalently, for any edge  $e = (x, y)$  in  $G$  there exists a path between  $x$  and  $y$  of distance at most  $t/w_e$  in  $H$ .

the quantum sampling routine (as remarked after Theorem 4.2, we can assume that the probabilities  $p_e$  have an  $O(\log n)$ -bit description). Since

$$\sum_e p_e \leq \frac{C \log(n)}{\epsilon^2} \sum_e w_e \tilde{R}_e^H \leq \frac{(1 + 1/100)^2 C \log(n)}{\epsilon^2} \sum_e w_e R_e^G,$$

and  $\sum_e w_e R_e^G = n - 1$  [19, Theorem 25], we have that  $\sum_e p_e \in \tilde{O}(n/\epsilon^2)$ , and so the expected runtime of the sampling routine is  $\tilde{O}(\sqrt{mn}/\epsilon)$ . Moreover, by the Chernoff bound the runtime of the latter routine will indeed be  $\tilde{O}(\sqrt{mn}/\epsilon)$  with probability at least  $1 - 1/n$ . Hence we can abort the algorithm whenever the runtime exceeds this bound, and the algorithm will still succeed with high probability, while the total runtime becomes  $\tilde{O}(\sqrt{mn}/\epsilon)$  in the worst case.  $\square$

**5. Quantum algorithm for building spanners.** The Koutis–Xu sparsification algorithm identifies “important” edges by growing spanners inside the graph. In this section we propose a quantum algorithm for growing spanners, speeding up the best classical algorithms.

Recall from section 2 that a  $t$ -spanner of a graph  $G = (V, E, w)$  is a subgraph  $H = (V, E_H \subseteq E, w)$  that preserves all pairwise distances between nodes up to a stretch factor  $t$ . For every pair  $u, v \in V$ , it should hold that

$$\delta_G(u, v) \leq \delta_H(u, v) \leq t\delta_G(u, v),$$

where we recall that  $\delta_G(u, v) = \min_{u-v \text{ path } P} \sum_{e \in P} 1/w_e$ . We illustrate this in Figure 2. A spanner preserves the original weights on its edges. This is in contrast to spectral sparsifiers which are necessarily reweighted. A classic result by Althöfer et al. [4] shows that, for any parameter  $k > 0$ , any  $n$ -node graph has a  $(2k - 1)$ -spanner with  $O(n^{1+1/k})$  edges. We refer the interested reader to the classic book by Peleg [83] or the very recent survey by Ahmed et al. [2].

There exists a range of classical algorithms for constructing spanners. We will make use of one by Thorup and Zwick [99], which follows from their work on “approximate distance oracles.” The main bottleneck of their algorithm is the growth of shortest-path trees in subgraphs. We speed up this bottleneck by using the quantum algorithm of Dürr et al. [41] for growing a shortest-path tree in time  $\tilde{O}(\sqrt{mn})$ .

**5.1. Thorup–Zwick algorithm.** The spanner algorithm from [99] makes use of *shortest-path trees* (SPTs). An SPT  $T(v)$  from a node  $v$  spanning a subset  $C$  is defined as a tree, rooted at  $v$  and spanning  $C$ , so that the distance in this tree from  $v$  to any node in  $C$  is the same as their distance in the original graph  $G$ .

---

**Algorithm 5.1**  $H = \text{Spanner}(G, k)$ .

---

- 1: let  $A_0 = V$  and  $A_k = \emptyset$
  - 2: **for**  $i = 1, 2, \dots, k$  **do**
  - 3:   if  $i < k$ , let  $A_i$  contain each element of  $A_{i-1}$ , independently, with probability  $n^{-1/k}$
  - 4:   **for**  $v \in A_{i-1} - A_i$  **do**
  - 5:     grow shortest-path tree  $T(v)$  from  $v$  spanning  $C(v) = \{w \in V \mid \delta(w, v) < \delta(w, A_i)\}$
  - 6:     add  $T(v)$  to  $H$
  - 7:   **end for**
  - 8: **end for**
- 

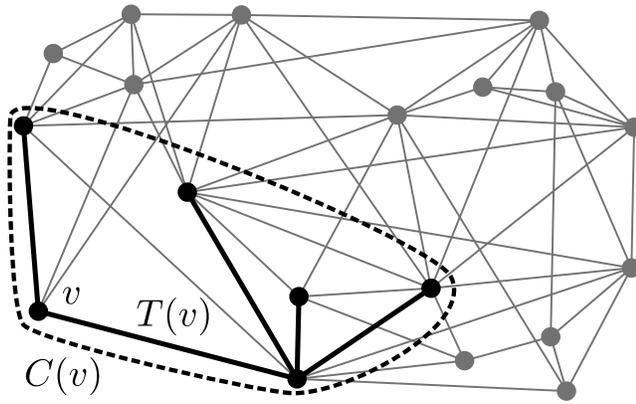


FIG. 3. The spanner algorithm by Thorup and Zwick [99]. In each iteration, the algorithm grows a shortest-path tree  $T(v)$  from a vertex  $v$  until it spans some subset  $C(v)$ .

The Thorup–Zwick algorithm, presented in Algorithm 5.1, randomly partitions the node set into  $k$  layers  $\{A_i\}$ , which are increasingly sparsified. The nodes in these layers function as “hubs” for the nearby nodes. SPTs are then grown that allow efficient routing along these hubs. The resulting spanner consists of the union of these SPTs. In the algorithm below, we set  $\delta(w, \emptyset) = \infty$  for any  $w \in V$ .

Figure 3 illustrates the use of SPTs inside the algorithm. Apart from the correctness of the algorithm, we will require some additional bounds on the size of the intermediate clusters  $C(v)$ . We extract the following theorem from the analysis by Thorup and Zwick.

**THEOREM 5.1** ([99]).

- The output graph  $H$  of  $\text{Spanner}(G, k)$  is a  $(2k - 1)$ -spanner of  $G$ .
- The expected number of edges in  $H$  is  $O(\mathbb{E}(\sum_v |C(v)|)) \in O(kn^{1+1/k})$ .
- The expected number of edges with at least one node in the clusters is

$$\mathbb{E} \left( \sum_v |E(C(v))| \right) \in O(kmn^{1/k}).$$

Setting  $k = 1/2 + \log n$ , as we will do later on, this yields a  $2 \log n$ -spanner with an expected number of edges  $O(n \log n)$ .

Downloaded 01/06/23 to 192.16.191.136 . Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

**5.2. Quantum spanner algorithm.** We can use a quantum algorithm from Dürr et al. [41] to speed up the construction of the SPT  $T(v)$ , spanning  $C(v)$ . We slightly generalize their algorithm to deal with “forbidden edges,” which are encoded by associating a weight  $w_e = 0$  to them (which corresponds to an infinite resistance or cost). Such edges will correspond to edges going outside of  $C(v)$ , as well as edges that have already been discarded by our sparsification algorithm.

In Appendix A we prove the following statement. We define the connected component of a node  $v_0$  as the smallest subset  $C_{v_0} \subseteq V$  such that  $v_0 \in C_{v_0}$  and either  $E(C_{v_0}, V \setminus C_{v_0}) = \emptyset$  or  $\max\{w_e \mid e \in E(C_{v_0}, V \setminus C_{v_0})\} = 0$ . This implies that there is no path of finite distance between  $v_0$  and any node outside  $C_{v_0}$ .

**PROPOSITION 5.2.** *Assume adjacency-list access to a weighted and undirected graph  $G = (V, E, w)$ . Let  $v_0$  be a source node and  $C_{v_0}$  its connected component. Then there exists a quantum algorithm that outputs, with probability at least  $1 - \delta$ , an SPT from  $v_0$  that spans  $C_{v_0}$ . It has a runtime  $\tilde{O}(\sqrt{|C_{v_0}|} |E(C_{v_0})| \log(n/\delta))$  and requires  $O(\log n)$  qubits and a QRAM of  $\tilde{O}(|C_{v_0}|)$  bits.*

From this we can speed up the spanner construction rather straightforwardly. To see this, note that the runtime of the Thorup–Zwick algorithm is dominated by the task of growing the SPTs  $T(v)$ , spanning the local clusters  $C(v)$ , for all nodes  $v \in V$ . By setting  $w_e = 0$  for any edge reaching out of  $C(v)$ , this task corresponds to an SPT on the connected component of  $v$ . If we use the above quantum algorithm to accelerate this, the total runtime becomes

$$\tilde{O}\left(\sum_v \sqrt{|C(v)| |E(C(v))|}\right) \in \tilde{O}\left(\sqrt{\sum_v |C(v)|} \sqrt{\sum_v |E(C(v))|}\right),$$

where the containment follows from the Cauchy–Schwarz inequality. By Theorem 5.1 we know that  $\mathbb{E}(\sum_v |C(v)|) \in O(kn^{1+1/k})$  and  $\mathbb{E}(\sum_v |E(C(v))|) \in O(kmn^{1/k})$ . By Markov’s inequality this implies that with probability close to 1 the runtime is

$$\tilde{O}\left(\sqrt{kn^{1+1/k}} \sqrt{kmn^{1/k}}\right) \in \tilde{O}(kn^{1/k} \sqrt{mn}).$$

What remains to be shown is how we (implicitly) set  $w_e = 0$  for all edges reaching out of  $C(v)$ . To that end we follow the idea of Thorup and Zwick of connecting a new source node  $s$  to every node in  $A_i$ , with edges of infinite weight, and construct an SPT from  $s$  to  $V$ . It is easy to see that this returns the shortest path from any node  $w \notin A_i$  to  $A_i$ , allowing us to calculate  $\delta(w, A_i)$ . Using the standard quantum SPT algorithm of [41] we can construct this SPT in time  $O(\sqrt{mn})$ , and we do this whenever we construct a new  $A_i$ . Now assume that the quantum SPT algorithm at some point wishes to choose an edge  $(w, w')$ , with  $w$  part of the SPT constructed so far, and  $w'$  an adjacent node. Then by design this must be a cheapest border edge of the SPT constructed so far, and  $\delta(v, w') = \delta(v, w) + \delta(w, w')$ . Hence we know  $\delta(v, w')$  and can simply check whether  $\delta(v, w') < \delta(w, A_i)$ , setting the weight of the edge equal to zero if this is not the case. This proves the following theorem.

**THEOREM 5.3.** *There is a quantum algorithm that outputs in time  $\tilde{O}(kn^{1/k} \sqrt{mn})$  with high probability a  $(2k - 1)$ -spanner of  $G$  of size  $O(kn^{1+1/k})$ . The algorithm uses  $O(\log n)$  qubits and a QRAM of  $\tilde{O}(kn^{1+1/k})$  bits.*

Setting  $k = \log n + 1/2$ , we find an  $\tilde{O}(\sqrt{mn})$  quantum algorithm for constructing  $2 \log n$ -spanners, as required by our sparsification algorithm.

**6. Matching lower bound: A hidden sparsifier.** In this section we prove that the runtime of our quantum algorithm for spectral sparsification is optimal, up to polylog-factors. In fact, we show that even constructing a weaker *cut* sparsifier requires the same complexity. The following is a rephrasing of Theorem 1.2 from the introduction.

**THEOREM 6.1.** *Fix  $n, m$  and let  $\epsilon \geq \sqrt{n/m}$ . Consider the problem of outputting, with high probability, an explicit description of an  $\epsilon$ -cut sparsifier of a weighted, undirected graph  $G$  with  $n$  nodes and  $m$  edges, given adjacency-list access to  $G$ . The quantum query complexity of this problem is  $\tilde{\Omega}(\sqrt{mn}/\epsilon)$ .*

Note that sparsification is only meaningful under the constraint  $\epsilon \geq \sqrt{n/m}$ , since for  $\epsilon \in O(\sqrt{n/m})$  the sparsifier would have at least as many edges as the original graph. We prove this lower bound by “hiding” a sparsifier in a larger graph, and then proving a quantum lower bound for finding the sparsifier. More specifically, we use a random graph construction by Andoni et al. [6] such that any cut sparsifier must contain a constant fraction of the edges of the graph. We then hide a number of copies of this random graph by embedding it in a larger, denser graph. Finally we show that finding a constant fraction of the edges of the initial random graph requires  $\tilde{\Omega}(\sqrt{mn}/\epsilon)$  queries. To prove this lower bound, we combine a quantum lower bound for the OR-function with an information-theoretic lower bound for the problem of finding nonzero bits in a Boolean matrix. We can combine these separate lower bounds by using a composition theorem for adversary bounds, applicable to the composition of a relational problem with a function. Prompted by our question, such a composition theorem was very recently proven by Belovs and Lee [17].

**6.1. Hiding a sparsifier.** We use a random graph construction of Andoni et al. [6] for which a sparsifier must output a constant fraction of its edges. We then carefully hide a number of copies of this graph in a larger graph, which will later allow for the reduction of a query problem to the construction of a sparsifier.

**6.1.1. An unsparsifiable graph.** Andoni et al. [6] construct a communication problem that is described by a family of random graphs on  $2/\epsilon^2$  nodes with  $1/(2\epsilon^4)$  edges. They show that for a constant fraction of the inputs ( $> 3/5$ ), the communication problem requires the communication of  $\Omega(1/\epsilon^4)$  bits. On the other hand, for at least a  $2/3$ -fraction of the inputs, the communication problem can be solved by communicating an  $\epsilon$ -cut sparsifier. This shows that for at least a  $(3/5 - 1/3) > 1/4$ -fraction of the inputs, the description of an  $\epsilon$ -cut sparsifier requires  $\Omega(1/\epsilon^4)$  bits. Using a slightly more refined argument, they even show that the number of edges of the sparsifier for these instances must be  $\Omega(1/\epsilon^4)$ .

