

A Simple Framework for Finding Balanced Sparse Cuts via APSP*

Li Chen[†] Rasmus Kyng[‡] Maximilian Probst Gutenberg[§] Sushant Sachdeva[¶]

Abstract

We present a very simple and intuitive algorithm to find balanced sparse cuts in a graph via shortest-paths. Our algorithm combines a new multiplicative-weights framework for solving unit-weight multi-commodity flows with standard ball growing arguments.

Using Dijkstra’s algorithm for computing the shortest paths afresh every time gives a very simple algorithm that runs in time $\tilde{O}(m^2/\phi)$ and finds an $\tilde{O}(\phi)$ -sparse balanced cut, when the given graph has a ϕ -sparse balanced cut. Combining our algorithm with known deterministic data-structures for answering approximate All Pairs Shortest Paths (APSP) queries under increasing edge weights (decremental setting), we obtain a simple deterministic algorithm that finds $m^{o(1)}\phi$ -sparse balanced cuts in $m^{1+o(1)}/\phi$ time. Our deterministic almost-linear time algorithm matches the state-of-the-art in randomized and deterministic settings up to subpolynomial factors, while being significantly simpler to understand and analyze, especially compared to the only almost-linear time deterministic algorithm, a recent breakthrough by Chuzhoy-Gao-Li-Nanongkai-Peng-Saranurak (FOCS 2020).

1 Introduction

Graph partitioning is a fundamental algorithmic primitive that has been studied extensively. There are several ways to formalize the question. We focus on the question of finding balanced separators in a graph. More precisely, given an m -edge graph $G = (V, E)$, the conductance of a cut is defined by $\Phi_G(S) = \frac{|E_G(S, \bar{S})|}{\min\{\text{vol}(S), \text{vol}(V \setminus S)\}}$ where $E_G(S, \bar{S})$ is the set of edges with exactly one endpoint in S , and the volume of S , denoted $\text{vol}(S)$ is the sum of the degrees of vertices in S . We say that a cut $(S, V \setminus S)$ is b -balanced if $\text{vol}(S), \text{vol}(V \setminus S) \geq b \cdot \text{vol}(V)$. The objective in the Balanced Separator problem is

Given positive parameters $b, \phi, \alpha < 1$, either find a cut $(S, V \setminus S)$ that is b -balanced and has conductance $\Phi_G(S) \leq \phi$, or certify that every $\Omega(b)$ -balanced¹ cut has conductance at least $\alpha\phi$.

The Balanced Separator problem is a classic NP-hard problem and under the Small-Set-Expansion hypothesis, even NP-hard to approximate to within an arbitrary constant [29]. Thus, the above formulation allows for α -approximation for some $\alpha < 1$. This problem has been studied extensively due to its application to divide-and-conquer on graphs, and theoretical connections to random walks, spectral graph theory, and metric embeddings.

Our Results. In this paper, we present a very simple and intuitive algorithm for Balanced Separator. Our algorithm gives a simple framework based on (scalar) multiplicative weights that reduces the problem to computing approximate shortest paths in a graph under increasing lengths for the edges (decremental setting). Our framework either finds a balanced cut with small conductance, or certifies that every balanced cut has large conductance (Theorem 4.1).

If one simply uses Dijkstra’s algorithm to compute the necessary shortest paths afresh each time, our algorithm gives an $\tilde{O}(m^2/\phi)$ time algorithm that achieves approximation $\alpha = \Omega(1/\log^2 n)$ for cuts of constant balance, and

*The full version of the paper can be accessed at <https://arxiv.org/abs/2209.08845>

[†]Georgia Tech, lichen@gatech.edu. Li Chen was supported by NSF Grant CCF-2106444.

[‡]ETH Zurich, kyng@inf.ethz.ch. The research leading to these results has received funding from the grant “Algorithms and complexity for high-accuracy flows and convex optimization” (no. 200021 204787) of the Swiss National Science Foundation.

[§]ETH Zurich, maxprobst@ethz.ch. The research leading to these results has received funding from the grant “Algorithms and complexity for high-accuracy flows and convex optimization” (no. 200021 204787) of the Swiss National Science Foundation.

[¶]University of Toronto, sachdeva@cs.toronto.edu. Sushant Sachdeva’s research is supported by an NSERC (Natural Sciences and Engineering Research Council of Canada) Discovery Grant.

¹Note that we allow the algorithm to return an $\Omega(b)$ -balanced sparse cut when the graph has a b -balanced sparse cut. Such an algorithm is known as a pseudo-approximation algorithm. All known efficient algorithms for balanced cut find pseudo-approximations.

$\alpha = \Omega(1/\log n \cdot \log \log n)$ for cuts of constant balance and conductance (Theorem 2.3). If we instead use known $n^{o(1)}$ -approximate deterministic dynamic algorithms for decremental All-Pairs-Shortest-Paths (APSP), we obtain an algorithm that runs in $m^{1+o(1)}/\phi$ and achieves an approximation of $\alpha = 1/n^{o(1)}$ (Theorem 2.2). This algorithm has almost-optimal runtime in the size of the graph. We point out that algorithms are known that run in near-linear/almost-linear time in m with no dependency in ϕ , however, most applications consider the balanced cut problem for $b = 1/n^{o(1)}$ and $\phi = 1/n^{o(1)}$.

Our reduction also extends to directed graphs (see [8, 16] for the generalizations of conductance and sparse cuts to directed cuts; state-of-the-art and applications) where it reduces the problem to decremental All-Pairs Shortest Paths in a directed graph. Since Dijkstra’s algorithm works on directed graphs, the former result thus translates to directed graphs. The latter however relies on Data Structures that exploit the structure of the undirected graph and therefore do not extend. We focus henceforth on undirected graphs, unless specified explicitly.

Our algorithm can be described very simply. We attempt to embed an explicit expander H as a multi-commodity flow using paths of length $\tilde{O}(\phi^{-1})$ in G , while ensuring that the congestion on the edges in G is at most $\tilde{O}(\phi^{-1})$. If the endpoints of an edge $e \in H$ are connected in G using a short path, we use the path in G to route e . Further, we increase the length of each edge on this path by a multiplicative factor. This increased length makes it less likely that this path will be used in the future. A simple multiplicative-weights argument now allows us to bound the congestion over the course of entire algorithm. If our algorithm succeeds in embedding most edges of H in G , this provides us a certificate that all balanced cuts in G have conductance $\tilde{\Omega}(\phi)$. If our algorithm fails, we find several edges of H such that the endpoints of these edges are at distance $\tilde{\Omega}(\phi^{-1})$ as measured by the lengths of the edges computed by the algorithm. Now, we can apply a simple ball-growing argument to recover a balanced cut of conductance ϕ .

Applications. While finding the Balanced Sparsest Cut is a crucial ingredient in Divide-And-Conquer frameworks for many algorithms (see [33] for an introduction), and has various applications ranging from VLSI Design, Image Segmentation [32] to PRAM emulation, we want to point out in particular that our algorithm can be used to replace the use of the Cut-Matching framework [18] in the work of Saranurak-Wang [30] (see Remark 4.2 in Section 4.1). Together, this gives an elegant framework for computing expander decompositions which in turn have been pivotal in various recent breakthroughs in algorithmic graph theory with applications to computing Electric Flows [34], Maximum Flows and Min-Cost Flows [10], Gomory-Hu Trees [2, 1] for finding Global Min-Cuts deterministically [17, 21, 20], and many, many more.

Comparison to Previous Works. There has been a lot of work on algorithms for Balanced Separator. The celebrated work of Leighton and Rao [19] showed that one could achieve an $O(\log n)$ approximation to Balanced Separator by repeatedly solving a linear program that computes a fractional multi-commodity flow. Several works give a faster implementation of this approach via multiplicative-weights algorithms for multi-commodity flow [28, 35, 15, 13], and by using the Leighton-Rao result as a black-box to deduce that they compute an $O(\log n)$ approximation. However, the running time they achieved for Balanced Separator was $\Omega(nm^2)$ since they repeatedly find and remove low-conductance cuts, each of which might be highly unbalanced, possibly introducing a factor of n . In contrast, our algorithm works directly with balanced cuts, rather than multi-commodity flows. Our algorithm is in the same spirit as the Garg-Könemann, Fleischer framework from [15, 13], but directly incorporates the Leighton-Rao algorithm for finding low conductance cuts.

The groundbreaking work of Spielman and Teng on solving Laplacian linear systems [34] introduced the notion of local algorithms for finding low-conductance cuts, where the running time of the algorithm scales almost-linearly with the smaller size of the output cut. Thus the algorithm can be applied repeatedly to find balanced cuts in almost-linear time. Inspired by this work, multiple local algorithms were proposed [3, 4]. While all these algorithms are fast, and almost-linear in running time $m^{1+o(1)}/\phi$, they are inherently randomized, and the balanced cut found has conductance $\tilde{\Omega}(\sqrt{\phi})$. In contrast, our algorithm is deterministic, and finds a cut of conductance at most $\phi \cdot m^{o(1)}$.

