


Space Complexity of Minimum Cut Problems in Single-Pass Streams

Matthew Ding   

University of California, Berkeley, CA, USA

Alexandro Garces 

Massachusetts Institute of Technology, Cambridge, MA, USA

Jason Li   

Carnegie Mellon University, Pittsburgh, PA, USA

Honghao Lin   

Carnegie Mellon University, Pittsburgh, PA, USA

Jelani Nelson   

University of California, Berkeley, CA, USA

Vihan Shah   

University of Waterloo, Canada

David P. Woodruff   

Carnegie Mellon University, Pittsburgh, PA, USA

Abstract

We consider the problem of finding a minimum cut of a weighted graph presented as a single-pass stream. While graph sparsification in streams has been intensively studied, the specific application of finding minimum cuts in streams is less well-studied. To this end, we show upper and lower bounds on minimum cut problems in insertion-only streams for a variety of settings, including for both randomized and deterministic algorithms, for both arbitrary and random order streams, and for both approximate and exact algorithms. One of our main results is an $\tilde{O}(n/\varepsilon)$ space algorithm with fast update time for approximating a spectral cut query with high probability on a stream given in an arbitrary order. Our result breaks the $\Omega(n/\varepsilon^2)$ space lower bound required of a sparsifier that approximates all cuts simultaneously. Using this result, we provide streaming algorithms with near optimal space of $\tilde{O}(n/\varepsilon)$ for minimum cut and approximate all-pairs effective resistances, with matching space lower-bounds. The amortized update time of our algorithms is $\tilde{O}(1)$, provided that the number of edges in the input graph is at least $(n/\varepsilon^2)^{1+o(1)}$. We also give a generic way of incorporating sketching into a recursive contraction algorithm to improve the post-processing time of our algorithms. In addition to these results, we give a random-order streaming algorithm that computes the *exact* minimum cut on a simple, unweighted graph using $\tilde{O}(n)$ space. Finally, we give an $\Omega(n/\varepsilon^2)$ space lower bound for deterministic minimum cut algorithms which matches the best-known upper bound up to polylogarithmic factors.

2012 ACM Subject Classification Theory of computation \rightarrow Streaming, sublinear and near linear time algorithms

Keywords and phrases minimum cut, approximate, random order, lower bound

Digital Object Identifier 10.4230/LIPIcs.ITCS.2025.43

Related Version *Full Version:* <https://arxiv.org/abs/2412.01143> [17]

Funding *Matthew Ding:* Supported in part by NSF CCF-1951384.

Alexandro Garces: Supported in part by NSF CNS-2150186.

Jason Li: This work was done in part as a Research Fellow at the Simons Institute for the Theory of Computing.

Honghao Lin: Supported in part by a Simons Investigator Award, NSF CCF-2335412, and a CMU Paul and James Wang Sercomm Presidential Graduate Fellowship.



© Matthew Ding, Alexandro Garces, Jason Li, Honghao Lin, Jelani Nelson, Vihan Shah, and David P. Woodruff;

licensed under Creative Commons License CC-BY 4.0

16th Innovations in Theoretical Computer Science Conference (ITCS 2025).

Editor: Raghu Meka; Article No. 43; pp. 43:1–43:23



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Jelani Nelson: Supported in part by NSF CCF-1951384 and NSF CCF-2427808.

Vihan Shah: Supported in part by Sepehr Assadi’s Sloan Research Fellowship and NSERC Discovery Grant.

David P. Woodruff: Supported in part by a Simons Investigator Award and NSF CCF-2335412.

Acknowledgements The authors would like to thank the ITCS 2025 reviewers for their anonymous feedback. Alexandro Garces and Vihan Shah are extremely grateful to Sepehr Assadi for many helpful conversations throughout the project. They also thank the organizers of DIMACS REU in Summer 2023, in particular Lazaros Gallos, for initiating this collaboration and for all their help and encouragement along the way.

1 Introduction

We consider the graph streaming model, which is a key model for computations on massive graph datasets that has been extensively studied over the past couple of decades (see, e.g., [36] for a survey). We specifically study the problem of finding minimum cuts in adversarial and random-order streams, which has been less studied than the problem of general cut sparsification of graphs. Besides being theoretically interesting, finding minimum cuts is also a problem of practical interest. For example, they allow for the calculation of social network metrics such as influence [50] and are used to quantify the robustness of power networks [23] and road networks [48] to failures.

1.1 Adversarial Streams

The key method so far for solving minimum cut with low-memory is the usage of cut sparsifiers. The notion of cut sparsifiers was introduced by Benczur and Karger [7] and has been extremely influential. Given a weighted graph $G = (V, E, w)$ with $n = |V|$ vertices and $m = |E|$ edges, and polynomially bounded edge weights $w : E \rightarrow \mathbb{R}_+$, together with an accuracy parameter $\varepsilon > 0$, a cut sparsifier of G is a sparse subgraph H on the same vertex set V but with (possibly) different edge weights such that the weight of every cut in G is $(1 + \varepsilon)$ -approximated by the weight of the corresponding cut in H . For two sets $S, T \subseteq V$, we use $E(S, T) = \{(u, v) \in E : u \in S, v \in T\}$ to denote the set of edges between S and T in graph G and $w_G(S, T) = \sum_{e \in E(S, T)} w_e$ to denote the total weight of edges between S and T in graph G . Formally, we have the following definition,

► **Definition 1 (For-All Cut Sparsifier).** *H is a $(1 + \varepsilon)$ for-all cut sparsifier of G if and only if the following holds for all $\emptyset \subset S \subset V$:*

$$w_H(S, V \setminus S) \approx_\varepsilon w_G(S, V \setminus S)$$

where $a \approx_\varepsilon b$ is defined as $(1 - \varepsilon) \cdot b \leq a \leq (1 + \varepsilon) \cdot b$.

Specifically, [7] shows that a cut sparsifier always exists with $O(n \log n / \varepsilon^2)$ edges. This bound was improved to $O(n / \varepsilon^2)$ edges by [6] and [1, 2, 11] proved a lower bound of $\Omega(n(\log n) / \varepsilon^2)$ bits (even in the case when the cut-sparsifier H is a not necessarily a subgraph of G). The existence bound was also extended to the stronger notion of spectral sparsifiers [46, 47, 45, 6], where the quadratic form associated with the Laplacian of graph H provides a $(1 + \varepsilon)$ -approximation to that of G . Let L_G denote the Laplacian matrix of graph G . For every subset $S \subseteq V$, let x_S be the binary indicator vector of S such that $x_i = 1$ if $i \in S$ and $x_i = 0$ if $i \notin S$. Then we have $x_S^\top L_G x_S = w_G(S, V \setminus S)$. The definition of the spectral sparsifier extends the assumption that x is a binary vector to an arbitrary vector.

► **Definition 2** (For-All Spectral Sparsifier). *Let G and H be two weighted undirected graphs. Fix $0 < \varepsilon < 1$. We say H is a $(1 + \varepsilon)$ for-all spectral sparsifier of G iff the following holds:*

$$(1 - \varepsilon) \cdot L_G \preceq L_H \preceq (1 + \varepsilon) \cdot L_G$$

where we use $A \preceq B$ to denote $\forall x \in \mathbb{R}^n, x^\top Ax \leq x^\top Bx$.

This construction has had a tremendous impact on cut problems in graphs, see, e.g., [7, 8, 30, 44, 35]. However, for very small values of ε , the $1/\varepsilon^2$ dependence in cut sparsifiers may be prohibitive on large-scale graphs. Motivated by this, the work of [2] relaxed the cut sparsification problem to that of outputting a data structure sk_G such that given any *fixed* cut $S \subseteq V$, the value of $\text{sk}_G(S)$ is within a $(1 + \varepsilon)$ factor of the cut value of S in G , with probability at least $2/3$ ¹. This relaxed notion of the problem is called the “for-each” model, which should be contrasted with the previous “for-all” model of cut and spectral sparsification.

Surprisingly, [2] showed that such a data structure exists for cut sparsification with $\text{poly}(n)$ -bounded integer edge weights of size $\tilde{O}(n/\varepsilon)^2$, which is optimal up to polylogarithmic factors. In a follow-up work of [24], the authors extended the $\tilde{O}(n/\varepsilon)$ upper bound to the for-each model for spectral sparsifiers. Another very interesting work is [14], which shows that such a sketch can be chosen as a reweighted subgraph of G . We specifically deal with these graphical for-each sparsifiers:

► **Definition 3** (Graphical For-Each Spectral Sparsifier). *Let G and H be two weighted undirected graphs. Fix $0 < \varepsilon < 1$. We say H is a $(1 + \varepsilon)$ graphical for-each spectral sparsifier of G if and only if for each $x \in \mathbb{R}^n$, with probability at least $2/3$:*

$$x^\top L_H x \approx_\varepsilon x^\top L_G x$$

In a line of work, see, e.g., [36, 27, 28], efficient algorithms for sparsifiers in the for-all model in a graph stream were constructed. The state-of-the-art work [28] gives a single-pass algorithm for spectral sparsification in dynamic streams which uses space $\tilde{O}(n/\varepsilon^2)$.

► **Lemma 4** ([28]). *Given an input graph G in a stream, there exists a one-pass streaming algorithm that outputs a subgraph H of G such that with high probability H is a $(1 + \varepsilon)$ -spectral sparsifier of G . Moreover, the algorithm uses $\tilde{O}(n/\varepsilon^2)$ space and $\tilde{O}(m + n/\varepsilon^2)$ time.*

However, it was not known how to construct a for-each sketch for cut sparsification in a stream in better than $\tilde{O}(n/\varepsilon^2)$ space, which follows from computing a for-all sparsifier. A natural question is whether we can implement a streaming algorithm that computes a for-each sparsifier with space matching the $\Omega(n/\varepsilon)$ offline lower bound that holds for for-each sparsification [2].

If we can answer this question in the affirmative, we can also get a better space bound for the approximate minimum cut problem. It is well-known that there are only $\text{poly}(n)$ $O(1)$ -approximate minimum cuts³ in a graph. So we can run a for-all cut sparsification algorithm in a stream with $\varepsilon = \Theta(1)$ in $\tilde{O}(n)$ space and obtain a candidate set of $\text{poly}(n)$

¹ This can be amplified to high probability, which we define to mean with probability at least $1 - 1/\text{poly}(n)$. This is done by independently repeating the data structure $O(\log n)$ times and outputting the median estimate. This also lets us estimate the cut value for $\text{poly}(n)$ cuts simultaneously.

² Throughout we use $\tilde{O}(\cdot)$ to hide polylog(n) factors.