Fix any  $\epsilon > 0$ . Let  $B_\epsilon$  be any bipartite graph with  $1/\epsilon^2$  nodes on each side, and every left node connected to a corresponding subset of half of the right nodes. From Andoni et al. [6, Theorem 3.3] we can extract the following theorem. Indeed, if the claim would not hold, then we could violate their communication lower bound by communicating an  $\epsilon$ -cut sparsifier.

**THEOREM 6.2 ([6]).** *For at least a  $1/4$ -fraction of all graphs  $B_\epsilon$ , any  $\epsilon$ -cut sparsifier must contain  $\Omega(1/\epsilon^4)$  edges.*

It follows that at least a  $1/4$ -fraction of all graphs  $B_\epsilon$  cannot be significantly sparsified. Similarly to [6], we will also consider larger families of disjoint copies of  $B_\epsilon$ . We can easily prove the following corollary using the Chernoff bound.

**COROLLARY 6.3.** *Consider the disjoint union of  $\ell$  distinct copies of  $B_\epsilon$ . There exists a constant  $\eta > 0$ , independent of  $\ell$ , such that for at least a  $9/10$ -fraction of all such graphs it holds that any  $\epsilon$ -cut sparsifier must contain at least  $\eta\ell/\epsilon^4$  edges.*

**6.1.2. Embedding the sparsifier.** Fix  $n, m \leq n^2/4$  and  $\epsilon \geq \sqrt{n/m}$ . Consider  $\ell = \epsilon^2 n/2$  independent copies  $B^{(k)}$  of  $B_\epsilon$ , yielding a graph with  $n$  nodes. We wish to “hide” this graph in a larger, denser graph. To this end, we use an  $m$ -bit string  $x$  to (redundantly) describe the resulting graph, which we denote by  $B(x)$ . The description  $x$  consists of  $\ell = \epsilon^2 n/2$  matrices  $x^{(k)}$  of dimension  $1/\epsilon^2 \times 1/\epsilon^2$ ,

$$x^{(k)} = \{x_{i,j}^{(k)} \mid i, j \in [1/\epsilon^2]\}, \quad k \in [\epsilon^2 n/2].$$

Every matrix  $x^{(k)}$  is used to describe the bipartite adjacency matrix of the corresponding copy  $B^{(k)}$ . Rather than bits, however, the entries  $x_{i,j}^{(k)}$  correspond to smaller strings of  $N = 2\epsilon^2 m/n$  bits each, with at most one nonzero bit per string. We say that an input  $x$  is *valid* if it is of this form. The presence of an edge between the  $i$ th left node and the  $j$ th right node in  $B^{(k)}$  is determined by the presence of a nonzero bit in the string  $x_{i,j}^{(k)}$ , i.e., by  $\text{OR}(x_{i,j}^{(k)})$ . The bipartite adjacency list of the  $k$ th copy  $B^{(k)}$  is hence described as

$$\text{adj}(B^{(k)}) = \text{OR}(x^{(k)}) = \text{OR} \left( \begin{bmatrix} 00010 & 00000 & 00000 & 10000 \\ 00100 & 00000 & 00100 & 00000 \\ 00000 & 01000 & 00001 & 00000 \\ 00000 & 00000 & 00010 & 01000 \end{bmatrix} \right) = \begin{bmatrix} 1 & 0 & 0 & 1 \\ 1 & 0 & 1 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 0 & 1 & 1 \end{bmatrix},$$

where we give concrete values to the bits in  $x^{(k)}$  for illustration. In the next section we prove a lower bound on the identification of a constant fraction of the 1-bits in a valid input  $x$ . In this section we show how to embed the corresponding graph  $B(x)$  in a larger, denser graph  $G(x)$ , so that an  $\epsilon$ -cut sparsifier of  $G(x)$  must identify a constant fraction of the edges of  $B(x)$ . This identifies a constant fraction of the 1-bits of the input  $x$ , so that the aforementioned lower bound effectively yields a lower bound for the construction of a sparsifier. We must take particular care in embedding  $B(x)$  in  $G(x)$  so as not to reveal additional information about  $x$ . For example, we must prevent the degrees of  $G(x)$  from revealing anything about the location of 1-bits in the input. We do this essentially by ensuring that a query to the adjacency list of  $G(x)$  can be performed using a single query to  $x$ .

Initially, the oblivious (input-independent) mother graph  $G = G_1 \cup G_2$  is the disjoint union of two bipartite graphs  $G_1 = (L_1 \cup R_1, E_1)$  and  $G_2 = (L_2 \cup R_2, E_2)$ . The first graph  $G_1$  is a multigraph consisting of  $\ell = \epsilon^2 n/2$  disjoint copies  $B^{(k)}$  of the *complete* bipartite graph on  $1/\epsilon^2$  nodes, containing  $N = 2\epsilon^2 m/n$  parallel copies of each edge (we will get rid of these multiedges later). As such,  $G_1$  has exactly  $n$  nodes and  $m$  (sometimes parallel) edges. We match every edge of  $G_1$  to a unique input bit by matching the parallel copies of edge  $(i, j)$  in  $B^{(k)}$  to the input bits in the string  $x_{i,j}^{(k)}$ . The second bipartite graph  $G_2$  has at most  $n$  nodes and exactly  $m$  edges, but no multiedges (we do not need to further specify this graph at this point). We formally match every single copy of an edge in  $G_1$  to a unique edge in  $G_2$ , so that every input bit now corresponds to a unique edge in  $G_1$  and a unique edge in  $G_2$ . While doing so, we ensure that all edges leaving a left or right node of  $G_1$  are matched to edges in  $G_2$  whose left (resp., right) ends are distinct. We will later clarify why this is important and defer a proof that this is possible for some choice of  $G_2$  to Appendix B. At this point, all edges in  $G$  have zero weight.

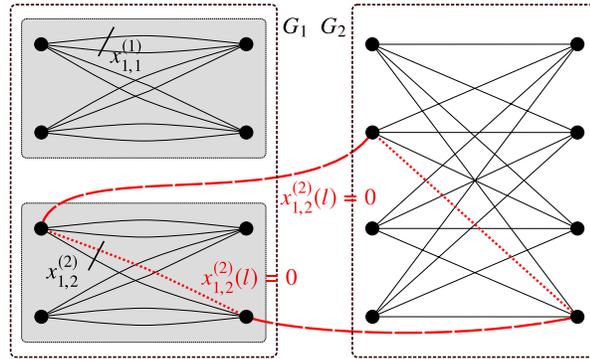


FIG. 4. Matching input bits to edges in  $G_1$  and  $G_2$  for  $n = 8$ ,  $m = 16$ , and  $\epsilon = 1/\sqrt{2}$  (we set  $G_2$  to be the complete bipartite graph for illustration only). The dashed red edges depict a pair of matched edges, which correspond to input bit  $x_{1,2}^{(2)}(l)$ . If  $x_{1,2}^{(2)}(l) = 1$ , these edges are kept in  $G(x)$ ; if  $x_{1,2}^{(2)}(l) = 0$ , they are “flipped” with the dotted edges.

Next we take the input  $x$  into account, turning  $G$  into  $G(x)$ . To this end, we “flip” edge pairs conditioned on the input bit. Specifically, consider a bit  $x(i)$  that corresponds to edges  $(l, r)$  in  $G_1$  and  $(l', r')$  in  $G_2$ . If  $x(i) = 1$ , then we keep these edges as they are, except that we give  $(l, r)$  a unit weight. If  $x(i) = 0$ , then we “flip” the edges: we replace  $(l, r)$  and  $(l', r')$  by edges  $(l, l')$  and  $(r, r')$ . This is illustrated in Figure 4. We see now that if two outgoing edges from  $l$  were matched to edges with the same left end  $l'$  in  $L_2$ , then this could create the edge  $(l, l')$  twice. Similarly if two incoming edges from  $r$  were matched to the same right ends  $r'$ , this could create the edge  $(r, r')$  twice. Our matching ensures that this can never happen.

Now consider a pair  $l \in L_1$  and  $r \in R_1$ , corresponding to the edge  $(i, j)$  in  $B^{(k)}$ . Then  $G(x)$  will contain a unique, unit-weight edge between  $l$  and  $r$  if and only if the string  $x_{i,j}^{(k)}$  has a unique nonzero bit. As a consequence,  $G(x)$  restricted to the node set  $L_1 \cup R_1$  exactly describes  $B(x)$ . Moreover, we can perform a single query to the adjacency list of  $G(x)$  using a single query to  $x$ , as we prove in the lemma below.

LEMMA 6.4. Fix  $n, m \leq n^2/4$  and  $\epsilon \geq \sqrt{n/m}$ , and consider a valid input  $x \in \{0, 1\}^m$ . Then  $G(x)$  has at most  $2n$  nodes and exactly  $2m$  edges, and any query to its adjacency list can be simulated using a single query to  $x$ .  $G(x)$  has  $B(x)$  as a subgraph with unit edge weights, and all remaining edges have zero weight.

*Proof.* We only prove the claim about the query access, as all other claims follow easily from the construction. A degree query is trivially simulated, since  $G(x)$  has the same degrees as  $G$ . To simulate a neighbor query, say that we wish to query the  $k$ th entry of the adjacency list of node  $l \in L_1$  in  $G_1$  for some  $k \in [2m/n]$ . By construction, this corresponds to a known edge  $(l, r)$  in  $G_1$ , a corresponding edge  $(l', r')$  in  $G_2$ , and a unique input bit  $x(i)$ . We query the input bit  $x(i)$ . If it equals 1, then  $G(x)$  has the unit-weight edge  $(l, r)$ , and hence we return  $r \in R_1$  and weight 1. If it equals 0, then we flipped the edges and  $G(x)$  has the edge  $(l, l')$ , and so we return the node  $l' \in L_1$  and weight 0. The same reasoning applies to queries from all other nodes of  $G(x)$ .  $\square$

We will also use the claim below, which follows easily from the fact that only the subgraph  $B(x)$  of  $G(x)$  has unit-weight edges, and all remaining edges have weight zero.

CLAIM 6.5. Any  $\epsilon$ -cut sparsifier of  $G(x)$  contains an  $\epsilon$ -cut sparsifier of  $B(x)$ .

*Proof.* Let  $H$  be any  $\epsilon$ -cut sparsifier of  $G(x)$ . Then we can simply remove all edges from  $H$  that are not in  $B(x)$  (and necessarily have zero weight) to yield an  $\epsilon$ -cut sparsifier for  $B(x)$ .  $\square$

Combining this claim with Corollary 6.3, we can deduce that for at least  $9/10$  of all valid inputs  $x$ , any  $\epsilon$ -cut sparsifier of  $G(x)$  must identify a constant fraction (specifically, at least  $\eta n/\epsilon^2$ ) of the edges of  $B(x)$ . Since the presence of an edge in  $B(x)$ , say  $(l_i, r_j)$  in copy  $B^{(k)}$ , reveals that string  $x_{i,j}^{(k)}$  has a nonzero entry, we find the following corollary.

**COROLLARY 6.6.** *For at least a  $9/10$ -fraction of all valid inputs  $x$ , it holds that any  $\epsilon$ -cut sparsifier of  $G(x)$  must identify a constant  $2\eta$ -fraction of the nonzero strings.*

**6.2. Finding the sparsifier.** Now we wish to prove a lower bound on identifying nonzero strings, thereby proving a lower bound on the complexity of sparsifying  $G(x)$ . To this end, we first formalize the new problem.

**DEFINITION 6.7.** *The problem  $\text{FindBits}_{r,c}$  takes as input an  $r \times c$  Boolean matrix, with each row containing exactly  $c/2$  nonzero bits. A correct output consists of the indices of a  $2\eta$ -fraction of the nonzero bits.*

Note that this is a *relational* problem: for every input there are many different possible correct outputs. Formally, a relational problem  $f$  corresponds to a set  $f \subseteq S \times T$ , where  $S \subseteq \{0,1\}^M$  is the set of possible inputs and  $T$  denotes the set of possible outputs. Output  $t$  is deemed a correct output for  $f$  on input  $x \in S$  if and only if  $(x,t) \in f$ . For our relational problem  $\text{FindBits}_{r,c}$ ,  $S$  is the allowed set of  $r \times c$  Boolean input matrices, and each  $t \in T$  corresponds to a set of  $\eta rc$  indices of that matrix. Output  $t$  is correct for input matrix  $x$  if all indices in  $t$  correspond to 1-bits in  $x$ .

Using information theory, we can prove a lower bound on the bounded-error quantum query complexity of this problem, i.e., on the number of queries to  $x$  required by a quantum algorithm that returns a correct output with probability at least  $2/3$ .