Another line of work develops fast SDP algorithms based on matrix-multiplicative weights. The most popular of these is the Cut-Matching framework of Khandekar-Rao-Vazirani [18] which achieves near-optimal running time $\tilde{O}(m)$ on undirected graphs by using [27] as a subroutine, and running time $m^{1+o(1)}$ on directed graphs via [23, 10]. Note that these algorithms do not have a dependency in ϕ . Inspired by [18], several works [6, 25, 26, 24] obtained almost-linear time algorithms for Balanced Separator building on the matrix-multiplicative weights framework. While the cut-matching framework and the resulting algorithms are elegant, they rely on rather involved techniques that are non-intuitive. The celebrated work of Arora-Rao-Vazirani [7] obtained an $O(\sqrt{\log n})$

approximation for Balanced Separator via an SDP based algorithm. Faster algorithms built on their ideas [5, 31] achieved almost-linear running time with $O(\sqrt{\log n})$ approximation. However, these algorithms are very involved, based on matrix-multiplicative weights, randomized, and rely on near-linear time (approximate) max-flow. Our algorithm and analysis work with scalar multiplicative weights and are very simple to understand. Further, our algorithm only need to invoke approximate shortest-path oracles under increasing edge weights.

The only previous deterministic, almost-linear time approximation algorithm for Balanced Separator was given recently by Chuzhoy-Gao-Li-Nanongkai-Peng-Saranurak [12] achieving runtime $m^{1+o(1)}$ which is again independent of ϕ . Their algorithm relies on a rather intricate recursive scheme that implicitly uses at each recursion level a reduction to decremental APSP. But even the analysis on a single level relies on the rather involved expander pruning framework. In contrast to their work, the simplicity of our algorithm and analysis stands out.

We also point out that a generalization of [12] to weighted graphs was given by Li and Saranurak [22]. This algorithm implicitly uses [12], and is therefore even more involved.

2 Main Result

We formally state our results in this section. Our main result is the following theorem.

THEOREM 2.1. *Given an n -vertex, m -edge graph G , an α_{APSP} -approx decremental APSP algorithm and conductance parameter ϕ and balance parameter $b \in [1/n, 1/4]$, the algorithm `LOWCONDUCTANCECUTCERTIFY`(G, ϕ, b) either*

1. *Returns a cut (S, \bar{S}) with $\text{vol}_G(S), \text{vol}_G(\bar{S}) \geq b \cdot \text{vol}(G)$ with conductance $\Phi_G(S) \leq \phi$, or*
2. *Certifies that every cut (X, \bar{X}) with $\text{vol}_G(X), \text{vol}_G(\bar{X}) = \Omega(b \cdot \text{vol}_G(G))$ has conductance at least $\phi \cdot \Omega\left(\frac{1}{\alpha_{APSP} \log n \cdot \log(1/b) \cdot \log(\log(n) \alpha_{APSP} / (b\phi))}\right) = \phi \cdot \Omega\left(\frac{1}{\alpha_{APSP} \log^3(n)}\right)$.*

The algorithm is deterministic and requires the APSP data structure to undergo $O(\alpha_{APSP} \cdot m\phi^{-1} \log^3 n)$ updates, queries it $O(m)$ times and spends an additional $O(\alpha_{APSP} \cdot m\phi^{-1} \log^3 n)$ time.

REMARK 2.1. *Our algorithm requires the APSP data structure to work against an adaptive adversary, i.e. the generated update sequence relies on the information obtained from queries to the data structure.*

REMARK 2.2. *APSP data structures often answer queries in time proportional to the number of edges on the approximate shortest path that they return. Our algorithm ensures that the number of such edges on all paths is bounded by $O(\alpha_{APSP} \cdot m\phi^{-1} \log^3 n)$.*

We note that for computing balanced cuts (i.e. cuts where b is constant) which is arguably the most interesting case, our approximation guarantee becomes $\Omega(1/\alpha_{APSP} \log n \cdot \log(\alpha_{APSP} \phi^{-1} \cdot \log n))$. For a decremental APSP data structure with constant-approximation and $\phi \geq \Omega(1/\log^{O(1)} n)$, this further simplifies to $\Omega(1/\log n \log \log n)$.

Using the efficient $n^{o(1)}$ -approximate decremental APSP data structure from [9] or [11], we obtain the following result²:

THEOREM 2.2. *Given an n -vertex, m -edge graph G , a conductance parameter ϕ and balance parameter $b \in [1/n, 1/4]$, there is an algorithm `LOWCONDUCTANCECUTCERTIFY`(G, ϕ, b) that can either*

1. *Find a cut (S, \bar{S}) with $\text{vol}_G(S), \text{vol}_G(\bar{S}) \geq b \cdot \text{vol}(G)$ with conductance $\Phi_G(S) \leq \phi$, or*
2. *Certify that every cut (X, \bar{X}) with $\text{vol}_G(X), \text{vol}_G(\bar{X}) = \Omega(b \cdot \text{vol}_G(G))$ has conductance $\phi/n^{o(1)}$.*

The algorithm is deterministic and runs in $m^{1+o(1)}/\phi$ time.

On the other hand, one can run Dijkstra's shortest path algorithm for every query and obtain the following:

²We remark that both data structures [9, 11] implicitly rely on the framework of Chuzhoy-Gao-Li-Nanongkai-Peng-Saranurak [12], thus, our reduction in combination with these data structures does not yield a simpler algorithm in itself. We are however optimistic that simpler data structures for the decremental APSP problem are available in the future that do not necessarily rely on expander techniques.

THEOREM 2.3. *Given an n -vertex, m -edge graph G , a conductance parameter ϕ and balance parameter $b \in [1/n, 1/4]$, there is a deterministic algorithm $\text{LOWCONDUCTANCECUTCERTIFY}(G, \phi, b)$ that can either*

1. *Find a cut (S, \bar{S}) with $\text{vol}_G(S), \text{vol}_G(\bar{S}) \geq b \cdot \text{vol}(G)$ with conductance $\Phi_G(S) \leq \phi$, or*
2. *Certify that every cut (X, \bar{X}) with $\text{vol}_G(X), \text{vol}_G(\bar{X}) = \Omega(b \cdot \text{vol}_G(G))$ has conductance $\phi \cdot \Omega\left(\frac{1}{\log n \cdot \log(1/b) \cdot \log(\log(n)/(\phi b))}\right)$.*

The algorithm is deterministic and runs in $\tilde{O}(m^2/\phi)$ time.

3 Preliminaries

Sparsity and Expanders. In this article, we consider an undirected n -vertex graph $G = (V, E)$. For such a graph, we define the sparsity of a cut $\emptyset \subsetneq S \subsetneq V$ by $\Psi_G(S) = \frac{|E_G(S, \bar{S})|}{\min\{|S|, |\bar{S}|\}}$ where $E_G(A, B)$ is defined to be the set of edges with exactly one endpoint in A and B respectively. For any set $S \subseteq V(G)$, we sometimes use the shorthand $E_G(S)$ in place of $E_G(S, \bar{S} = V \setminus S)$. The sparsity of a graph G is defined $\Psi(G) = \min_{\emptyset \subsetneq S \subsetneq V} \Psi(S)$. If G contains no ψ -sparse cut, we say that G is a ψ -expander.

Conductance vs. Sparsity. Via a simple reduction replacing each vertex of degree d with an explicit expander graph on d vertices (see appendix A), we can reduce to the case where every vertex has degree at most 10. In such a graph, for any set $S \subseteq V$, $|S| \leq \text{vol}(S) \leq 10|S|$, and thus, instead of conductance $\Phi_G(S) = \frac{|E_G(S, \bar{S})|}{\min\{\text{vol}(S), \text{vol}(V \setminus S)\}}$, we can work with sparsity $\Psi_G(S) = \frac{|E_G(S, \bar{S})|}{\min\{|S|, |V \setminus S|\}}$. Throughout the rest of the article, we will therefore work with sparsity instead of conductance.

Expander Constructions. Given any n , there is a deterministic construction of a $\Omega(1)$ -expander on n vertices of bounded degree. This will be an essential tool used in our proof and we use ψ_0 to denote the universal lower bound on the sparsity of such family of expanders.