³ When we say a cut is an α -approximate minimum cut, we mean that its cut value is at most α times the minimum cut value.

cuts containing the minimum cut. In parallel, we estimate the value of each cut using a for-each sparsifier with failure probability $1/\text{poly}(n)$. The smallest cut out of the candidate cuts will be a $(1 + \varepsilon)$ -approximate minimum cut. Thus, the space complexity is just $\tilde{O}(n)$ plus that of maintaining a for-each sparsifier in a stream. However, the existing algorithms for computing for-each sparsifiers are not streaming algorithms. Motivated by the existence of a for-each sketch in $\tilde{O}(n/\varepsilon)$ bits of space, one may wonder if it is possible to obtain a $(1 + \varepsilon)$ -approximation to the minimum cut in $\tilde{O}(n/\varepsilon)$ space in a single-pass stream.

► **Question 5.** What is the space and time complexity of obtaining a $(1 + \varepsilon)$ -approximate minimum cut in a single-pass insertion-only stream?

We also note that an exact global minimum cut streaming algorithm exists in $\tilde{O}(n)$ space with two passes for simple graphs [5] (see also [43]) and $\log n$ passes for weighted graphs [38].

1.2 Random-order streams

We know that finding the minimum cut exactly in single-pass adversarial streams needs $\Omega(n^2)$ space [52]. We need an additional pass to get the exact minimum cut in $\tilde{O}(n)$ space. However, many streaming applications require solving the problem in exactly one pass because the stream cannot be stored.

Thus, to surpass the $\Omega(n^2)$ barrier of adversarial streams in one pass, we consider a relaxation known as random-order streams and ask whether we can find the minimum cut exactly in $o(n^2)$ space in this setting. In this model, the graph can still be chosen adversarially, but its edges arrive in the stream in a random order.

The motivation for studying this model stems from the fact that real-world data is typically not adversarial. Adversarial streams are adversarial in both their input and arrival order, while random-order streams relax the arrival order, making them more representative of real-world data. Additionally, the random-order streaming model may help explain the empirical performance of certain heuristics.

This model was first studied by [39] for sorting in limited space. Many subsequent works have studied problems in random-order streams [21, 26, 12, 37, 42, 16, 13], some of which show a clear separation between the adversarial and random-order models.

One notable line of work focuses on maximum matching within this model [33, 19, 32, 3, 18, 9, 4, 22]. The best-known upper bound in the random order model [4] gets a slightly better than $2/3$ -approximation in semi-streaming space, which is better than the best possible $\frac{1}{1+\ln 2} \approx 0.59$ -approximation in the adversarial model due to the lower bound of [25], showing another separation between adversarial and random-order streams.

Motivated by this, we study the minimum cut problem in the random-order model and ask the following question:

► **Question 6.** What is the space complexity of obtaining the exact minimum cut in a single-pass random-order stream?

This work resolves the above questions with matching algorithms and lower bounds. We also carefully optimize the running time and update time of our algorithms. We also obtain several additional results, which we outline in Section 3.

2 Preliminaries

2.1 All-Pairs Effective Resistances

Given $G(V, E, w)$ as an electrical network on n nodes in which each edge e corresponds to a link of conductance w_e (i.e., a resistor of resistance $1/w_e$), the effective resistance of an edge e is the potential difference induced across it when a unit current is injected at one end of e and extracted at the other end of e . Equivalently it is equal to $R_{\text{eff}}^G = \chi_{u,v}^\top L_G^\dagger \chi_{u,v}$, where L_G^\dagger denotes the pseudoinverse of the Laplacian of graph G and $\chi_{u,v} := \mathbf{1}_u - \mathbf{1}_v$.

In the all-pairs effective resistance problem, we are required to output a data structure, which with high probability can generate effective resistance values for each of the n^2 pairs of vertices, or a $(1+\varepsilon)$ -approximation of each of the n^2 pairs of values in the $(1+\varepsilon)$ -approximate version of the problem. Note that the data structure implicitly represents these n^2 values and responds with one such value on a given query pair.

2.2 Short-Cycle Decomposition

Short cycle decompositions are a recent algorithmic tool for graph sketching problems introduced by Chu et al. [14].

► **Definition 7** (Short-Cycle Decomposition). *An (\hat{m}, L) short-cycle decomposition of a graph G is a decomposition of the graph into edge-disjoint cycles of at most length L and an additional \hat{m} edges outside of these cycles.*

The main contribution of [14] to spectral sparsification is in proving the following claim:

▷ **Claim 8.** Given an undirected weighted graph G and an (\hat{m}, L) short-cycle decomposition routine `CYCLEDECOMP`, there exists an algorithm which returns with high probability a $(1+\varepsilon)$ -spectral sparsifier *graph* with $\tilde{O}(\hat{m} + nL\varepsilon^{-1})$ edges and run-time

$$\tilde{O}(m) + T_{\text{CYCLEDECOMP}}(O(m \log n), n)$$

where $T_{\text{CYCLEDECOMP}}(m, n)$ is the running time of `CYCLEDECOMP` on a graph with m edges and n vertices.

As a high-level overview, the short-cycle decomposition is key as a degree-preserving sparsification method. Given each cycle, we can label them numerically in order and sample either all the odd or all the even labeled edges, each with $1/2$ probability. Previous works on “for-each” sparsifiers [2, 24] used a recursive expander decomposition and subsampling at each level, with the main issue being that the degrees were not well-preserved. This issue was solved by explicitly storing the degrees of the graph, causing their spectral sparsifier data structure to no longer be a graph. The work of [14] circumvents this issue with the short-cycle decomposition by constructing subsampled graphs that exactly preserve their vertex degree.

Note that [14] already gives a graphical spectral sparsifier with $\tilde{O}(n/\varepsilon)$ edges and bits of working memory using a basic brute-force search for cycle decomposition. However, its runtime for construction is $O(mn)$, which is prohibitively slow as a subroutine. On the other hand, later work by Parter and Yogev [41] shows a deterministic algorithm to compute $(O(n \log n), O(\log^2 n))$ short-cycle decompositions. This combined with Claim 8 will allow us to design an algorithm in $m^{1+o(1)}$ time which returns a $(1+\varepsilon)$ -spectral sparsifier which is a reweighted graph with $\tilde{O}(n/\varepsilon)$ edges. This gives near-optimal parameters except for a required $m^{1+o(1)}$ bits of working memory. However, in Section 4.1, we show that we can implement this subroutine so that it is guaranteed to use only $\tilde{O}(m)$ working memory, achieving near-optimal results in both working memory usage and running time.

■ **Table 1** Minimum Cut Space Complexity in the Single-Pass Insertion-Only Streaming Setting.

Stream Type	Exact/Approx ($1 + \varepsilon$)	Upper Bound	Lower Bound
Adversarial	Exact	$O(n^2)$ (full graph)	$\Omega(n^2)$ [52]
Adversarial	Approx, Deterministic	$\tilde{O}(n/\varepsilon^2)$ [6]	$\Omega(n/\varepsilon^2)^*$ (Theorem 12)
Adversarial	Approx, Randomized	$\tilde{O}(n/\varepsilon)$ (Theorem 9)	$\Omega(n/\varepsilon)$ (Theorem 12)
Random-Order	Exact	$\tilde{O}(n)$ (Theorem 14)	$\Omega(n)$ [12]

3 Overview of Results

We list the current state-of-the-art results for single-pass minimum cut streaming algorithms in Table 1. Our main result is an $\tilde{O}(n/\varepsilon)$ space streaming algorithm that finds a $(1 + \varepsilon)$ -approximate minimum cut in a single-pass stream.

► **Theorem 9.** *There is a one-pass insertion-only streaming algorithm that, with high probability, computes an $(1 + \varepsilon)$ -approximation of the minimum cut on weighted graphs using $\tilde{O}(n/\varepsilon)$ bits of space. Moreover, our algorithm takes $\tilde{O}(m) + (n/\varepsilon^2)^{1+o(1)}$ total update time and $\tilde{O}(n^2/\varepsilon^2)$ post-processing time.*

This shows, somewhat surprisingly, that estimating the minimum cut is easier than computing a for-all cut sparsifier in a data stream. Our algorithm is mainly based on the following new for-each spectral sparsifier in a graph stream.

► **Lemma 10.** *A one-pass insertion-only streaming algorithm exists that, with high probability, constructs a $(1 + \varepsilon)$ for-each spectral sparsifier of weighted graphs with $\tilde{O}(n/\varepsilon)$ edges using $\tilde{O}(n/\varepsilon)$ bits of space. This algorithm also has total runtime $\tilde{O}(m) + (n/\varepsilon^2)^{1+o(1)}$.*

Our algorithm's total runtime is $\tilde{O}(m) + (n/\varepsilon^2)^{1+o(1)}$, which notably implies its amortized update time is $\tilde{O}(1)$ when $m \geq (n/\varepsilon^2)^{1+o(1)}$. With the $\Omega(n/\varepsilon)$ data structure lower bound in [2], our algorithm is tight in space complexity up to polylogarithmic factors. Based on our Lemma 10, we can obtain an $\tilde{O}(n/\varepsilon)$ space streaming algorithm that can $(1 \pm \varepsilon)$ -estimate all-pairs effective resistances.

► **Corollary 11.** *There exists a one-pass insertion-only streaming algorithm that constructs a data structure that calculates $(1 + \varepsilon)$ -approximations to the effective resistances between every pair of vertices on weighted graphs using $\tilde{O}(n/\varepsilon)$ bits of space. Our algorithm has $\tilde{O}(m) + (n/\varepsilon^2)^{1+o(1)}$ total time during the stream and $\tilde{O}(n^2/\varepsilon)$ post-processing time.*

Notice that we are able to match the sketching bounds of [24] up to polylogarithmic factors in the space and time required.

We also show an $\Omega(n/\varepsilon)$ lower bound for minimum cut and all-pairs effective resistances in Section 4.4 and Section 4.6 respectively, showing that our results of Theorem 9 and Corollary 11 are space-optimal up to polylogarithmic factors.

► **Theorem 12.** *Fix $\varepsilon > 1/n$. Any randomized algorithm that outputs a $(1 + \varepsilon)$ -approximation to the minimum cut of a simple, undirected graph in a single pass over a stream with probability at least $2/3$ requires $\Omega(n/\varepsilon)$ bits of space. If the algorithm is deterministic and $\varepsilon \geq 1/n^{1/4}$, then the algorithm requires $\Omega(n/\varepsilon^2)$ bits of space.⁴*

⁴ This latter assumption on ε is the same as used in [11].

► **Theorem 13.** Fix $\varepsilon > 1/n$. Suppose $sk(\cdot)$ is a sketching algorithm that outputs at most $s = s(n, \varepsilon)$ bits, and f is an estimation algorithm such that,

$$\forall a, b \in V, \quad \Pr[f(a, b, sk(G)) \in (1 \pm \varepsilon)r_{a,b}] \geq \frac{2}{3},$$

where $r_{a,b}$ is the effective resistance of nodes a, b . Then we have $s \geq \Omega(n/\varepsilon)$.