**CLAIM 6.8.** *The bounded-error quantum query complexity of  $\text{FindBits}_{r,c}$  on any constant fraction of its valid inputs is  $\tilde{\Omega}(rc)$ .*

*Proof.* This lower bound follows from combining an information-theoretic lower bound on how much information the algorithm extracts about its input, with Holevo's quantum information bound. More precisely, we examine the mutual information  $I(A; B)$  between a uniformly random input matrix  $A$  and an output  $B$  for  $\text{FindBits}_{r,c}$  that is correct with probability at least  $2/3$ . We refer the reader to the textbook of Cover and Thomas [36] for precise definitions. By Holevo's theorem [54], this mutual information lower bounds the quantum query complexity times  $O(\log(rc))$ .<sup>3</sup>

We lower bound  $I(A; B) = H(A) - H(A|B)$  by lower bounding  $H(A)$  and upper bounding  $H(A|B)$ . To lower bound  $H(A)$ , let  $S$  denote the set of all inputs, consisting of all  $r \times c$  Boolean matrices with all row sums equal to  $c/2$ , and let  $S' \subseteq S$  denote any constant fraction subset of  $S$  (say  $|S'| = \beta|S|$  for some  $\beta > 0$ ). Then  $A$  corresponds to a uniformly random element of  $S'$ , and we have

<sup>3</sup> This is a fairly standard argument. We can consider a communication protocol where the  $T$ -query quantum algorithm implements each quantum query using one round of  $O(\log(rc))$  qubits of communication to and from another party that holds  $A$  (by sending the query and answer registers back and forth), thus learning  $I(A; B)$  bits of information about  $A$  while only communicating  $O(T \log(rc))$  qubits. It now follows from [54] that  $T \in \Omega(I(A; B)/\log(rc))$ .

$$\begin{aligned}
 H(A) &= \log |S'| = \log(\beta|S|) \\
 &= \log \left( \beta \prod_{j=1}^r \binom{c}{c/2} \right) \\
 &= r \log \binom{c}{c/2} + \log \beta \geq r(c - O(\log c)).
 \end{aligned}$$

To upper bound  $H(A|B)$ , we use that with probability at least  $2/3$  the output  $B$  is correct, revealing a  $2\eta$ -fraction of the nonzero entries of  $A$  and hence significantly reducing the entropy of  $A$ . Let  $E$  denote an indicator bit which equals 1 if  $B$  is a correct output, and 0 otherwise. Then we can bound

$$\begin{aligned}
 H(A|B) &\leq H(A|B, E) + H(E) \\
 &= \Pr(E = 1) H(A|B, E = 1) + \Pr(E = 0) H(A|B, E = 0) + H(E) \\
 &\leq \frac{2}{3} H(A|B, E = 1) + \frac{1}{3} rc + 1.
 \end{aligned}$$

We now want to upper bound  $H(A|B, E = 1)$ . Fix integers  $\{d(j)\}_{j \in [r]}$  such that  $\sum_j d(j) = \eta rc$ , and condition on the event  $D_{\{d(j)\}}$  that  $B$  reveals  $d(j)$  entries of the  $j$ th row of  $A$ . Let  $S_{\{d(j)\}} \subseteq S'$  denote the set of inputs which are compatible with  $\{d(j)\}$ . We can bound

$$\begin{aligned}
 H(A|B, E = 1, D_{\{d(j)\}}) &\leq \log |S_{\{d(j)\}}| \leq \log \left( \prod_{j=1}^r \binom{c - d(j)}{c/2 - d(j)} \right) \\
 &= \sum_{j=1}^r \log \binom{c - d(j)}{c/2 - d(j)} \leq \sum_{j=1}^r (c - d(j)) = r(1 - \eta)c.
 \end{aligned}$$

This upper bound holds irrespective of the specific setting of  $\{d(j)\}$ . Hence, taking the expectation over all possible settings of  $\{d(j)\}$ , we obtain

$$H(A|B, E = 1) = \mathbb{E}_{\{d(j)\}} [H(A|B, E = 1, D_{\{d(j)\}})] \leq r(1 - \eta)c.$$

Our lower bound on  $H(A)$  and upper bound on  $H(A|B)$  together imply

$$I(A; B) \geq \frac{2}{3} r(\eta c - O(\log c)) - 1 \in \Omega(rc),$$

which finalizes our proof. □

Next we will compose the relational problem  $\text{FindBits}_{r,c}$  with  $r \times c$  copies of the  $\text{OR}_N$ -function on  $N$  bits each, so that every single input bit of  $\text{FindBits}_{r,c}$  is now described by an  $\text{OR}_N$ . We denote this composed problem as  $\text{FindBits}_{r,c} \circ \text{OR}_N^{rc}$ . To ensure that the input of the composed problem is a valid input for  $\text{FindBits}_{r,c}$ , we must restrict it to  $r \times c$  matrices of strings, each carrying  $N$  bits, so that exactly  $c/2$  strings per row have one nonzero entry, and the remaining strings only have zeros.

The composed problem  $\text{FindBits}_{r,c} \circ \text{OR}_N^{rc}$  corresponds to the composition of a relational problem (multiple outputs are correct for  $\text{FindBits}_{r,c}$ ) and a function. Bounds on the quantum query complexity of the composition of two functions are well understood [55]. When the outer problem is a relational problem, however, no such result seemed to be known. Fortunately, prompted by our question, Belovs and Lee [17] very recently proved the following theorem.

**THEOREM 6.9** (Corollary 27 of [17]). *Let  $f \subseteq S \times T$ , with  $S \subseteq \{0,1\}^M$ , be a relational problem with bounded-error quantum query complexity  $L$ . Assume that  $f$  is efficiently verifiable, in the sense that there is a bounded-error quantum algorithm that, given some  $t \in T$  and oracle access to  $x \in S$ , verifies whether  $(x,t) \in f$  using  $o(L)$  queries to  $x$ . Let  $g : D \rightarrow \{0,1\}$ , with  $D \subseteq \{0,1\}^N$ , be a Boolean function whose bounded-error quantum query complexity is  $Q$ . Then the bounded-error quantum query complexity of the relational problem  $f \circ g^M$ , restricted to inputs  $x \in \{0,1\}^{NM}$  such that  $g^M(x) \in S$ , is  $\Theta(LQ)$ .*

We instantiate this by setting  $f$  to the relational problem  $\text{FindBits}_{r,c}$ , with  $S$  a  $9/10$ -fraction of its valid inputs. By Claim 6.8 its bounded-error quantum query complexity is  $L \in \tilde{\Omega}(rc)$ . Using Grover's algorithm, we can efficiently verify an output, using  $O(\sqrt{rc}) \in o(L)$  queries to  $f$ 's  $M$ -bit input. We let  $g$  be the function  $\text{OR}_N$  restricted to the set  $D$  of all  $N$ -bit inputs of Hamming weight 0 or 1; its bounded-error quantum query complexity is  $\Theta(\sqrt{N})$ . Plugging this into Theorem 6.9 yields the following corollary.

**COROLLARY 6.10.** *Solving the problem  $\text{FindBits}_{r,c} \circ \text{OR}_N^{rc}$  has bounded-error quantum query complexity  $\tilde{\Omega}(rc\sqrt{N})$ . This holds even when the inputs are restricted to a constant fraction of the valid inputs.*

Finally, using the graph embedding of the previous section, we show in the claim below how to solve this composed relational problem by constructing a cut sparsifier of an associated graph. Combining this with the lower bound from Corollary 6.10 then yields our lower bound for finding a cut sparsifier (Theorem 6.1).

**CLAIM 6.11.** *Fix  $n$ , mlet  $\epsilon \geq \sqrt{n/m}$ , and set  $r = n/2$ ,  $c = 1/\epsilon^2$  (number of potential neighbors) and  $N = 2\epsilon^2 m/n$ . For at least a  $9/10$ -fraction of all valid inputs  $x$ , we can reduce the problem  $(\text{FindBits}_{r,c} \circ \text{OR}_N^{rc})(x)$  to finding an  $\epsilon$ -cut sparsifier of  $G(n, m, \epsilon, x)$ . Specifically, any  $T$ -query quantum algorithm that sparsifies  $G(n, m, \epsilon, x)$  with success probability at least  $2/3$  can solve  $(\text{FindBits}_{r,c} \circ \text{OR}_N^{rc})(x)$  with success probability at least  $2/3$  on at least a  $9/10$ -fraction of its valid inputs using  $O(T)$  queries.*

*Proof.* Consider a valid input  $x \in \{0,1\}^m$  and the associated graph  $G(x) = G(n, m, \epsilon, x)$  as defined in section 6.1.2. For a  $9/10$ -fraction of these inputs, we know by Corollary 6.6 that we can solve  $(\text{FindBits}_{r,c} \circ \text{OR}_N^{rc})(x)$  by constructing an  $\epsilon$ -cut sparsifier of  $G(x)$ . Moreover, by Lemma 6.4 we can query the adjacency list of  $G(x)$  using a single query to  $x$ . Hence, if we output with probability at least  $2/3$  an  $\epsilon$ -cut sparsifier of  $G(x)$  using  $T$  queries to its adjacency list, then we also solve  $(\text{FindBits}_{r,c} \circ \text{OR}_N^{rc})(x)$  with probability at least  $2/3$  using  $T$  queries to  $x$ .  $\square$

**7. Applications.** In this section we nonexhaustively list some applications of our quantum sparsification algorithm, proving speedups in cut approximation and Laplacian solving. Given the broad applicability of sparsification in classical algorithms, we expect that more applications will follow.

**7.1. Cut approximation.** A range of cut approximation algorithms have a near-linear runtime in the number of edges in the graph. In the following we discuss a number of these algorithms, and show how our quantum algorithm for cut sparsification allows us to speed them up.

**7.1.1. max cut.** The MAX CUT problem for a weighted graph  $G = (V, E, w)$  asks for a cut  $(S, S^c)$  that maximizes the total weight  $\text{val}_G(S)$  of the cut. Its decision

version is one of the 21 problems famously shown to be NP-complete by Karp [64]. The best current approximation factor of this problem is roughly .8785, as was famously proven by Goemans and Williamson [45] using an SDP relaxation. Khot et al. [69] showed that this approximation factor is optimal under the unique games conjecture.

In later work, Arora and Kale [10] showed how to solve the Goemans–Williamson SDP in time  $\tilde{O}(m)$ , using, among other things, the cut sparsification algorithm by Benczúr and Karger [18]. We find a quantum speedup by replacing the Benczúr–Karger algorithm by our quantum sparsification algorithm. Specifically we construct an  $\epsilon$ -spectral sparsifier  $H$  for some small but constant  $\epsilon > 0$  in time  $\tilde{O}(\sqrt{mn}/\epsilon)$ . Since all cuts are preserved up to a multiplicative error  $\epsilon$ , we can apply the Arora–Kale algorithm on  $H$  to retrieve a MAX CUT approximation factor of at least  $.8785(1 - \epsilon)$ . Choosing  $\epsilon$  sufficiently small we find the following claim.

CLAIM 7.1. *There exists a quantum algorithm that outputs a .878-approximate max cut of a weighted graph in time  $\tilde{O}(\sqrt{mn})$ .*

We note that for *unweighted* graphs, MAX CUT can be approximately solved classically in time  $\tilde{O}(n)$ . If we wish to output the MAX CUT bipartition, then this is trivially optimal both for classical and quantum algorithms, and hence no quantum speedup is possible. The  $\tilde{O}(n)$  classical algorithm follows from the fact that for unweighted graphs a trivial sparsification procedure suffices to approximately preserve the MAX CUT value (pick  $\tilde{O}(n)$  edges uniformly at random—see, for instance, [100, section 2]). In the adjacency list model, this can be done classically in time  $\tilde{O}(n)$ . For weighted graphs this approach no longer works, as in this case the edges have to be sampled with probability proportional to their edge weights. This cannot be done classically in time  $o(m)$  since we could, for instance, hide a single heavy edge among  $m$  light edges. From all the cut problems we consider, this is the only one for which a classical sublinear algorithm with multiplicative error exists—albeit only for the unweighted case.

**7.1.2. min cut.** Given a weighted graph  $G = (V, E, w)$ , the MIN CUT problem asks for a cut  $(S, S^c)$  with minimum weight  $\text{val}_G(S)$ . Up to polylog-factors in the runtime, the current best algorithm is by Karger [63]. It builds on cut sparsification to output an exact min cut of the graph with high probability in time  $\tilde{O}(m)$ . Running this algorithm on our  $\epsilon$ -cut sparsifier with  $\tilde{O}(n/\epsilon^2)$  edges thus requires time  $\tilde{O}(n/\epsilon^2)$  and returns an  $\epsilon$ -approximate min cut.

CLAIM 7.2. *There exists a quantum algorithm that outputs an  $\epsilon$ -approximate min cut of a weighted graph in time  $\tilde{O}(\sqrt{mn}/\epsilon)$ .*

We leave as an open question whether this algorithm can be improved to also output an *exact* min cut, similar to the algorithm in [63].

**7.1.3. min *st*-cut.** Given a weighted graph  $G = (V, E, w)$ , the MIN *st*-CUT problem requires one to output a cut  $C = (S, S^c)$  with minimum value  $\text{val}_G(S)$ , such that  $s \in S$  and  $t \notin S$ . The current best algorithms for approximate MIN *st*-CUT build on the max-flow min-cut theorem, which states that the value of the min *st*-cut equals the value of a maximum *st*-flow. Combining classical cut sparsification with the recent  $\tilde{O}(m/\epsilon^3)$  solver for  $\epsilon$ -approximate max flows by Peng [84], this yields an  $\tilde{O}(m + n/\epsilon^5)$  time algorithm for MIN *st*-CUT. We obtain the following claim by running this algorithm on our cut sparsifier with  $\tilde{O}(n/\epsilon^2)$  edges.

CLAIM 7.3. *There exists a quantum algorithm that outputs an  $\epsilon$ -approximate minimum *st*-cut of a weighted graph in time  $\tilde{O}(\sqrt{mn}/\epsilon + n/\epsilon^5)$ .*

**7.1.4. sparsest cut and balanced separator.** Given a weighted graph  $G = (V, E, w)$ , the SPARSEST CUT problem asks for a cut  $C = (S, S^c)$  which minimizes the ratio  $\text{val}_G(S)/(|S| |S^c|)$ . The BALANCED SEPARATOR problem asks in addition that the cut be “balanced,” i.e.,  $\mu \leq |S|/|V| \leq 1/2$  for some constant  $\mu > 0$ . Exactly solving either of these problems is NP-hard, and optimally trying to approximate them has led to an interesting line of research [11, 75]. Currently the best polynomial-time classical algorithm for both problems achieves an  $O(\sqrt{\log n})$ -approximation factor in time  $\tilde{O}(m + n^{1+\delta})$  for an arbitrary positive constant  $\delta$ . This is achieved by combining cut sparsification, work by Sherman [88] which shows that an  $O(\sqrt{\log n})$ -approximation can be calculated by solving  $\tilde{O}(n^\delta)$  max flows, and the  $\tilde{O}(m/\epsilon^3)$  max flow algorithm by Peng [84] for constant  $\epsilon > 0$ . Applying these algorithms to our  $\epsilon$ -cut sparsifier for constant  $\epsilon > 0$ , we get the claim below.