THEOREM 3.1. (SEE THM. 2.4 OF [12] BASED ON THM 2 OF [14].) *There is an universal constant $\psi_0 \in (0, 1)$ and an algorithm $\text{CONSTDEGEXPANDER}(n)$ that returns a ψ_0 -expander H on a vertex set of size n with maximum degree 9. The algorithm runs in time $O(n)$.*

REMARK 3.1. *While deterministic algorithms to construct a constant-degree, constant sparsity expander require rather involved proof techniques, we prove in Appendix C a simple randomized algorithm to construct a $O(\log n)$ -degree $\Omega(\log n)$ -expander H in $O(n \log n)$ time. Using this randomized algorithm in place of the above theorem only affects guarantees of our overall algorithm by polylogarithmic factors.*

Graph Embeddings. Given graphs H and G that are defined over the same vertex set, then we say that a function $\Pi_{H \rightarrow G}$ is an *embedding* if it maps each edge $(u, v) \in H$ to a u -to- v path $P_{u,v} = \Pi_{H \rightarrow G}(u, v)$ in G . We say that the *congestion* of $\Pi_{H \rightarrow G}$ is the maximum number of times that any edge $e \in E(G)$ appears on any embedding path:

$$\text{cong}(\Pi_{H \rightarrow G}) = \max_{e \in E(G)} |\{e' \in E(H) \mid e \in \Pi_{H \rightarrow G}(e')\}|.$$

Certifying Expander Graphs via Embeddings. Graph embeddings are useful since they allow us to argue that if we can embed a graph H that is known to be an expander into a graph G , then we can reason about the sparsity of G , as shown below.

LEMMA 3.1. *Given a ψ -expander graph H and an embedding of H into G with congestion C , then G must be an $\Omega\left(\frac{\psi}{C}\right)$ -expander.*

Proof. Consider any cut $(S, V \setminus S)$ with $|S| \leq |V \setminus S|$. Since H is a ψ -expander, we have that $|E_H(S, V \setminus S)| \geq \psi|S|$. We also know by the embedding of H into G , that for each edge $(u, v) \in E_H(S, V \setminus S)$, we can find path a $P_{u,v}$ in G that also has to cross the cut $(S, V \setminus S)$ at least once. But since each edge in G is on at most C such paths, we can conclude that at least $|E_H(S, V \setminus S)|/C \geq \psi|S|/C$ edges in G cross the cut $(S, V \setminus S)$. \square

We use the following generalization of this Folklore result to balanced sparse cuts.

LEMMA 3.2. *Given a ψ -expander graph H , a subgraph $H' \subseteq H$ with $|E(H \setminus H')| \leq \frac{\psi}{2}bn$ for some $b \in [0, 1]$ and an embedding $\Pi_{H' \rightarrow G}$ of H' into G with congestion C , then for all cuts (S, \bar{S}) where $bn \leq |S| \leq n/2$, we have $\Psi_G(S) = \Omega\left(\frac{\psi}{C}\right)$.*

Proof. Observe that for each such (S, \bar{S}) , we have $|E_{H'}(S, \bar{S})| \geq |E_H(S, \bar{S})| - |E(H \setminus H')| \geq \psi|S| - \frac{\psi}{2}bn \geq \frac{\psi}{2}|S|$. Using the same argument as above, the cut size of S in G is at least $|E_G(S, \bar{S})| \geq |E_{H'}(S, \bar{S})|/C \geq \psi|S|/2C$. \square

Decremental All-Pairs Shortest-Paths (APSP). A decremental α_{APSP} -approximate All-Pairs Shortest-Paths (APSP) data structure (abbreviated α_{APSP} -APSP) is a data structure that is initialized to an m -edge n -vertex graph G and supports the following operations:

- **INCREASEEDGEWEIGHT** (u, v, Δ) : increases the edge weight of (u, v) by Δ .
- **QUERYDISTANCE** (u, v) : for any $u, v \in V$ returns a distance estimate $\tilde{d}(u, v)$ that α_{APSP} -approximates the distance from u to v in the current graph G denoted $d_G(u, v)$, i.e. $\tilde{d}(u, v) \in [d_G(u, v), \alpha_{\text{APSP}} \cdot d_G(u, v)]$.
- **QUERYPATH** (u, v) : returns a path π from u to v in the current graph G of total weight $\tilde{d}(u, v)$ (that is the value of the distance estimate if queried).

We denote the total time required by the data structure to execute a series of q queries and u update operations on an n -vertex constant-degree graph by $T_{\text{APSP}}(q, u)$.

Recently, deterministic $n^{o(1)}$ -approximate APSP data structures have been developed (see [11, 9]) that process any sequence of $\tilde{O}(m)$ edge weight increases in total time $m^{1+o(1)}$ while answering distance queries in time $n^{o(1)}$ time and for a path query, returns paths in time near-linear in the number of edges on the path (i.e. if it returns a path P , it takes at most time $|P|n^{o(1)}$). We conjecture that in the near-future, $O(\log n)$ -APSP data structures are found that implement edge weight increases in time $\tilde{O}(m)$ and answers distance queries in time $\tilde{O}(1)$ and path queries in time $\tilde{O}(|P|)$.

4 Our Algorithm

In this section, we present an algorithm to find sparse cuts with respect to sparsity or embed an expander into a constant-degree graph G . By standard reductions (given in Appendix A and Appendix B), one can translate between sparsity and conductance and remove the bounded-degree assumption, both with only a constant loss in quality. Thus, by proving the theorem below, we directly establish our main result, Theorem 2.1.

THEOREM 4.1. *Given a graph G of degree at most 10, an α_{APSP} -approx decremental APSP algorithm and sparsity parameter ψ and balance parameter $b \in [1/n, 1/4]$, there is an algorithm $\text{SPARSECUTCERTIFY}(G, \psi, b)$ (Algorithm 2) that can either*

1. Find a cut (S, \bar{S}) with $|S|, |\bar{S}| \geq bn$ of sparsity $\leq \psi$, or
2. Certify that every cut (X, \bar{X}) with $|X|, |\bar{X}| = \Omega(bn)$ has sparsity $\psi \cdot \Omega\left(\frac{1}{\alpha_{\text{APSP}} \log n \cdot \log(1/b) \cdot \log(\log(n) \alpha_{\text{APSP}} / (b\psi))}\right)$.

The algorithm is deterministic and requires the APSP data structure to undergo $O(\alpha_{\text{APSP}} \cdot n/\psi \log^3 n)$ updates, queries it $O(n)$ times and spends an additional $O(\alpha_{\text{APSP}} \cdot n/\psi \log^3 n)$ time.

REMARK 4.1. *Our algorithm ensures that the total number of edges summed across all queried paths is bounded by $O(\alpha_{\text{APSP}} \cdot n/\psi \log^3 n)$.*

The algorithm contains two phases. The first phase tries to embed an $\Omega(1)$ -expander into the input graph G with congestion $\tilde{O}(1/\psi)$. Let F be the subset of expander-edges the algorithm cannot embed. If $|F| = O(bn)$, i.e. the algorithm embed all but $O(bn)$ edges, Lemma 3.2 ensures that every b -balanced cut has sparsity $\tilde{\Omega}(\psi)$. Otherwise, $|F| = \Omega(bn)$ and the algorithm outputs an edge weight \mathbf{w} such that every $(u, v) \in F$ are far apart w.r.t. \mathbf{w} . In this case, the second phase is initiated to extract a sparse $\Omega(b)$ -balanced cut from these far-apart pairs of vertices.

4.1 An Algorithm to Separate Or Certify First, we present the algorithm for the first phase that either embeds a large portion of an expander or finds a large set of far-apart vertex-pairs w.r.t. some edge weights \mathbf{w} .

LEMMA 4.1. *Given an α_{APSP} -APSP data structure, two graphs G and H over the same vertex set V where G has maximum degree 10, a congestion parameter $C \in [1, n]$, and a balance parameter $b \in [1/n, 1/2]$. The algorithm SEPARATEORCERTIFY(G, H, C, b) (Algorithm 1) outputs either*

1. A set of weights $\mathbf{w} \in \mathbb{R}_{\geq 1}^{E(G)}$ with $\|\mathbf{w}\|_1 \leq 20n$, a number $b' \in [b, 1/2]$, and a subset of edges $F \subseteq E(H)$ with $|F| > 10b'n$ such that

$$\forall (u, v) \in F, \quad \text{dist}_{\mathbf{w}}(u, v) > \frac{C}{b'}, \text{ or}$$

2. A graph $H' \subseteq H$ with $|E(H) \setminus E(H')| \leq 10bn$ and an embedding $\Pi_{H' \rightarrow G}$ that maps each edge (u, v) in H' to a uv -path in G with congestion $O(C \cdot \alpha_{APSP} \cdot \log(1/b) \cdot \log(C \cdot \alpha_{APSP}/b))$.

The algorithm is deterministic and requires the APSP data structure to undergo $O(C\alpha_{APSP}n \log^2 n)$ edge updates and $O(n)$ distance queries along with additional $O(C\alpha_{APSP}n \log^2 n)$ time.