We also study the minimum cut problem in the random-order streaming model, where edges arrive in a random order instead of an arbitrary worst-case order. We prove the following result for *simple* unweighted graphs, which are graphs with at most one edge between any two vertices.

► **Theorem 14.** There exists a one-pass insertion-only streaming algorithm in the random-order model that outputs all minimum cuts $(S, V \setminus S)$, along with their corresponding edges $E(S, V \setminus S)$, for a simple, unweighted graph and with high probability using $\tilde{O}(n)$ space. The algorithm has $\tilde{O}(n)$ update time and $\tilde{O}(n^2k)$ post-processing time, where k is the value of the minimum cut.

We note that Theorem 14 provides the exact minimum cut, which we find surprisingly possible in a single pass. This further adds to the surprises regarding minimum cut in the streaming model, as a 2-pass exact algorithm in arbitrary order streams was known [5] while computing the exact minimum cut in a single pass in an arbitrary order stream is known to require $\Omega(n^2)$ memory [52].

Given that determining whether a graph is connected from a random order stream requires $\Omega(n)$ space [12], our $\tilde{O}(n)$ space algorithm for finding the exact minimum cut in random-order streams is optimal up to polylogarithmic space factors, even for just outputting the value. Note that we compute the minimum cut and all edges crossing the cut in addition to the value.

3.1 Our Techniques

Approximate Minimum Cut

Suppose that H is a $(1 + \varepsilon)$ -for-all sparsifier of G . Then the minimum cut of H is a $(1 + \varepsilon)$ -approximation to that of G . However, such a conclusion will not hold if H is instead a $(1 + \varepsilon)$ for-each sparsifier, as the total number of cuts is exponential. As discussed earlier, fortunately the number of 1.1-approximate minimum cuts is $O(n^2)$ [29]. Hence, if we run a for-all sparsifier algorithm in parallel with accuracy $\varepsilon' = O(1)$, we can obtain a list of $O(n^2)$ candidate minimum cuts in G . Using the the $(1 + \varepsilon)$ for-each sparsifier, we can then find a $(1 + \varepsilon)$ -approximation to the minimum cut by union bounding over the $O(n^2)$ candidates.

For the lower bound, we consider the k -edge-connectivity problem. The *edge-connectivity* of a graph is the minimum number of edges that need to be deleted to disconnect the graph. The k -edge-connectivity problem asks whether the edge connectivity of a graph is $< k$ or $\geq k$. Suppose that there is an algorithm that solves the $(1 + \varepsilon)$ -approximate minimum cut value problem. Then, we can use it to solve the k -edge-connectivity problem with $k = O(1/\varepsilon)$. Combining with the $\Omega(kn)$ bit space lower bound in the work of [49], it follows that there is an $\Omega(n/\varepsilon)$ bit space lower bound.

For-Each Spectral Sparsifier

The main difficulty in implementing the algorithms of [2] and [24] in the streaming setting is that both algorithms require careful graph decomposition. Take the algorithm in [2] as an example: it partitions the graph into several components, where each component is

well-connected and does not have a sparse cut smaller than $1/\varepsilon$. In the worst case, such a procedure may take $\log n$ levels recursively, which seems unachievable in a single-pass stream. Besides, since the for-each sketch of both algorithms is not a graph, it is unclear how to merge two sketches directly. Indeed, natural ways of merging such sketches may destroy the decomposition into sparse cuts. To address these issues, we instead consider the work of [14] in which the authors show how to construct a for-each sparsifier which is, in fact, a re-weighted subgraph of the input graph, with $n^{1+o(1)}/\varepsilon$ edges. Combining this and the merge-and-reduce framework gives a for-each spectral sparsifier with $n^{1+o(1)}/\varepsilon$ edges in a single-pass stream. The extra $n^{o(1)}$ factor in the space is an issue here. We notice that the extra $n^{o(1)}$ factor in the work of [14] is due to their near-linear time short cycle decomposition algorithm. Leveraging ideas from [41], we give a new short-cycle decomposition algorithm that trades off running time for space, namely, it achieves $m^{1+o(1)}$ time and $\tilde{O}(m)$ space. While this translates to $\tilde{O}(n/\varepsilon)$ space, it unfortunately gives $n^{o(1)}$ update time per edge. To fix this, we use online leverage score sampling to produce a virtual stream of only $\tilde{O}(n/\varepsilon^2)$ edges that we instead run our algorithm on. By doing this, we reduce the time to $\tilde{O}(1)$ per edge provided the number m of edges in the input satisfies $m \geq (n/\varepsilon^2)^{1+o(1)}$.

Approximate All-Pairs Effective Resistance

In [24], the authors show if we can get a for-each sparsifier, then we can use it to generate a data structure in near-linear time, and such a data structure can approximate all pairs of the effective resistance in $\tilde{O}(n^2/\varepsilon)$ time. Combining this and our Lemma 10 yields our upper bound.

For the lower bound, we use communication complexity and consider a similar graph construction in [2], though we use new arguments based on random walks. Specifically, given a random binary string $s \in \{0, 1\}^{n/\varepsilon}$, we encode it into a graph G where G is divided into $O(\varepsilon n)$ disjoint bipartite graphs G_i , and in each G_i , the existence of an edge between each pair corresponds to one random bit in s . We show that for every such pair (a, b) , with high probability there is a $(1 + \varepsilon)$ -separation between the two cases (there is an edge between a, b or not), which yields an $\Omega(n/\varepsilon)$ lower bound. To prove this, in particular, we use the connection between the effective resistance and the hitting time on a graph and a recent concentration result about the hitting time on a random graph $G(n, p)$ (each possible edge on n vertices is included independently with probability p).

Exact Minimum Cut in Random-order Streams

Recall that we assume that the graph is simple and unweighted, and each edge arrives in a random order stream. We also assume the minimum cut size is $\Omega(\log n)$. If the size is smaller, a for-all cut sparsifier with accuracy parameter $\varepsilon' = 1/\log^2 n$ gives the exact value of the minimum cut. The main idea here is that by looking at the prefix of edges in a stream, which we will call the graph H , we roughly learn the sizes of all the cuts up to a small constant factor. Using this information, we learn all 1.1-approximate minimum cuts in G . The next question is how to get the exact value of all these cuts (so we can find the minimum one). For the cut edges in H , note that at this time, the cut values of all these cuts in H is $\Theta(\log n)$, and hence a for-all sparsifier of H with accuracy parameter $\varepsilon' = 1/\log^2 n$ gives the exact value of these cuts in H . We note that the total number of edges that participate in at least one of the 1.1-approximate non-singleton cuts is $O(n)$ [43], which is interesting because there could be as many as $O(n^2)$ such cuts. Hence, we can store all these $O(n)$ edges to get the exact cut values of the approximate minimum cuts in $G \setminus H$. Putting the two things together, we obtain the exact value of all 1.1-approximate minimum cuts and thus obtain the exact minimum cut value in G (along with all the minimum cuts).

Faster Runtime

In the above discussion, our algorithm for $(1 + \varepsilon)$ min-cut in worst-case streams has an $\tilde{O}(n^3)$ post-processing time, and the algorithm for exact min-cut in a random-order stream has an $\tilde{O}(n^3)$ update time. The difficulty here is related: when enumerating all the $O(1)$ -approximate minimum cuts, we need $O(n)$ time to evaluate the cost for each specific cut. We thus propose a general algorithmic framework to overcome this, which is based on the recursive contraction algorithm in [31] (Section 4.4). Namely, when enumerating all the approximate minimum cuts in the recursion tree, we simultaneously maintain a low-space sketch of the columns of the edge-vertex incidence matrix of the corresponding graph. Specifically, we apply a Johnson-Lindenstrauss sketch in our first case and a sparse recovery sketch in our second case, which reduces the evaluation time $O(n)$ to $O(\log n/\varepsilon^2)$ and $\text{polylog}(n)$, respectively.

3.2 Open Problems

While we resolve several gaps between the upper and lower bounds, we discuss a remaining open problem. There exist fully dynamic streaming algorithms for approximate minimum cut using $\tilde{O}(n/\varepsilon^2)$ space where the stream is allowed to both add or delete edges. From Theorem 12, we know that algorithms solving approximate minimum cut in insertion-only streams must use $\Omega(n/\varepsilon)$ space. Therefore, the space complexity of approximate minimum cut in a dynamic stream lies somewhere within these two values, and we believe resolving this gap is a problem of theoretical interest.

► **Open Question 15.** What is the exact space complexity of calculating a $(1 + \varepsilon)$ approximation to the minimum cut of simple weighted graphs in a one-pass dynamic stream?

We do not have an exact conjecture about this, but we provide an alternative and fully self-contained proof to Theorem 12 in the full version of the paper [17], for which the techniques could help prove a lower bound for the fully dynamic case.

4 Spectral Sparsification and Minimum Cut in Worst-Case Streams

A key result of the section is our construction of Lemma 10, a single-pass insertion-only streaming algorithm for a $(1 + \varepsilon)$ -for-each cut sparsifier in undirected weighted graphs in $\tilde{O}(n/\varepsilon)$ space and $n^{o(1)}$ update time. In Section 4.1 and 4.2, we first give our description of graphical spectral sketches (i.e., spectral sparsifiers that are reweighted subgraphs) and online leverage score sampling, which we use as subroutines in our final algorithm. In Section 4.3, we give our full algorithm and prove the correctness of our algorithm. In Section 4.4 and 4.6, we prove the two applications of our algorithm: streaming approximate minimum cut and all-pairs effective resistances.

4.1 Graphical Spectral Sketches

Our algorithm uses the following result in [14] as a subroutine.

► **Lemma 16** ([14]). *For $\varepsilon \in (0, 1]$, there is an algorithm, which given G , runs in time $m^{1+o(1)}$, and with high probability returns a $(1 + \varepsilon)$ -for-each spectral sparsifier with $n^{1+o(1)}/\varepsilon$ edges.*

As pointed out by the work of [41], in the above algorithm, the extra $n^{o(1)}$ factor in the number of edges is due to the nearly-linear time short cycle decomposition used in the whole algorithm. We use the simpler and improved short-cycle decomposition algorithm of [41],

which outputs in $m^{1+o(1)}$ time a collection of edge-disjoint cycles of length $O(\log^2 n)$ that cover all but $O(n \log n)$ edges of the graph. Naïvely, the algorithm is also implemented in $m^{1+o(1)}$ space, but we show that a simple modification achieves near-linear space, as follows:

► **Lemma 17.** *There is an $m^{1+o(1)}$ time, $\tilde{O}(m)$ space algorithm that outputs a collection of edge-disjoint cycles of length $O(\log^2 n)$ that cover all but $O(n \log n)$ edges of the graph.*

Proof. The algorithm in Section 3 of [41] first computes a *low-congestion cycle cover*, which is a collection of short cycles that cover most of the edges such that each edge belongs to a small number of cycles. The precise parameters are specified in the statement of Lemma 2 of the full version of [41]: compute a collection of cycles of length $d = O(2^{1/\varepsilon} \log n)$ that cover all but $O(n \log n)$ edges such that each edge appears on at most $c = 1/\varepsilon \cdot n^{O(\varepsilon)}$ cycles. For $\varepsilon = 1/\log \log n$, we obtain $d = O(\log^2 n)$ and $c = n^{o(1)}$. To obtain the short cycle decomposition, the proof of Lemma 2 describes a simple procedure of greedily selecting a subset of edge-disjoint cycles that cover an $\Omega(1/(dc))$ fraction of the edges in the cycle cover and then iterating on the uncovered edges.