CLAIM 7.4. *There is a quantum algorithm that outputs an  $O(\sqrt{\log n})$ -approximate sparsest cut or balanced separator of a weighted graph in time  $\tilde{O}(\sqrt{mn} + n^{1+\delta})$  for an arbitrary constant  $\delta > 0$ .*

**7.2. Quantum Laplacian solver.** The complexity of the best known linear system solver is  $\tilde{O}(n^\omega)$ , with  $\omega < 2.373$  the matrix multiplication coefficient. Building on a long line of work, Spielman and Teng [94] famously showed that the special case of Laplacian systems can be solved in time  $\tilde{O}(m) \in \tilde{O}(n^2)$ , with  $m$  the number of nonzero entries of the Laplacian. To this end they exploit the connection of Laplacians with graphs, allowing for a combinatorial interpretation and treatment of the linear system.

More specifically, a Laplacian solver aims to solve a linear system  $L_G x = b$ , where  $L_G$  is the Laplacian associated to some graph  $G$ . We denote the solution by  $x = L_G^+ b$ . The original motivation for spectral sparsification was in fact to create better Laplacian solvers. Indeed, as follows from the lemma below, we can solve the Laplacian system in the sparsifier and retrieve an approximation of the original system. Here we use the  $A$ -induced norm  $\|v\|_A = \sqrt{v^\dagger A v} = \|A^{1/2} v\|$  for a positive semidefinite matrix  $A$ , with  $v^\dagger$  the Hermitian transpose of vector  $v$ .

CLAIM 7.5. *Consider a linear system  $L_G x = b$ , where  $L_G$  is the Laplacian of a weighted, undirected graph  $G$ . If  $H$  is an  $\epsilon$ -spectral sparsifier of  $G$ , with Laplacian  $L_H$ , then solving  $L_H x = b$  yields an approximate solution to the original system:*

$$\|L_H^+ b - x\|_{L_G} \leq 2\epsilon \|x\|_{L_G}.$$

*Proof.* Since  $H$  is an  $\epsilon$ -spectral sparsifier of  $G$ , we have that  $(1 - \epsilon)L_G \preceq L_H \preceq (1 + \epsilon)L_G$ . This implies that the nonzero eigenvalues of  $L_G L_H^+$  (and hence also of  $L_G^{1/2} L_H^+ L_G^{1/2}$ ) lie between  $1/(1 + \epsilon)$  and  $1/(1 - \epsilon)$ . With  $I$  the identity matrix restricted to the image of  $L_G$  and  $L_H$ , this implies that  $\|L_G^{1/2} L_H^+ L_G^{1/2} - I\| \leq \epsilon/(1 - \epsilon) \leq 2\epsilon$  for  $\epsilon \leq 1/2$ . From this we can bound

$$\begin{aligned} \|L_H^+ b - x\|_{L_G} &= \|L_G^{1/2} L_H^+ b - L_G^{1/2} x\| \\ &= \|L_G^{1/2} L_H^+ L_G x - L_G^{1/2} x\| \\ &= \|(L_G^{1/2} L_H^+ L_G^{1/2} - I) L_G^{1/2} x\| \leq \|L_G^{1/2} L_H^+ L_G^{1/2} - I\| \|L_G^{1/2} x\|. \end{aligned}$$

Since  $\|L_G^{1/2} L_H^+ L_G^{1/2} - I\| \leq 2\epsilon$  and  $\|L_G^{1/2} x\| = \|x\|_{L_G}$ , this proves the lemma.  $\square$

This observation allowed Spielman and Teng to reduce the task of solving a potentially dense Laplacian system to solving a very sparse one in the Laplacian of the

sparsifier, having  $\tilde{O}(n/\epsilon^2)$  edges. They then invoke other methods to efficiently solve the sparse Laplacian system in additional time  $\tilde{O}(n/\epsilon^2)$ .

An immediate consequence is that we can use our quantum sparsification algorithm to speed up Laplacian solving. We can create a sparsifier  $H$  with  $\tilde{O}(n/\epsilon^2)$  edges in time  $\tilde{O}(\sqrt{mn}/\epsilon)$ , and then use a classical Laplacian solver to solve the system  $L_H x = b$  in an additional time  $\tilde{O}(n/\epsilon^2)$ . This proves the proposition below.

**PROPOSITION 7.6** (quantum Laplacian solver). *There exists a quantum algorithm that, given adjacency-list access to a weighted and undirected graph  $G$ , outputs with high probability an approximate solution  $\tilde{x}$  to the linear system  $L_G x = b$  such that  $\|\tilde{x} - x\|_{L_G} \leq \epsilon \|x\|_{L_G}$  in time  $\tilde{O}(\sqrt{mn}/\epsilon)$ .*

The use of the  $L_G$ -induced norm  $\|\cdot\|_{L_G}$  is common in the study of Laplacian solvers. If, however, an  $\epsilon$ -approximate solution in the regular 2-norm is desired, then one can use that  $\|\tilde{x} - x\|_{L_G} \leq \epsilon \|x\|_{L_G}$  implies that  $\|\tilde{x} - x\| \leq \epsilon \sqrt{\kappa} \|x\|$ , with  $\kappa = \lambda_{\max}/\lambda_{\min}$  the condition number of  $L_G$ . Setting  $\epsilon = \delta/\sqrt{\kappa}$  hence yields a  $\delta$ -approximation in the 2-norm, and the runtime of the quantum Laplacian solver becomes  $\tilde{O}(\sqrt{mn\kappa}/\delta)$ . This factor- $\sqrt{\kappa}$  overhead would typically be too large, apart from the case where the Laplacian is well-conditioned. A resolution to this issue follows from improving the error-dependence of our quantum Laplacian solver down to polylogarithmic, which is one of the open questions that we mentioned in the introduction.

**7.2.1. Solving SDD systems.** Laplacian systems are special cases of the more general class of *real, symmetric, weakly diagonally dominant* (SDD) systems  $Ax = b$ , where  $A$  is such that

$$A = A^T \quad \text{and} \quad A_{ii} \geq \sum_{j \neq i} |A_{ij}| \quad \forall i \in [n].$$

A Laplacian system  $Lx = b$  is the special case of an SDD system where the off-diagonals of  $L$  are nonpositive, and its row sums are equal to zero. It is well known [48] that classically solving SDD systems can be reduced in time  $O(m)$  to solving a Laplacian system. This implies that the near-linear time Laplacian solver in [97] yields a near-linear time solver for the more general case of SDD systems as well.

Since we are aiming for a sublinear runtime in the quantum case, this classical reduction is too costly. However, taking inspiration from the classical reduction, we show in Appendix C that our quantum sparsification algorithm can approximately reduce an SDD system to a much sparser SDD system in time  $\tilde{O}(\sqrt{mn}/\epsilon)$ . Using a classical near-linear time SDD solver [97] on this sparser SDD system then yields the following proposition. As is customary in quantum linear algebra routines [44], we assume *sparse access* to the matrix  $A$ : given a query  $(i, k) \in [n]^2$ , this returns the index and value  $(j, A_{ij}) \in [n] \times \mathbb{R}$  of the  $k$ th nonzero element of the  $i$ th row  $A_i$ . (and an error symbol if there are less than  $k$  nonzero elements).

**PROPOSITION 7.7** (quantum SDD solver). *There exists a quantum algorithm that, given sparse access to an SDD matrix  $A$ , outputs with high probability an approximate solution  $\tilde{x}$  to the linear system  $Ax = b$  such that  $\|\tilde{x} - x\|_A \leq \epsilon \|x\|_A$  in time  $\tilde{O}(\sqrt{mn}/\epsilon)$ .*

**7.2.2. Effective resistances and commute times.** Electrical networks, consisting of nodes  $\{v \in V\}$  and resistors  $\{r_e | e \in E\}$ , are conveniently described by a weighted graph  $G = (V, E, \{w_e = 1/r_e\})$  [19]. Certain quantities of the electrical network can then be expressed using the Laplacian of  $G$ . One example is the effective resistance  $R_{s,t}$  between a pair of nodes  $s$  and  $t$ , which we already encountered in section 4. This can be expressed as a quadratic form in the inverse of the Laplacian:

$$R_{s,t} = (\chi_s - \chi_t)^T L^+ (\chi_s - \chi_t).$$

This effectively measures the *dissipated power*  $\mathcal{E}(j_{s,t})$  of the electric flow that results from injecting a unit current in  $s$ , and extracting it from  $t$ , as is described by the demand vector  $j_{s,t} = \chi_s - \chi_t$ . For a more general demand vector  $j \in \mathbb{R}^n$ , with  $1^T j = \sum_{v \in V} j(v) = 0$ , the dissipated power is defined analogously:

$$\mathcal{E}(j) = j^T L^+ j.$$

Closely related to the effective resistance is the *commute time* of a random walk. The random walk commute time  $C_{s,t}$  between nodes  $s$  and  $t$  is defined as the expected number of steps the walk must take from  $s$  to reach  $t$ , and then return to  $s$ . By a result of Chandra et al. [26] we know that  $C_{s,t} = 2WR_{s,t}$ , with  $W = \sum_e w_e$  the total edge weight in the graph.

Since the effective resistance and the dissipated power correspond to quadratic forms in the inverse of the Laplacian, they can be  $\epsilon$ -approximated by calculating the corresponding quantity in an  $\epsilon$ -spectral sparsifier. In addition, the total edge weight  $W_H$  of an  $\epsilon$ -cut sparsifier  $\epsilon$ -approximates the original edge weight  $W$ , so that  $(1-\epsilon)^2 C_{s,t}^G \leq C_{s,t}^H \leq (1+\epsilon)^2 C_{s,t}^G$  and hence also the commute times are approximated. Using our quantum sparsification algorithm, together with a classical Laplacian solver, we can approximate any of these quantities in time  $\tilde{O}(\sqrt{mn}/\epsilon)$ .

**CLAIM 7.8.** *Let  $j \in \mathbb{R}^n$  be a current demand vector, with  $1^T j = 0$ . Then there exists a quantum algorithm that outputs an  $\epsilon$ -approximation to the dissipated power  $\mathcal{E}(j)$  in time  $\tilde{O}(\sqrt{mn}/\epsilon)$ . In particular, if  $j = \chi_s - \chi_t$ , then this yields an  $\epsilon$ -approximation of the effective resistance  $R_{s,t}$  and the commute time  $C_{s,t}$ .*

By a similar argument, we can create an  $\epsilon$ -spectral sparsifier and construct the approximate resistance oracle of Spielman and Srivastava (see section 4.1.1) on this sparsifier. By Theorem 4.1 we can construct this oracle in time  $\tilde{O}(n/\epsilon^4)$  (which corresponds to  $\tilde{O}(m/\epsilon^2)$  when we input a sparsifier with  $m \in \tilde{O}(n/\epsilon^2)$  edges). This proves the following claim.

**CLAIM 7.9.** *There exists an  $\tilde{O}(\sqrt{mn}/\epsilon + n/\epsilon^4)$ -time quantum algorithm that outputs a  $(24 \log(n)/\epsilon^2) \times n$  matrix  $Z$  such that with high probability, for any  $s, t \in V$  it holds that*

$$(1-\epsilon)R_{s,t} \leq \|Z(\chi_s - \chi_t)\|^2 \leq (1+\epsilon)R_{s,t}.$$

After creating the matrix  $Z$ , we can hence  $\epsilon$ -approximate *any* effective resistance  $R_{s,t}$  in time  $\tilde{O}(1/\epsilon^2)$  simply by calculating the norm of the difference between two  $\tilde{O}(1/\epsilon^2)$ -length columns of  $Z$ . Apart from its use for spectral sparsification, as we demonstrated in section 4, such an oracle can also be used to obtain geometric embeddings of the graph; see, for instance, [103].

**7.2.3. Cover time.** The cover time  $\tau_{\text{cov}}(G)$  of a weighted graph  $G$  denotes the expected number of steps before a random walk has visited all nodes, starting from the worst initial node. A classic bound on the cover time called *Matthew's bound* [79] states that

$$\max_{s,t} H_{s,t} \leq \tau_{\text{cov}}(G) \leq (1 + \log(n)) \max_{s,t} H_{s,t},$$

with  $H_{s,t}$  the hitting time from node  $s$  to node  $t$ . Tighter characterizations were later proven by Kahn et al. [59] and Ding, Lee, and Peres [39]. We extract the following claims from the latter.

THEOREM 7.10 ([39, Theorems 1.6 and 4.14]).

- If  $H$  is an  $O(1)$ -spectral sparsifier of  $G$ , then  $\tau_{\text{cov}}(H) \in \Theta(\tau_{\text{cov}}(G))$ .
- There is an algorithm that outputs in time  $\tilde{O}(m)$  with high probability an  $O(1)$ -approximation of  $\tau_{\text{cov}}(G)$ .

We can easily derive the following claim by using our quantum sparsification algorithm to construct an  $O(1)$ -spectral sparsifier, and then approximating the cover time on this sparsifier using the algorithm from [39].