Algorithm 1: SEPARATEORCERTIFY(G, H, C, b)

```

1  $H' = (V, \emptyset); \Pi_{H' \rightarrow G} \leftarrow \emptyset; \mathbf{w} \leftarrow \mathbf{1}^{|E(G)|}; \eta \leftarrow \frac{1}{4C\alpha_{APSP} \log_2(10/b)}$ .
2 Maintain an  $\alpha_{APSP}$ -approximate APSP data structure on  $G$  weighted by  $\mathbf{w}$ .
3 for  $i = 0, 1, \dots, \lfloor \log_2(1/b) \rfloor$  do
4   foreach  $e = (u, v) \in E(H) \setminus E(H')$  do
5     if APSP.QUERYDIST( $u, v$ )  $\leq 2^i \cdot C\alpha_{APSP}$  then
6       Add  $e$  to  $H'$ ;  $\Pi_{H' \rightarrow G}(e) \leftarrow$  APSP.QUERYPATH( $u, v$ ).
7       foreach  $f \in \Pi_{H' \rightarrow G}(e)$  do
8         APSP.INCREASEWEIGHT( $e, \eta\mathbf{w}_e$ );  $\mathbf{w}_e \leftarrow (1 + \eta)\mathbf{w}_e$ .
9       end
10    end
11  end
12  if  $|E(H) \setminus E(H')| > 10n/2^i$  then return  $(\mathbf{w}, 2^{-i}, E(H) \setminus E(H'))$ .
13 end
14 return  $(H', \Pi_{H' \rightarrow G})$ .

```

The Algorithm. Algorithm 1 implements SEPARATEORCERTIFY(G, H, C, b). Here, the task of finding an embedding of H into G is interpreted as a multicommodity flow problem, that is each edge $(u, v) \in H$ gives rise to the demand to route one unit of flow from u to v . Later, we use a ψ_0 -expander in place of H .

The goal of the algorithm is to find such an embedding/ multicommodity flow with small congestion which combined with our choice of H certifies that G is a good (almost) expander (i.e. contains no balanced sparse cut). Here, we guess the congestion to be roughly C and want to enforce $\text{cong}(\Pi_{H \rightarrow G}) \leq C$. In fact, we even provide a slightly tighter analysis.

To achieve this goal, we use a technique which is an instance of the *Multiplicative Weight Update (MWU)* framework. Initially, we define a uniform weight function \mathbf{w} with weights over G . We try to embed each edge $(u, v) \in E(H)$ using a short uv -path P_{uv} in G with respect to \mathbf{w} . Whenever we embed an edge (u, v) in such a way and the path P_{uv} contains an edge $e \in E(G)$, we increase the weight \mathbf{w}_e by a multiplicative factor $(1 + \eta)$. Naturally, after t edges have been embedded by using the edge e , we have scaled up the weight of e by a factor of $(1 + \eta)^t$. Using $e^x \leq (1 + 2x), x \in [0, 1]$, and setting $\eta \approx C$ ensures that the weight \mathbf{w}_e approaches a large polynomial in n for $t \gg 2\eta \log n$ (which again is $\approx C$).

At the same time, the algorithm only embeds edges $(u, v) \in E(H)$ if the distance between the endpoints in G w.r.t. \mathbf{w} is small. This ensures that $\|\mathbf{w}\|_1 = O(n \log(1/b))$ and that we never use an edge e into which many embedding paths are already routed.

More precisely, we proceed in rounds to embed edges in H . At later rounds (i.e. when i large), we have already embed a large number of edges in H . Since the number of remaining edges is small, we allow for them

to be embed with slightly longer paths which still lets us argue that $\|\mathbf{w}\|_1$ is increased by at most $O(n)$ in the current round. If in any round, it is not possible to embed many of the remaining edges with paths of weight at most the current threshold, we can simply return these edges and end up in the first scenario.

Correctness (Returning in Line 12). We start by proving the following claim which then immediately establishes correctness if Algorithm 1 terminates at Line 12 (i.e. in the second scenario).

INVARIANT 4.1. *After the i -th iteration of the for-loop in Line 3, we have $\|\mathbf{w}\|_1 \leq 10n(1+2\eta C\alpha_{\text{APSP}}(i+1)) \leq 20n$.*

Proof. Initially, $\|\mathbf{w}\|_1 = \|\mathbf{1}^{E(G)}\|_1 \leq 10n$.

To gauge the increase in $\|\mathbf{w}\|_1$ during the i -th iteration of the for-loop, consider the effect of embedding a new edge e in the foreach-loop starting in Line 4 (we only consider such iterations if the if-statement in Line 5 evaluates true as otherwise \mathbf{w} does not change). Letting \mathbf{w}^{OLD} denote \mathbf{w} just before the foreach-loop iteration and \mathbf{w}^{NEW} right after. We clearly have that $\|\mathbf{w}^{NEW}\|_1 = \|\mathbf{w}^{OLD}\|_1 + \eta \cdot \mathbf{w}^{OLD}(\Pi_{H' \rightarrow G}(e))$ from Line 8. But since the if-statement was true, we have that $\mathbf{w}^{OLD}(\Pi_{H' \rightarrow G}(e)) \leq 2^i \cdot C\alpha_{\text{APSP}}$. We conclude that each edge that is newly embed increases $\|\mathbf{w}\|_1$ by at most $\eta \cdot 2^i \cdot C\alpha_{\text{APSP}}$.

At the beginning of the i -th iteration of the for-loop, there are at most $10n/2^i$ edges in $E(H) \setminus E(H')$. At the very first iteration $i = 0$, $|E(H)| \leq 20n$ as the max degree of H is at most 10. Later, $|E(H) \setminus E(H')| \leq 10n/2^{i-1}$ holds or otherwise the algorithm would terminate after the $(i-1)$ -th iteration in Line 12. Thus, during the i -th iteration, the foreach-loop in Line 4 iterates over at most $10n/2^{i-1}$ edges as well. We can bound the total increase of $\|\mathbf{w}\|_1$ during the i -th iteration by

$$\frac{10n}{2^{i-1}} \cdot \eta \cdot 2^i \cdot C\alpha_{\text{APSP}} = 20n \cdot \eta C\alpha_{\text{APSP}}.$$

The total number of iterations is at most $\lceil \log_2(1/b) \rceil + 1$. This establish the second inequality using the definition of η . \square

Note that for every edge (u, v) that is in $E(H) \setminus E(H')$ when the algorithm returns in Line 12, the preceding foreach-loop iterated over (u, v) and found that $\text{APSP.QUERYDIST}(u, v) > 2^i \cdot C\alpha_{\text{APSP}}$ (as otherwise (u, v) would have been added to $E(H')$). But this implies that $\text{dist}_{\mathbf{w}}(u, v) > 2^i \cdot C = C/b'$ by our choice of b' . To establish correctness, it only remains to use the if-condition preceding Line 12 and observe that the condition does not hold when $i = 0$.

Correctness (Returning in Line 14). It is straight-forward to see from Algorithm 1 that $\Pi_{H' \rightarrow G}$ is a correct embedding from H' to G and that $|E(H) \setminus E(H')| \leq 10bn$. It thus only remains to bound the congestion of $\Pi_{H' \rightarrow G}$.

LEMMA 4.2. *The congestion of $\Pi_{H' \rightarrow G}$ is at most $\frac{2 \log(2C\alpha_{\text{APSP}}/b)}{\eta}$.*

Proof. Let us fix any edge $e \in E(G)$. Note that each time we add an embedding path in the foreach-loop starting in Line 4 that contains e , we increase the weight \mathbf{w}_e to $(1 + \eta)\mathbf{w}_e$. Since initially, $\mathbf{w}_e = 1$, we have that after t times that the edge e was used to embed an edge in the foreach-loop, we have that $\mathbf{w}_e = (1 + \eta)^t \geq e^{t\eta/2}$ since $e^x \leq 1 + 2x$ for $x \in [0, 1]$. In particular, if the algorithm embeds t times into e for $t > \frac{2 \log(2C\alpha_{\text{APSP}}/b)}{\eta}$, then at the end of the algorithm, we would have $\mathbf{w}_e > \frac{2C\alpha_{\text{APSP}}}{b}$.

However, note that by the if-condition in Line 5, we never embed into an edge e that has weight more than $2^{\log_2(1/b)} \cdot C\alpha_{\text{APSP}} = \frac{C\alpha_{\text{APSP}}}{b}$ since otherwise the path using this edge has higher weight. We can thus conclude that at the end of the algorithm, $\mathbf{w}_e \leq (1 + \eta) \frac{C\alpha_{\text{APSP}}}{b} \leq \frac{2C\alpha_{\text{APSP}}}{b}$, which leads to a contradiction. \square

Run time Analysis. The for-loop of the algorithm runs at most $O(\log(1/b))$ iterations and in the i^{th} iteration at most $O(n/2^i)$ edges are iterated over in the foreach-loop starting in Line 4. Thus, the total number of queries to the APSP data structure can be bounded by $O(\sum_i n/2^i) = O(n)$.