Note that the low-congestion cycle cover consists of up to cm edges since each edge can appear in c cycles. Since $c = n^{o(1)}$, storing the entire cycle cover takes up too much space. Our key insight is to terminate the collection of cycles in each iteration of algorithm IMPROVEDSHORTCYCLEDECOMP (see Figure 3 in the full version of [41]) once the cycle cover has a total of $O(m)$ edges. Terminating this subroutine early cannot increase the congestion compared to not terminating, thus we maintain the guarantee that the total set of cycles has congestion c . Using the above greedy selection of edge-disjoint cycles, we cover an $\Omega(1/(dc))$ fraction of the edges in our cycle cover, which is $\Omega(1/(dc))$ of all uncovered edges in the graph. Iterating on the uncovered edges gives the same iterations as before up to logarithmic factors. ◀

With our improved short-cycle decomposition, we get the following lemma:

► **Lemma 18.** *Given $\varepsilon \in (0, 1]$, there is an algorithm we denote $\text{SPECTRALSKETCH}(G, \varepsilon)$, which given G , runs in time $m^{1+o(1)}$ and space $\tilde{O}(m)$, and with high probability returns a $(1 + \varepsilon)$ -for-each spectral sparsifier of G with $\tilde{O}(n/\varepsilon)$ edges.*

4.2 Online Leverage Score Sampling

For the purposes of a faster runtime, we consider only sampling (and reweighting) a fraction of edges during the stream. Existing online leverage score sampling methods allow us to choose edges in an online stream without retracting our choices such that our final graph has $\tilde{O}(n/\varepsilon^2)$ edges and is a spectral-sparsifier of the original graph. Indeed, let $B_n \in \mathbb{R}^{\binom{n}{2} \times n}$ be the vertex edge incidence matrix of an undirected, unweighted complete graph on n vertices, where the e -th row b_e for edge $e = (u, v)$ has a 1 in column u , a (-1) in column v , and zeroes elsewhere. Then for an arbitrary undirected graph G , we can write its vertex edge incidence matrix $B = SB_n$ where $S \in \mathbb{R}^{\binom{n}{2} \times \binom{n}{2}}$ is a diagonal matrix with entry $\sqrt{w_e}$ in the e -th diagonal entry where w_e is the weight of edge e . It is well-known that the Laplacian matrix L of graph G can be written as $L = B^\top B$. Hence, if we can sample a subset of the rows of B to form a new matrix C (which corresponds to a subset of the edges in G) such that for every x , $\|Bx\|_2 \approx_\varepsilon \|Cx\|_2$, we then have $x^\top Lx \approx_\varepsilon x^\top C^\top Cx$ so that the subgraph H corresponding to C is a $(1 + \varepsilon)$ -spectral sparsifier of G . For more details, we refer the reader to [27, 15]. Formally, we have the following lemma.

► **Lemma 19** (Corollary 2.4 from [15]). *Let G be an undirected simple graph with n vertices and $\text{poly}(n)$ bounded edge weights, and $\varepsilon \in (0, 1)$. We can construct a $(1 + \varepsilon)$ -spectral sparsifier of G as a reweighted subgraph with $O(n \log^2 n / \varepsilon^2)$ edges, using only $O(n \log^2 n)$ bits of working memory in the online model. Additionally, the total running time is near-linear in the number of edges of G .*

4.3 Main Algorithm

Our algorithm is presented in Algorithm 1, which is based on the merge-and-reduce paradigm (see, for example, [10]) and uses an additional factor of $\text{polylog}(n)$ space. At a high level, our algorithm maintains a number of blocks of edges $\mathbf{B}_0, \mathbf{B}_1, \dots, \mathbf{B}_{\log(n/\varepsilon)}$, each with size $m_{\text{space}} = \tilde{O}(n/\varepsilon)$. The most recent edges are stored in \mathbf{B}_0 ; whenever \mathbf{B}_0 is full, the successive non-empty blocks $\mathbf{B}_0, \dots, \mathbf{B}_i$ are merged and reduced to a new graph with m_{space} edges, which will be stored in \mathbf{B}_{i+1} . Next, we show the correctness of our algorithm. Following a similar argument in [10], we first show the following lemma.

■ **Algorithm 1** STREAMINGSPECTRALSPARSIFIER(G, ε).

```

1 Input: Undirected graph  $G(V, E)$  in a stream with  $n$  vertices and  $m$  edges, accuracy
   parameter  $\varepsilon \in (\frac{1}{n}, 1)$ , offline spectral-sparsifier subroutine SPECTRALSKETCH
2 Output: A graph  $H$  with weights  $w$ .
3 Initialize  $m_{\text{space}} = n \log^c(n)/\varepsilon$  for some constant  $c$ ,  $\mathbf{B}_i \leftarrow \emptyset$ 
4 foreach sampled edge  $e_t$  (Lemma 19) with rescaled weight  $u_t$  do
5   if  $\mathbf{B}_0$  does not contain  $m_{\text{space}}$  edges then
6      $\mathbf{B}_0 \leftarrow e_t \cup \mathbf{B}_0$ ;
7   else
8     Let  $i > 0$  be the minimal index such that  $\mathbf{B}_i = \emptyset$ ;
9      $\mathbf{B}_i, w_i \leftarrow \text{SPECTRALSKETCH} \left( \mathbf{M}, \frac{\varepsilon}{\log(n/\varepsilon)} \right)$ , where  $\mathbf{M} = \mathbf{B}_0 \cup \dots \cup \mathbf{B}_{i-1}$ ;
10    for  $j = 0$  to  $j = i - 1$  do
11       $\mathbf{B}_j \leftarrow \emptyset$ ;
12     $\mathbf{B}_0 \leftarrow e_t$ ;
13  $\mathbf{B} \leftarrow \text{SPECTRALSKETCH}(\mathbf{B}_{\log(n/\varepsilon)} \cup \dots \cup \mathbf{B}_0, \varepsilon)$ 
14 return  $\mathbf{B}$ .
```

► **Lemma 20** (see Lemma 5.2 in [10]). *Suppose that $\mathbf{B}_0, \dots, \mathbf{B}_{i-1}$ are all empty while \mathbf{B}_i is non-empty. Then \mathbf{B}_i is a $(1 + \frac{\varepsilon}{\log(n/\varepsilon)})^i$ -for-each spectral sparsifier for the last $2^{i-1} m_{\text{space}}$ edges.*

Proof. We prove this by induction on $i \geq 0$. Recall that \mathbf{B}_i can only be non-empty if at some point \mathbf{B}_0 contains m_{space} and $\mathbf{B}_0, \mathbf{B}_1, \dots, \mathbf{B}_{i-1}$ are all non-empty. By induction, suppose that for every $1 \leq j < i$, \mathbf{B}_j is a $(1 + \frac{\varepsilon}{\log(n/\varepsilon)})^j$ -for-each spectral sparsifier for $2^{j-1} m_{\text{space}}$ edges. Then we have that \mathbf{B}_i is a $(1 + \frac{\varepsilon}{\log(n/\varepsilon)})$ -for-each spectral sparsifier for $\mathbf{B}_0 \cup \mathbf{B}_1 \cup \dots \cup \mathbf{B}_{i-1}$. From the mergeability property of spectral sparsifier graphs, we get that \mathbf{B}_i is a $(1 + \frac{\varepsilon}{\log(n/\varepsilon)})(1 + \frac{\varepsilon}{\log(n/\varepsilon)})^{i-1} = (1 + \frac{\varepsilon}{\log(n/\varepsilon)})^i$ -for-each spectral sparsifier for the $m_{\text{space}} + \sum_{j=0}^{i-2} 2^j m_{\text{space}} = 2^{i-1} m_{\text{space}}$ edges, as needed. ◀

Proof of Lemma 10. The success probability of each call to subroutine SPECTRALSKETCH is at least $1 - 1/\text{poly}(n)$, hence we can assume that each call to this subroutine and the online leverage score sampling procedure is successful with high probability after taking a union

43:12 Space Complexity of Minimum Cut Problems in Single-Pass Streams

bound. It then follows from Lemma 20 that $\mathbf{B}_{\log(n/\varepsilon)} \cup \dots \cup \mathbf{B}_0$ is a $(1 + \frac{\varepsilon}{\log(n/\varepsilon)})^{\log(n/\varepsilon)} \leq (1 + \varepsilon)$ -for-each spectral sparsifier of the original graph G . Thus \mathbf{B} is a $(1 + O(\varepsilon))(1 + \varepsilon) = (1 + O(\varepsilon))$ -for-each spectral sparsifier of the original graph G .

Space Complexity. Next, we analyze the space complexity of our algorithm. As stated in Lemma 19, the online leverage score sampling procedure can be implemented in $O(n \log^2 n)$ bits of working memory. We maintain $\log(n/\varepsilon)$ blocks \mathbf{B}_i , each taking at most $O(m_{\text{space}})$ words of space. From Lemma 16 we have that each call to the subroutine SPECTRALSKETCH takes at most $\tilde{O}(m_{\text{space}} \cdot \log(n/\varepsilon))$ words of space, as the total number of edges never exceeds $m_{\text{space}} \cdot \log(n/\varepsilon)$. Hence, the total space of the algorithm is $\tilde{O}(n \log^2 n + m_{\text{space}} \log(n/\varepsilon)) = \tilde{O}(n/\varepsilon)$.

Time Complexity. We finally analyze the time complexity of our algorithm. As stated in Lemma 19, the online leverage score sampling procedure can be done in time $\tilde{O}(m)$, and after the sampling procedure, there are at most $O(n \log^2 n / \varepsilon^2)$ edges. For the merge-and-reduce procedure, each edge participates in at most $\log(n/\varepsilon)$ different spectral sparsifiers, one in block \mathbf{B}_i . Additionally, from Lemma 16 we have that SPECTRALSKETCH runs in $m^{1+o(1)}$ time where $m = \tilde{O}(n/\varepsilon)$. We deduce that our total runtime is $\tilde{O}(m + (n \log^2 n / \varepsilon^2)^{1+o(1)}) = \tilde{O}(m) + (n/\varepsilon^2)^{1+o(1)}$. ◀

4.4 Approximate Minimum Cut Streaming Algorithm

Our main application of Lemma 10 is the one-pass streaming algorithm for a $(1 + \varepsilon)$ -approximation to the minimum cut. To achieve this, first recall that we can get a $(1 + \varepsilon)$ -for-all sparsifier of G in a one-pass stream in nearly-linear time, using $\tilde{O}(n/\varepsilon^2)$ bits of space (Lemma 4).