CLAIM 7.11. *Let  $G$  be an unweighted, undirected graph. There exists a quantum algorithm that outputs a constant-factor approximation to the cover time  $\tau_{\text{cov}}(G)$  in time  $\tilde{O}(\sqrt{mn})$ .*

**7.2.4. Eigenvalues and spectral clustering.** The bottom eigenvalues and eigenvectors of a graph Laplacian provide useful information about the graph, as is shown by, e.g., Cheeger’s inequality and spectral clustering [80], the PageRank algorithm [21], and in fact the entire field of spectral graph theory [32, 37]. Laplacian solvers allow one to efficiently approximate these eigenvalues and eigenvectors. Spielman and Teng [97] (with a later refinement by Koutis, Levin, and Peng [70]) showed that in time  $\tilde{O}(m + kn/\epsilon^2)$  it is possible to compute (i) an  $\epsilon$ -approximation to the  $k$  smallest eigenvalues  $\lambda_1, \dots, \lambda_k$  of a Laplacian, and (ii) a set of  $k$  orthogonal unit vectors  $v_1, \dots, v_k$  such that

$$(7.1) \quad v_\ell^T L v_\ell \leq (1 + \epsilon)\lambda_\ell, \quad 1 \leq \ell \leq k.$$

This set of vectors approximates the subspace spanned by the  $k$  bottom eigenvectors of the Laplacians. Already for constant  $\epsilon > 0$ , such a set can be used for spectral clustering. For the case of two clusters, this is explicitly discussed in [95] and [97, section 7] for RatioCut. For  $k$  clusters, the most common approach is spectral  $k$ -means clustering [80, 103]. Using the same analysis as in [85], one can show that a set obeying (7.1) can be used to obtain the same performance [106].

Using our quantum sparsification algorithm, we find a direct speedup for this task. To that end, note that it suffices to calculate the  $k$  smallest eigenvalues and approximate eigenvectors of an  $\epsilon'$ -spectral sparsifier for, say,  $\epsilon' = \epsilon/10$ . This will yield  $\epsilon$ -approximate eigenvalues and eigenvectors of the original graph; see, e.g., [97, Proposition 7.3]. We can hence construct an  $(\epsilon/10)$ -spectral sparsifier with  $\tilde{O}(n/\epsilon^2)$  edges in time  $\tilde{O}(\sqrt{mn}/\epsilon)$ , and then use a classical algorithm to solve the problem in the sparsifier, taking additional time  $\tilde{O}(kn/\epsilon^2)$ . This proves the following claim.

CLAIM 7.12. *There exists an  $\tilde{O}(\sqrt{mn}/\epsilon + kn/\epsilon^2)$ -time quantum algorithm that outputs with high probability an  $\epsilon$ -approximation of each of the  $k$  smallest eigenvalues and a set of orthogonal unit vectors  $v_1, \dots, v_k$  such that  $v_\ell^T L v_\ell \leq (1 + \epsilon)\lambda_\ell$  for all  $1 \leq \ell \leq k$ .*

This directly yields a speedup for the aforementioned spectral clustering algorithms, provided that we are given adjacency-list access to some similarity graph of the data [103]. Consider, for instance, the spectral  $k$ -means clustering algorithm in [80]. Given the similarity graph, its classical time complexity is dominated by (i) the time to construct a set of  $k$  vectors obeying (7.1) for constant  $\epsilon > 0$ , which is  $\tilde{O}(m + nk)$ , and (ii) the time to perform  $k$ -means clustering on these vectors, which is  $\tilde{O}(n\text{poly}(k))$ . This yields a total classical complexity  $\tilde{O}(m + n\text{poly}(k))$ . From the above claim, we immediately find the following corollary.

COROLLARY 7.13. *There exists a quantum algorithm that, given adjacency-list access to the similarity graph of a data set, performs spectral  $k$ -means clustering on this data set in time  $O(\sqrt{mn} + \text{npoly}(k))$ .*

### Appendix A. Quantum algorithm for shortest-path trees.

In this section we describe a quantum algorithm for constructing a single-source shortest-path tree (SPT), based<sup>4</sup> on the algorithm by Dürr et al. [41]. We slightly generalize the algorithm to the case where there is a set of “forbidden edges,” which we encode by associating a weight  $w_e = 0$  to these edges. The algorithm should then return an SPT over the connected component of the source. We recall that the cost of traversing an edge is described by its resistance  $1/w_e$ , and the distance  $\delta_G(u, v)$  between nodes  $u$  and  $v$  is

$$\delta_G(u, v) = \min_{u-v \text{ path } P} \sum_{e \in P} \frac{1}{w_e}.$$

The key routine that is used is a quantum algorithm for minimum-finding. It assumes oracle access to a “value function”  $f : [N] \rightarrow \mathbb{R} \cup \{\infty\}$  and a “type function”  $g : [N] \rightarrow \mathbb{N}$ , with a total number of types

$$M = |\text{Im}(g)| = |\{g(j) : j \in [N]\}|.$$

Given an integer  $d \in \mathbb{N}$  with  $d \leq N/2$ , the problem MINFIND  $(d, f, g)$  is to output a subset  $I \subseteq [N]$  of size  $|I| = \min\{d, M\}$ , such that

- $g(i) \neq g(j)$  for all distinct  $i, j \in I$ , and
- for all  $j \in [N] \setminus I$  and  $i \in I$ , if  $f(j) < f(i)$ , then  $f(i') \leq f(j)$  for some  $i' \in I$  with  $g(i') = g(j)$ .

PROPOSITION A.1 ([41, Theorem 3.4]). *There is a quantum algorithm to solve minfind  $(d, f, g)$  in time  $\tilde{O}(\sqrt{Nd} \log(1/\delta))$  with success probability at least  $1 - \delta$ . The algorithm requires  $O(\log N)$  qubits and a QRAM of  $\tilde{O}(d)$  bits.*

The quantum algorithm for constructing an SPT effectively implements Dijkstra’s algorithm, which we recall below for completeness. The algorithm assigns a “cost” to any (ordered) pair of nodes  $u, v$  by setting

$$\text{cost}(u, v) = \delta_G(v_0, u) + w(u, v).$$

We define the *component* of a node  $v_0$  as the smallest subset  $S \subseteq V$  such that  $v_0 \in S$  and either  $|E(S, V \setminus S)| = 0$  or  $\min\{w(u, v) \mid (u, v) \in E(S, V \setminus S)\} = \infty$ .

Starting from a node  $v_0$ , Dijkstra’s algorithm (Algorithm A.1) returns an SPT from  $v_0$  that spans the component of  $v_0$ . A key feature of the algorithm is that, at any time, the tree  $T$  is an SPT for the nodes in  $T$  (i.e.,  $\delta_T(v_0, u) = \delta_G(v_0, u)$  for all  $u \in T$ ). It follows that for any node  $u \in T$  we can efficiently evaluate  $\delta_G(v_0, u)$  and hence  $\text{cost}(u, v) = \delta_G(v_0, u) + 1/w(u, v)$ .

Now we describe a quantum SPT algorithm that is based on Dijkstra’s algorithm. It will have the same feature that at any time, the tree  $T$  is an SPT for the nodes in  $T$ .

<sup>4</sup> It was brought to our attention by J. van Apeldoorn and T. de Vos that there is an error in the original algorithm from [41]. Namely, in the proof of [41, Theorem 7.1] the cost function  $f((u, v)) = 1/w(u, v)$  is used, while this should have been  $f((u, v)) = \delta_G(v_0, u) + 1/w(u, v)$ . We here describe the corrected algorithm, which requires a more elaborate proof of correctness.

---

**Algorithm A.1 Dijkstra’s algorithm.**

---

```

1: let  $T = (V_T = \{v_0\}, E_T = \emptyset)$ 
2: while  $|V_T| < n - 1$  do
3:   find an edge  $(u, v) \in E(V_T, V \setminus V_T)$  with minimal  $\text{cost}(u, v)$ 
4:   if  $w(u, v) = \infty$  then
5:     abort and output  $T$ 
6:   else
7:     add  $v$  to  $V_T$  and  $(u, v)$  to  $E_T$ 
8:   end if
9: end while

```

---



---

**Algorithm A.2  $T = \text{SPT}(G, v_0)$ .**

---

```

1: let  $T = (V_T = \{v_0\}, E_T = \emptyset)$ ,  $L = 1$  and  $P_1 = \{v_0\}$ 
2: set  $\text{dist}(v_0) = 0$  and  $\text{dist}(u) = \infty$  for all  $u \neq v_0$ 
3: while  $|V_T| < n - 1$  do
4:   let  $B_L$  be the output of  $\text{MINFIND}(|P_L|, f, g)$  on  $E(P_L)$  where
      (i) type function  $g$  returns the end node  $g((u, v)) = v$  of an edge  $(u, v)$ ,
      (ii) value function  $f((u, v)) = \text{cost}(u, v) = \text{dist}(u) + 1/w(u, v)$  if  $u \in P_L$ ,
            $v \notin T$ , and  $f((u, v)) = \infty$  otherwise
5:   let  $(u, v) \in B_1 \cup \dots \cup B_L$  have minimal  $\text{cost}(u, v)$  with  $v \notin P_1 \cup \dots \cup P_L$ 
6:   if  $w(u, v) = 0$  then
7:     abort and output  $T$ 
8:   else
9:     add  $v$  to  $V_T$  and  $(u, v)$  to  $E_T$ , set  $P_{L+1} = \{v\}$  and  $L = L + 1$ 
10:    set  $\text{dist}(v) = \text{dist}(u) + 1/w(u, v)$ 
11:   end if
12:   as long as  $L \geq 2$  and  $|P_L| = |P_{L-1}|$ , merge  $P_L$  into  $P_{L-1}$  (removing  $P_L$ ) and
      set  $L = L - 1$ 
13: end while

```

---

First we prove correctness of the algorithm.

**PROPOSITION A.2.** *Let  $G = (V, E, w)$  be an undirected graph with weights  $w : E \rightarrow \mathbb{R}_{\geq 0}$ , and  $C_{v_0}$  the connected component of  $v_0 \in V$ . Then Algorithm **SPT** $(G, v_0)$  outputs an SPT from  $v_0$  that spans  $C_{v_0}$ .*

*Proof.* We choose the success probability of each call to **MINFIND** to be  $1 - \delta/n$ , so that the total success probability is at least  $(1 - \delta/n)^n \geq 1 - \delta$ . Assuming this, correctness of the algorithm follows from Dijkstra’s algorithm, provided that we can show that each iteration adds a least-cost border edge to  $T$ , if it exists. To this end, notice that in every iteration we have the following invariants directly after step 3:

1. The nodes of  $T$  are partitioned into subsets  $P_1, \dots, P_L$  whose sizes are non-increasing powers of 2, and such that  $|P_k| > \sum_{i=k+1}^L |P_i|$ .
2. For every  $1 \leq k \leq L$ , the set  $B_k$  contains up to  $|P_k|$  least-cost border edges from  $P_k$  to distinct nodes outside of  $P_k$ .

Now let  $T$  be the SPT constructed so far, let  $(u, v)$  be a least-cost border edge of  $T$ , and let  $P_1, \dots, P_L$  denote the node subsets directly after step 3. Without loss of

generality, assume that  $u \in P_k$ . Since the edge  $(u, v)$  is a least-cost border edge of  $T$ , it must also be a least-cost border edge of  $P_k$  from the subset

$$E_{\neq k} = \{(u', v') \mid u' \in P_k, v' \notin V(T) = P_1 \cup \dots \cup P_L\}.$$

Indeed, if there were another edge  $(\hat{u}, \hat{v}) \in E_{\neq k}$  with  $\text{cost}(\hat{u}, \hat{v}) < \text{cost}(u, v)$ , then  $(\hat{u}, \hat{v})$  would be a border edge of  $T$  with lower cost than  $(u, v)$ , which is a contradiction.

Now it suffices to show that there exists an edge  $(u', v') \in B_k$  with  $v' \notin T$  and  $\text{cost}(u', v') = \text{cost}(u, v)$  (potentially  $(u', v') = (u, v)$ ), because this would imply that step 5 indeed adds a least-cost border edge to  $T$ . Recall that  $B_k$  was constructed as the set of up to  $|P_k|$  least-cost border edges with distinct endpoints from the larger subset<sup>5</sup>

$$E_{>k} = \{(u', v') \mid u' \in P_k, v' \notin P_1 \cup \dots \cup P_k\}.$$

By the invariant property we have that  $|P_{k+1} \cup \dots \cup P_L| < |P_k|$ , and so the set  $B_k$  must contain at least one edge  $(u', v')$  with  $v' \notin P_{k+1} \cup \dots \cup P_L$  (and so  $v' \notin T$ ). This must be a least-cost edge from  $E_{\neq k}$ , and so  $\text{cost}(u', v') = \text{cost}(u, v)$ .  $\square$

It remains to prove that there is an efficient quantum algorithm for implementing **SPT**.