The time the algorithm spends updating the weights in Line 8 can be bounded by observing that each edge e has its weight increased only after an additional embedding path was added through e ; but the congestion is bounded by $O(\log n/\eta)$ by Lemma 4.2, thus the foreach-loop is executed at most $O(n \log n/\eta)$ times over the entire course of the algorithm. This concludes our analysis of the number of updates to the APSP data structure. The runtime analysis of the algorithm follows along the same line of reasoning.

REMARK 4.2. Our algorithm can be extended to compute expander decompositions, following the approach of [30]. We refer the reader to this paper for additional background and the necessary definitions. For readers familiar with [30], we briefly describe the key step we need to implement: When $\text{SEPARATEORCERTIFY}(G, H, C, b)$ certifies that most edges in the expander H can be embedded into G (and hence by Lemma 3.2 there are no sparse balanced cuts in G) then we need to be able to extract a large expander from G so that we only need to recurse on a small (potentially) non-expanding part. To find an induced subgraph with large expansion, we first produce a new graph G' by adding the edges $E(H) \setminus E(H')$ to G . This ensures that G' is a good expander. We then use the expander pruning of [30] to delete the same edges $E(H) \setminus E(H')$ from G' , resulting in a large leftover expander G'' with vertex set V'' . By construction $G[V'']$ is now a large expander.

4.2 Extracting the Sparsest Cut In order to prove Theorem 4.1, we now have to show how to extract a sparsest cut from the weight function that is returned in case no embedding is found. We point out that in order to do so it is significantly more convenient to work with an integral weight function \mathbf{w} . We therefore round the weight function that we obtain Lemma 4.1 up which might result in $\|\mathbf{w}\|_1$ being at most twice as large as stated.

We use the following auxiliary algorithm that finds a cut with few edges crossing given any two vertices at large distance.

CLAIM 4.1. The procedure $\text{FINDTHINLAYER}(G, \mathbf{w}, u, v, D)$ takes a graph G weighted by $\mathbf{w} \in \mathbb{N}_{\geq 1}^{E(G)}$ and two vertices u, v such that $\text{dist}_{\mathbf{w}}(u, v) > D$ for some integer $D > 4 \log_2 \|\mathbf{w}\|_1$. It returns a set of vertices $S \neq \emptyset$ such that $|S| \leq |V|/2$ and $|E_G(S, V \setminus S)| \leq \frac{4\mathbf{w}(E_G(S)) \log_2 \|\mathbf{w}\|_1}{D}$. The algorithm runs in time $O(|E_G(S)| \log |E_G(S)|)$.

Given this auxiliary algorithm, we can state the final algorithm and prove our main result, Theorem 4.1. As described before, we use the algorithm $\text{SEPARATEORCERTIFY}(G, H, C, \hat{b})$ with a constant degree, constant sparsity expander H . It is straight-forward to conclude that G contains no balanced sparse cuts, if the procedure can embed H .

Otherwise, we take the weight function and repeatedly find a separator between the endpoints of edges in F that are far from each other (using the auxiliary algorithm). Note that if there are roughly $b'n$ edges in F at distance roughly C/b' , then using the auxiliary algorithm repeatedly with $D \approx C/b'$, produces a cut where the smaller side has $\Omega(|F|) = \Omega(b'n)$ vertices. Using the guarantees from the auxiliary procedure, we further have that the number of edges in the induced cut are at most $\tilde{O}(b'n/C)$. Thus, the sparsity of the cut must be $\tilde{O}(1/C)$ where $C \approx 1/\psi$ by our choice of parameters.

Algorithm 2: $\text{SPARSECUTORCERTIFY}(G, \psi, b)$

```

1  $H \leftarrow \text{CONSTDEGEXPANDER}(|V(G)|)$ ;  $C \leftarrow 320 \log n / \psi$ ;
2 if  $\text{SEPARATEORCERTIFY}(G, H, C, 2b)$  returns  $(H', \Pi_{H' \rightarrow G})$  then
3   | return  $(H', \Pi_{H' \rightarrow G})$ .
4 end
5 else // i.e. if it returns  $(\mathbf{w}, b', F)$ 
6   |  $\hat{\mathbf{w}} \leftarrow \lceil \mathbf{w} \rceil$ .
7   |  $X \leftarrow V(G)$ .
8   |  $D \leftarrow 2C/b'$ .
9   | while  $\exists (u, v) \in H[X] \cap F$  and  $|V \setminus X| \leq n/4$  do
10  | | //  $\text{dist}_{\hat{\mathbf{w}}}(u, v) > D$ 
11  | |  $S \leftarrow \text{FINDTHINLAYER}(G[X], \hat{\mathbf{w}}, u, v, D)$ .
12  | |  $X \leftarrow X \setminus S$ .
13  | end
14  | return  $V \setminus X$ .
15 end

```

THEOREM 4.2. Given a graph G of degree at most 10, an α_{APSP} -approx decremental APSP algorithm and sparsity parameter ψ and balance parameter $b \in [1/n, 1/4]$, there is an algorithm $\text{SPARSECUTORCERTIFY}(G, \psi, b)$ (Algorithm 2) that can either

1. Find a cut (S, \bar{S}) with $|S|, |\bar{S}| \geq bn$ of sparsity $\leq \psi$, or
2. Certify that every cut (X, \bar{X}) with $|X|, |\bar{X}| = \Omega(bn)$ has sparsity $\psi \cdot \Omega\left(\frac{1}{\alpha_{APSP} \log n \cdot \log(1/b) \cdot \log(\log(n)^{\alpha_{APSP}/(b\psi)})}\right)$.

The algorithm is deterministic and requires the APSP data structure to undergo $O(\alpha_{APSP} \cdot n/\psi \log^3 n)$ updates, queries it $O(n)$ times and spends an additional $O(\alpha_{APSP} \cdot n/\psi \log^3 n)$ time.

Proof. The case where Algorithm 2 returns in Line 3 follows directly from Lemma 4.1, Theorem 3.1 and Lemma 3.2. Let us therefore analyze the remaining case where the algorithm returns in Line 13 (the while-loop can be seen to terminate since each iteration shrinks the set X by Claim 4.1 and $X = \emptyset$ trivially has no two vertices at far distance).

We first prove that the final set $V \setminus X$ has size $b'n \leq |V \setminus X| \leq \frac{3}{4}n$:

- $b'n \leq |V \setminus X|$: Initially, $H[X] = H$ and $F \subseteq H$ contains more than $10b'n$ edges by Lemma 4.1. Every edge $(u, v) \in F$ has $\text{dist}_{\hat{w}}(u, v) \geq \text{dist}_w(u, v) > C/b'$. Since the maximum degree of H is 10, as long as $|V \setminus X| < b'n$, $H[X]$ contains all but $10b'n$ edges from H . Thus, $H[X] \cap F$ is not empty and the while-loop continues. We conclude that $b'n \leq |V \setminus X|$ holds.
- $|V \setminus X| \leq \frac{3}{4}n$: Since the while-loop condition allows only invocations of FINDTHINLAYER if $|V \setminus X| \leq n/4$, and since this procedure returns the smaller side of the cut it produces by Claim 4.1 (which is found on $G[X]$), we can conclude that at the end of the algorithm $|V \setminus X| \leq n/4 + n/2 \leq \frac{3}{4}n$.

This indicates that $|X| \geq n/4 \geq b'n/2 \geq bn$ since $2b \leq b' \leq \frac{1}{2}$.

Next, we bound the sparsity of the cut $V \setminus X$. Let S_1, S_2, \dots, S_k be the sets returned by procedure FINDTHINLAYER one after another over the course of the while-loop, such that $V \setminus X = \cup S_i$. We first observe that these sets are vertex-disjoint since after the i -th iteration, the procedure FINDTHINLAYER is invoked on the graph $G_i = G[V \setminus (S_1 \cup \dots \cup S_i)]$ to find S_{i+1} . Further, the final cut $(X, V \setminus X)$ contains only edges that were previously in a thin layer, i.e.

$$E_G(X, V \setminus X) \subseteq \bigcup_i E_{G_i}(V \setminus (S_1 \cup \dots \cup S_i), S_i).$$

It remains to use the guarantee of Claim 4.1 that for each S_i , we have $|E_{G_i}(S_i, V \setminus (S_1 \cup \dots \cup S_i))| \leq \frac{4\hat{w}(S_i) \log_2 \|\mathbf{w}\|_1}{D}$ and by the vertex-disjointness of S_1, S_2, \dots, S_k , we thus have that

$$\begin{aligned} |E_G(X, V \setminus X)| &\leq \left| \bigcup_i E_{G_i}(S_i, V \setminus (S_1 \cup \dots \cup S_i)) \right| \leq \sum_i \frac{4\hat{w}(S_i) \log \|\hat{\mathbf{w}}\|_1}{D} \\ &\leq \frac{4\|\hat{\mathbf{w}}\|_1 \log \|\hat{\mathbf{w}}\|_1}{D} = \frac{8n \cdot b' \log n}{C} \end{aligned}$$

where we use $\|\mathbf{w}\| \leq 20n$ from Theorem 4.1 and $\hat{\mathbf{w}}$ is obtained from rounding up \mathbf{w} , and our choice of D . Since we have shown that $|X|, |V \setminus X| \geq b'n/2 \geq bn$, choosing $C = 320 \log n/\psi$, we have $\Psi(V \setminus X) = \Psi(X) \leq \psi$, as desired.