The next lemma bounds the number of approximate minimum cuts.

► **Lemma 21** ([29]). *For constant $\alpha > 0$, the number of α -approximate minimum cuts is $O(n^{\lfloor 2\alpha \rfloor})$.*

Motivated by [2], we run our algorithm in Lemma 10 along with the streaming for-all sparsifier algorithm ALG in Lemma 4 with accuracy parameter $\varepsilon' = O(1)$. Suppose that the output of ALG is H . We then use H to find the cuts that have a size less than 1.5 times that of the minimum cut value (which can be done in $\text{poly}(n)$ time, see, e.g., [29]). Then, since the number of these approximate minimum cuts is $O(n^3)$, using our $(1 + \varepsilon)$ -for-each sparsifier, we can get a $(1 + \varepsilon)$ -approximation to each of these cut values with high probability after taking a union bound. Picking the cut with the minimum value gives us a $(1 + \varepsilon)$ -approximate minimum cut.

Faster Post-processing Time

The above algorithm is simple but, unfortunately, takes $O(n^3)$ post-processing time. To achieve a faster post-processing time, we consider the recursive contraction algorithm introduced by [31], which enumerates all α -approximate minimum cuts in time $O(n^{2\alpha} \log^2 n)$. We combine this with a Johnson-Lindenstrauss sketch to achieve $\tilde{O}(n^2/\varepsilon^2)$ post-processing time, proving Theorem 9.

First, we state the lemma from [31]:

► **Lemma 22.** *There is an algorithm that can find all the α -approximate minimum cuts with high probability in $O(n^{2\alpha} \log^2 n)$ time and $\tilde{O}(m)$ space.*

Below, we first give a brief explanation of the algorithm. At a high level, in each step, the algorithm randomly samples an edge proportional to its weights and then contracts the two nodes of the edge. When there are only two remaining nodes, we get a partition of the nodes corresponding to a specific cut of the graph. The work of [31] shows that if we repeat the algorithm $\text{poly}(n)$ times, we can finally enumerate all approximate minimum cuts with high probability. The above process is simple but needs $\text{poly}(n)$ runtime. To make the algorithm faster, [31] then proposes a recursive implementation of the contraction process. Basically, at each level of the recursion, the algorithm reduces the number of nodes by contraction by a constant factor and repeats the algorithm twice in each recursion level if the number of remaining nodes is $\Omega(1)$. Then, each leaf node of the recursion tree corresponds to a specific cut of the graph.

The technical difficulty in our case is evaluating the value of the corresponding cut in a short time when we enter a leaf node in the above process. To achieve this, recall that for a given vertex set $S \subseteq V$, the cut value between S and $V \setminus S$ is $x_S^\top Lx_S = x_S^\top B^\top Bx_S = \|Bx_S\|_2^2$ where x_S is the binary indicator vector of S and B is the edge-vertex matrix of the graph (see Section 4.2 for more details). The key observation here is since we only care about a $(1 + \varepsilon)$ -approximate value, we can apply a JL matrix T to Bx_S , where the matrix T has only $O(\log(n/\varepsilon)/\varepsilon^2)$ rows where for every $x \in \mathbb{R}^{O(n/\varepsilon)}$, $\|Tx\|_2^2 \approx_\varepsilon \|x\|_2^2$ with high probability.

Our algorithm procedure is described as follows. As before, after the stream, we first get a for-all sparsifier K with the accuracy parameter $\varepsilon' = 1/\log n$ (Lemma 4) and a $(1 + \varepsilon)$ for-each sparsifier H of the original graph G (Lemma 10). We then use the algorithm in Lemma 22 to enumerate all the approximate minimum cuts of K with $\alpha = 1 + c/\log n$. When maintaining the recursion process of the contraction algorithm, we also maintain a sketch of the columns of the edge-vertex matrix B of H . Specifically, we initially compute a sketch of TB where T is a JL matrix with $O(\log n/\varepsilon^2)$ rows. Then, in the contraction process, when we contract a node pair (u, v) , we also merge the u, v -th columns of the sketch with their sums. Note that in this procedure, when we enter a leaf node, there will only remain two columns of our sketch, and the two vectors are exactly TBx_S and $TBx_{V \setminus S}$, for which we can directly compute the cut value in $O(\log n/\varepsilon^2)$ time.

Next, we analyze the time and space complexity in the above procedure. Since the sketch has $O(\log n/\varepsilon^2)$ rows at each recursion level, we need $\tilde{O}(N \log n/\varepsilon^2)$ words of space to save the sketch, where N is the number of the remaining nodes in the current recursion level. Combining two rows also takes $\tilde{O}(\log n/\varepsilon^2)$ time. Recall that we need $O(n)$ space in each level and $O(N)$ time to contract two nodes in the original recursive contraction algorithm. This implies after the modification, the space complexity of the above algorithm procedure remains the same of $\tilde{O}(m) = \tilde{O}(n/\varepsilon)$, and the time complexity becomes $\tilde{O}(n^2/\varepsilon^2)$. Putting everything together, we get the correctness of our Theorem 9.

4.5 Approximate Minimum Cut Streaming Lower Bound

To prove this lower bound, we consider the k -edge-connectivity problem as follows. The *edge-connectivity* of a graph is the minimum number of edges that need to be deleted to disconnect the graph. The k -edge-connectivity problem asks whether the edge connectivity of a graph is $< k$ or $\geq k$. The work of [49] shows that any deterministic algorithm that solves the k -edge-connectivity problem in insertion-only streams needs $\Omega(kn \log n)$ bits of space. We note that the proof in [49] also works for randomized algorithms with success probability at least $1 - 1/n$, which implies a lower bound of $\Omega(kn)$ bits of space for any randomized algorithm with constant probability of success.

43:14 Space Complexity of Minimum Cut Problems in Single-Pass Streams

▷ Claim 23 ([49]). Any randomized algorithm solving the k -edge-connectivity problem in an insertion-only stream with probability at least $2/3$ requires $\Omega(kn)$ bits of space. Moreover, any deterministic algorithm solving the k -edge-connectivity problem in an insertion-only stream requires $\Omega(kn \log n)$ bits.

We now prove the randomized and deterministic lower bounds of Theorem 12.

Randomized Lower Bound

► **Lemma 24.** Fix $\varepsilon > 1/n$. Any randomized algorithm that outputs a $(1 + \varepsilon)$ -approximation to the minimum cut of a simple, undirected graph in a single pass over a stream with probability at least $2/3$ requires $\Omega(n/\varepsilon)$ bits of space

Proof. Consider the lower bound for k -edge-connectivity in [49] with $k = \frac{1}{10\varepsilon}$ (Claim 23). Distinguishing between whether the graph is $< k$ -connected or $\geq k$ -connected needs $\Omega(kn) = \Omega(n/\varepsilon)$ space. This implies that deciding whether the min-cut has value $< k$ or $\geq k$ also requires $\Omega(kn) = \Omega(n/\varepsilon)$ space. Suppose that there is a $(1 + \varepsilon)$ -approximation algorithm for the minimum cut value, in the case where the edge connectivity $\lambda < k$, the approximate value is at most $(1 + \varepsilon)\lambda \leq (1 + 1/10k) \cdot (k - 1) < k - 9/10$. In the case where the connectivity $\lambda \geq k$, the approximate value is at least $(1 - \varepsilon)\lambda \geq (1 - 1/10k) \cdot k > k - 1/10$. This implies that an algorithm that gives a $(1 + \varepsilon)$ -approximation to minimum cut value, solves k -edge-connectivity for $k = \frac{1}{10\varepsilon}$. Thus, we obtain a lower bound of $\Omega(n/\varepsilon)$ bits of space for $(1 + \varepsilon)$ -approximating the value of the minimum cut. ◀

In addition to the proof above that considers the k -edge-connectivity problem, we also provide an alternative self-contained construction for the randomized result using the Index communication problem in the full version of the paper [17].

Deterministic Lower Bound

If the minimum cut value algorithm is deterministic, then we deterministically solve k -edge-connectivity for $k = \frac{1}{10\varepsilon}$. This gives us a lower bound of $\Omega(kn \log n) = \Omega(n \log n/\varepsilon)$ bits of space for deterministic algorithms. We show how to improve this lower bound to $\Omega(n/\varepsilon^2)$ bits for $\varepsilon \geq n^{-1/4}$. Note that this assumption on ε is consistent with the previous work of [11]. They also use this assumption on ε and prove a lower bound of $\Omega(n \log n/\varepsilon^2)$ bits for cut sparsifiers.

► **Lemma 25.** Fix $\varepsilon > 1/n^{1/4}$. Any deterministic algorithm that outputs a $(1 + \varepsilon)$ -approximation to the minimum cut value of a simple, undirected graph in a single pass over a stream requires $\Omega(n/\varepsilon^2)$ bits of space

Proof. A randomized insertion-only streaming algorithm for a for-all cut sparsifier needs $\Omega(n/\varepsilon^2)$ bits of space [2] (since any data structure that stores this information needs that much space). If we carefully look at the proof of [2] in section 3.1, their hard instance is a disjoint union of $\varepsilon^2 n/2$ bipartite graphs with $2/\varepsilon^2$ vertices each. They use the cut sparsifier to query cuts that have vertices only in one of these disjoint graphs, which implies that the cut sparsifier only needs to preserve the cut values for cuts that have at most $1/\varepsilon^4$ edges. Thus, a $(1 + \varepsilon)$ cut sparsifier that preserves the cut values of all cuts of size at most $1/\varepsilon^4$ edges needs $\Omega(n/\varepsilon^2)$ bits of space.

We will now show how to simulate a $(1 + \varepsilon)$ cut sparsifier that preserves the cut values of all cuts of size at most $1/\varepsilon^4$ using an algorithm for $(1 + \varepsilon)$ -approximate minimum cut. This will give us the desired lower bound. Let A be a deterministic insertion-only streaming algorithm for $(1 + \varepsilon)$ -approximate minimum cut value we run during the stream for input graph G .

After the stream, consider any cut S whose value we want to approximate within a $(1 + \varepsilon)$ factor (note that the value is at most $1/\varepsilon^4$). We now add some extra edges to the stream for A . We first construct two cliques C_1 and C_2 on $2n$ vertices each and add their edges to the stream. We then add edges between all vertices of S and C_1 and similarly for all vertices in \bar{S} and C_2 . Call this new graph G' . We now end the stream and look at the output of A . We claim this is the $(1 + \varepsilon)$ -approximate value of cut S in the original graph.