**PROPOSITION A.3.** *There is a quantum algorithm that implements **SPT**( $G, v_0$ ) with success probability  $1 - \delta$  in time  $\tilde{O}(\sqrt{|C_{v_0}|} |E(C_{v_0})| \log(n/\delta))$  and by making  $\tilde{O}(\sqrt{|C_{v_0}|} |E(C_{v_0})| \log(n/\delta))$  quantum queries to the adjacency list of  $G$ . The algorithm requires  $O(\log n)$  qubits and a QRAM of  $\tilde{O}(|C_{v_0}|)$  bits.*

*Proof.* The bound on the runtime follows from bounding the runtime of the calls to **MINFIND**, which are dominant. It follows from Proposition A.1 that a call of **MINFIND** ( $|P_L|, f, g$ ) requires time  $O(\sqrt{|E(P_L)|} |P_L| \log(n/\delta))$ . The total runtime is hence given by

$$O\left(\sum_{P_L} \sqrt{|E(P_L)|} |P_L| \log(n/\delta)\right),$$

where the sum runs over the sets  $P_L$  in step 3 of each iteration. We can bound this by noting that the merging procedure in step 9 ensures that in every iteration  $|P_L| = 2^{r_L}$  for some integer  $r_L \leq \log n$ , and any two  $P_L$  of the same size are necessarily disjoint. Since necessarily  $P_L \subseteq C_{v_0}$ , there are at most  $|C_{v_0}|/2^r$  (disjoint) sets of size  $2^r$ , and we can bound

$$\begin{aligned} \sum_{P_L: |P_L|=2^r} \sqrt{|E(P_L)|} |P_L| &= 2^{r/2} \sum_{P_L: |P_L|=2^r} \sqrt{|E(P_L)|} \\ &\leq 2^{r/2} \sqrt{2^{-r} |C_{v_0}| \sum_{P_L: |P_L|=2^r} |E(P_L)|} \leq \sqrt{2 |C_{v_0}| |E(C_{v_0})|}, \end{aligned}$$

where the second inequality follows from Cauchy–Schwarz and the third inequality from the fact that  $\sum_{P_L: |P_L|=2^r} |E(P_L)| \leq 2|E(C_{v_0})|$ . Summing this over all  $r \leq \log n$ , we find the claimed runtime.

Finally we bound the space complexity. The only quantum routine is running **MINFIND** ( $|P_L|, f, g$ ) on the set of edges, with  $|P_L| \leq n$ . By Proposition A.1 this requires  $O(\log n)$  qubits and a QRAM of  $\tilde{O}(n)$  bits. All remaining routines are classical

<sup>5</sup>Note that  $B_k$  was constructed at the iteration where  $P_k$  was “formed” (potentially by merging other subsets), and that the subsets  $P_1, \dots, P_k$  do not change between this iteration and the iteration that we are examining.

operations on the  $\tilde{O}(n)$  bits describing (i) the nodes and edges of the tree  $T$ , (ii) the sets  $B_1, \dots, B_L$ , and (iii) the string  $\text{dist}(\cdot)$ .  $\square$

**Appendix B. Existence of disjoint matching.**

Fix  $n, m \leq n^2/4$  and let  $\epsilon \geq \sqrt{n/m}$ . Let  $G_1 = (L_1 \cup R_1, E_1)$  consist of  $\epsilon^2 n/2$  disjoint copies  $B^{(k)}$  of the complete bipartite graph on  $1/\epsilon^2$  left and right nodes, containing  $2\epsilon^2 m/n$  parallel copies of every edge. In this way,  $G_1$  has  $n$  nodes and  $m$  edges and is  $2m/n$ -regular. We index the  $i$ th left and  $j$ th right node of  $B^{(k)}$  as  $l_i^{(k)}$ , respectively,  $r_j^{(k)}$ . Let  $G_2 = (L_2 \cup R_2, E_2)$  be the complete bipartite graph on  $2m/n$  left nodes and  $n/2$  right nodes. In this appendix we prove that every edge in  $G_1$  can be matched to a unique edge in  $G_2$  such that

1. all edges leaving a left node  $l \in L_1$  are matched to edges in  $G_2$  with distinct left ends;
2. all edges leaving a right node  $r \in R_1$  are matched to edges in  $G_2$  with distinct right ends.

We do this by considering maximum bipartite matchings in  $G_2$  of the form

$$M_j = \{(i, i + j) \mid i \in [2m/n]\} \subset E_2, \quad 0 \leq j < n/2,$$

where the sum is modulo  $n/2$ . These matchings form a partition of the edges set  $E_2$ . We interpret every  $M_j$  as a set of  $2m/n$  ordered edges. Now for every node  $l_i^{(k)} \in L_1$  we will match the  $2m/n$  (lexicographically ordered) outgoing edges  $E(l_i^{(k)})$  from that node to the (lexicographically ordered) edges in some  $M_j$ . This ensures that all edges leaving  $l_i^{(k)}$  are matched to edges in  $G_2$  whose left ends are distinct, so that condition 1 is satisfied. In order to satisfy condition 2, we specify the matching as follows:

$$E(l_i^{(k)}) \Leftrightarrow M_{k-1+(i-1)\epsilon^2 n/2} \quad \forall k \in [\epsilon^2 n/2], i \in [1/\epsilon^2].$$

Indeed, one can check that the  $2m/n$  edges matched to, for instance, the incoming edges of node  $r_1^{(1)}$  are described by

$$(\alpha, \beta), \quad \text{with } \alpha \in [1, 2\epsilon^2 m/n] \text{ and } \beta = \ell \frac{\epsilon^2 n}{2} + \alpha, \quad 0 \leq \ell < 1/\epsilon^2.$$

Since we assumed that  $m \leq n^2/4$ , and hence  $2m/n \leq n/2$ , the right ends of these edges are indeed disjoint. The same reasoning applies to all other nodes  $r_j^{(k)}$ . We illustrate this matching in Figure 5 below.

**Appendix C. Quantum SDD solver.**

In this appendix we describe a quantum algorithm for approximately solving an SDD linear system  $Ax = b$ . The blueprint of the algorithm is given below. An *SDDM matrix* is an SDD matrix whose off-diagonals are nonpositive.

---

**Algorithm C.1 qSDD( $A, x, b$ ).**

---

- 1: reduce the SDD system  $Ax = b$  to an SDDM system  $\hat{A}\hat{x} = \hat{b}$  using a classical trick [48] (see section C.1)
  - 2: construct an  $\epsilon$ -spectral sparsifier  $\tilde{A}$  of  $\hat{A}$  using our quantum sparsification algorithm (see section C.2)
  - 3: solve the sparse SDDM linear system  $\tilde{A}\hat{x} = \hat{b}$  using a *classical* SDD solver [97]
-

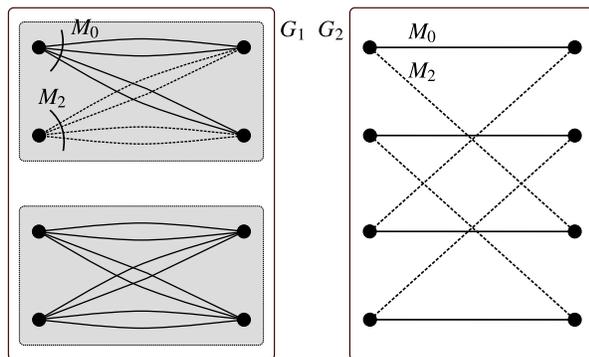


FIG. 5. Matching edges between  $G_1$  and  $G_2$  for  $n = 8$ ,  $m = 16$ , and  $\epsilon = 1/\sqrt{2}$ . The left ends of matchings  $M_0$  are distinct, ensuring condition 1 for node  $l_1^{(1)}$ . Similarly, the right ends of all edges matched to  $r_1^{(1)}$  are distinct, ensuring condition 2 for this node. We note that in general  $G_2$  is a bipartite graph between  $2m/n$  left nodes and  $n/2$  right nodes, with  $2m/n \leq n/2$ .

By the same reasoning as in Claim 7.5, it follows that an  $\epsilon$ -approximate solution of the SDDM system  $\hat{A}\hat{x} = \hat{b}$  yields a  $2\epsilon$ -approximate solution to the system  $\hat{A}\hat{x} = \hat{b}$ . This proves correctness of the algorithm. The runtime is dominated by the second step, which takes time  $\tilde{O}(\sqrt{mn}/\epsilon)$  (see section C.2). Hence this describes an  $\epsilon$ -approximate quantum SDD solver with runtime  $\tilde{O}(\sqrt{mn}/\epsilon)$ , which proves Proposition 7.7.

**C.1. Classical SDD to SDDM reduction.**

Gremban [48] showed how to reduce an SDD system  $Ax = b$  to an SDDM system  $\hat{A}\hat{x} = \hat{b}$ . Let  $A = D + P + N$  be an SDD matrix, with the diagonal matrix  $D$  containing the diagonal elements, and  $N$  and  $P$  containing the negative and positive off-diagonal elements, respectively. To this matrix we can associate an SDDM matrix  $\hat{A} \in \mathbb{R}^{2n \times 2n}$  and a vector  $\hat{b} \in \mathbb{R}^{2n \times 1}$ , which we define as

$$\hat{A} = \begin{bmatrix} D + N & -P \\ -P & D + N \end{bmatrix}, \quad \hat{b} = \begin{bmatrix} b \\ -b \end{bmatrix}.$$

Now Gremban showed that if the  $2n$ -dimensional vector  $\hat{z} = [\hat{z}_1 \hat{z}_2]^T$  is an  $\epsilon$ -approximate solution of the linear system  $\hat{A}\hat{x} = \hat{b}$ , then the  $n$ -dimensional vector  $z = (\hat{z}_1 - \hat{z}_2)/2$  is an  $\epsilon$ -approximate solution of the original system  $Ax = b$ . More precisely,

$$\|\hat{z} - \hat{A}^+ \hat{b}\|_{\hat{A}} \leq \epsilon \|\hat{A}^+ \hat{b}\|_{\hat{A}} \quad \Rightarrow \quad \|z - A^+ b\|_A \leq \epsilon \|A^+ b\|_A.$$

This implies that it suffices to solve the SDDM system  $\hat{A}\hat{x} = \hat{b}$ .

Given sparse access to  $A$ , we now show how to simulate sparse access to  $\hat{A}$ . The latter takes as input a pair  $(i, k) \in [2n]^2$ . If  $i \leq n$ , let  $(j, A_{ij})$  denote the output of an  $(i, k)$ -query to the original matrix  $A$ . If this yields an error symbol, output the error symbol. Otherwise, do the following:

$$(i, k) \mapsto \begin{cases} (j, A_{ij}) & \text{if } j = i \text{ or } A_{ij} < 0, \\ (j + n, -A_{ij}) & \text{otherwise.} \end{cases}$$

If  $i \geq n + 1$ , let  $(j, A_{ij})$  denote the output of an  $(i - n, k)$ -query to  $A$ . If this yields an error symbol, return the error symbol, otherwise do the following:

$$(i, k) \mapsto \begin{cases} (j + n, A_{ij}) & \text{if } j = i - n \text{ or } A_{ij} < 0, \\ (j, -A_{ij}) & \text{otherwise.} \end{cases}$$

This proves that we can simulate sparse access to  $\hat{A}$  using a single query to  $A$ . Note that the relative ordering between elements in the induced adjacency list for  $\hat{A}$  may be different from the ordering in the adjacency list for  $A$ ; that is not a problem because we are not assuming the adjacency lists to be ordered in any particular way.

**C.2. Sparsifying SDDM matrices.**

By the previous section, it suffices to solve an SDDM system  $\hat{A}\hat{x} = \hat{b}$  to which we have sparse access. We will do so by using our quantum algorithm to construct a much sparser SDDM matrix  $\tilde{A}$  with respect to which we can solve linear systems classically. If we can ensure that<sup>6</sup>

$$\tilde{A} \approx_\epsilon \hat{A},$$

then an  $\epsilon$ -approximate solution of  $\tilde{A}\hat{x} = \hat{b}$  will be an  $O(\epsilon)$ -approximate solution to  $\hat{A}\hat{x} = \hat{b}$ .

Since  $\hat{A}$  is SDDM, we can rewrite it as  $\hat{A} = \hat{L} + \hat{D}$ , with  $\hat{L}$  a Laplacian and  $\hat{D} = \text{diag}(\hat{A}) - \text{diag}(\hat{L})$  a nonnegative, diagonal “excess”-matrix. Given sparse access to  $\hat{A}$ , we can easily simulate adjacency-list access to the graph associated to  $\hat{L}$ . Using our sparsification algorithm, this allows us to explicitly output a sparsified Laplacian  $\tilde{L}$  with  $\tilde{O}(n/\epsilon^2)$  nonzero entries such that

$$\tilde{L} \approx_\epsilon \hat{L}.$$

Now we define the matrix

$$\tilde{A} = \tilde{L} + \text{diag}(\hat{A}) - \text{diag}(\tilde{L}).$$

This matrix has  $\tilde{O}(n/\epsilon^2)$  nonzero entries, of which we have an explicit description. We can prove that  $\tilde{A}$  closely approximates  $\hat{A}$ , as follows.

LEMMA C.1. *Let  $\tilde{A}$  be as described above. Then  $\tilde{A} \approx_{2\epsilon} \hat{A}$ .*

*Proof.* First we use the fact that if  $A, B$ , and  $C$  are positive semidefinite matrices and  $A \approx_\epsilon B$ , then  $A + C \approx_\epsilon B + C$ . Since  $\hat{L} \approx_\epsilon \tilde{L}$ , this shows that  $\hat{A} = \hat{L} + \hat{D} \approx_\epsilon \tilde{L} + \hat{D}$ . As a special case, this implies that  $\text{diag}(\hat{A}) = \text{diag}(\hat{L}) + \hat{D} \approx_\epsilon \text{diag}(\tilde{L}) + \hat{D}$ . Now we can simply rewrite

$$\tilde{A} = \tilde{L} + \hat{D} + \text{diag}(\hat{L}) - \text{diag}(\tilde{L}) \approx_\epsilon \tilde{L} + \hat{D} + \text{diag}(\tilde{L}) - \text{diag}(\tilde{L}) = \tilde{L} + \hat{D} \approx_\epsilon \hat{A}.$$

Now it suffices to note that if  $\tilde{A} \approx_\epsilon B \approx_\epsilon \hat{A}$ , then  $\tilde{A} \approx_{2\epsilon} \hat{A}$ . □

We can hence solve the sparse linear system  $\tilde{A}\hat{x} = \hat{b}$  to yield an approximation of the original system. Since  $\tilde{A}$  has  $\tilde{O}(n/\epsilon^2)$  nonzero entries, we can do this efficiently using a classical near-linear time SDD solver [97].