We use the disjointness of S_1, S_2, \dots, S_k to argue that the total time spend in procedure FINDTHINLAYER can be bounded by $O(n \log n)$. The remainder of the runtime analysis is trivial given Lemma 4.1. \square

It remains to provide an implementation of FINDTHINLAYER(G, \mathbf{w}, u, v, D) and prove Claim 4.1. The algorithm follows a simple ball-growing procedure. It grows balls from both endpoints u and v . Because the distance between u and v are guaranteed to be large, the procedure takes longer time. However, these two balls cannot be larger than the entire graph. There must be a moment that one of the ball grows only by a thin layer.

CLAIM 4.1. *The procedure FINDTHINLAYER(G, \mathbf{w}, u, v, D) takes a graph G weighted by $\mathbf{w} \in \mathbb{N}_{\geq 1}^{E(G)}$ and two vertices u, v such that $\text{dist}_w(u, v) > D$ for some integer $D > 4 \log_2 \|\mathbf{w}\|_1$. It returns a set of vertices $S \neq \emptyset$ such that $|S| \leq |V|/2$ and $|E_G(S, V \setminus S)| \leq \frac{4\mathbf{w}(E_G(S)) \log_2 \|\mathbf{w}\|_1}{D}$. The algorithm runs in time $O(|E_G(S)| \log |E_G(S)|)$.*

Proof. Since $\text{dist}_w(u, v) > D$ by assumption, we have that at least one of u and v have their ball to radius $D/2$ contain at most half the vertices in G . More formally, for some $z \in \{u, v\}$, $|B_{G,w}(z, D/2)| \leq |V|/2$. We claim that there is a radius $0 < r \leq D/2$, such that taking $S = B(z, r)$ satisfies the above guarantees. For this proof, it is convenient to define the following auxiliary function $\Phi(z, r) = \sum_{e \in E} \Phi(z, r, e)$ where the latter functions are defined for all edges $e = (x, y) \in E$ by

$$\Phi(z, r, e) = \begin{cases} |\text{dist}_w(z, x) - \text{dist}_w(z, y)| & \text{if } \text{dist}_w(z, x) \leq r \text{ and } \text{dist}_w(z, y) \leq r \\ r - \text{dist}_w(z, x) & \text{if } \text{dist}_w(z, x) \leq r < \text{dist}_w(z, y) \\ r - \text{dist}_w(z, y) & \text{if } \text{dist}_w(z, y) \leq r < \text{dist}_w(z, x) \\ 0 & \text{otherwise} \end{cases}$$

Here, an edge $e = (x, y) \in E(G)$ contributes the distance between its two endpoints x and y (which is at most w_e) to $\Phi(z, r, e)$ if both endpoints are fully contained in the ball $B(z, r)$. If neither of the endpoints are contained it contributes 0. Otherwise, $e = (x, y)$ contributes the distance of the endpoint closer to z to the boundary of the ball. In both cases, $0 \leq \Phi(z, r, e) \leq w_e$. This means in particular that the weight of edges incident to $B(z, r)$ denoted by $w(E(B(z, r)))$ is always greater-equal to $\Phi(z, r)$, i.e. $w(E(B(z, r))) \geq \Phi(z, r)$ for all r .

Note further that $\Phi(z, r+1) - \Phi(z, r)$ is exactly $|E_G(B(z, r), V \setminus B(z, r))|$, the number of edges that leave $B(z, r)$. To see this, observe that an edge $e = (x, y)$ contributes 1 to the difference if $\text{dist}_w(x, z) \leq r < r+1 \leq \text{dist}_w(y, z)$ holds, i.e. e leaves $B(z, r)$. Otherwise, the contribution of e are identical in both $\Phi(z, r)$ and $\Phi(z, r+1)$. Here we use that w is integral and so are distances in G .

Given this set-up, assume for contradiction that for all $0 < r < D/2$, we have

$$\Phi(z, r+1) > \left(1 + \frac{4 \log_2 \|w\|_1}{D}\right) \Phi(z, r).$$

By induction we have that

$$\Phi(z, D/2) \geq \left(1 + \frac{4 \log_2 \|w\|_1}{D}\right)^{D/2-1} \Phi(z, 1) > \|w\|_1$$

where we use that $1 + x \geq 2^x$ for $x \in [0, 1]$. This would give a contradiction since $\|w\|_1 \geq w(E(B(z, D/2))) \geq \Phi(z, D/2) > \|w\|_1$.

Therefore, there must be some radius $0 < r < D/2$ such that

$$\Phi(z, r+1) \leq \left(1 + \frac{4 \log_2 \|w\|_1}{D}\right) \Phi(z, r).$$

Combining with our previous discussion yields that

$$\begin{aligned} |E(B(z, r), V \setminus B(z, r))| &= \Phi(z, r+1) - \Phi(z, r) \\ &\leq \frac{4 \log_2 \|w\|_1}{D} \Phi(z, r) \\ &\leq \frac{4 w(E(B(z, r))) \log_2 \|w\|_1}{D}. \end{aligned}$$

We can therefore take $S = B(z, r)$, as desired.

Finally, to compute this cut, we run Dijkstra's algorithm from u and v in parallel and check for the earliest radius r for either of them such that the inequality holds. Thus, the algorithm runs in time $O(|E_G(S)| \log |E_G(S)|)$. \square

References

- [1] Amir Abboud, Robert Krauthgamer, Jason Li, Debmalya Panigrahi, Thatchaphol Saranurak, and Ohad Trabelsi. Breaking the cubic barrier for all-pairs max-flow: Gomory-hu tree in nearly quadratic time. FOCS, 2022.

- [2] Amir Abboud, Robert Krauthgamer, and Ohad Trabelsi. Subcubic algorithms for gomory–hu tree in unweighted graphs. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 1725–1737, 2021.
- [3] Reid Andersen, Fan R. K. Chung, and Kevin J. Lang. Using pagerank to locally partition a graph. *Internet Mathematics*, 4(1):35–64, 2007.
- [4] Reid Andersen and Yuval Peres. Finding sparse cuts locally using evolving sets. In *Proceedings of the Forty-first Annual ACM Symposium on Theory of Computing*, STOC '09, pages 235–244, New York, NY, USA, 2009. ACM.
- [5] Sanjeev Arora, Elad Hazan, and Satyen Kale. $O(\sqrt{\log(n)})$ approximation to SPARSEST CUT in $\tilde{o}(n^2)$ time. *SIAM J. Comput.*, 39(5):1748–1771, 2010.
- [6] Sanjeev Arora and Satyen Kale. A combinatorial, primal-dual approach to semidefinite programs. In David S. Johnson and Uriel Feige, editors, *Proceedings of the 39th Annual ACM Symposium on Theory of Computing, San Diego, California, USA, June 11-13, 2007*, pages 227–236. ACM, 2007.
- [7] Sanjeev Arora, Satish Rao, and Umesh V. Vazirani. Expander flows, geometric embeddings and graph partitioning. *J. ACM*, 56(2):5:1–5:37, 2009. Announced at STOC'04.
- [8] Aaron Bernstein, Maximilian Probst Gutenberg, and Thatchaphol Saranurak. Deterministic decremental reachability, scc, and shortest paths via directed expanders and congestion balancing. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1123–1134. IEEE, 2020.
- [9] Aaron Bernstein, Maximilian Probst Gutenberg, and Thatchaphol Saranurak. Deterministic decremental sssp and approximate min-cost flow in almost-linear time. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1000–1008. IEEE, 2022.
- [10] Li Chen, Rasmus Kyng, Yang P Liu, Richard Peng, Maximilian Probst Gutenberg, and Sushant Sachdeva. Maximum flow and minimum-cost flow in almost-linear time. *Accepted to FOCS'2022*, 2022.
- [11] Julia Chuzhoy. Decremental all-pairs shortest paths in deterministic near-linear time. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 626–639, 2021.
- [12] Julia Chuzhoy, Yu Gao, Jason Li, Danupon Nanongkai, Richard Peng, and Thatchaphol Saranurak. A deterministic algorithm for balanced cut with applications to dynamic connectivity, flows, and beyond. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1158–1167. IEEE, 2020.
- [13] Lisa K Fleischer. Approximating fractional multicommodity flow independent of the number of commodities. *SIAM Journal on Discrete Mathematics*, 13(4):505–520, 2000.
- [14] Ofer Gabber and Zvi Galil. Explicit constructions of linear-sized superconcentrators. *Journal of Computer and System Sciences*, 22(3):407–420, 1981.
- [15] Naveen Garg and Jochen Könemann. Faster and simpler algorithms for multicommodity flow and other fractional packing problems. *SIAM Journal on Computing*, 37(2):630–652, 2007.
- [16] Yiding Hua, Rasmus Kyng, Maximilian Probst Gutenberg, and Zihang Wu. Maintaining expander decompositions via sparse cuts. *SODA 2023*, 2022.
- [17] Ken-ichi Kawarabayashi and Mikkel Thorup. Deterministic edge connectivity in near-linear time. *Journal of the ACM (JACM)*, 66(1):1–50, 2018.
- [18] Rohit Khandekar, Satish Rao, and Umesh V. Vazirani. Graph partitioning using single commodity flows. *J. ACM*, 56(4):19:1–19:15, 2009.
- [19] Tom Leighton and Satish Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *J. ACM*, 46(6):787–832, November 1999.
- [20] Jason Li. Deterministic mincut in almost-linear time. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 384–395, 2021.
- [21] Jason Li and Debmalya Panigrahi. Deterministic min-cut in poly-logarithmic max-flows. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 85–92. IEEE, 2020.
- [22] Jason Li and Thatchaphol Saranurak. Deterministic weighted expander decomposition in almost-linear time. *arXiv preprint arXiv:2106.01567*, 2021.
- [23] Anand Louis. Cut-matching games on directed graphs. *arXiv preprint arXiv:1010.1047*, 2010.
- [24] 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 *STOC*, 2012.
- [25] Lorenzo Orecchia, Leonard J. Schulman, Umesh V. Vazirani, and Nisheeth K. Vishnoi. On partitioning graphs via single commodity flows. In Cynthia Dwork, editor, *Proceedings of the 40th Annual ACM Symposium on Theory of Computing, Victoria, British Columbia, Canada, May 17-20, 2008*, pages 461–470. ACM, 2008.
- [26] Lorenzo Orecchia and Nisheeth K. Vishnoi. Towards an sdp-based approach to spectral methods: a nearly-linear-time algorithm for graph partitioning and decomposition. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 532–545. SIAM, 2011.
- [27] Richard Peng. Approximate undirected maximum flows in $O(m \text{ polylog}(n))$ time. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1862–1867. SIAM, 2016. Available at