To prove this, we must show that S is the minimum cut in the modified graph G' . We know that the value of S is at most $1/\varepsilon^4 \leq n$ (since $\varepsilon \geq n^{-1/4}$). Consider any cut that separates vertices in $C_1 \cup S$. Such a cut has a size of at least $2n$. The same applies to any cut separating vertices in $C_2 \cup \bar{S}$. Thus, the minimum cut is $C_1 \cup S$. None of the new edges we added cross this cut, so the size of this cut remains unchanged in G' , implying that we get a $(1 + \varepsilon)$ -approximation to the cut value in G . Note that the algorithm A has to be deterministic for us to do this because we have to repeat this process for exponentially many cuts (potentially 2^n).

The space taken to build the sparsifier for insertion-only streams is the space taken for A when run on $5n$ vertices. This implies that any deterministic insertion-only streaming algorithm for $(1 + \varepsilon)$ -approximate minimum cut value needs $\Omega(n/\varepsilon^2)$ bits of space, proving the lower bound for deterministic algorithms. ◀

Proof of Theorem 12. The proof of the theorem directly results from combining Lemma 24 and Lemma 25. Notably, this shows that the spectral sparsifier result in Lemma 10 is optimal in space complexity up to polylogarithmic factors. ◀

4.6 Application: All-Pairs Effective Resistances

Another application of our streaming algorithm from Lemma 10 is calculating all-pairs effective resistances in graph streams. Suppose we have a $(1 + \varepsilon)$ for-each spectral sparsifier H of the original graph G . Then, we may hope it will be helpful to approximate the pseudoinverse of G . In fact, in the work of [24], it has been shown that if we have a $(1 + \varepsilon)$ for-each spectral sparsifier, we can use it to generate an effective resistance sketch in near-linear time, which can compute the all-pairs effective resistances in $\tilde{O}(n^2/\varepsilon)$ additional time (Algorithm 8 in [24]). From this, we immediately see the correctness of our Corollary 11.

Lower Bound

We next consider lower bounds. In Theorem 13, we show that if a sketch can approximate each effective resistance with constant probability, then such a sketch must have size at least $\Omega(n/\varepsilon)$. We need the following lemmas about effective resistances and random graphs.

▶ **Lemma 26** (see, e.g., [51]). *For a graph G with unit weight for each edge, we have $C(s, t) = 2mr_{\text{eff}}(s, t)$, where $C(s, t)$ is the commute time of s and t , which is defined as the expected number of steps of a random walk to go from s to t and back.*

► **Lemma 27** ([40]). *Let $G(n, p)$ be an Erdos-Renyi random graph. Then we have, with high probability that for any pair $w \neq v$,*

$$h(w, v) = \frac{2|E|}{\deg(v)} + \begin{cases} -1 & \text{if } (w, v) \in E \\ -1 + 1/p & \text{if } (w, v) \notin E \end{cases} + O\left(\frac{(\log n)^{3/2}}{\sqrt{n}}\right),$$

where $h(s, t)$ is the hitting time of s, t , defined as the expected number of steps of a random walk from s to t . Note that we have $C(s, t) = h(s, t) + h(t, s)$.

Our reduction is based on the following Index problem.

► **Lemma 28** (Index Lower Bound [34]). *Suppose that Alice has a random string $u \in \{0, 1\}^n$ and Bob has a random index $i \in [n]$. If Alice sends a single message to Bob from which Bob can recover u_i with probability at least $2/3$, where the probability is over both the randomness of the protocol and the input, then Alice must send $\Omega(n)$ bits to Bob.*

At a high level, we will reduce the Index problem to the effective resistance problem. Suppose Alice has a string $s \in \{0, 1\}^{\Theta(n/\varepsilon)}$. We will construct a graph G to encode s such that Bob can recover s_i with constant probability from a $(1 \pm \varepsilon)$ effective resistance sketch of G . By the communication complexity lower bound of the Index problem (Lemma 28), the cut sketch must use $\Omega(n/\varepsilon)$ bits.

Proof of Theorem 13. We first give our construction of the graph G .

Construction of G . We use a bipartite graph G to encode s . Let L and R be the left and right nodes of G where $|L| = |R| = n/2$. We partition L into $O(n\varepsilon)$ disjoint blocks $L_1, \dots, L_{O(n\varepsilon)}$ of the same size $\frac{1}{\varepsilon}$, and similarly, we partition R into $R_1, \dots, R_{O(n\varepsilon)}$. We divide s into $n\varepsilon$ disjoint strings $s_i \in \{0, 1\}^{\left(\frac{1}{\varepsilon}\right)^2}$ of the same length. We will encode s_i using the edges from L_i to R_i , where we use G_i to denote the subgraph (L_i, R_i) . In particular, if the $s_i(j) = 1$, we form an edge that connects the corresponding node pair in L_i and R_i with unit weight. Note that each G_i is disconnected from the other G_j in this construction.

Recovering a bit in s from an effective resistance sketch of G . Suppose Bob wants to recover a specific bit of s , which belongs to the sub-string $z = s_i$ and has an index t in z . This coordinate corresponds to whether there is an edge between $u \in L_i$ and $v \in R_i$.

Next, we consider the effective resistance between u and v for these two cases. Note that since G_i is disconnected from the other G_j , we just need to consider the sub-graph G_i . By construction, G_i is a random graph with $p = \frac{1}{2}$. From Lemma 26 we get that we only need to show both the hitting time $h(u, v)$ and $h(v, u)$ has a $(1 + \varepsilon)$ -separation (then the $r_{\text{eff}}(u, v)$) will also have a $(1 + \varepsilon)$ -separation. Then, from Lemma 27, we get that both the $h(u, v)$ and $h(v, u)$ will have a $\Theta(1)$ gap for the two cases, and with high probability, we have that both of $h(u, v)$ and $h(v, u)$ are $\Theta(1/\varepsilon)$. From this, we can get that $r_{\text{eff}}(u, v)$ will have a $(1 + \varepsilon)$ -separation for the two cases (note that Alice can also send the degree of each node and the number of edges in each G_i to Bob, which only needs $O(n \log(1/\varepsilon))$ bits), which yields an $\Omega(n/\varepsilon)$ lower bound of the effective resistance sketch size. ◀

5 Exact Minimum Cut in Random-Order Streams

In this section, we give an $\tilde{O}(n)$ algorithm that outputs the minimum cut in a simple, unweighted graph in a single-pass random-order stream. Recall that in a simple graph, there is at most one edge between any pair of vertices. For a clearer demonstration, we will present an initial construction for an algorithm with an $\tilde{O}(n^2)$ update time when an edge arrives. Then, we shall show how to improve the update time to $\tilde{O}(n)$, proving Theorem 14.

5.1 Initial Construction

We will use the algorithm in Lemma 4 (for-all sparsifier in a stream) as a subroutine. The whole algorithm is given in Algorithm 2. Below, we first give a high-level explanation of our algorithm. We first consider the case when the minimum cut size is $s = \Omega(\log n)$. Since we are now considering the random-order model, a prefix of the edges gives us partial information about the cut sizes. Let H be the subgraph of G which is formed by a prefix of the edges of the stream where $|H| \approx \frac{|G| \log n}{s}$ and $|G|, |H|$ are the number of the edges in G and H respectively. By a Chernoff bound, one can show that with high probability, for every subset $S \subset V$, $w_H(S, V \setminus S)$ is a small constant approximation to $w_G(S, V \setminus S)$. H might be too large to store in memory, so we store H_1 , a $(1 + \varepsilon)$ -for-all sparsifier of H , during the stream (we will decide ε later). Therefore, using the graph H_1 , we know all cuts in G whose cut size is within a factor of 1.1 of the true minimum cut of G . After this step, when a new edge $e \in G \setminus H$ arrives, we can use graph H_1 to check whether e belongs to some 1.1-approximate minimum cut in G , and if so, we save this edge (from Lemma 30 we know that the total number of these edges saved is at most $O(n)$). The next question is how to estimate the exact value of these 1.1-approximate cuts in H . The crucial observation is since the minimum cut size of H is an integer and is at most $\Theta(\log n)$, we can set the approximation parameter $\varepsilon = 1/\log^2 n$ for H_1 . This gives the exact value of $w_H(S, V \setminus S)$ if S and $V \setminus S$ is a constant-approximate minimum cut in G (since the error is at most $O(\log n) \cdot 1/\log^2 n < 1$). Putting these things together, after the stream, we can enumerate all 1.1-approximate minimum cuts and obtain their exact values, thus getting the exact minimum cut value.

The remaining case to handle is when the minimum cut size is $O(\log n)$. Similarly, in this case, setting $\varepsilon = 1/\log^2 n$ will give us the exact value of the minimum cut. Finally, we do not know the value of s , so we will try all powers of 2 for it. Note that we do not need the exact value; a 2-approximation to s will suffice. The detailed description of the algorithm is given in Algorithm 2.

To show the correctness of our algorithm, we first prove the following key lemma.

► **Lemma 29.** *Suppose that H is a subgraph of G formed by a prefix of the edges in the random order model. We have that for every $S \subset V$, if $|H| = |G| \cdot \frac{\ell}{w_G(S, V \setminus S)}$, then with probability at least $1 - 2e^{-c_2 \cdot \ell}$, $|w_H(S, V \setminus S) - \ell| \leq 0.1\ell$.*

Proof. Consider each edge e that is a cut edge between S and $V \setminus S$ (let T denote the set of such edges). Let y_e denote the indicator random variable where $y_e = 1$ if $e \in H$ and $y_e = 0$ otherwise. Then we have that $\mathbb{E}[w_H(S, V \setminus S)] = \mathbb{E}[\sum_{e \in T} y_e] = \ell$ from the assumption in the random-order model. By a standard Chernoff bound, we have that

$$\Pr[|w_H(S, V \setminus S) - \ell| \geq 0.1\ell] \leq 2e^{-c_2 \cdot \ell}$$

for some constant c_2 . ◀

We also need the following structural result for the approximate minimum cuts of a graph.

► **Lemma 30** ([43]). *For any simple, unweighted graph and any constant $\varepsilon > 0$, the total number of edges that participate in non-singleton $(2 - \varepsilon)$ -approximate minimum cuts is at most $O(n)$.*

We first consider the case when the minimum cut size $s < c \log n$. In this case, the minimum cut of H will never exceed $c \log n$ since H is a subgraph of G . Since the cut value is an integer, setting $\varepsilon = 1/\log^2 n$ gives the exact value of the size of the minimum cut.

■ **Algorithm 2** MINCUTRANDOMORDER.

Input : Undirected and unweighted graph $G(V, E)$ in a random-order stream with n vertices and m edges.