**Acknowledgments.** We are very grateful to Aleksandrs Belovs and Troy Lee for proving a composition property of the adversary method for relational problems which we required for the proof of our lower bound, and to Joran van Apeldoorn and Tijn de Vos for spotting an error in the quantum shortest-path tree algorithm from [41], which we have fixed here. This work also benefited from discussions with André Chailloux, Shantanav Chakraborty, András Gilyén, Michael Kapralov, Robin

<sup>6</sup>We use the shorthand  $B \approx_\epsilon C$  to mean that  $(1 - \epsilon)B \preceq C \preceq (1 + \epsilon)\hat{B}$ .

Downloaded 01/06/23 to 192.16.191.136 . Redistribution subject to SIAM license or copyright; see https://pubs.siam.org/terms-privacy

Kothari, Anthony Leverrier, Christian Majenz, Christian Schaffner, Luca Trevisan, and Luca Zanetti. We also thank the anonymous FOCS referees for helpful comments.

## REFERENCES

- [1] D. AHARONOV AND L. ZHOU, *Hamiltonian sparsification and gap-simulation*, in Proceedings of the 10th Innovations in Theoretical Computer Science Conference (ITCS), ACM, 2019, 2.
- [2] R. AHMED, G. BODWIN, F. D. SAHNEH, K. HAMM, M. J. JEBELLI, S. KOBOUROV, AND R. SPENCE, *Graph Spanners: A Tutorial Review*, preprint, <https://arxiv.org/abs/1909.03152>, 2019.
- [3] E. AÏMEUR, G. BRASSARD, AND S. GAMBS, *Quantum clustering algorithms*, in Proceedings of the 24th International Conference on Machine Learning (ICML), ACM, 2007, pp. 1–8, <https://doi.org/10.1145/1273496.1273497>.
- [4] I. ALTHÖFER, G. DAS, D. DOBKIN, D. JOSEPH, AND J. SOARES, *On sparse spanners of weighted graphs*, *Discrete Comput. Geom.*, 9 (1993), pp. 81–100, <https://doi.org/10.1007/BF02189308>.
- [5] A. AMBAINIS, *Variable time amplitude amplification and quantum algorithms for linear algebra problems*, in Proceedings of the 29th Symposium on Theoretical Aspects of Computer Science (STACS), Leibniz International Proceedings in Informatics (LIPIcs), 2012, pp. 636–647.
- [6] A. ANDONI, J. CHEN, R. KRAUTHGAMER, B. QIN, D. P. WOODRUFF, AND Q. ZHANG, *On sketching quadratic forms*, in Proceedings of the 2016 Innovations in Theoretical Computer Science Conference (ITCS), ACM, 2016, pp. 311–319.
- [7] A. ANDONI, R. KRAUTHGAMER, AND Y. POGROW, *On solving linear systems in sublinear time*, in Proceedings of the 10th Innovations in Theoretical Computer Science Conference (ITCS), ACM, 2019, 3.
- [8] S. APERS, A. AUZA, AND T. LEE, *A Sublinear Time Quantum Algorithm for  $s$ - $t$  Minimum Cut on Dense Simple Graphs*, preprint, <https://arxiv.org/abs/2110.15587>, 2021.
- [9] S. APERS AND T. LEE, *Quantum complexity of minimum cut*, in Proceedings of the 36th IEEE Conference on Computational Complexity (CCC), Leibniz International Proceedings in Informatics (LIPIcs), 2021, 28.
- [10] S. ARORA AND S. KALE, *A combinatorial, primal-dual approach to semidefinite programs*, *J. ACM*, 63 (2016), 12, <https://doi.org/10.1145/2837020>.
- [11] S. ARORA, S. RAO, AND U. VAZIRANI, *Expander flows, geometric embeddings and graph partitioning*, *J. ACM*, 56 (2009), 5, <https://doi.org/10.1145/1502793.1502794>.
- [12] B. AXELROD, Y. P. LIU, AND A. SIDFORD, *Near-optimal approximate discrete and continuous submodular function minimization*, in Proceedings of the 14th ACM-SIAM Symposium on Discrete Algorithms (SODA), SIAM, 2020, pp. 837–853, <https://doi.org/10.1137/1.9781611975994.51>.
- [13] N. BANSAL, O. SVENSSON, AND L. TREVISAN, *New notions and constructions of sparsification for graphs and hypergraphs*, in Proceedings of the 60th IEEE Symposium on Foundations of Computer Science (FOCS), IEEE, 2019, pp. 910–928.
- [14] J. BATSON, D. A. SPIELMAN, AND N. SRIVASTAVA, *Twice-Ramanujan sparsifiers*, *SIAM J. Comput.*, 41 (2012), pp. 1704–1721, <https://doi.org/10.1137/090772873>.
- [15] J. BATSON, D. A. SPIELMAN, N. SRIVASTAVA, AND S.-H. TENG, *Spectral sparsification of graphs: Theory and algorithms*, *Commun. ACM*, 56 (2013), pp. 87–94, <https://doi.org/10.1145/2492007.2492029>.
- [16] R. BEALS, H. BUHRMAN, R. CLEVE, M. MOSCA, AND R. DE WOLF, *Quantum lower bounds by polynomials*, *J. ACM*, 48 (2001), pp. 778–797.
- [17] A. BELOVS AND T. LEE, *The Quantum Query Complexity of Composition with a Relation*, preprint, <https://arxiv.org/abs/2004.06439>, 2020.
- [18] A. A. BENCZÚR AND D. R. KARGER, *Approximating  $s - t$  minimum cuts in  $\tilde{O}(n^2)$  time*, in Proceedings of the 28th ACM Symposium on Theory of Computing (STOC), ACM, 1996, pp. 47–55, <https://doi.org/10.1145/237814.237827>.
- [19] B. BOLLOBÁS, *Modern Graph Theory*, Grad. Texts in Math. 184. Springer Science & Business Media, 2013, <https://doi.org/10.1007/978-1-4612-0619-4>.
- [20] F. G. S. L. BRANDÃO, R. KUENG, AND D. STILCK FRANÇA, *Faster Quantum and Classical SDP Approximations for Quadratic Binary Optimization*, preprint, <https://arxiv.org/abs/1909.04613>, 2019.
- [21] S. BRIN AND L. PAGE, *The anatomy of a large-scale hypertextual web search engine*, *Comput. Networks ISDN*, 30 (1998), pp. 107–117, [https://doi.org/10.1016/S0169-7552\(98\)00110-X](https://doi.org/10.1016/S0169-7552(98)00110-X).

- [22] H. BUHRMAN, L. FORTNOW, I. NEWMAN, AND H. RÖHRIG, *Quantum property testing*, SIAM J. Comput., 37 (2008), pp. 1387–1400, <https://doi.org/10.1137/S0097539704442416>.
- [23] C. CADE, F. LABIB, AND I. NIESEN, *Quantum Motif Clustering*, preprint, <https://arxiv.org/pdf/2111.13222.pdf>, 2021.
- [24] L. J. CARTER AND M. N. WEGMAN, *Universal classes of hash functions*, J. Comput. Syst. Sci., 18 (1979), pp. 143–154, [https://doi.org/10.1016/0022-0000\(79\)90044-8](https://doi.org/10.1016/0022-0000(79)90044-8).
- [25] S. CHAKRABORTY, A. GILYÉN, AND S. JEFFERY, *The power of block-encoded matrix powers: Improved regression techniques via faster Hamiltonian simulation*, in Proceedings of the 46th International Colloquium on Automata, Languages, and Programming (ICALP), Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2019, 33.
- [26] A. K. CHANDRA, P. RAGHAVAN, W. L. RUZZO, R. SMOLENSKY, AND P. TIWARI, *The electrical resistance of a graph captures its commute and cover times*, Comput. Complex, 6 (1996), pp. 312–340, <https://doi.org/10.1007/BF01270385>.
- [27] L. CHEN, R. KYNG, Y. P. LIU, R. PENG, M. PROBST GUTENBERG, AND S. SACHDEVA, *Maximum flow and minimum-cost flow in almost-linear time*, in Proceedings of the 63rd IEEE Symposium on Foundations of Computer Science (FOCS), IEEE, 2022.
- [28] Y. CHEN AND R. DE WOLF, *Quantum Algorithms and Lower Bounds for Linear Regression with Norm Constraints*, preprint, <https://arxiv.org/abs/2110.13086>, 2021.
- [29] A. M. CHILDS, R. KOTHARI, AND R. D. SOMMA, *Quantum algorithm for systems of linear equations with exponentially improved dependence on precision*, SIAM J. Comput., 46 (2017), pp. 1920–1950, <https://doi.org/10.1137/16M1087072>.
- [30] T. CHRISTIANI, R. PAGH, AND M. THORUP, *From independence to expansion and back again*, in Proceedings of the 47th ACM Symposium on Theory of Computing (STOC), ACM, 2015, pp. 813–820.
- [31] P. CHRISTIANO, J. A. KELNER, A. MADRY, D. A. SPIELMAN, AND S.-H. TENG, *Electrical flows, Laplacian systems, and faster approximation of maximum flow in undirected graphs*, in Proceedings of the 43rd ACM Symposium on Theory of Computing (STOC), ACM, 2011, pp. 273–282.
- [32] F. R. CHUNG, *Spectral Graph Theory*, American Mathematical Society, 1997, <https://doi.org/10.1090/cbms/092>.
- [33] M. B. COHEN, J. A. KELNER, J. PEEBLES, R. PENG, A. SIDFORD, AND A. VLADU, *Faster algorithms for computing the stationary distribution, simulating random walks, and more*, in Proceedings of the 57th IEEE Symposium on Foundations of Computer Science (FOCS), IEEE, 2016, pp. 583–592.
- [34] M. B. COHEN, Y. T. LEE, C. M. MUSCO, C. P. MUSCO, R. PENG, AND A. SIDFORD, *Uniform sampling for matrix approximation*, in Proceedings of the 6th Innovations in Theoretical Computer Science Conference (ITCS), ACM, 2015, pp. 181–190.
- [35] M. B. COHEN, A. MADRY, D. TSIPRAS, AND A. VLADU, *Matrix scaling and balancing via box constrained Newton’s method and interior point methods*, in Proceedings of the 58th IEEE Symposium on Foundations of Computer Science (FOCS), IEEE, 2017, pp. 902–913.
- [36] T. M. COVER AND J. A. THOMAS, *Elements of Information Theory*, John Wiley & Sons, 2012.
- [37] D. M. CVETKOVIC, M. DOOB, AND H. SACHS, *Spectra of Graphs*, Academic Press, New York, 1980.
- [38] A. DASKIN, *Quantum spectral clustering through a biased phase estimation algorithm*, TWMS J. Appl. Engrg. Math., 1 (2020), pp. 24–33.
- [39] J. DING, J. R. LEE, AND Y. PERES, *Cover times, blanket times, and majorizing measures*, in Proceedings of the 43rd annual ACM symposium on Theory of computing, ACM, 2011, pp. 61–70, <https://doi.org/10.4007/annals.2012.175.3.8>.
- [40] P. DRINEAS, M. W. MAHONEY, AND S. MUTHUKRISHNAN, *Sampling algorithms for  $l_2$  regression and applications*, in Proceedings of the 17th ACM-SIAM Symposium on Discrete Algorithms (SODA), SIAM, 2006, pp. 1127–1136.
- [41] C. DÜRR, M. HEILIGMAN, P. HØYER, AND M. MHALLA, *Quantum query complexity of some graph problems*, SIAM J. Comput., 35 (2006), pp. 1310–1328, <https://doi.org/10.1137/050644719>.
- [42] T. EDEN AND W. ROSENBAUM, *Lower bounds for approximating graph parameters via communication complexity*, in Proceedings of APPROX-RANDOM, 2018, 11.
- [43] W.-S. FUNG, R. HARIHARAN, N. J. A. HARVEY, AND D. PANIGRAHI, *A general framework for graph sparsification*, SIAM J. Comput., 48 (2019), pp. 1196–1223, <https://doi.org/10.1137/16M1091666>.
- [44] A. GILYÉN, Y. SU, G. H. LOW, AND N. WIEBE, *Quantum singular value transformation and beyond: Exponential improvements for quantum matrix arithmetics*, in Proceedings of the 51st ACM Symposium on Theory of Computing (STOC), ACM, 2019, pp. 193–204.