<http://arxiv.org/abs/1411.7631>.

- [28] Serge A Plotkin, David B Shmoys, and Éva Tardos. Fast approximation algorithms for fractional packing and covering problems. *Mathematics of Operations Research*, 20(2):257–301, 1995.
- [29] Prasad Raghavendra, David Steurer, and Madhur Tulsiani. Reductions between expansion problems. In *2012 IEEE 27th Conference on Computational Complexity*, pages 64–73, 2012.
- [30] Thatchaphol Saranurak and Di Wang. Expander decomposition and pruning: Faster, stronger, and simpler. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019*, SODA '19, 2019.
- [31] Jonah Sherman. Breaking the multicommodity flow barrier for $o(\log n)$ -approximations to sparsest cut. In *FOCS*, pages 363–372. IEEE Computer Society, 2009.
- [32] Jianbo Shi and Jitendra Malik. Normalized cuts and image segmentation. *IEEE Transactions on pattern analysis and machine intelligence*, 22(8):888–905, 2000.
- [33] David B Shmoys. Cut problems and their application to divide-and-conquer. *Approximation algorithms for NP-hard problems*, pages 192–235, 1997.
- [34] Daniel A Spielman and Shang-Hua Teng. Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, pages 81–90, 2004.
- [35] Neal E. Young. Randomized rounding without solving the linear program. In *Proceedings of the Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '95*, page 170–178, USA, 1995. Society for Industrial and Applied Mathematics.

A Reducing Conductance to Sparsity

Here, we prove Theorem 2.1. The proof is an adaption of Lemma 5.4 and Theorem 5.5 of [12].

The Transformation Algorithm. Our algorithm is essentially a wrapper function around our main result Theorem 4.1. That is, we first construct a bounded degree graph \widehat{G} from G , then run the algorithm from Theorem 4.1 on \widehat{G} . If the algorithm certifies that \widehat{G} has no balanced sparse cuts, we prove that G has no balanced low-conductance cuts. Otherwise, if the algorithm returns a sparse cut in \widehat{G} , we recover a balanced low-conductance cut in G .

We first describe the construction of \widehat{G} given $G = (V, E)$. Let us assume an arbitrary ordering of the edges incident to each vertex $v \in V$. $\widehat{G} = (\widehat{V}, \widehat{E})$ is constructed as follows:

1. For each vertex $v \in V$, create a set of vertices $X_v = \{v_1, v_2, \dots, v_{\deg(v)}\}$, and an ψ_0 -expander H_v on X_v using Theorem 3.1. Add H_v to \widehat{G} .
2. For each edge $e = (u, v) \in E$, we add (u_i, v_j) to \widehat{E} if e is the i^{th} (and j^{th}) edge incident to u (and v , respectively).

Clearly, \widehat{G} has $\text{vol}(\widehat{G}) = 2m$ vertices and each vertex has at most 10 incident edges, 9 from the expander and 1 from the corresponding edge in G .

We now run the algorithm from our main result, Theorem 4.1, on the graph \widehat{G} . If the algorithm certifies that no b -balanced ψ -sparse cut exists in \widehat{G} , we return the same result for G . Otherwise, we run Algorithm 3 on the returned cut (A, \overline{A}) in \widehat{G} to obtain a $\Omega(b)$ -balanced $\Omega(\psi)$ -sparse cut (S, \overline{S}) in G . It is straight-forward to check that Algorithm 3 is deterministic and runs in time linear in the number of edges of G and thus the runtimes stated in Theorem 4.1 are asymptotically not affected.

Algorithm 3: TRANSFORM($G, \widehat{G}, A \subseteq V(\widehat{G})$)

1 return $S = \{u \in V \mid |X_u \cap A| \geq |X_u \setminus A|\}$.

Certifying G . We start by showing that if no $\Omega(b)$ -balanced $O(\psi)$ -sparse cut is found on \widehat{G} , then no such cut exists in G either.

LEMMA A.1. *Given a balance parameter $b \in (0, 1/4)$, if every cut (X, \overline{X}) in \widehat{G} with $|X|, |\overline{X}| \geq b \cdot |V(\widehat{G})|$ has $\Psi_{\widehat{G}}(S) \geq \psi$, then every cut (S, \overline{S}) in G with $\text{vol}_G(S), \text{vol}_G(\overline{S}) \geq b \cdot \text{vol}(G)$ has $\Phi_G(X) \geq \psi$.*

Proof. Let (S, \overline{S}) be any cut in G with $\text{vol}_G(S), \text{vol}_G(\overline{S}) \geq b \cdot \text{vol}(G)$. Define $X_S = \cup_{u \in S} X_u$ and $\overline{X}_S = \cup_{u \notin S} X_u =$

$X_{\bar{S}}$. Observe that $|E_G(S, \bar{S})| = |E_{\hat{G}}(X_S, \bar{X}_S)|$ because the ψ_0 -expander edges in \hat{G} do not appear in the cut and every cut edge (u_i, v_j) in \hat{G} corresponds to the cut edge $(u, v) \in G$.

By construction of \hat{G} , we have that $|X_S| = \text{vol}_G(S)$ and $|\bar{X}_S| = \text{vol}_G(\bar{S})$ and therefore $|X_S|, |\bar{X}_S| \geq b \cdot \text{vol}(G) = b|V(\hat{G})|$ by assumption on (S, \bar{S}) . Thus (X_S, \bar{X}_S) is balanced in \hat{G} and we can use the guarantee that $\Phi_{\hat{G}}(X_S) \geq \psi$. This yields

$$|E_G(S, \bar{S})| = |E_{\hat{G}}(X_S, \bar{X}_S)| \geq \psi \cdot \min\{|X_S|, |\bar{X}_S|\} = \psi \cdot \min\{\text{vol}(S), \text{vol}(\bar{S})\}.$$

□

Returning a Sparse Cut. It remains to prove that the above algorithm transforms any balanced sparse cut in \hat{G} to a balanced low conductance cut in G . We prove this claim in two steps. We first show that the number of edges in the cut (S, \bar{S}) in G is comparable to the number of edges in (A, \bar{A}) in \hat{G} .

CLAIM A.1. $|E_G(S, \bar{S})| = O(|E_{\hat{G}}(A, \bar{A})|)$.

Proof. Define $X_S = \cup_{u \in S} X_u$. Consider any vertex $u \in V$, we have that the graph H_u contributes at least $\psi_0 \cdot \min\{|X_u \cap A|, |X_u \setminus A|\}$ edges to the cut $|E_{\hat{G}}(A, \bar{A})|$. But in \hat{G} , the number of edges incident to X_u that are in the cut (S, \bar{S}) but were previously not in the cut (A, \bar{A}) can be at most $\min\{|X_u \cap A|, |X_u \setminus A|\}$ since H_u is contained entirely in S or \bar{S} and only one additional edge is incident to each vertex in X_u .