- 1 ALG_1 is an instance of Lemma 4 with $\varepsilon = 1/\log^2 n$.
- 2 Maintain the degree of each node d_i during the stream.
- 3 **foreach** edge e in the graph stream **do**
- 4 Feed e to ALG_1 ;
- 5 **if** e is the 2^i -th edge for some $i \in \mathbb{N}$ **then**
- 6 $H_1 \leftarrow$ output of ALG_1 (Here H_1 is a for-all sparsifier of the prefix graph H);
- 7 **if** the minimum cut size of H_1 is larger than $c \log n$ **then**
- 8 Save H_1 and break the for loop.
- 9 $H_1 \leftarrow$ output of ALG_1 ;
- 10 **if** there is no new edge in the stream **then**
- 11 **return** the minimum cut value of H_1 .
- 12 $T \leftarrow \emptyset$.
- 13 **foreach** new edge e during the stream **do**
- 14 **if** e is the cut edge between some S and $V \setminus S$ where S and $V \setminus S$ is a non-singleton 1.1-approximate minimum cut in H_1 **then**
- 15 Add e to T
- 16 **foreach** S where S and $V \setminus S$ is a non-singleton 1.1-approximate minimum cut in H_1 **do**
- 17 $v_S \leftarrow w_{H_1}(S, V \setminus S) + |\{e \in T : e \text{ is a cut edge between } S \text{ and } V \setminus S\}|$
- 18 **return** the minimum value of $\min v_S$ and $\min d_i$

We next consider the other case when $s \geq c \log n$. We will first show that, as the number of edges in H increases, the minimum cut size of H will increase to $\Theta(\log n)$ when $|H| = \Theta(\frac{|G| \log n}{s})$. We next analyze the value of $w_H(S, V \setminus S)$. Let \mathcal{S}_i ($1 \leq i \leq \log n$) denote the set of nodes such that

$$\mathcal{S}_i = \{S \subset V \mid s \cdot 2^{i-1} \leq w_G(S, V \setminus S) \leq s \cdot 2^i\}$$

For each \mathcal{S}_i , from Lemma 21 we know that $|\mathcal{S}_i| \leq O(n^{2^{i+1}})$ and for every $S \in \mathcal{S}_i$, from Lemma 29 we know that with probability at least $1 - 2e^{-c_2 c \cdot \log n \cdot 2^{i-1}} = 1 - 2n^{-c_2 c \cdot 2^{i-1}}$, $w_H(S, V \setminus S) \cdot \frac{|G|}{|H|}$ is a 1.1-approximation of $w_G(S, V \setminus S)$. Note that the constant c here can be sufficiently large, and after taking a union bound over all $S \in \mathcal{S}_i$ and all $\log n$ \mathcal{S}_i we get that with probability at least 9/10, for every $S \subset V$, $w_H(S, V \setminus S) \cdot \frac{|G|}{|H|}$ is a 1.1-approximation of $w_G(S, V \setminus S)$. From the above analysis we immediately know that when $\frac{|G|}{|H|} = \Theta(\frac{s}{\log n})$ from the subgraph H , we can learn the set

$$\mathcal{S} = \{S \subset V \mid s \leq w_G(S, V \setminus S) \leq 1.1^2 s\},$$

and the next step is choosing the true minimum cut among them. For every non-single node set $S \in \mathcal{S}$, as mentioned, since $w_H(S, V \setminus S)$ is $\Theta(\log n)$, setting $\varepsilon = 1/\log^2 n$ we can get its exact value. The remaining step is to estimate the value of $w_{G \setminus H}(S, V \setminus S)$. As described in Algorithm 2, after the previous step, we save every edge $e \in G \setminus H$, where e belongs to at least one of the non-singleton 1.1²-approximate minimum cuts in \mathcal{S} . We know the exact

value of $w_{G \setminus H}(S, V \setminus S)$ from this edge set. Putting the two things together, for every $S \in \mathcal{S}$ and $|S| \geq 2$, we know the value of $w_G(S, V \setminus S)$ after taking a sum of the two parts. For every singleton cut, we simply maintain the degree of each node during the stream. Thus, after taking the minimum value of the two parts (singleton and non-singleton cuts), we get the exact value of the minimum cut.

Space Complexity

In the first phase of the algorithm, we use one instance of the algorithm in Lemma 4 with $\varepsilon = 1/\log^2 n$, which takes space $\tilde{O}(n)$. In the second phase of the algorithm, we save all the edges that belong to the approximate minimum cut in H . From Lemma 30, we know that there are at most $O(n)$ such edges, and hence this part takes $O(n)$ words of space. We also maintain the degree of each node during the stream, which takes $O(n)$ words of space. Putting everything together, the space usage of our algorithm is $\tilde{O}(n)$.

Time Complexity

In the first phase of the algorithm, when one edge e comes, the update and query time of ALG_1 is polynomial, and we can also find the minimum cut in polynomial time. Hence, the update time here is still polynomial. In the second phase, when one edge e comes, we can enumerate all 1.1-approximate minimum cuts to check whether e belongs to one of the approximate minimum cuts (we can use $O(n^2)$ time to enumerate all approximate minimum cuts, see, e.g., [29]). Hence, we can do the update of this step in polynomial time. After all edges come, we can similarly enumerate all minimum cuts, from which the overall algorithm can be implemented in polynomial time.

5.2 Faster Update and Post-Processing Time

In the above algorithm, the runtime bottleneck is when a non-prefix edge e comes; we need to enumerate all approximate minimum cuts in the prefix graph to check whether to keep this edge. To get a faster runtime, we instead do a check when we collect n edges and consider a similar procedure to what we did in Section 4.4 where we use the recursive contraction algorithm with parameter $\alpha = 1 + \frac{1}{\log n}$. The difference is when doing the recursive contraction algorithm, we maintain the sketch SB where S is a k -sparse recovery matrix (see, e.g., [20]) with $k \log(n/k)$ rows where $k = O(\log n)$ and B is the edge-vertex matrix of the n edges we currently collected. Particularly, during the recursive contraction process, when we contract the nodes u, v , we merge the columns that u, v corresponds to in the sketch SB and replace them with their sum. Next, we consider each leaf node in the recursion, corresponding to one specific cut for the prefix graph. We can check whether it is a non-singleton and approximate minimum cut in $O(1)$ time from the information the algorithm keeps. Note that if without the sketch matrix S , for each edge e we want to check, it belongs to this specific cut if and only if the e -th coordinates of the remaining two columns of the edge-vertex matrix are 1 and -1 , and otherwise these two coordinates are both 0. Since the minimum cut of the prefix graph is $\Theta(\log n)$, we can get that both of the remaining columns are $O(\log n)$ -sparse before multiplying the sketch matrix S . This means that a k -sparse approximate recovery algorithm with $k = O(\log n)$ is sufficient to recover the indices of the non-zero coordinates of the remaining columns, which helps us to find the corresponding edges.

Time and Space Complexity

We analyze the time and space complexity of the above procedure. At each level of the recursion, since the sketch has $k \log(n/k)$ rows, we need $\tilde{O}(Nk)$ words of space to save the sketch, and it takes $\tilde{O}(k)$ time to combine two rows, where N is the number of the remaining nodes in the current level of the recursion. Recall that we need $O(n)$ space in each level and $O(N)$ time to contract two nodes in the original recursive contraction algorithm. This implies the modification only increases the time and space complexity by a factor of $\tilde{O}(k)$. Since the decoding time of the k -sparse recovery sketch is $k \cdot \text{polylog}(n)$ and $k = O(\log n)$, we have that the procedure has time complexity $\tilde{O}(n^2)$ and space complexity $\tilde{O}(n)$. Finally, note that when we collect a set of n edges in the suffix graph, we can do our above checking procedure while collecting the following n edges in the graph stream, which results in a strictly $O(n)$ update time of our algorithm. Similarly, suppose the minimum cut of the graph is $\Theta(c)$. When we do the post-processing, we can use the c -sparse recovery sketch when enumerating all the approximate minimum cuts, which results in a $\tilde{O}(n^2c)$ post-processing time. Putting everything together, we have the correctness of Theorem 14.

Lastly, we note that our algorithm can not only return the exact min-cut value but also collect all the edges that participate in any of the minimum cuts. When computing the for-all sparsifier of the prefix graph, for an edge that participates in at least one approximate min-cut, the probability that it will be sampled is $\Omega(1)$ as otherwise, the estimated cut value in such a cut will have two different values each with a constant probability, contradicting the guarantees of the for-all sparsifier. Hence, we can collect all of these edges from oversampling by a constant factor, which results in finding all the minimum cuts of the input graph and the edges crossing them.

References

- 1 Noga Alon. On the edge-expansion of graphs. *Combinatorics, Probability and Computing*, 6(2):145–152, 1997. URL: <http://journals.cambridge.org/action/displayAbstract?aid=46517>.
- 2 Alexandr Andoni, Jiecao Chen, Robert Krauthgamer, Bo Qin, David P. Woodruff, and Qin Zhang. On sketching quadratic forms. In Madhu Sudan, editor, *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science (ITCS)*, pages 311–319, 2016. doi:10.1145/2840728.2840753.
- 3 Sepehr Assadi, MohammadHossein Bateni, Aaron Bernstein, Vahab Mirrokni, and Cliff Stein. Coresets meet edcs: algorithms for matching and vertex cover on massive graphs. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1616–1635. SIAM, 2019. doi:10.1137/1.9781611975482.98.
- 4 Sepehr Assadi and Soheil Behnezhad. Beating Two-Thirds For Random-Order Streaming Matching. In Nikhil Bansal, Emanuela Merelli, and James Worrell, editors, *48th International Colloquium on Automata, Languages, and Programming (ICALP 2021)*, volume 198 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 19:1–19:13, Dagstuhl, Germany, 2021. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ICALP.2021.19.
- 5 Sepehr Assadi and Aditi Dudeja. *A Simple Semi-Streaming Algorithm for Global Minimum Cuts*, pages 172–180. Society for Industrial and Applied Mathematics, January 2021. doi:10.1137/1.9781611976496.19.
- 6 Joshua D. Batson, Daniel A. Spielman, and Nikhil Srivastava. Twice-Ramanujan sparsifiers. *SIAM J. Comput.*, 41(6):1704–1721, 2012. doi:10.1137/090772873.
- 7 András A. Benczúr and David R. Karger. Approximating s - t minimum cuts in $\tilde{O}(n^2)$ time. In Gary L. Miller, editor, *Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing, Philadelphia, Pennsylvania, USA, May 22-24, 1996*, pages 47–55. ACM, 1996. doi:10.1145/237814.237827.