- [45] M. X. GOEMANS AND D. P. WILLIAMSON, *Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming*, J. ACM, 42 (1995), pp. 1115–1145, <https://doi.org/10.1145/227683.227684>.
- [46] O. GOLDBREICH *Introduction to testing graph properties*, in Property Testing, Springer, 2010, pp. 105–141.
- [47] O. GOLDBREICH AND D. RON, *Property testing in bounded degree graphs*, Algorithmica, 32 (2002), pp. 302–343, <https://doi.org/10.1007/s00453-001-0078-7>.
- [48] K. D. GREMBAN, *Combinatorial Preconditioners for Sparse, Symmetric, Diagonally Dominant Linear Systems*, Ph.D. thesis, Carnegie Mellon University, 1996.
- [49] S. GRIBLING AND H. NIEUWBOER, *Improved quantum lower and upper bounds for matrix scaling*, in Proceedings of the 39th Symposium on Theoretical Aspects of Computer Science (STACS), 2021, 35.
- [50] L. K. GROVER, *A fast quantum mechanical algorithm for database search*, in Proceedings of the 28th ACM Symposium on Theory of Computing (STOC), ACM, 1996, pp. 212–219.
- [51] Y. HAMOUDI, P. REBENTROST, A. ROSMANIS, AND M. SANTHA, *Quantum and classical algorithms for approximate submodular function minimization*, Quant. Inform. Comput., 19 (2019), pp. 1325–1349.
- [52] A. W. HARROW, A. HASSIDIM, AND S. LLOYD, *Quantum algorithm for linear systems of equations*, Phys. Rev. Lett., 103 (2009), 150502.
- [53] S. HERBERT AND S. SUBRAMANIAN, *Spectral Sparsification of Matrix Inputs as a Preprocessing Step for Quantum Algorithms*, preprint, <https://arxiv.org/abs/1910.02861>, 2019.
- [54] A. S. HOLEVO, *Bounds for the quantity of information transmitted by a quantum communication channel*, Probl. Peredachi Inform.; English translation in Probl. Inform. Transmiss., 9 (1973), pp. 3–11.
- [55] P. HØYER, T. LEE, AND R. SPALEK, *Negative weights make adversaries stronger*, in Proceedings of the 39th ACM Symposium on Theory of Computing (STOC), ACM, 2007, pp. 526–535.
- [56] T. ITO AND S. JEFFERY, *Approximate span programs*, Algorithmica, 81 (2019), pp. 2158–2195.
- [57] A. JAMBULAPATI AND A. SIDFORD, *Ultrasparse ultrasparsifiers and faster Laplacian system solvers*, in Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA), SIAM, 2021, pp. 540–559, <https://doi.org/10.1137/1.9781611976465.33>.
- [58] M. JARRET, S. JEFFERY, S. KIMMEL, AND A. PIEDRAFITA, *Quantum algorithms for connectivity and related problems*, in Proceedings of the 26th European Symposium on Algorithms (ESA), Springer, 2018, 49.
- [59] J. KAHN, J. H. KIM, L. LOVASZ, AND V. H. VU, *The cover time, the blanket time, and the Matthews bound*, in Proceedings of the 41st IEEE Symposium on Foundations of Computer Science (FOCS), IEEE, 2000, pp. 467–475.
- [60] M. KAPRALOV, Y. T. LEE, C. M. MUSCO, C. P. MUSCO, AND A. SIDFORD, *Single pass spectral sparsification in dynamic streams*, SIAM J. Comput., 46 (2017), pp. 456–477, <https://doi.org/10.1137/141002281>.
- [61] M. KAPRALOV AND R. PANIGRAHY, *Spectral sparsification via random spanners*, in Proceedings of the 3rd Innovations in Theoretical Computer Science Conference (ITCS), ACM, 2012, pp. 393–398, <https://doi.org/10.1145/2090236.2090267>.
- [62] D. R. KARGER, *Using randomized sparsification to approximate minimum cuts*, in Proceedings of the 5th ACM-SIAM Symposium on Discrete Algorithms (SODA), SIAM, 1994, pp. 424–432.
- [63] D. R. KARGER, *Minimum cuts in near-linear time*, J. ACM, 47 (2000), pp. 46–76, <https://doi.org/10.1145/331605.331608>.
- [64] R. M. KARP, *Reducibility among combinatorial problems*, in Complexity of computer computations, Springer, 1972, pp. 85–103, [https://doi.org/10.1007/978-1-4684-2001-2\\_9](https://doi.org/10.1007/978-1-4684-2001-2_9).
- [65] J. A. KELNER, Y. T. LEE, L. ORECCHIA, AND A. SIDFORD, *An almost-linear-time algorithm for approximate max flow in undirected graphs, and its multicommodity generalizations*, in Proceedings of the 2014 ACM-SIAM Symposium on Discrete Algorithms (SODA), SIAM, 2014, pp. 217–226, <https://doi.org/10.1137/1.9781611973402.16>.
- [66] J. A. KELNER AND A. LEVIN, *Spectral sparsification in the semi-streaming setting*, Theoret. Comput. Syst., 53 (2013), pp. 243–262, <https://doi.org/10.1007/s00224-012-9396-1>.
- [67] I. KERENIDIS AND J. LANDMAN, *Quantum Spectral Clustering*, preprint, <https://arxiv.org/abs/2007.00280>, 2020.
- [68] I. KERENIDIS, J. LANDMAN, A. LUONGO, AND A. PRAKASH, *q-means: A quantum algorithm for unsupervised machine learning*, in Advances in Neural Information Processing Systems, pp. 4136–4146, 2019.
- [69] S. KHOT, G. KINDLER, E. MOSSEL, AND R. O’DONNELL, *Optimal inapproximability results for MAX-CUT and other 2-variable CSPs?*, SIAM J. Comput., 37 (2007), pp. 319–357, <https://doi.org/10.1137/S0097539705447372>.

- [70] I. KOUTIS, A. LEVIN, AND R. PENG, *Faster spectral sparsification and numerical algorithms for SDD matrices*, ACM Trans. Algorithms, 12 (2016), 17.
- [71] I. KOUTIS, G. L. MILLER, AND R. PENG, *Approaching optimality for solving SDD linear systems*, SIAM J. Comput., 43 (2014), pp. 337–354, <https://doi.org/10.1137/110845914>.
- [72] I. KOUTIS, G. L. MILLER, AND D. TOLLIVER, *Combinatorial preconditioners and multilevel solvers for problems in computer vision and image processing*, Comput. Vis. Image Understand., 115 (2011), pp. 1638–1646, [https://doi.org/10.1007/978-3-642-10331-5\\_99](https://doi.org/10.1007/978-3-642-10331-5_99).
- [73] I. KOUTIS AND S. C. XU, *Simple parallel and distributed algorithms for spectral graph sparsification*, ACM Trans. Parallel Comput., 3 (2016), pp. 1–14.
- [74] Y. T. LEE, *Probabilistic Spectral Sparsification in Sublinear Time*, preprint, <https://arxiv.org/abs/1401.0085>, 2013.
- [75] T. LEIGHTON AND S. RAO, *Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms*, J. ACM, 46 (1999), pp. 787–832, <https://doi.org/10.1145/331524.331526>.
- [76] M. LI, G. L. MILLER, AND R. PENG, *Iterative row sampling*, in Proceedings of the 54th IEEE Symposium on Foundations of Computer Science (FOCS), IEEE, 2013, pp. 127–136.
- [77] S. LLOYD, M. MOHSENI, AND P. REBENTROST, *Quantum Algorithms for Supervised and Unsupervised Machine Learning*, preprint, <https://arxiv.org/abs/1307.0411>, 2013.
- [78] A. W. MARCUS, D. A. SPIELMAN, AND N. SRIVASTAVA, *Interlacing families II: Mixed characteristic polynomials and the Kadison-Singer problem*, Ann. of Math. (2), 182 (2015), pp. 327–350.
- [79] P. MATTHEWS, *Covering problems for Markov chains*, Ann. Probab., 16 (1988), pp. 1215–1228, <https://doi.org/10.1214/aop/1176991686>.
- [80] A. Y. NG, M. I. JORDAN, AND Y. WEISS, *On spectral clustering: Analysis and an algorithm*, in Advances in Neural Information Processing Systems, 2002, pp. 849–856.
- [81] M. A. NIELSEN AND I. CHUANG, *Quantum Computation and Quantum Information*, Cambridge University Press, 2002.
- [82] L. ORECCHIA, S. SACHDEVA, AND N. K. VISHNOI, *Approximating the exponential, the Lanczos method and an  $\tilde{O}(m)$ -time spectral algorithm for balanced separator*, in Proceedings of the 44th ACM Symposium on Theory of Computing (STOC), ACM, 2012, pp. 1141–1160.
- [83] D. PELEG, *Distributed Computing*, SIAM Monogr. Discrete Math. Appl. 5, SIAM, 2000, <https://doi.org/10.1137/1.9780898719772>.
- [84] R. PENG, *Approximate undirected maximum flows in  $O(m \text{ polylog}(n))$  time*, in Proceedings of the 2016 ACM-SIAM Symposium on Discrete Algorithms (SODA), SIAM, 2016, pp. 1862–1867, <https://doi.org/10.1137/1.9781611974331.ch130>.
- [85] R. PENG, H. SUN, AND L. ZANETTI, *Partitioning well-clustered graphs: Spectral clustering works!*, SIAM J. Comput., 46 (2017), pp. 710–743, <https://doi.org/10.1137/15M1047209>.
- [86] J.-C. PICARD AND M. QUEYRANNE, *Selected applications of minimum cuts in networks*, INFOR Inform. Syst. Oper. Res., 20 (1982), pp. 394–422, <https://doi.org/10.1080/03155986.1982.11731876>.
- [87] S. PIDDOCK, *Quantum Walk Search Algorithms and Effective Resistance*, preprint, <https://arxiv.org/abs/1912.04196>, 2019.
- [88] J. SHERMAN, *Breaking the multicommodity flow barrier for  $O(\sqrt{\log n})$ -approximations to sparsest cut*, in Proceedings of the 50th IEEE Symposium on Foundations of Computer Science (FOCS), IEEE, 2009, pp. 363–372.
- [89] J. SHERMAN, *Nearly maximum flows in nearly linear time*, in Proceedings of the 54th IEEE Symposium on Foundations of Computer Science (FOCS), IEEE, 2013, pp. 263–269.
- [90] D. B. SHMOYS, *Cut Problems and Their Application to Divide-and-conquer*, in Approximation Algorithms for NP-Hard Problems, PWS Publishing, 1997, pp. 192–235.
- [91] M. K. SILVA, N. J. A. HARVEY, AND C. M. SATO, *Sparse sums of positive semidefinite matrices*, ACM Trans. Algorithms, 12 (2016), 9.
- [92] T. SOMA AND Y. YOSHIDA, *Spectral sparsification of hypergraphs*, in Proceedings of the 2019 ACM-SIAM Symposium on Discrete Algorithms (SODA), SIAM, 2019, pp. 2570–2581, <https://doi.org/10.1137/1.9781611975482.159>.
- [93] D. A. SPIELMAN AND N. SRIVASTAVA, *Graph sparsification by effective resistances*, SIAM J. Comput., 40 (2011), pp. 1913–1926, <https://doi.org/10.1137/080734029>.
- [94] D. A. SPIELMAN AND S.-H. TENG, *Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems*, in Proceedings of the 36th ACM Symposium on Theory of Computing (STOC), ACM, 2004, pp. 81–90.
- [95] D. A. SPIELMAN AND S.-H. TENG, *Spectral partitioning works: Planar graphs and finite element meshes*, Linear Algebra Appl., 421 (2007), pp. 284–305, <https://doi.org/10.1016/j.laa.2006.07.020>.

- [96] D. A. SPIELMAN AND S.-H. TENG, *Spectral sparsification of graphs*, SIAM J. Comput., 40 (2011), pp. 981–1025, <https://doi.org/10.1137/08074489X>.
- [97] D. A. SPIELMAN AND S.-H. TENG, *Nearly linear time algorithms for preconditioning and solving symmetric, diagonally dominant linear systems*, SIAM J. Matrix Anal. Appl., 35 (2014), pp. 835–885, <https://doi.org/10.1137/090771430>.
- [98] S.-H. TENG, *The Laplacian paradigm: Emerging algorithms for massive graphs*, in Theory and Applications of Models of Computation (TAMC 2010), Springer, 2010, pp. 2–14, [https://doi.org/10.1007/978-3-642-13562-0\\_2](https://doi.org/10.1007/978-3-642-13562-0_2).
- [99] M. THORUP AND U. ZWICK, *Approximate distance oracles*, J. ACM, 52 (2005), pp. 1–24, <https://doi.org/10.1145/1044731.1044732>.
- [100] L. TREVISAN, *Max cut and the smallest eigenvalue*, SIAM J. Comput., 41 (2012), pp. 1769–1786, <https://doi.org/10.1137/090773714>.
- [101] J. VAN APELDOORN, S. GRIBLING, Y. LI, H. NIEUWBOER, M. WALTER, AND R. DE WOLF, *Quantum algorithms for matrix scaling and matrix balancing*, in Proceedings of the 48th International Colloquium on Automata, Languages, and Programming (ICALP), Leibniz International Proceedings in Informatics (LIPIcs), Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021, 110.
- [102] N. K. VISHNOI,  $Lx = b$ , Found. Trends Theoret. Comput. Sci., 8 (2013), pp. 1–141, <https://doi.org/10.1561/04000000054>.
- [103] U. VON LUXBURG, *A tutorial on spectral clustering*, Stat. Comput., 17 (2007), pp. 395–416.
- [104] G. WANG, *Efficient quantum algorithms for analyzing large sparse electrical networks*, Quantum Inform. Comput., 17 (2017), pp. 987–1026.
- [105] N. WIEBE, A. KAPOOR, AND K. SVORE, *Quantum algorithms for nearest-neighbor methods for supervised and unsupervised learning*, Quantum Inform. Comput., 15 (2015), pp. 316–356.
- [106] L. ZANETTI, *Private communication*, 2019.
- [107] M. ZHANDRY, *Secure identity-based encryption in the quantum random oracle model*, Internat. J. Quantum Inform., 13 (2015), 1550014, [https://doi.org/10.1007/978-3-642-32009-5\\_44](https://doi.org/10.1007/978-3-642-32009-5_44).
- [108] D. ZHOU, J. HUANG, AND B. SCHÖLKOPF, *Learning from labeled and unlabeled data on a directed graph*, in Proceedings of the 22nd International Conference on Machine Learning (ICML), ACM, 2005, pp. 1036–1043, <https://doi.org/10.1145/1102351.1102482>.
- [109] X. ZHU, Z. GHAHRAMANI, AND J. D. LAFFERTY, *Semi-supervised learning using Gaussian fields and harmonic functions*, in Proceedings of the 20th International Conference on Machine Learning (ICML), 2003, pp. 912–919.