Thus, we can charge each edge in $E_{H_u}(A, \bar{A})$ with at most $1/\psi_0$ edges from $E_{\hat{G}}(S, \bar{S}) \setminus E_{\hat{G}}(A, \bar{A})$ incident on u and cover all such edges. We conclude that $|E_{\hat{G}}(S, \bar{S})| \leq |E_{\hat{G}}(A, \bar{A})| + |E_{\hat{G}}(A, \bar{A})|/\psi_0$, and finally use that $|E_G(S, \bar{S})| = |E_{\hat{G}}(X_S, \bar{X}_S)|$ as observed in Lemma A.1. □

Next, we prove that (S, \bar{S}) is a balanced cut.

CLAIM A.2. *If $\Psi_G(A) \leq \psi_0/2$, we have $\text{vol}_G(S) \geq \frac{1}{2}|A|$ and $\text{vol}_G(\bar{S}) \geq \frac{1}{2}|\bar{A}|$.*

Proof. We prove $\text{vol}_G(S) \geq \frac{1}{2}|A|$ (the proof of $\text{vol}_G(\bar{S}) \geq \frac{1}{2}|\bar{A}|$ is symmetric). Let us assume for the sake of contradiction that $\text{vol}_G(S) < \frac{1}{2}|A|$. We argued before that for every $u \in V$, we have $|E_{H_u}(A, \bar{A})| \geq \psi_0 \cdot \min\{|A \cap X_u|, |A \setminus X_u|\}$. We again define $X_S = \cup_{u \in S} X_u$ and observe that the fact that $\sum_{u \in S} |A \cap X_u| \leq |X_S| = \text{vol}_G(S) < \frac{1}{2}|A|$ implies that $\sum_{u \in S} |A \setminus X_u| \geq |A| - |X_S| > \frac{1}{2}|A|$. Definition of S also yields that $|A \setminus X_u| \leq |A \cap X_u|$.

Combining insights, we conclude

$$|E_{\hat{G}}(A, \bar{A})| \geq \sum_u |E_{H_u}(A, \bar{A})| \geq \sum_u \psi_0 \cdot \min\{|A \cap X_u|, |A \setminus X_u|\} \geq \sum_{u \in S} \psi_0 \cdot |A \setminus X_u| > \frac{\psi_0}{2}|A|$$

which implies that $\Psi_G(A) > \psi_0/2$ which contradicts our assumption, as desired. □

Finally, we combine our insights to prove Theorem 2.1.

Proof. [Proof of Theorem 2.1.] We have from the algorithm that $|A|, |\bar{A}| \geq b \cdot |V(\hat{G})| = 2bm$. Therefore, by Claim A.2, we produce a cut (S, \bar{S}) in G with $\text{vol}_G(S), \text{vol}_G(\bar{S}) \geq b/2 \cdot \text{vol}(G)$. By Claim A.1, we further have that $|E_G(S, \bar{S})| \leq O(|E_{\hat{G}}(A, \bar{A})|)$ and therefore $\Phi_G(S) = \frac{|E_G(S, \bar{S})|}{\text{vol}_G(S), \text{vol}_G(\bar{S})} = O\left(\frac{|E_{\hat{G}}(A, \bar{A})|}{\min\{|A|, |\bar{A}|\}}\right) = O(\phi)$ where the last equality stems from the fact that (A, \bar{A}) had $\Psi_{\hat{G}}(A) \leq \phi$ by Theorem 4.1. □

B The Constant-Degree Assumption

In this section, we prove that the following assumptions are without loss of generality.

ASSUMPTION B.1. *When computing a sparse cut with respect to sparsity, we may assume at a cost of a constant factor in the output quality that the input graph G has maximum degree 10.*

Proof. Consider obtaining the graph \widehat{G} from G by adding $\lceil m/n \rceil$ self-loops to each vertex in G . We then invoke Theorem 2.1 on \widehat{G} with ψ and parameter b .

Note first that in a connected graph G , we have that $\widehat{G} \leq 4m$. Further, note that since self-loops do not appear in cuts, we have $E_G(S, \overline{S}) = E_{\widehat{G}}(S, \overline{S})$ for all S .

Now, if the algorithm certifies low conductance of \widehat{G} , we have for each (S, \overline{S}) in G where $|S|, |\overline{S}| \geq 4b \cdot n$ that $\text{vol}_{\widehat{G}}(S) \geq |S| \lceil m/n \rceil \geq 4bm \geq b \cdot \text{vol}(\widehat{G})$. Since $|E_G(S, \overline{S})| = |E_{\widehat{G}}(S, \overline{S})| \geq \psi \min\{\text{vol}_{\widehat{G}}(S), \text{vol}_{\widehat{G}}(\overline{S})\} \geq \frac{1}{3}\psi \min\{|S|, |\overline{S}|\}$. Thus every $4b$ -balanced sparse cut has sparsity at least $\frac{1}{3}\psi$. Otherwise, the algorithm returns a cut (S, \overline{S}) of conductance at most ϕ in \widehat{G} . But, we have $\Phi_{\widehat{G}}(S) \geq \Psi_{\widehat{G}}(S) = \Psi_G(S)$ for all S . \square

C A Simple Randomized Algorithm to Construct Low-Degree Expanders

Algorithm 4: RANDCONSTDEGEXPANDER(n)

```

1 Construct an empty graph  $H$  on  $n$  vertices.
2 foreach  $v \in V(H)$  do
3   for  $i = 1, 2, \dots, k = 80 \log n$  do
4     | Sample a vertex  $u$  from  $V(H)$  uniformly and i.i.d. at random.
5     | Add edge  $(u, v)$  to  $H$ .
6   end
7 end
8 return  $H$ 

```

The Algorithm. Here, we provide Algorithm 4 which implements the algorithm mentioned in Remark 3.1.

Analysis. Before we start our analysis, we recall the following Chernoff bound.

THEOREM C.1. *Given i.i.d. $\{0, 1\}$ -random variables X_1, X_2, \dots, X_k , $X = \sum_i X_i$ and any $\delta \geq 0$, we have $P[X \geq (1 + \delta)\mathbb{E}[X]] \leq e^{-\frac{\delta^2 \mathbb{E}[X]}{(2+\delta)}}$ and $P[X \leq (1 - \delta)\mathbb{E}[X]] \leq e^{-\frac{\delta^2 \mathbb{E}[X]}{2}}$.*

Let us first prove that H has bounded degree.

CLAIM C.1. *Algorithm 4 returns H such that w.h.p., the maximum degree is $O(\log n)$.*

Proof. Each vertex u is selected as the second endpoint of an edge added to H in the inner for-loop with probability $1/n$ per iteration. As there are nk iterations of this for-loop, and each iteration is independent, we have by the Chernoff bound that each vertex u is at most k times selected with probability at least $1 - e^{-\frac{4k}{n}} = 1 - e^{-k} = 1 - n^{-32}$.

Since each vertex u has degree equal to k plus the number of times it is sampled, we have that its degree is at most $2k$ with probability at most $1 - n^{-32}$. We obtain our result over all vertices in H by applying a union bound. \square

CLAIM C.2. *Algorithm 4 returns a $\Omega(\log n)$ -expander H w.h.p.*

Proof. Consider any set S with $|S| \leq n/2$. Then, we have that $\mathbb{E}[E_H(S, \overline{S})] \geq \frac{1}{2}|S|k$ since each edge (u, v) sampled when the foreach-loop iterates over a vertex $v \in S$ has $u \notin S$ with probability at least $\frac{1}{2}$ and there are $|S|k$ such sampling events. Since they are independent, we further have from the Chernoff bound that $P[|E_H(S, \overline{S})| \leq \frac{1}{4}|S|k] \leq e^{-\frac{|S|k}{16}} = n^{-5|S|}$. It is clear that if $|E_H(S, \overline{S})| > \frac{1}{4}|S|k$ then $\Psi_H(S) \geq \frac{k}{4} = \Omega(\log n)$.

The remaining difficulty is that there are an exponential number of cuts so a union bound seems at first hard to apply. However, we observe that there are at most $\binom{\alpha}{n} \leq \left(\frac{n\epsilon}{\alpha}\right)^\alpha \leq n^{3\alpha}$ for $\alpha \geq 1$ cuts where the smaller half contains α vertices. As we have proven that a cut is ψ -sparse with probability at most $n^{-5\alpha}$, we can thus conclude by a simple union bound argument that H is not $\Omega(\log n)$ -expander with probability at most $\sum_{\alpha \geq 1} \binom{\alpha}{n} \cdot n^{-5\alpha} \leq 1/n$. \square