- 8 András A. Benczúr and David R. Karger. Randomized approximation schemes for cuts and flows in capacitated graphs. *SIAM J. Comput.*, 44(2):290–319, 2015. doi:10.1137/070705970.
- 9 Aaron Bernstein. Improved bounds for matching in random-order streams. *Theory of Computing Systems*, pages 1–15, 2023.
- 10 Vladimir Braverman, Petros Drineas, Cameron Musco, Christopher Musco, Jalaj Upadhyay, David P. Woodruff, and Samson Zhou. Near optimal linear algebra in the online and sliding window models. In Sandy Irani, editor, *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020*, pages 517–528. IEEE, 2020. doi:10.1109/FOCS46700.2020.00055.
- 11 Charles Carlson, Alexandra Kolla, Nikhil Srivastava, and Luca Trevisan. Optimal lower bounds for sketching graph cuts. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2565–2569. SIAM, 2019. doi:10.1137/1.9781611975482.158.
- 12 Amit Chakrabarti, Graham Cormode, and Andrew McGregor. Robust lower bounds for communication and stream computation. *Theory of Computing*, 12(10):1–35, 2016. doi:10.4086/toc.2016.v012a010.
- 13 Amit Chakrabarti, Prantar Ghosh, Andrew McGregor, and Sofya Vorotnikova. Vertex ordering problems in directed graph streams. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1786–1802. SIAM, 2020. doi:10.1137/1.9781611975994.109.
- 14 Timothy Chu, Yu Gao, Richard Peng, Sushant Sachdeva, Saurabh Sawlani, and Junxing Wang. Graph sparsification, spectral sketches, and faster resistance computation, via short cycle decompositions. In Mikkel Thorup, editor, *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7-9, 2018*, pages 361–372. IEEE Computer Society, 2018. doi:10.1109/FOCS.2018.00042.
- 15 Michael B. Cohen, Cameron Musco, and Jakub Pachocki. Online row sampling. *Theory of Computing*, 16(15):1–25, 2020. doi:10.4086/toc.2020.v016a015.
- 16 Artur Czumaj, Hendrik Fichtenberger, Pan Peng, and Christian Sohler. Testable properties in general graphs and random order streaming. *arXiv preprint*, 2019. arXiv:1905.01644.
- 17 Matthew Ding, Alexandro Garces, Jason Li, Honghao Lin, Jelani Nelson, Vihan Shah, and David P. Woodruff. Space complexity of minimum cut problems in single-pass streams, 2024. arXiv:2412.01143.
- 18 Alireza Farhadi, Mohammad Taghi Hajiaghayi, Tung Mah, Anup Rao, and Ryan A Rossi. Approximate maximum matching in random streams. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1773–1785. SIAM, 2020.
- 19 Buddhima Gamlath, Sagar Kale, Slobodan Mitrovic, and Ola Svensson. Weighted matchings via unweighted augmentations. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*, pages 491–500, 2019. doi:10.1145/3293611.3331603.
- 20 Anna C. Gilbert, Yi Li, Ely Porat, and Martin J. Strauss. Approximate sparse recovery: Optimizing time and measurements. *SIAM J. Comput.*, 41(2):436–453, 2012. doi:10.1137/100816705.
- 21 Sudipto Guha and Andrew McGregor. Stream order and order statistics: Quantile estimation in random-order streams. *SIAM Journal on Computing*, 38(5):2044–2059, 2009. doi:10.1137/07069328X.
- 22 Diba Hashemi and Weronika Wrzoz-Kaminska. Weighted matching in the random-order streaming and robust communication models. *arXiv preprint*, 2024. doi:10.48550/arXiv.2408.15434.
- 23 Julien M. Hendrickx, Karl Henrik Johansson, Raphaël M. Jungers, Henrik Sandberg, and Kin Cheong Sou. Efficient computations of a security index for false data attacks in power networks. *IEEE Transactions on Automatic Control*, 59(12):3194–3208, 2014. doi:10.1109/TAC.2014.2351625.

- 24 Arun Jambulapati and Aaron Sidford. Efficient $\tilde{o}(n/\epsilon)$ spectral sketches for the laplacian and its pseudoinverse. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2487–2503. SIAM, 2018.
- 25 Michael Kapralov. Space lower bounds for approximating maximum matching in the edge arrival model. In *Proceedings of the Thirty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '21, pages 1874–1893, USA, 2021. Society for Industrial and Applied Mathematics. doi:10.1137/1.9781611976465.112.
- 26 Michael Kapralov, Sanjeev Khanna, and Madhu Sudan. Approximating matching size from random streams. In *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms*, pages 734–751. SIAM, 2014. doi:10.1137/1.9781611973402.55.
- 27 Michael Kapralov, Yin Tat Lee, Cameron Musco, Christopher Musco, and Aaron Sidford. Single pass spectral sparsification in dynamic streams. *SIAM J. Comput.*, 46(1):456–477, 2017. doi:10.1137/141002281.
- 28 Michael Kapralov, Aida Mousavifar, Cameron Musco, Christopher Musco, Navid Nouri, Aaron Sidford, and Jakab Tardos. Fast and space efficient spectral sparsification in dynamic streams. In *Proceedings of the Thirty-First Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '20, pages 1814–1833, USA, 2020. Society for Industrial and Applied Mathematics. doi:10.1137/1.9781611975994.111.
- 29 David R. Karger. Minimum cuts in near-linear time. *J. ACM*, 47(1):46–76, 2000. doi:10.1145/331605.331608.
- 30 David R. Karger and Matthew S. Levine. Random sampling in residual graphs. In John H. Reif, editor, *Proceedings on 34th Annual ACM Symposium on Theory of Computing, May 19-21, 2002, Montréal, Québec, Canada*, pages 63–66. ACM, 2002. doi:10.1145/509907.509918.
- 31 David R. Karger and Clifford Stein. A new approach to the minimum cut problem. *J. ACM*, 43(4):601–640, 1996. doi:10.1145/234533.234534.
- 32 Christian Konrad. A simple augmentation method for matchings with applications to streaming algorithms. In *43rd International Symposium on Mathematical Foundations of Computer Science, MFCS 2018*, pages 74–1. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, 2018.
- 33 Christian Konrad, Frédéric Magniez, and Claire Mathieu. Maximum matching in semi-streaming with few passes. In *International Workshop on Approximation Algorithms for Combinatorial Optimization*, pages 231–242. Springer, 2012. doi:10.1007/978-3-642-32512-0_20.
- 34 Eyal Kushilevitz and Noam Nisan. *Communication Complexity*. Cambridge University Press, 1996. doi:10.1017/CB09780511574948.
- 35 Aleksander Madry. Fast approximation algorithms for cut-based problems in undirected graphs. In *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, October 23-26, 2010, Las Vegas, Nevada, USA*, pages 245–254. IEEE Computer Society, 2010. doi:10.1109/FOCS.2010.30.
- 36 Andrew McGregor. Graph stream algorithms: A survey. *SIGMOD Rec.*, 43(1):9–20, May 2014. doi:10.1145/2627692.2627694.
- 37 Morteza Monemizadeh, Shan Muthukrishnan, Pan Peng, and Christian Sohler. Testable bounded degree graph properties are random order streamable. *arXiv preprint*, 2017. arXiv:1707.07334.
- 38 Sagnik Mukhopadhyay and Danupon Nanongkai. Weighted min-cut: sequential, cut-query, and streaming algorithms. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2020, Chicago, IL, USA, June 22-26, 2020*, pages 496–509. ACM, 2020. doi:10.1145/3357713.3384334.
- 39 J.I. Munro and M.S. Paterson. Selection and sorting with limited storage. *Theoretical Computer Science*, 12(3):315–323, 1980. doi:10.1016/0304-3975(80)90061-4.
- 40 Andrea Ottolini and Stefan Steinerberger. Concentration of hitting times in erdős-rényi graphs. *Journal of Graph Theory*, 2023.

- 41 Merav Parter and Eylon Yogev. Optimal short cycle decomposition in almost linear time. In *46th International Colloquium on Automata, Languages, and Programming, ICALP 2019*, Leibniz International Proceedings in Informatics, LIPIcs. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, July 2019. doi:10.4230/LIPIcs.ICALP.2019.89.
- 42 Pan Peng and Christian Sohler. Estimating graph parameters from random order streams. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2449–2466. SIAM, 2018. doi:10.1137/1.9781611975031.157.
- 43 Aviad Rubinfeld, Tselil Schramm, and S. Matthew Weinberg. Computing exact minimum cuts without knowing the graph. In Anna R. Karlin, editor, *9th Innovations in Theoretical Computer Science Conference, ITCS 2018, January 11-14, 2018, Cambridge, MA, USA*, volume 94 of *LIPIcs*, pages 39:1–39:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2018. doi:10.4230/LIPIcs.ITCS.2018.39.
- 44 Jonah Sherman. Breaking the multicommodity flow barrier for $o(\log n)$ -approximations to sparsest cut. In *50th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2009, October 25-27, 2009, Atlanta, Georgia, USA*, pages 363–372. IEEE Computer Society, 2009. doi:10.1109/FOCS.2009.66.
- 45 Daniel A. Spielman and Nikhil Srivastava. Graph sparsification by effective resistances. *SIAM J. Comput.*, 40(6):1913–1926, 2011. doi:10.1137/080734029.
- 46 Daniel A. Spielman and Shang-Hua Teng. Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In László Babai, editor, *Proceedings of the 36th Annual ACM Symposium on Theory of Computing, Chicago, IL, USA, June 13-16, 2004*, pages 81–90. ACM, 2004. doi:10.1145/1007352.1007372.
- 47 Daniel A. Spielman and Shang-Hua Teng. Spectral sparsification of graphs. *SIAM J. Comput.*, 40(4):981–1025, 2011. doi:10.1137/08074489X.
- 48 Satoshi Sugiura and Fumitaka Kurauchi. Isolation vulnerability analysis in road network: Edge connectivity and critical link sets. *Transportation Research Part D: Transport and Environment*, 119:103768, 2023.
- 49 Xiaoming Sun and David P. Woodruff. Tight Bounds for Graph Problems in Insertion Streams. In Naveen Garg, Klaus Jansen, Anup Rao, and José D. P. Rolim, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2015)*, volume 40 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 435–448, Dagstuhl, Germany, 2015. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.APPROX-RANDOM.2015.435.
- 50 Xinjue Wang, Ke Deng, Jianxin Li, Jeffery Xu Yu, Christian S Jensen, and Xiaochun Yang. Efficient targeted influence minimization in big social networks. *World Wide Web*, 23(4):2323–2340, 2020. doi:10.1007/S11280-019-00748-Z.
- 51 David P. Williamson. Lecture notes in spectral graph theory, October 2016. URL: <https://people.orie.cornell.edu/dpw/orie6334/Fall2016/lecture13.pdf>.
- 52 Mariano Zelke. Intractability of min- and max-cut in streaming graphs. *Information Processing Letters*, 111(3):145–150, 2011. doi:10.1016/J.IPL.2010.10.017.